



FortiOS™ Handbook
for FortiOS 5.0



FortiOS™ Handbook for FortiOS 5.0

June 15, 2015

01-5010-99686-20150615

Copyright© 2015 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., in the U.S. and other jurisdictions, and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. In no event does Fortinet make any commitment related to future deliverables, features or development, and circumstances may change such that any forward-looking statements herein are not accurate. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.

Technical Documentation	docs.fortinet.com
Knowledge Base	kb.fortinet.com
Customer Service & Support	support.fortinet.com
Training Services	training.fortinet.com
FortiGuard	fortiguard.com
Document Feedback	techdocs@fortinet.com

Contents Quick Look

	Change Log.....	87
	Introduction.....	88
Chapter 1:	What's New for FortiOS 5.0	90
	New Features in FortiOS 5.0 Patch 7.....	92
	New features in FortiOS 5.0 Patch 6.....	94
	New features in FortiOS 5.0 Patch 5.....	97
	New features in FortiOS 5.0 Patch 4.....	111
	New features in FortiOS 5.0 Patch 3.....	121
	New features in FortiOS 5.0 Patch 2.....	132
	Security Features	138
	Authentication: users and devices.....	159
	FortiOS and BYOD.....	173
	Client Reputation.....	181
	Wireless	187
	IPv6	201
	Logging and reporting	214
	Firewall	218
	WAN optimization and Web Caching.....	230
	Usability enhancements	243
	SSL VPN	251
	Other new features.....	253
Chapter 2:	Advanced Routing for FortiOS 5.0	268
	Advanced Static Routing.....	269
	Dynamic Routing Overview	304
	Routing Information Protocol (RIP)	319
	Border Gateway Protocol (BGP)	358
	Open Shortest Path First (OSPF)	396
	Intermediate System to Intermediate System Protocol (IS-IS).....	438
Chapter 3:	Authentication for FortiOS 5.0.....	456
	Introduction to authentication	457
	Authentication servers.....	466
	Users and user groups.....	488

	Managing Guest Access	509
	Configuring authenticated access.....	514
	Certificate-based authentication	532
	SSO using a FortiAuthenticator unit.....	550
	Single Sign-On to Windows AD.....	554
	Agent-based FSSO	563
	SSO using RADIUS accounting records.....	602
	Monitoring authenticated users.....	609
	Examples and Troubleshooting	613
Chapter 4:	FortiOS Carrier.....	631
	Overview of FortiOS Carrier features	632
	Carrier web-based manager settings.....	661
	MMS Security features.....	701
	Message flood protection.....	720
	Duplicate message protection	731
	Configuring GTP on FortiOS Carrier.....	738
	GTP message type filtering	745
	GTP identity filtering.....	752
	Troubleshooting	760
Chapter 5:	Compliance	767
	Configuring FortiGate units for PCI DSS compliance.....	768
Chapter 6:	Deploying Wireless Networks for FortiOS 5.0	783
	Introduction to wireless networking.....	784
	Configuring a WiFi LAN.....	793
	Access point deployment	814
	Wireless Mesh.....	828
	WiFi-Ethernet Bridge Operation.....	835
	Protecting the WiFi Network	843
	Wireless network monitoring	846
	Configuring wireless network clients.....	852
	Wireless network examples	863
	Using a FortiWiFi unit as a client	879
	Support for location-based services	881
	Reference.....	883
Chapter 7:	Firewall for FortiOS 5.0	888

	Firewall concepts	891
	Firewall objects	909
	Security policies	952
	Network defense	974
	GUI & CLI - What You May Not Know.....	983
	Building firewall objects and policies.....	989
	Multicast forwarding	1030
Chapter 8:	Hardware Acceleration	1067
	Hardware acceleration overview	1068
	NP6 Acceleration.....	1082
	FortiGate NP6 architectures.....	1089
	NP4 Acceleration.....	1093
	FortiGate NP4 architectures.....	1105
Chapter 9:	High Availability for FortiOS 5.0	1116
	Solving the High Availability problem	1117
	An introduction to the FGCP	1121
	Configuring and connecting HA clusters	1152
	Virtual clusters.....	1217
	Full mesh HA.....	1239
	Operating a cluster.....	1252
	HA and failover protection.....	1290
	HA and load balancing	1343
	HA with FortiGate-VM and third-party products	1360
	VRRP.....	1364
	FortiGate Session Life Support Protocol (FGSP).....	1370
	Configuring FRUP.....	1380
	1385
Chapter 10:	Install and System Administration for FortiOS 5.0	1385
	Differences between Models and Firmware	1386
	Using the web-based manager	1387
	Using the CLI	1405
	Basic Administration.....	1426
	Best practices.....	1456
	FortiGuard	1461
	FortiCloud.....	1474
	Interfaces	1480

Central management.....	1502
Monitoring	1506
VLANs	1523
PPTP and L2TP	1555
Advanced concepts.....	1568
Session helpers	1602
Chapter 11: IPsec VPN for FortiOS 5.0.....	1612
IPsec VPN concepts.....	1614
IPsec VPN Overview.....	1620
IPsec VPN in the web-based manager	1624
Auto Key phase 1 parameters	1637
Phase 2 parameters	1653
Defining VPN security policies	1659
Gateway-to-gateway configurations	1665
Hub-and-spoke configurations	1679
Dynamic DNS configuration	1695
FortiClient dialup-client configurations.....	1709
FortiGate dialup-client configurations	1724
Supporting IKE Mode config clients.....	1732
Internet-browsing configuration	1737
Redundant VPN configurations.....	1741
Transparent mode VPNs.....	1766
IPv6 IPsec VPNs	1772
L2TP and IPsec (Microsoft VPN)	1786
GRE over IPsec (Cisco VPN).....	1798
Protecting OSPF with IPsec	1808
Hardware offloading and acceleration.....	1816
Monitoring and troubleshooting	1822
IPv6 for FortiOS 5.0	1828
IPv6 Features	1830
IPv6 Configuration.....	1855
Chapter 12: Load Balancing for FortiOS 5.0	1880
Configuring load balancing	1881
Load balancing configuration examples	1911
Chapter 13: Logging and Reporting	1928
.....	1929

Logging and reporting overview	1930
Logging and reporting for small networks	1955
Logging and reporting for large networks	1960
Advanced logging.....	1967
Troubleshooting and logging	1985
Appendix: FortiGate report charts	1989
Chapter 14: Managing Devices for FortiOS 5.0	1995
Managing “bring your own device”	1996
Endpoint Protection	2003
Vulnerability Scan.....	2011
Chapter 15: Unified Threat Management for FortiOS 5.0.....	2018
Security Profiles overview	2019
Client Reputation.....	2025
AntiVirus	2032
Email filter	2049
Intrusion protection.....	2062
Custom Application & IPS Signatures	2080
Web filter	2095
Data leak prevention	2130
Application control	2152
ICAP	2164
Other Security Profiles considerations	2169
Chapter 16: SSL VPN for FortiOS 5.0.....	2189
Introduction to SSL VPN	2190
Basic Configuration.....	2195
The SSL VPN client.....	2219
Setup examples	2221
Chapter 17: Traffic Shaping for FortiOS 5.0.....	2233
The purpose of traffic shaping.....	2234
Traffic shaping methods.....	2243
Examples.....	2259
Troubleshooting traffic shaping.....	2266
Chapter 18: Troubleshooting	2270
Life of a Packet.....	2271
Verifying FortiGate admin access security	2285

Troubleshooting resources	2290
Troubleshooting tools	2292
Troubleshooting methodologies	2321
Technical Support Organization Overview	2325
Chapter 19: Virtual Domains	2331
Virtual Domains	2332
Virtual Domains in NAT/Route mode.....	2361
Virtual Domains in Transparent mode.....	2380
Inter-VDOM routing.....	2400
Troubleshooting Virtual Domains	2440
Chapter 20: Virtual FortiGate Units for FortiOS 5.0.....	2445
FortiGate VM Overview	2446
Deployment example: VMware	2451
Deployment example: MS Hyper-V.....	2457
Deployment example: KVM	2470
Deployment example: OpenXen.....	2473
Deployment example: Citrix XenServer.....	2477
FortiGate VM Initial Configuration	2482
Chapter 21: VoIP Solutions: SIP for FortiOS 5.0.....	2488
FortiGate VoIP solutions: SIP	2489
Chapter 22: WAN Optimization, Web Cache, Explicit Proxy, and WCCP for FortiOS 5.0..	2577
Example network topologies.....	2580
Configuring WAN optimization.....	2590
Peers and authentication groups.....	2610
Configuration examples.....	2616
Web caching and SSL offloading.....	2643
FortiClient WAN optimization	2661
The FortiGate explicit web proxy	2665
The FortiGate explicit FTP proxy	2689
FortiGate WCCP	2703
Storage	2716
Diagnose commands	2719
Index	2728

Table of Contents

Change Log	87
Introduction	88
Chapter 1: What's New for FortiOS 5.0	90
New Features in FortiOS 5.0 Patch 7	92
OpenSSL Vulnerability (Heartbleed) Fixed.....	92
New features in FortiOS 5.0 Patch 6	94
Endpoint Control Daemon Improvement.....	94
IPS Hardware Acceleration.....	94
802.11g Protection Mode	95
Miglogd Child Processes.....	95
IPv6 in CRL/SCEP	95
Extended IPS Database for D-series Desktop Models.....	95
Logging Options for 3000 and 5000 Series Models	96
Wireless Controller on FortiGate-30D.....	96
New features in FortiOS 5.0 Patch 5	97
Improvements to Endpoint Control	97
New menu options.....	97
Default profile.....	98
FortiClient Monitor	98
FortiAP LAN port support	98
Bridging with the FortiAP's SSID(s)	98
Bridging with the WAN port	99
Configuring bridging	99
Restrictions	100
Automatically allowing basic applications.....	100
Pre-authorizing a FortiAP unit.....	101
Preventing IP fragmentation of packets in CAPWAP tunnels.....	102
Limiting access for unauthenticated users	103
Use case - allowing limited access for unauthenticated users.....	103
Use case - multiple levels of authentication	104
LDAP browser to import users into a user group	104
Dedicated management CPU	105
Improvements to the Traffic History and Threat History widgets.....	105
Assigning an IP address to a dynamic IPsec VPN interface	105
SSL VPN History widget	106
Port Block Allocation (PBA) for CGN to reduce logs.....	106

Neighbor cache table for IPv6	106
Improved HA diagnose commands	107
Secure disk erasing	107
Anonymize user names in logs	107
VLAN interface traffic statistics.....	107
Preserving the Class of Service bit.....	107
Front panel illustration	108
USB entropy token support.....	108
Station locate for FortiWiFi units.....	108
Switch Controller added models 200D, 240D, 600C, 800C, and 1000C	109
Diagnose command for 5000 series FortiGate units	109
New platforms for FortiGate-VM.....	109
Supported RFCs	109
New features in FortiOS 5.0 Patch 4.....	111
FortiSandbox	111
Wireless Health Dashboard	111
IPsec VPN.....	112
Dial-up IPsec VPN Creation Wizard.....	112
Show or Hide policy-based IPsec VPN	112
Managing FortiAP units	112
Units remain online when their WiFi Controller goes offline	113
Assigning the same profile to multiple FortiAP units	113
Dynamic VLANs for SSIDs.....	113
NAT46 & NAT64.....	114
Enhancements to Tables	114
Policy Table.....	114
Member Display	114
Fortinet Top Bar.....	115
FortiAnalyzer and FortiManager log encryption.....	115
FortiToken Mobile.....	115
Load balancing for explicit web proxy forwarding server groups	115
Server load balancing enhancements	116
SNMP traps.....	116
HTTP redirects	116
Additional filters for IPS and Application Control	117
Blocking IPv6 packets by extension headers.....	117
Distinguishing between HTTP GET and POST in DLP	117
RADIUS Accounting.....	118
H3C Compatibility.....	118
Web filter administrative overrides	118
Configurable idle timeout for console admin login sessions	118

TCP reset	119
Log Volume Monitor	119
Invalid Packet log.....	119
Server limits	119
PoE Power Management display.....	119
Other new features	120
New features in FortiOS 5.0 Patch 3.....	121
Security Features.....	122
Exempting IP addresses from IPS.....	122
DLP Watermarking Client	122
Predefined Device Groups.....	122
Client Reputation Configuration	122
Feature Select.....	122
Changes to Endpoint Control	122
Endpoint control for Android	122
Assigning endpoint profiles to specific users and user groups.....	123
Endpoint profile portal pages.....	123
Managing FortiAP units	123
Firmware Auto-detection	123
Wireless Device Locating Service.....	123
More Wireless Controller MIB Support.....	124
Normal or Remote WTP mode parameter	125
FortiGuard Subscription Services.....	125
Adding Explicit Web Proxy services	125
SSO Authentication failover for the Explicit Web Proxy	126
User Creation Wizard.....	127
FortiClient Registration	127
DSS and ECDSA Certificates for FortiGate SSL-related features	127
LDAP Servers.....	127
User Monitor	127
Web Filter Profiles.....	127
CAPWAP Administrative Access	128
IPS Algorithms	128
NAC-Quarantine Traffic Logs	128
New System Report Charts	128
Memory Logging.....	128
URL-based Web Proxy Forwarding.....	128
Changes to Routing	129
RADIUS Support for Dynamic VLANs.....	129
Dedicated Management Port.....	129
URL Filtering	129

URL Source Tracking.....	129
IPv6 Denial of Service Policies	130
Support for NAT46, VIP64 and VIP46.....	130
Packet Capture Filters	130
Configure hosts in an SNMP v1/2c community to send queries or receive traps	130
IP in IP tunneling support (RFC 1853)	131
GTP-u acceleration on FortiGate units with SP3 processors.....	131
New features in FortiOS 5.0 Patch 2.....	132
Endpoint Profile Changes	132
Client Reputation Changes.....	132
Changes to logging in security policies.....	132
Configuring the FortiGate unit to be an NTP Server.....	133
Customizing and viewing the local FortiGate UTM Security Analysis Report	133
Wireless changes: Custom mesh downlink SSIDs and new identifier for local bridge SSIDs.....	134
SSL-VPN Realm Support (multiple custom SSL VPN logins).....	135
Automatically add devices found by device identification to the vulnerability scanner configuration.....	136
The SIP ALG can receive SIP traffic on multiple TCP and UDP ports.....	136
IPv6 PIM sparse mode multicast routing.....	137
Wireless RADIUS-Based MAC Authentication	137
Security Features	138
FortiSandbox	138
Configuration	138
Sending files to FortiSandbox.....	139
Tracking submitted files.....	139
Botnet and phishing protection	139
Windows file sharing (CIFS) flow-based antivirus scanning	140
Advanced Application Control and IPS sensor creation	142
Custom Application Control signatures and IPS signatures	143
Exempting IP addresses from IPS.....	144
Flow-based inspection improvements	145
Configuring SSL inspection for flow-based and proxy protection	145
Explicit web Proxy Extensions – SSL inspection, IPS, Application Control, and flow-based antivirus, web filtering and DLP.....	146
Replacement messages for flow-based web filtering of HTTPS traffic.....	146
DNS web filtering	146
FortiGuard Web Filter quotas can be set based on traffic volume.....	147
Customizing the authentication replacement message for a FortiGuard web filter category.....	148
YouTube Education Filter implemented in Web Filtering Profiles	148

IPS hardware acceleration.....	149
New SIP ALG features	149
Inspecting SIP over SSL/TLS (secure SIP)	150
Opening and closing SIP via and record-route pinholes	152
Adding the original IP address and port to the SIP header after NAT	152
DLP watermarking	153
Fortinet watermarking utility	154
SSH inspection	156
Optimizing SSL encryption/decryption performance	157
Authentication: users and devices.....	159
User authentication menu changes	159
User identity policy changes.....	159
Authentication-based routing	160
Secondary and tertiary RADIUS, LDAP, and TACAS+ servers.....	161
FortiToken two-factor authentication and FortiToken Mobile	162
Configuring FortiToken mobile soft token support	162
SSO using a FortiAuthenticator unit	164
User's view of FortiAuthenticator SSO authentication	164
Administrator's view of FortiAuthenticator SSO authentication	165
SSO with Windows AD or Novell	165
Citrix Agent support for Single Sign On.....	166
Installing Citrix/Terminal Service Support Agent (TS Agent)	166
Installing the FSSO collector	167
To enable single sign-on using polling mode	167
Verifying the configuration	167
Configuring guest access	167
User's view of guest access	167
Administrator's view of guest access	167
Creating guest management administrators	168
Creating guest user groups	168
Creating guest user accounts.....	169
Batch guest account creation.....	170
Vulnerability Scanning	170
Running and configuring scans and viewing scan results.....	171
FortiOS and BYOD.....	173
Device monitoring.....	173
Device Groups	175
Creating a custom device group.....	175
Controlling access with a MAC Address Access Control List.....	176
Device policies.....	176
Device policy portal options	178
Creating the WiFi SSID	178

Configuring Internet access for guests with mobile devices	179
Client Reputation.....	181
Setting the client reputation profile/definition.....	182
Applying client reputation monitoring to your network.....	183
Viewing client reputation results	184
Expanding client reputation to include more types of behavior	184
Client reputation execute commands.....	186
Client reputation diagnose commands.....	186
Wireless	187
Wireless IDS.....	187
WiFi performance improvements.....	190
FortiAP web-based manager and CLI	190
WiFi guest access provisioning	192
Adding guest access to a WiFi network	193
FortiAP local bridging (Private Cloud-Managed AP).....	193
WiFi data channel encryption	195
Configuring DTLS on the FortiGate unit	196
Configuring encryption on the FortiAP unit	196
Wireless client load balancing for high-density deployments	197
Access point hand-off.....	197
Frequency hand-off or band-steering.....	197
Configuration	197
Bridge SSID to FortiGate wired network	198
IPv6	201
IPv6 Policy routing	201
IPv6 security policies	202
IPv6 Explicit web proxy	203
Restricting the IP address of the explicit IPv6 web proxy	204
Restricting the outgoing source IP address of the IPv6 explicit web proxy ..	204
IPv6 NAT – NAT64, DNS64, NAT66.....	205
NAT64 and DNS64.....	205
NAT66	208
NAT66 destination address translation.....	209
IPv6 Forwarding Policies - IPS, Application Control, and flow-based antivirus, web filtering and DLP.....	209
New Fortinet FortiGate IPv6 MIB fields	210
New OIDs.....	211
EXAMPLE SNMP get/walk output	212
IPv6 Per-IP traffic shaper.....	212
DHCPv6 relay	212
FortiGate interfaces can get IPv6 addresses from an IPv6 DHCP server	213

Logging and reporting	214
Log message reorganization	214
Log Viewer Improvements	214
The FortiGate Security Analysis Report.....	215
Viewing the current report	216
Viewing the saved (historical) security analysis reports.....	216
Customizing the security analysis report.....	216
Converting compact log format.....	217
Firewall	218
Choosing the policy type	218
Creating a basic security policy.....	218
Creating a security policy to authenticate users.....	219
Creating a security policy to authenticate devices for BYOD.....	220
Creating a policy-based IPsec VPN security policy	220
Creating a route-based IPsec VPN security policy.....	221
Creating an SSL VPN security policy.....	222
Reorganized Firewall Services.....	223
Editing and deleting services	223
Adding an address to a service	224
Adding a new service.....	224
Adding a new service category.....	224
Local in policies	224
Multicast Policies.....	225
Adding DoS Anomaly protection to a FortiGate interface	226
Changes to security proxy options.....	227
Protocol port mapping.....	227
Common options, web options and email options.....	228
SSL and SSH inspection	228
SSL inspection options.....	229
SSH inspection options	229
WAN optimization and Web Caching.....	230
Configuring WAN optimization profiles.....	230
Dynamic data chunking for WAN optimization byte caching	233
Policy-based WAN optimization configuration changes summary	234
On the client side	234
On the server side.....	234
Client side configuration summary	235
Server Side configuration summary.....	237
Combining web caching for HTTP traffic with WAN optimization.....	238
Turning on web caching and SSL offloading for HTTPS traffic.....	238
Changing the ports on which to look for HTTP and HTTPS traffic to cache.....	239
Web proxy URL debugging	240
Debugging caching of a specific web page.....	240

Debugging caching of multiple web pages	241
FortiOS Web Caching now caches Windows/MS-Office software updates	242
Usability enhancements	243
Feature Select.....	243
Security Features Presets.....	244
Improved list editing	244
Dynamic comment fields	245
Setup Wizard enhancements.....	245
Fortinet Top Bar.....	245
VDOM Mode GUI changes	246
Enhanced Top Sessions dashboard widget.....	246
Top Sources.....	246
Top Destinations	247
Top Applications	248
Identifying Skype sessions	248
Customizing the Top Sessions dashboard widget	249
Improved CLI syntax for multi-value fields	249
SSL VPN	251
New default SSL VPN portals.....	251
SSL VPN user groups no longer required.....	251
SSL VPN policy interface name change	251
Support SSL VPN push configuration of DNS suffix	251
Other new features.....	253
New FortiGuard features	253
FortiGate Auto-config using DHCP	254
FortiGate Session Life Support Protocol (FGSP).....	254
HA failover supports more features.....	255
New HA mode: Fortinet redundant UTM protocol (FRUP)	255
ICAP and the explicit web proxy	256
Example ICAP sequence for an ICAP server performing web URL filtering on web proxy HTTP requests.....	256
Example ICAP configuration	256
Adding ICAP to a web proxy security policy - web-based manager.....	257
Adding ICAP to a web proxy security policy - CLI	257
New interface features - DHCP server and authentication.....	258
Adding a DHCP server to an interface.....	258
Reserving, assigning and blocking MAC addresses	259
Authentication - Captive Portal.....	259
Replacement Message Improvements	260
Acceleration of Inter-VDOM Traffic (by NP4).....	261
Virtual Hardware Switch	262
FortiExplorer for iOS devices.....	263

Connecting to and logging into a FortiGate unit.....	264
Updating firmware and configuring network settings.....	264
Inter-VDOM links between NAT mode and Transparent mode VDOMs.....	264
About inter-VDOM links between NAT and Transparent mode VDOMs	265
Sniffer modes: one-armed and normal.....	265
Configuring an interface to operate as a one-arm sniffer	265
Integrated switch fabric (ISF) access control list (ACL) short-cut path	266
Generalized TTL Security Mechanism (GTSM) support	267
Firewall services.....	267

Chapter 2: Advanced Routing for FortiOS 5.0 268

Advanced Static Routing 269

Routing concepts.....	269
Routing in VDOMs	269
Default route	270
Adding a static route.....	270
Routing table.....	270
Building the routing table.....	277
Static routing security	277
Multipath routing and determining the best route	279
Route priority	280
Troubleshooting static routing	281
Static routing tips.....	283
Policy routing	284
Adding a policy route.....	285
Moving a policy route	287
Transparent mode static routing	287
Static routing example.....	288
Network layout and assumptions	288
General configuration steps.....	289
Get your ISP information such as DNS, gateway, etc.	290
Configure FortiGate unit	290
Configure Admin PC and Dentist PCs	295
Testing network configuration	296
Advanced static example: ECMP failover and load balancing	297
Equal-Cost Multi-Path (ECMP)	297
Configuring interface status detection for gateway load balancing	298
Configuring spillover or usage-based ECMP.....	300
Configuring weighted static route load balancing	302

Dynamic Routing Overview 304

What is dynamic routing?	304
Comparing static and dynamic routing.....	304
Dynamic routing protocols.....	305

Minimum configuration for dynamic routing	307
Comparison of dynamic routing protocols	307
Features of dynamic routing protocols	307
When to adopt dynamic routing	310
Choosing a routing protocol	312
Dynamic routing terminology	313
IPv6 in dynamic routing	318
Routing Information Protocol (RIP)	319
RIP background and concepts	319
Background	319
Parts and terminology of RIP	320
How RIP works	325
Troubleshooting RIP	330
Routing Loops.....	330
Holddowns and Triggers for updates	333
Split horizon and Poison reverse updates	333
Debugging IPv6 on RIPng.....	334
Simple RIP example.....	334
Network layout and assumptions	335
General configuration steps.....	336
Configuring the FortiGate units system information	336
Configuring FortiGate unit RIP router information	346
Configuring other networking devices	350
Testing network configuration	351
RIPng – RIP and IPv6.....	351
Network layout and assumptions	351
Configuring the FortiGate units system information	352
Configuring RIPng on FortiGate units	355
Configuring other network devices.....	356
Testing the configuration	356
Border Gateway Protocol (BGP)	358
BGP background and concepts	358
Background	358
Parts and terminology of BGP	358
How BGP works.....	367
Troubleshooting BGP	370
Clearing routing table entries.....	371
Route flap.....	371
Dual-homed BGP example	375
Network layout and assumptions	376
Configuring the FortiGate unit	378
Configuring other networking devices	387
Testing this configuration.....	387

Redistributing and blocking routes in BGP	389
Network layout and assumptions	389
Configuring the FortiGate unit	390
Testing network configuration	395
Open Shortest Path First (OSPF)	396
OSPF Background and concepts	396
Background	396
The parts and terminology of OSPF	396
How OSPF works.....	403
Troubleshooting OSPF	408
Clearing OSPF routes from the routing table.....	408
Checking the state of OSPF neighbors	409
Passive interface problems.....	409
Timer problems	409
Bi-directional Forwarding Detection (BFD)	409
Authentication issues.....	410
DR and BDR election issues.....	410
Basic OSPF example	410
Network layout and assumptions	411
Configuring the FortiGate units.....	412
Configuring OSPF on the FortiGate units	415
Configuring other networking devices	422
Testing network configuration	422
Advanced inter-area OSPF example	423
Network layout and assumptions	423
Configuring the FortiGate units.....	425
Configuring OSPF on the FortiGate units	429
Configuring other networking devices	433
Testing network configuration	433
Controlling redundant links by cost	433
Adjusting the route costs	435
Verifying route redundancy	437
Intermediate System to Intermediate System Protocol (IS-IS)	438
IS-IS background and concepts.....	438
Background	438
How IS-IS works	439
Parts and terminology of IS-IS.....	440
Troubleshooting IS-IS	445
Routing loops.....	445
Split horizon and Poison reverse updates	448
Simple IS-IS example	448
Network layout and assumptions	449
Expectations	449
CLI configuration.....	450

Verification	452
Troubleshooting	455
Chapter 3: Authentication for FortiOS 5.0.....	456
Introduction to authentication	457
What is authentication?	457
Methods of authentication	457
Local password authentication.....	458
Server-based password authentication.....	458
Certificate-based authentication.....	458
Two-factor authentication.....	459
Types of authentication	459
Firewall authentication (identity-based policies).....	460
VPN authentication	461
Single Sign On authentication for users	462
User's view of authentication	463
Web-based user authentication.....	463
VPN client-based authentication	463
FortiGate administrator's view of authentication.....	464
General authentication settings	464
Authentication servers.....	466
FortiAuthenticator servers	466
RADIUS servers	466
Configuring the FortiGate unit to use a RADIUS server.....	470
LDAP servers	471
Components and topology	472
LDAP directory organization	473
Configuring the FortiGate unit to use an LDAP server.....	474
Example – wildcard admin accounts - CLI	476
Example of LDAP to allow Dial-in through member-attribute - CLI.....	478
Troubleshooting LDAP	479
TACACS+ servers.....	480
Configuring a TACACS+ server on the FortiGate unit	480
SSO servers.....	481
RSA ACE (SecurID) servers	483
Components	483
Configuring the SecurID system	483
Users and user groups.....	488
Users.....	488
Local users.....	489
PKI or peer users	493
Two-factor authentication.....	494
FortiToken.....	497

IM users	501
Monitoring users	502
User groups	503
Firewall user groups.....	503
SSO user groups.....	507
Configuring Peer user groups.....	507
Viewing, editing and deleting user groups.....	508
Managing Guest Access	509
Introduction.....	509
User's view of guest access	509
Administrator's view of guest access	509
Configuring guest user access	509
Creating guest management administrators	509
Creating guest user groups	510
Creating guest user accounts.....	511
Guest access in a retail environment.....	512
Implementing email harvesting	512
Configuring authenticated access.....	514
Authentication timeout.....	514
Security authentication timeout	514
SSL VPN authentication timeout	514
Password policy	515
Authentication protocols.....	517
Authentication in security policies	517
Enabling authentication protocols	518
Authentication replacement messages.....	518
Access to the Internet.....	520
Configuring authentication security policies.....	521
Identity-based policy	523
NTLM authentication.....	524
Certificate authentication	525
Restricting number of concurrent user logons	526
Limited access for unauthenticated users.....	526
Use case - allowing limited access for unauthenticated users.....	527
Use case - multiple levels of authentication	527
VPN authentication	528
Configuring authentication of SSL VPN users.....	528
Configuring authentication of remote IPsec VPN users	528
Configuring authentication of PPTP VPN users and user groups	530
Configuring authentication of L2TP VPN users/user groups.....	531
Certificate-based authentication	532
What is a security certificate?.....	532
Certificates overview	533

Certificates and protocols.....	533
IPsec VPNs and certificates.....	534
Certificate types on the FortiGate unit.....	534
Certificate signing	535
Managing X.509 certificates	535
Generating a certificate signing request	536
Generating certificates with CA software	538
Obtaining and installing a signed server certificate from an external CA	538
Installing a CA root certificate and CRL to authenticate remote clients	539
Troubleshooting certificates	540
Online updates to certificates and CRLs	541
Backing up and restoring local certificates.....	542
Configuring certificate-based authentication	543
Authenticating administrators with security certificates	544
Authenticating SSL VPN users with security certificates	544
Authenticating IPsec VPN users with security certificates	545
Example — Generate a CSR on the FortiGate unit	545
Example — Generate and Import CA certificate with private key pair on OpenSSL..	546
Assumptions	546
Generating and importing the CA certificate and private key	546
Example — Generate an SSL certificate in OpenSSL	547
Assumptions	548
Generating a CA signed SSL certificate	548
Generating a self-signed SSL certificate	548
Import the SSL certificate into FortiOS.....	549
SSO using a FortiAuthenticator unit.....	550
User's view of FortiAuthenticator SSO authentication	550
Administrator's view of FortiAuthenticator SSO authentication	551
Configuring the FortiAuthenticator unit.....	551
Configuring the FortiGate unit	552
Adding a FortiAuthenticator unit as an SSO agent.....	552
Configuring an FSSO user group.....	552
Configuring security policies.....	552
Configuring the FortiClient SSO Mobility Agent	553
Viewing SSO authentication events on the FortiGate unit.....	553
Single Sign-On to Windows AD.....	554
Introduction to Single Sign-On with Windows AD.....	554
Configuring Single Sign On to Windows AD.....	554
Configuring LDAP server access	555
Creating Fortinet Single Sign-On (FSSO) user groups	557
Configuring the LDAP Server as a Single Sign-On server	557
Creating security policies.....	557

Enabling guest access through FSSO security policies	559
FortiOS FSSO log messages	559
Enabling authentication event logging	559
Testing FSSO	561
Troubleshooting FSSO	561
General troubleshooting tips for FSSO	561
Users on a particular computer (IP address) can not access the network....	562
Guest users do not have access to network	562
Agent-based FSSO	563
Introduction to agent-based FSSO	563
Introduction to FSSO agents	564
FSSO for Windows AD	565
FSSO for Citrix	567
FSSO for Novell eDirectory	568
FSSO security issues	569
FSSO NTLM authentication support	569
NTLM in a multiple domain environment	570
Agent installation	571
Collector agent installation	572
DC agent installation	573
Citrix TS agent installation	575
Novell eDirectory agent installation	575
Updating FSSO agents on Windows AD	576
Configuring the FSSO Collector agent for Windows AD	576
Configuring Windows AD server user groups	577
Configuring Collector agent settings	577
Selecting Domain Controllers and working mode for monitoring	580
Configuring Directory Access settings	581
Configuring the Ignore User List	582
Configuring FortiGate group filters	583
Configuring FSSO ports	584
Configuring alternate user IP address tracking	585
Viewing FSSO component status	585
Configuring the FSSO TS agent for Citrix	586
Configuring the FSSO eDirectory agent for Novell eDirectory	587
Configuring FSSO on FortiGate units	589
Configuring LDAP server access	589
Specifying your Collector agents or Novell eDirectory agents	590
Creating Fortinet Single Sign-On (FSSO) user groups	591
Creating security policies	591
Enabling guest access through FSSO security policies	594
FortiOS FSSO log messages	594
Enabling authentication event logging	594
Testing FSSO	595

Troubleshooting FSSO	596
General troubleshooting tips for FSSO.....	597
User status “Not Verified” on the Collector agent	597
After initial configuration, there is no connection to the Collector agent.....	597
Collector Agent service freezing and shutting down	598
FortiGate performance is slow on a large network with many users.....	598
Users from the Windows AD network are not able to access the network ...	599
Users on a particular computer (IP address) can not access the network....	599
Guest users do not have access to network	600
Can’t find the DCagent service.....	600
User logon events not received by FSSO Collector agent	600
User list from Windows AD is empty	600
Mac OS X users can’t access external resources after waking from sleep mode	601
601	
SSO using RADIUS accounting records.....	602
User’s view of RADIUS SSO authentication	602
Configuration Overview	602
Configuring the RADIUS server	603
Creating the FortiGate RADIUS SSO agent.....	603
Selecting which RADIUS attributes are used for RSSO	604
Configuring logging for RSSO	604
Defining local user groups for RADIUS SSO	605
Creating security policies	605
Example: webfiltering for student and teacher accounts	607
Monitoring authenticated users.....	609
Monitoring firewall users.....	609
Monitoring SSL VPN users	609
Monitoring IPsec VPN users	610
Monitoring banned users.....	610
Monitoring IM users	611
Examples and Troubleshooting	613
Firewall authentication example	613
Overview	613
Creating a locally-authenticated user account	614
Creating a RADIUS-authenticated user account	614
Creating user groups	615
Defining policy addresses.....	617
Creating security policies.....	618
LDAP Dial-in using member-attribute.....	620
RADIUS SSO example.....	621
Assumptions	621
Topology	622
General configuration.....	622

Configuring RADIUS	622
Configuring FortiGate interfaces.....	622
Configuring a RADIUS SSO Agent on the FortiGate unit	624
Creating a RADIUS SSO user group.....	624
Configuring FortiGate regular and RADIUS SSO security policies.....	625
Testing	628
Troubleshooting.....	629
Chapter 4: FortiOS Carrier.....	631
Overview of FortiOS Carrier features	632
Overview	632
MMS.....	632
GTP	632
Registering FortiOS Carrier.....	633
MMS background	633
MMS content interfaces.....	633
How MMS content interfaces are applied	634
How FortiOS Carrier processes MMS messages.....	636
FortiOS Carrier and MMS content scanning.....	637
FortiOS Carrier and MMS duplicate messages and message floods.....	642
MMS protection profiles	644
Bypassing MMS protection profile filtering based on carrier endpoints	645
Applying MMS protection profiles to MMS traffic	645
GTP basic concepts	645
PDP Context	645
GPRS security.....	647
Parts of a GTPv1 network.....	648
Radio access	649
Transport.....	649
Billing and records	652
GPRS network common interfaces	653
Packet flow through the GPRS network.....	654
SCTP.....	655
Overview	656
SCTP Firewall.....	658
SCTP example scenario	658
Carrier web-based manager settings.....	661
MMS profiles.....	661
MMS Content Checksum	674
Notification List.....	675
Message Flood	678
Duplicate Message	680
Carrier Endpoint Filter Lists	681

GTP Profile.....	683
MMS Security features.....	701
Why scan MMS messages for viruses and malware?	701
Example: COMMWARRIOR.....	701
MMS virus scanning	702
MMS virus monitoring.....	703
MMS virus scanning blocks messages (not just attachments).....	703
Scanning MM1 retrieval messages.....	703
Configuring MMS virus scanning.....	703
Removing or replacing blocked messages.....	703
Carrier Endpoint Block.....	704
MMS Content Checksum	706
Passing or blocking fragmented messages.....	707
Client comforting	707
Server comforting	708
Handling oversized MMS messages	708
MM1 sample messages.....	708
MMS file filtering	710
Built-in patterns and supported file types	711
MMS file filtering blocks messages (not just attachments)	713
Configuring MMS file filtering	713
Sender notifications and logging	713
MMS notifications	714
Replacement messages.....	715
Logging and reporting	715
MMS logging options.....	715
SNMP.....	715
MMS content-based Antispam protection	716
Overview	716
Scores and thresholds.....	717
Configuring content-based antispam protection.....	717
Configuring sender notifications.....	717
MMS DLP archiving	718
Configuring MMS DLP archiving	718
Viewing DLP archives	719
Message flood protection.....	720
Setting message flood thresholds	720
Example	721
Flood actions	722
Notifying administrators of floods.....	722
Example – three flood threshold levels with different actions for each threshold ...	722
Notifying message flood senders and receivers.....	725
Responses to MM1 senders and receivers	725

Forward responses for MM4 message floods	726
Viewing DLP archived messages.....	726
Order of operations: flood checking before duplicate checking	726
Bypassing message flood protection based on user's carrier endpoints	727
Configuring message flood detection.....	727
Sending administrator alert notifications.....	728
Configuring how and when to send alert notifications	728
Configuring who to send alert notifications to.....	730
Duplicate message protection.....	731
Using message fingerprints to identify duplicate messages	731
Messages from any sender to any recipient.....	732
Setting duplicate message thresholds	732
Duplicate message actions.....	732
Notifying duplicate message senders and receivers	733
Responses to MM1 senders and receivers	733
Forward responses for duplicate MM4 messages	734
Viewing DLP archived messages.....	734
Order of operations: flood checking before duplicate checking	735
Bypassing duplicate message detection based on user's carrier endpoints.....	735
Configuring duplicate message detection	735
Sending administrator alert notifications.....	735
Configuring how and when to send alert notifications	736
Configuring who to send alert notifications to.....	737
Select the duplicate thresholds at which to send alert notifications to the MSISDN.....	737
Configuring GTP on FortiOS Carrier.....	738
GTP support on the Carrier-enabled FortiGate unit	738
Packet sanity checking	739
GTP stateful inspection.....	739
Protocol anomaly detection and prevention.....	739
HA	739
Virtual domain support.....	740
Configuring General Settings on the Carrier-enabled FortiGate unit	740
Configuring Encapsulated Filtering in FortiOS Carrier.....	740
Configuring Encapsulated IP Traffic Filtering	740
Configuring Encapsulated Non-IP End User Address Filtering	741
Configuring the Protocol Anomaly feature in FortiOS Carrier.....	742
Configuring Anti-overbilling in FortiOS Carrier	742
Overbilling in GPRS networks.....	742
Anti-overbilling with FortiOS Carrier	742
Logging events on the Carrier-enabled FortiGate unit	743

GTP message type filtering	745
Common message types on carrier networks	745
GTP-C messages.....	745
GTP-U messages.....	746
Unknown Action messages	747
Configuring message type filtering in FortiOS Carrier	747
Message Type Fields	748
GTP identity filtering.....	752
IMSI on carrier networks.....	752
Other identity and location based information elements.....	752
When to use APN, IMSI, or advanced filtering	754
Configuring APN filtering in FortiOS Carrier	755
Configuring IMSI filtering in FortiOS Carrier	756
Configuring advanced filtering in FortiOS Carrier.....	757
Troubleshooting	760
FortiOS Carrier diagnose commands	760
GTP related diagnose commands	760
Applying IPS signatures to IP packets within GTP-U tunnels	761
GTP packets are not moving along your network	762
Attempt to identify the section of your network with the problem	762
Ensure you have an APN configured	762
Check the logs and adjust their settings if required	762
Check the routing table.....	763
Perform a sniffer trace	764
Generate specific packets to test the network	766
Chapter 5: Compliance	767
Configuring FortiGate units for PCI DSS compliance	768
Introduction to PCI DSS	768
What is PCI DSS?	768
What is the Customer Data Environment.....	768
PCI DSS objectives and requirements.....	768
Network topology	772
Internet.....	772
The CDE wired LAN	773
The CDE wireless LAN.....	773
Other internal networks.....	773
Security policies for the CDE network.....	773
Controlling the source and destination of traffic.....	773
Controlling the types of traffic in the CDE	774
The default deny policy.....	774
Wireless network security	774
On-wire detection of rogue APs	774

Setting up rogue access point scanning	774
Securing a CDE network WAP.....	775
Protecting stored cardholder data.....	776
Protecting communicated cardholder data.....	776
Configuring IPsec VPN security	777
Configuring SSL VPN security	777
Protecting the CDE network from viruses	777
Enabling FortiGate antivirus protection	777
Configuring antivirus updates	778
Enforcing firewall use on endpoint PCs.....	778
Monitoring the network for vulnerabilities.....	778
FortiGate logs	778
Using the FortiOS Network Vulnerability Scan feature	778
Monitoring with other Fortinet products	779
Restricting access to cardholder data.....	779
Controlling access to the CDE network.....	779
Password complexity and change requirements.....	779
Password non-reuse requirement	780
Administrator lockout requirement	781
Administrator timeout requirement	781
Administrator access security.....	781
Remote access security.....	781

Chapter 6: Deploying Wireless Networks for FortiOS 5.0 783

Introduction to wireless networking..... 784

Wireless concepts	784
Bands and channels	784
Power	784
Antennas	785
Security.....	785
Whether to broadcast SSID	785
Encryption.....	785
Separate access for employees and guests.....	786
Captive portal.....	786
Power	786
Monitoring for rogue APs.....	786
Authentication.....	787
Wireless networking equipment	787
FortiWiFi units	787
FortiAP units	788
Deployment considerations	789
Types of wireless deployment	789
Deployment methodology.....	789
Single access point networks.....	791

Multiple access point networks	791
Automatic Radio Resource Provisioning	792
Configuring a WiFi LAN.....	793
Overview of WiFi controller configuration.....	793
About SSIDs on FortiWiFi units	794
About automatic AP profile settings	794
Process to create a wireless network.....	795
Setting your geographic location.....	795
Creating a custom AP Profile.....	796
Defining a wireless network interface (SSID)	798
Configuring DHCP for WiFi clients.....	801
Configuring security	801
Adding a MAC filter	804
Multicast enhancement.....	805
Dynamic VLAN assignment	805
Configuring user authentication.....	807
WPA-Enterprise authentication.....	807
MAC-based authentication	808
Authenticating guest WiFi users	808
Configuring firewall policies for the SSID	809
Customizing captive portal pages	811
Modifying the login page	811
Modifying the login failed page.....	812
Configuring the built-in access point on a FortiWiFi unit.....	813
Access point deployment	814
Overview	814
Network topology for managed APs.....	814
Discovering and authorizing APs	815
Configuring the network interface for the AP unit.....	816
Pre-authorizing a FortiAP unit.....	816
Enabling and configuring a discovered AP.....	817
Assigning the same profile to multiple FortiAP units	818
Checking and updating FortiAP unit firmware	819
Advanced WiFi controller discovery	820
Controller discovery methods.....	820
Connecting to the FortiAP CLI	821
Wireless client load balancing for high-density deployments	822
Access point hand-off.....	822
Frequency hand-off or band-steering.....	822
Configuration	823
LAN port options.....	823
Bridging a LAN port with a FortiAP SSID.....	823
Bridging a LAN port with the WAN port.....	824

Configuring FortiAP LAN ports	824
Preventing IP fragmentation of packets in CAPWAP tunnels.....	826
Wireless Mesh.....	828
Overview of Wireless Mesh	828
Wireless mesh deployment modes.....	829
Firmware requirements	829
Types of wireless mesh	829
Configuring a meshed WiFi network.....	831
Creating custom AP profiles	831
Configuring the mesh root AP	831
Configuring the mesh branches or leaves	833
Authorizing mesh branch/leaf APs.....	833
Viewing the status of the mesh network.....	834
Configuring a point-to-point bridge.....	834
WiFi-Ethernet Bridge Operation.....	835
Bridge SSID to FortiGate wired network	835
VLAN configuration	838
Additional configuration.....	838
FortiAP local bridging (Private Cloud-Managed AP).....	839
Continued FortiAP operation when WiFi controller connection is down	841
Using bridged FortiAPs to increase scalability	842
Protecting the WiFi Network	843
Wireless IDS.....	843
WiFi data channel encryption	845
Configuring encryption on the FortiGate unit	845
Configuring encryption on the FortiAP unit	845
Wireless network monitoring	846
Monitoring wireless clients	846
Monitoring rogue APs	847
On-wire rogue AP detection technique	847
Rogue AP scanning as a background activity	848
Configuring rogue scanning.....	848
Using the Rogue AP Monitor	850
Suppressing rogue APs	851
Monitoring wireless network health	851
Configuring wireless network clients.....	852
Windows XP client	852
Windows 7 client.....	856
Mac OS client	857
Linux client.....	859
Troubleshooting.....	861
Checking that the client has received IP address and DNS server information...	

Wireless network examples	863
Basic wireless network	863
Configuring authentication for wireless users.....	863
Configuring the SSID	864
Configuring firewall policies.....	865
Connecting the FortiAP units.....	866
A more complex example	868
Scenario	868
Configuration	868
Configuring authentication for employee wireless users.....	869
Configuring authentication for guest wireless users.....	869
Configuring the SSIDs	871
Configuring the custom AP profile.....	873
Configuring firewall policies.....	874
Connecting the FortiAP units.....	876
Using a FortiWiFi unit as a client	879
Use of client mode.....	879
Configuring client mode.....	880
Support for location-based services	881
Overview	881
Configuring location tracking.....	881
Viewing device location data on the FortiGate unit.....	882
Reference	883
Wireless radio channels.....	883
IEEE 802.11a/n channels	883
FortiAP CLI.....	885
FortiAP web-based manager.....	886
Chapter 7: Firewall for FortiOS 5.0	888
FortiGate Firewall Components.....	888
How does a FortiGate Protect Your Network.....	889
Firewall concepts	891
What is a Firewall?	891
Network Layer or Packet Filter Firewalls	891
Application Layer Firewalls	892
Proxy Servers.....	892
Security Profiles	893
IPv6.....	894
What is IPv6?	894
IPv6 in FortiOS.....	896
Dual Stack routing configuration	896
IPv6 Tunnelling.....	897

Tunnelling IPv6 through IPSec VPN.....	898
NAT	898
What is NAT?	898
The Origins of NAT.....	898
Static NAT.....	899
Dynamic NAT	899
Benefits of NAT.....	901
NAT in Transparent Mode.....	902
Central NAT Table.....	902
NAT 64 and NAT46	903
NAT 66	903
How Packets are handled by FortiOS	904
FortiGate Modes	905
NAT/Route Mode	906
Transparent Mode.....	906
Quality of Service.....	906
Traffic policing	906
Traffic Shaping.....	907
Queuing.....	907
Interfaces and Zones	907
Firewall objects	909
Addresses	909
IPv4 Address and Net Mask	910
FQDN Addressing	911
Geography Based Addressing.....	911
Address Groups.....	912
Wildcard Addressing.....	913
Virtual IP Addresses (VIPs).....	915
Virtual IP Groups.....	915
IP Pools.....	916
Fixed Port.....	918
Match-VIP	918
Services and TCP ports.....	918
Categories.....	919
Protocol Types.....	919
Service Groups	935
Example Scenario: Using FortiGate services to support Audio/Visual Conferencing	
936	
VIP	936
Creating an address for the subnet	937
Configuring the services	937
Creating the Service Group	939
Creating the IPS Security Profile	940
Policies.....	941

Firewall schedules	942
Schedule Groups	943
Schedule Expiration	944
Security profiles	944
AntiVirus	945
Web Filtering	945
Application Control	945
Intrusion Protection (IPS)	945
Email Filtering	946
Data Leak Prevention (DLP)	946
VoIP	946
ICAP	946
EndPoint Control	946
Proxy Option Components	947
The use of different proxy profiles and profile options	947
SSL/SSH Inspection	950
Creating a new SSL/SSH Inspection profile	950
Security policies	952
Firewall policies	952
What is not expressly allowed is denied	953
Policy order	954
Viewing Firewall Policies	957
How “Any” policy can remove the Section View	958
Security policy configuration extensions	958
Identity Based Policies	959
Identity-based policy positioning	959
Identity-based sub-policies	960
Identity policies an unauthenticated users	960
Device Identity Policies	960
VPN Policies	961
IPSec Policies	961
SSL VPN Policies	961
Interface Policies	962
DoS Protection	962
One-Arm IDS	966
IPv6 IPS	966
Traffic Destined to the FortiGate unit	966
Dropped, Flooded, Broadcast, Multicast and L2 packets	966
GUI and CLI	967
Local-In Policies	967
Security Policy 0	968
Deny Policies	969
Accept Policies	969
IPv6 Policies	969

Fixed Port	969
Endpoint Security	970
Traffic Logging.....	970
Quality of Service.....	972
Queuing.....	972
Policy Monitor.....	973
Upper Pane.....	973
Lower Pane.....	973
Network defense	974
Monitoring.....	974
Blocking external probes.....	974
Address sweeps	974
Port scans.....	975
Probes using IP traffic options.....	975
Evasion techniques.....	976
Defending against DoS attacks	979
The “three-way handshake”	979
SYN flood.....	979
SYN spoofing.....	980
DDoS SYN flood	980
Configuring the SYN threshold to prevent SYN floods.....	981
SYN proxy.....	981
Other flood types	981
DoS policies.....	982
GUI & CLI - What You May Not Know.....	983
Mouse Tricks	983
Changing the default column setting on the policy page.....	984
Example:.....	985
Naming Rules and Restrictions	985
Character Restrictions	985
Length of Fields Restrictions	986
Object Tagging and Coloring.....	986
Tags	987
Coloring.....	987
Numeric Values.....	987
Selecting options from a list	988
Enabling or disabling options	988
To Enable or Disable Optionally Displayed Features.....	988
Building firewall objects and policies.....	989
IPv4 Firewall Addresses.....	990
Scenario: Mail Server.....	990
Scenario: First Floor Network	990

Scenario: Marketing Department	991
Verification	992
IPv6 Firewall Addresses	992
Scenario: Mail Server	992
Scenario: First Floor Network	993
Verification	993
FQDN address	994
Verification	994
Changing the TTL of a FQDN address	995
New Geography-based Address	995
Wildcard Address	996
IPv4 Address Group	997
IPv6 Address Group	998
Multicast Address	998
Service Category	999
TCP/UDP/SCTP Service	1000
ICMP Service	1002
ICMPv6 Service	1003
Service Group	1004
Virtual IP address	1006
VIP Group	1007
IP Pool	1008
Central NAT Table	1009
Firewall Schedule - Recurring	1010
Firewall Schedule - One-time	1012
Schedule Group	1013
Proxy Option	1014
Oversized Files	1018
Firewall Address Policy	1019
Firewall User Identity Policy	1021
Firewall Device Identity Policy	1024
DoS Policy	1026
Multicast forwarding	1030
Sparse mode	1030
Dense mode	1031
Multicast IP addresses	1032
PIM Support	1033
Multicast forwarding and FortiGate units	1033
Multicast forwarding and RIPv2	1034
Configuring FortiGate multicast forwarding	1035
Adding multicast security policies	1036

Enabling multicast forwarding	1036
Multicast routing examples.....	1038
Example FortiGate PIM-SM configuration using a static RP	1038
FortiGate PIM-SM debugging examples	1047
Example multicast destination NAT (DNAT) configuration	1053
Example PIM configuration that uses BSR to find the RP	1055
Chapter 8: Hardware Acceleration	1067
Hardware acceleration overview	1068
Content processors (CP4, CP5, CP6 and CP8).....	1068
Determining the content processor in your FortiGate unit	1070
Viewing SSL acceleration status.....	1070
Disabling CP offloading	1070
Security processors (SPs).....	1071
SP Processing Flow	1071
Displaying information about security processing modules	1072
Network processors (NP1, NP2, NP3, NP4 and NP6)	1073
Determining the network processors installed on your FortiGate unit	1074
How NP hardware acceleration alters packet flow.....	1074
NP processors and traffic logging and monitoring	1075
NP session offloading in HA active-active configuration.....	1075
Configuring NP HMAC check offloading	1076
Offloading NP pre-IPS anomaly detection.....	1076
Software switch interfaces and NP processors.....	1077
Configuring NP accelerated VPN encryption/decryption offloading	1078
Checking that traffic is offloaded by NP processors	1078
Using the packet sniffer	1079
Checking the firewall session offload tag	1079
Verifying IPsec VPN traffic offloading	1080
Controlling IPS NPx and CPx acceleration.....	1081
NP6 Acceleration.....	1082
NP6 session fast path requirements.....	1083
Packet fast path requirements.....	1084
Mixing fast path and non-fast path traffic.....	1084
Viewing your FortiGate NP6 processor configuration	1084
Increasing NP6 offloading capacity using link aggregation groups (LAGs).....	1084
Configuring Inter-VDOM link acceleration with NP6 processors.....	1085
Using VLANs to add more accelerated Inter-VDOM links	1086
Confirm that the traffic is accelerated.....	1087
FortiGate NP6 architectures.....	1089
FortiGate-1500D fast path architecture	1089
FortiGate-3700D fast path architecture	1091

NP4 Acceleration	1093
Viewing your FortiGate's NP4 configuration.....	1094
NP4lite CLI commands (disabling NP4Lite offloading).....	1094
Configuring NP4 traffic offloading	1094
NP4 session fast path requirements.....	1094
Packet fast path requirements.....	1095
Mixing fast path and non-fast path traffic.....	1095
NP4 traffic shaping offloading	1096
NP4 IPsec VPN offloading	1096
NP4 IPsec VPN offloading configuration example	1097
Accelerated policy mode IPsec configuration	1098
Accelerated interface mode IPsec configuration.....	1099
Configuring Inter-VDOM link acceleration with NP4 processors.....	1101
Using VLANs to add more accelerated Inter-VDOM links	1101
Confirm that the traffic is accelerated.....	1102
FortiGate NP4 architectures	1105
FortiGate-600C	1105
FortiGate-800C	1106
FortiGate-1000C	1106
FortiGate-1240B	1107
FortiGate-3040B	1107
FortiGate-3140B	1108
FortiGate-3140B — load balance mode.....	1109
FortiGate-3240C	1110
FortiGate-3600C	1111
XAUI interfaces	1111
FortiGate-3950B and FortiGate-3951B	1112
FortiGate-3950B and FortiGate-3951B — load balance mode.....	1113
FortiGate-5001C	1114
FortiGate-5001B	1114
Setting switch-mode mapping on the ADM-XD4	1115
Chapter 9: High Availability for FortiOS 5.0	1116
Solving the High Availability problem	1117
FortiGate Cluster Protocol (FGCP)	1117
FortiGate Session Life Support Protocol (FGSP).....	1118
VRRP.....	1119
Fortinet redundant UTM protocol (FRUP).....	1119
An introduction to the FGCP	1121
About the FGCP.....	1122
FGCP failover protection	1123
Session Failover.....	1123

Load Balancing	1123
Virtual Clustering.....	1123
Full Mesh HA.....	1124
Cluster Management.....	1124
Synchronizing the configuration (and settings that are not synchronized).....	1124
Configuring FortiGate units for FGCP HA operation	1125
Connecting a FortiGate HA cluster	1127
Active-passive and active-active HA.....	1128
Active-passive HA (failover protection).....	1128
Active-active HA (load balancing and failover protection).....	1129
Identifying the cluster and cluster units.....	1129
Group name	1129
Password	1130
Group ID.....	1130
Device failover, link failover, and session failover.....	1130
Primary unit selection	1131
Primary unit selection and monitored interfaces	1132
Primary unit selection and age	1133
Primary unit selection and device priority.....	1136
Primary unit selection and the FortiGate unit serial number.....	1137
Points to remember about primary unit selection.....	1138
HA override	1138
Override and primary unit selection	1139
Controlling primary unit selection using device priority and override.....	1140
Points to remember about primary unit selection when override is enabled	1141
Configuration changes can be lost if override is enabled.....	1141
Override and disconnecting a unit from a cluster.....	1142
FortiGate HA compatibility with PPPoE and DHCP.....	1142
HA and distributed clustering	1143
Hard disk configuration and HA	1143
FGCP high availability best practices	1143
Heartbeat interfaces	1144
Interface monitoring (port monitoring)	1145
Troubleshooting	1145
FGCP HA terminology	1145
HA web-based manager options.....	1149
Configuring and connecting HA clusters	1152
About the procedures in this chapter	1152
Example: NAT/Route mode active-passive HA configuration	1152
Example NAT/Route mode HA network topology	1153
General configuration steps.....	1153
Configuring a NAT/Route mode active-passive cluster of two FortiGate-620B	

units - web-based manager	1154
Configuring a NAT/Route mode active-passive cluster of two FortiGate-620B units - CLI	1158
Example: Transparent mode active-active HA configuration	1164
Example Transparent mode HA network topology.....	1164
General configuration steps.....	1165
Configuring a Transparent mode active-active cluster of two FortiGate-620B units - web-based manager	1166
Configuring a Transparent mode active-active cluster of two FortiGate-620B units - CLI	1170
Example: advanced Transparent mode active-active HA configuration	1176
Example Transparent mode HA network topology.....	1177
Configuring a Transparent mode active-active cluster of three FortiGate-5005FA2 units - web-based manager.....	1177
Configuring a Transparent mode active-active cluster of three FortiGate-5005FA2 units - CLI	1180
Example: converting a standalone FortiGate unit to a cluster.....	1184
Example: adding a new unit to an operating cluster	1186
Example: replacing a failed cluster unit	1187
Example: HA and 802.3ad aggregated interfaces	1188
HA interface monitoring, link failover, and 802.3ad aggregation.....	1188
HA MAC addresses and 802.3ad aggregation	1188
Link aggregation, HA failover performance, and HA mode	1189
General configuration steps.....	1189
Configuring active-passive HA cluster that includes aggregated interfaces - web-based manager	1190
Configuring active-passive HA cluster that includes aggregate interfaces - CLI .	1194
Example: HA and redundant interfaces	1200
HA interface monitoring, link failover, and redundant interfaces.....	1200
HA MAC addresses and redundant interfaces	1201
Connecting multiple redundant interfaces to one switch while operating in active-passive HA mode.....	1201
Connecting multiple redundant interfaces to one switch while operating in active-active HA mode	1201
General configuration steps.....	1201
Configuring active-passive HA cluster that includes redundant interfaces - web-based manager	1202
Configuring active-passive HA cluster that includes redundant interfaces - CLI .	1206
Troubleshooting HA clusters	1212
Ignoring hardware revisions.....	1212
Before you set up a cluster	1212
Troubleshooting the initial cluster configuration.....	1213
More troubleshooting information	1215

Virtual clusters	1217
Virtual clustering overview	1217
Virtual clustering and failover protection	1217
Virtual clustering and heartbeat interfaces	1217
Virtual clustering and HA override	1218
Virtual clustering and load balancing or VDOM partitioning	1218
Configuring HA for virtual clustering	1219
Example: virtual clustering with two VDOMs and VDOM partitioning	1221
Example virtual clustering network topology	1221
General configuration steps	1222
Configuring virtual clustering with two VDOMs and VDOM partitioning - web-based manager	1223
Configuring virtual clustering with two VDOMs and VDOM partitioning - CLI.....	1228
Example: inter-VDOM links in a virtual clustering configuration.....	1236
Configuring inter-VDOM links in a virtual clustering configuration	1237
Troubleshooting virtual clustering.....	1238
Full mesh HA	1239
Full mesh HA overview	1239
Full mesh HA and redundant heartbeat interfaces	1240
Full mesh HA, redundant interfaces and 802.3ad aggregate interfaces	1240
Example: full mesh HA configuration.....	1241
FortiGate-620B full mesh HA configuration.....	1242
Full mesh switch configuration	1242
Full mesh network connections	1242
How packets travel from the internal network through the full mesh cluster and to the Internet	1242
Configuring FortiGate-620B units for HA operation - web-based manager	1243
Configuring FortiGate-620B units for HA operation - CLI	1247
Troubleshooting full mesh HA	1251
Operating a cluster	1252
Operating a cluster	1252
Operating a virtual cluster.....	1253
Managing individual cluster units using a reserved management interface.....	1254
Configuring the reserved management interface and SNMP remote management of individual cluster units.....	1255
The primary unit acts as a router for subordinate unit management traffic	1259
Cluster communication with RADIUS and LDAP servers	1260
Clusters and FortiGuard services	1260
FortiGuard and active-passive clusters	1260
FortiGuard and active-active clusters.....	1260
FortiGuard and virtual clustering	1261
Clusters and logging.....	1261
Viewing and managing log messages for individual cluster units	1261

HA log messages	1262
Fortigate HA message "HA master heartbeat interface <intf_name> lost neighbor information"	1262
Formatting cluster unit hard disks (log disks)	1264
Clusters and SNMP	1264
SNMP get command syntax for the primary unit	1264
SNMP get command syntax for any cluster unit	1266
Getting serial numbers of cluster units	1267
SNMP get command syntax - reserved management interface enabled....	1267
Clusters and file quarantine	1268
Cluster members list	1268
Virtual cluster members list	1270
Viewing HA statistics	1271
Changing the HA configuration of an operating cluster	1273
Changing the HA configuration of an operating virtual cluster	1273
Changing the subordinate unit host name and device priority	1273
Upgrading cluster firmware	1274
Changing how the cluster processes firmware upgrades	1275
Synchronizing the firmware build running on a new cluster unit	1275
Downgrading cluster firmware	1275
Backing up and restoring the cluster configuration.....	1276
Monitoring cluster units for failover	1277
Viewing cluster status from the CLI	1277
Examples	1279
About the HA cluster index and the execute ha manage command	1282
Managing individual cluster units	1284
Disconnecting a cluster unit from a cluster	1285
Adding a disconnected FortiGate unit back to its cluster	1286
HA diagnose commands	1287
all-xdb	1288
all-vcluster.....	1289
stat	1289
HA and failover protection.....	1290
About active-passive failover.....	1290
Device failure	1291
Link failure.....	1291
Session failover.....	1291
Primary unit recovery	1291
About active-active failover	1292
Device failover	1292
HA heartbeat and communication between cluster units.....	1293
Heartbeat interfaces	1293
Connecting HA heartbeat interfaces.....	1295

Heartbeat packets and heartbeat interface selection.....	1295
Interface index and display order	1296
HA heartbeat interface IP addresses	1296
Heartbeat packet Ethertypes	1297
Modifying heartbeat timing	1298
Enabling or disabling HA heartbeat encryption and authentication	1299
Cluster virtual MAC addresses	1300
Changing how the primary unit sends gratuitous ARP packets after a failover ...	1301
Disabling gratuitous ARP packets after a failover	1302
How the virtual MAC address is determined	1302
Displaying the virtual MAC address.....	1304
Diagnosing packet loss with two FortiGate HA clusters in the same broadcast domain.....	1305
Synchronizing the configuration	1307
Configuration settings that are not synchronized	1307
Disabling automatic configuration synchronization	1308
Incremental synchronization	1308
Periodic synchronization.....	1309
Console messages when configuration synchronization succeeds	1310
Console messages when configuration synchronization fails	1310
Comparing checksums of cluster units	1312
How to diagnose HA out of sync messages.....	1313
Recalculating the checksums to resolve out of sync messages	1315
Synchronizing kernel routing tables.....	1315
Configuring graceful restart for dynamic routing failover	1315
Controlling how the FGCP synchronizes kernel routing table updates	1316
Synchronizing IPsec VPN SAs.....	1318
Link failover (port monitoring or interface monitoring).....	1319
If a monitored interface on the primary unit fails	1320
If a monitored interface on a subordinate unit fails	1320
How link failover maintains traffic flow	1321
Recovery after a link failover and controlling primary unit selection (controlling falling back to the prior primary unit).....	1322
Preventing a primary unit change after a failed link is restored.....	1322
Testing link failover	1322
Updating MAC forwarding tables when a link failover occurs.....	1323
Multiple link failures	1323
Example link failover scenarios.....	1323
Subsecond failover	1324
Remote link failover	1325
Adding HA remote IP monitoring to multiple interfaces	1327
Changing the ping server failover threshold	1328
Monitoring multiple IP addresses from one interface	1329
Flip timeout	1329

Detecting HA remote IP monitoring failovers.....	1330
Session failover (session pick-up)	1330
If session pickup is not selected.....	1330
Improving session synchronization performance	1331
Session failover not supported for all sessions	1332
IPv6, NAT64, and NAT66 session failover	1333
SIP session failover.....	1333
Explicit web proxy, WCCP, and WAN optimization session failover	1333
SSL offloading and HTTP multiplexing session failover	1333
IPsec VPN session failover	1333
SSL VPN session failover and SSL VPN authentication failover	1333
PPTP and L2TP VPN sessions	1334
UDP, ICMP, multicast and broadcast packet session failover	1334
FortiOS Carrier GTP session failover	1334
Active-active HA subordinate units sessions can resume after a failover ...	1335
WAN optimization and HA	1335
Failover and attached network equipment	1335
Monitoring cluster units for failover	1336
NAT/Route mode active-passive cluster packet flow.....	1336
Packet flow from client to web server	1337
Packet flow from web server to client	1337
When a failover occurs	1338
Transparent mode active-passive cluster packet flow	1338
Packet flow from client to mail server.....	1339
Packet flow from mail server to client.....	1339
When a failover occurs	1340
Failover performance	1340
Device failover performance	1340
Link failover performance	1341
Reducing failover times	1342
HA and load balancing	1343
Load balancing overview	1343
Load balancing schedules	1344
Selecting which packets are load balanced	1345
More about active-active failover	1345
HTTPS sessions, active-active load balancing, and proxy servers	1345
Using FortiGate network processor interfaces to accelerate active-active HA performance	1346
Configuring load balancing settings	1347
Selecting a load balancing schedule	1347
Load balancing UTM sessions, TCP sessions, and UDP sessions	1347
Configuring weighted-round-robin weights.....	1348
Dynamically optimizing weighted load balancing according to how busy cluster units are	1350

NAT/Route mode active-active cluster packet flow	1354
Packet flow from client to web server	1354
Packet flow from web server to client	1355
When a failover occurs	1356
Transparent mode active-active cluster packet flow.....	1356
Packet flow from client to mail server.....	1357
Packet flow from mail server to client.....	1358
When a failover occurs	1359
HA with FortiGate-VM and third-party products	1360
FortiGate-VM for VMware HA configuration.....	1360
FortiGate VM for Hyper-V HA configuration	1361
Troubleshooting layer-2 switches.....	1361
Forwarding delay on layer 2 switches	1361
Failover issues with layer-3 switches	1361
Changing spanning tree protocol settings for some switches	1362
Spanning Tree protocol (STP).....	1362
Bridge Protocol Data Unit (BPDU)	1362
Failover and attached network equipment	1362
Ethertype conflicts with third-party switches	1363
LACP, 802.3ad aggregation and third-party switches	1363
VRRP.....	1364
Adding a VRRP virtual router to a FortiGate interface	1365
VRRP virtual MAC address	1365
Configuring VRRP	1366
Example VRRP configuration: two FortiGate units in a VRRP group	1366
Example VRRP configuration: VRRP load balancing two FortiGate units and two VRRP groups.....	1367
Optional VRRP configuration settings	1369
FortiGate Session Life Support Protocol (FGSP).....	1370
Synchronizing the configuration	1371
Synchronizing UDP and ICMP (connectionless) sessions.....	1372
Synchronizing NAT sessions	1372
Synchronizing expectation (asymmetric) sessions.....	1372
UTM Flow-based Inspection and Asymmetric Traffic	1373
Notes and limitations	1373
Configuring FGSP HA.....	1374
Configuring the session synchronization link	1374
Basic example configuration	1375
Verifying FGSP configuration and synchronization	1377
FGSP configuration summary and status.....	1378
Verifying that sessions are synchronized.....	1379

Configuring FRUP	1380
FRUP configuration example	1381
Configuring FGT-A	1381
Configuring FGT-B	1382
Connecting, testing and operating the FRUP cluster	1382
.....	1385
Chapter 10: Install and System Administration for FortiOS 5.0	1385
Differences between Models and Firmware	1386
Differences between Models	1386
Differences between Firmware Versions	1386
Using the web-based manager	1387
Web-based manager overview	1387
Web-based manager menus and pages	1387
Using information tables	1388
Using column settings	1389
Entering text strings	1389
Entering text strings (names)	1389
Entering numeric values	1390
Enabling or disabling options	1390
Dashboard	1390
Adding dashboards and widgets	1391
System Information widget	1391
License Information widget	1397
FortiGate unit Operation widget	1399
System Resources widget	1399
Alert Message Console widget	1399
CLI Console widget	1399
Session History widget	1400
Top Sessions widget	1400
USB Modem widget	1400
Advanced Threat Protection Statistics widget	1400
Features widget	1400
RAID monitor widget	1401
Basic configurations	1402
Changing your administrator password	1402
Changing the web-based manager language	1403
Changing administrative access	1403
Changing the web-based manager idle timeout	1403
Switching VDOMs	1403
Connecting to the CLI from the web-based manager	1403
Logging out	1404
Using the CLI	1405
Connecting to the CLI	1405

Connecting to the CLI using a local console	1405
Enabling access to the CLI through the network (SSH or Telnet)	1406
Connecting to the CLI using SSH	1407
Connecting to the CLI using Telnet	1408
Command syntax.....	1409
Terminology	1409
Indentation	1410
Notation	1410
Sub-commands	1412
Example of table commands	1414
Permissions	1415
Tips	1416
Help.....	1416
Shortcuts and key commands	1416
Command abbreviation	1417
Adding and removing options from lists	1417
Environment variables.....	1418
Special characters	1418
Using grep to filter get and show command output	1419
Language support and regular expressions	1420
Screen paging.....	1422
Baud rate	1423
Editing the configuration file on an external host	1423
Using Perl regular expressions	1423
Basic Administration	1426
Connecting to the FortiGate unit	1426
Connecting to the web-based manager	1426
Connecting to the CLI	1427
System configuration	1427
Setting the time and date.....	1427
Configuring FortiGuard	1428
Passwords	1429
Password considerations.....	1429
Password policy.....	1430
<i>Lost Passwords</i>	1431
Administrators.....	1431
Adding administrators.....	1431
LDAP Admin Access and Authorization.....	1432
Monitoring administrators.....	1433
Administrator profiles.....	1434
Regular (password) authentication for administrators	1435
Management access.....	1435
Security Precautions	1436
General Settings	1441

Administrative port settings	1441
Password policies	1441
Feature Select	1441
Configuration backups	1442
Backup and restore a configuration file using SCP	1443
Restoring a configuration	1445
Configuration revisions	1446
Restore factory defaults	1446
Firmware	1446
Downloading firmware	1447
Testing new firmware before installing	1447
Upgrading the firmware - web-based manager	1449
Upgrading the firmware - CLI	1449
Installing firmware from a system reboot using the CLI	1450
Reverting to a previous firmware version - web-based manager	1452
Reverting to a previous firmware version - CLI	1452
Configuration Revision	1453
Backup and Restore from a USB key	1453
Backup and Restore an encrypted config file from a USB key	1454
Controlled upgrade	1454
Best practices	1456
Hardware	1456
Environmental specifications	1456
Grounding	1457
Rack mount instructions	1457
Shutting down	1458
Performance	1458
Firewall	1458
Intrusion protection	1459
Antivirus	1459
Web filtering	1460
Antispam	1460
Security	1460
FortiGuard	1461
FortiGuard Services	1461
<i>Next Generation Firewall</i>	1461
Advanced Threat Protection	1462
Other Services	1462
Support Contract and FortiGuard Subscription Services	1463
FortiCloud	1463
Antivirus and IPS	1463
Detection during update	1463
Antivirus and IPS Options	1464

Manual updates	1464
Automatic updates.....	1465
Push updates.....	1465
Push IP override.....	1466
Web filtering.....	1467
Web Filtering and Email Filtering Options.....	1468
URL verification.....	1468
Email filtering	1469
Security tools	1469
URL lookup	1469
IP and signature lookup	1469
Online virus scanner	1470
Malware removal tools.....	1470
FortiSandbox	1470
Troubleshooting.....	1470
Web-based manager verification.....	1470
CLI verification	1472
Port assignment.....	1472
FortiCloud.....	1474
FortiCloud Features	1474
Simplified central management for your FortiGate network.....	1474
Hosted log retention with large default storage allocated	1474
Monitoring and alerting in real time	1474
Customized or pre-configured reporting and analysis tools	1474
Maintain important configuration information uniformly	1474
Service security.....	1474
Registration and Activation.....	1475
Registering with Support	1475
Registering and Activating your FortiCloud account.....	1475
Enabling logging to FortiCloud	1476
Logging into the FortiCloud portal.....	1477
Upgrading to a 200Gb subscription	1477
The FortiCloud Portal.....	1477
Using FortiCloud	1478
Cloud Sandboxing	1479
Interfaces	1480
Physical.....	1480
Interface settings	1482
Interface configuration and settings	1483
Software switch	1486
Soft switch example	1487
Virtual Switch	1488
Loopback interfaces	1489

Redundant interfaces.....	1489
One-armed sniffer.....	1490
Aggregate Interfaces	1491
DHCP addressing mode on an interface.....	1492
PPPoE addressing mode on an interface.....	1493
Administrative access.....	1495
Wireless	1495
Interface MTU packet size.....	1496
Secondary IP addresses to an interface.....	1497
Virtual domains	1497
Virtual LANs	1498
Zones	1499
Probing Interfaces.....	1500
Central management.....	1502
Adding a FortiGate to FortiManager	1502
FortiGate configuration	1502
FortiManager configuration.....	1503
Configuration through FortiManager	1503
Global objects.....	1504
Locking the FortiGate web-based manager	1504
Firmware updates	1504
FortiGuard.....	1504
Backup and restore configurations	1505
Administrative domains	1505
Monitoring	1506
Dashboard	1506
Widgets.....	1506
FortiClient software.....	1507
sFlow.....	1507
Configuration	1508
Monitor menus.....	1508
Logging.....	1508
FortiCloud	1509
FortiGate memory	1509
FortiGate hard disk	1509
Syslog server	1510
FortiAnalyzer	1510
Sending logs using a secure connection.....	1511
Packet Capture	1512
Alert email	1513
SNMP.....	1514
SNMP configuration settings.....	1515

Gigabit interfaces.....	1517
SNMP agent.....	1517
SNMP community.....	1518
Enabling on the interface.....	1519
Fortinet MIBs.....	1520
SNMP get command syntax.....	1521
VLANs.....	1523
VLAN ID rules.....	1524
VLAN switching and routing.....	1524
VLAN layer-2 switching.....	1524
VLAN layer-3 routing.....	1527
VLANs in NAT mode.....	1530
Adding VLAN subinterfaces.....	1530
Configuring security policies and routing.....	1532
Example VLAN configuration in NAT mode.....	1533
General configuration steps.....	1534
Configure the FortiGate unit.....	1535
Configure the VLAN switch.....	1540
Test the configuration.....	1541
VLANs in transparent mode.....	1541
VLANs and transparent mode.....	1541
Example of VLANs in transparent mode.....	1544
General configuration steps.....	1544
Configure the FortiGate unit.....	1545
Configure the Cisco switch and router.....	1548
Test the configuration.....	1550
Troubleshooting VLAN issues.....	1550
Asymmetric routing.....	1550
Layer-2 and Arp traffic.....	1551
Forward-domain solution.....	1552
NetBIOS.....	1553
STP forwarding.....	1553
Too many VLAN interfaces.....	1554
PPTP and L2TP.....	1555
How PPTP VPNs work.....	1555
FortiGate unit as a PPTP server.....	1557
Configuring user authentication for PPTP clients.....	1557
Enabling PPTP and specifying the PPTP IP address range.....	1558
Adding the security policy.....	1559
Configuring the FortiGate unit for PPTP VPN.....	1560
Configuring the FortiGate unit for PPTP pass through.....	1560
Configuring a virtual IP address.....	1560
Configuring a port-forwarding security policy.....	1561

Testing PPTP VPN connections	1562
Logging VPN events	1562
Configuring L2TP VPNs	1562
Network topology	1564
L2TP infrastructure requirements	1564
L2TP configuration overview	1564
Authenticating L2TP clients	1565
Enabling L2TP and specifying an address range	1565
Defining firewall source and destination addresses	1565
Adding the security policy	1566
Configuring a Linux client	1566
Monitoring L2TP sessions	1567
Testing L2TP VPN connections	1567
Logging L2TP VPN events	1567
Advanced concepts.....	1568
Dual internet connections (redundant Internet connections).....	1568
Redundant interfaces.....	1568
Load sharing	1571
Link redundancy and load sharing.....	1571
Single firewall vs. multiple virtual domains	1571
Single firewall vs. vdoms	1572
Modem.....	1574
USB modem port.....	1574
Modes	1574
Additional modem configuration.....	1576
Modem interface routing.....	1576
DHCP servers and relays.....	1577
DHCP Server configuration.....	1577
DHCP in IPv6	1578
Service	1578
Lease time.....	1578
DHCP options	1579
Exclude addresses in DHCP a range.....	1579
DHCP Monitor.....	1579
Breaking a address lease.....	1580
Assigning IP address by MAC address	1580
DNS services	1580
DNS settings	1580
Additional DNS CLI configuration	1581
DNS server.....	1581
Recursive DNS.....	1582
Dynamic DNS	1583
FortiClient discovery and registration	1583
FortiClient discovery	1584

FortiClient Registration	1584
IP addresses for self-originated traffic.....	1584
Administration for schools	1585
Security policies.....	1585
DNS.....	1586
Encrypted traffic (HTTPS)	1586
FTP.....	1586
Example security policies	1586
UTM security profiles	1587
Logging	1588
Tag management.....	1589
Adding and removing tags.....	1589
Reviewing tags.....	1590
Tagging guidelines	1590
Replacement messages list.....	1591
Replacement message images.....	1591
Adding images to replacement messages.....	1591
Modifying replacement messages	1592
Replacement message tags	1592
Administration replacement message	1594
Alert Mail replacement messages.....	1595
Authentication replacement messages.....	1595
Captive Portal Default replacement messages.....	1596
Device Detection Portal replacement message.....	1596
Email replacement messages	1596
Endpoint Control replacement message	1596
FTP replacement messages	1596
FortiGuard Web Filtering replacement messages	1596
HTTP replacement messages.....	1596
IM replacement messages.....	1597
Nntp replacement messages	1597
Spam replacement messages	1597
NAC quarantine replacement messages	1597
SSL VPN replacement message.....	1597
Web Proxy replacement messages	1597
Traffic quota control replacement messages	1598
MM1 replacement messages.....	1598
MM3 replacement messages.....	1598
MM4 replacement messages.....	1598
MM7 replacement messages.....	1598
MMS replacement messages	1598
Replacement message groups	1598
Disk.....	1599
Formatting the disk	1599
Setting space quotas	1599

CLI Scripts	1599
Uploading script files	1600
Rejecting PING requests	1600
Opening TCP 113	1601
Obfuscate HTTP responses.....	1601
Session helpers	1602
Viewing the session helper configuration	1602
Changing the session helper configuration	1603
Changing the protocol or port that a session helper listens on.....	1603
Disabling a session helper	1605
DCE-RPC session helper (dcerpc).....	1606
DNS session helpers (dns-tcp and dns-udp)	1606
File transfer protocol (FTP) session helper (ftp).....	1606
H.245 session helpers (h245I and h245O).....	1606
H.323 and RAS session helpers (h323 and ras)	1607
Alternate H.323 gatekeepers	1607
Media Gateway Controller Protocol (MGCP) session helper (mgcp).....	1607
ONC-RPC portmapper session helper (pmap)	1608
PPTP session helper for PPTP traffic (pptp).....	1608
Remote shell session helper (rsh)	1609
Real-Time Streaming Protocol (RTSP) session helper (rtsp)	1610
Session Initiation Protocol (SIP) session helper (sip).....	1610
Trivial File Transfer Protocol (TFTP) session helper (tftp)	1610
Oracle TNS listener session helper (tns).....	1611
Chapter 11: IPsec VPN for FortiOS 5.0.....	1612
IPsec VPN concepts.....	1614
VPN tunnels	1614
VPN gateways.....	1615
Clients, servers, and peers	1616
Encryption.....	1617
Authentication.....	1617
Preshared keys	1617
Additional authentication	1618
Phase 1 and Phase 2 settings	1618
Phase 1	1618
Phase 2	1618
Security Association	1619
IPsec VPN Overview.....	1620
Types of VPNs	1620
Route-based VPNs	1620

Policy-based VPNs	1621
Comparing policy-based or route-based VPNs.....	1621
Planning your VPN	1621
Network topologies	1622
General preparation steps	1623
How to use this guide to configure an IPsec VPN.....	1623
IPsec VPN in the web-based manager	1624
Auto Key (IKE)	1624
Phase 1 configuration	1625
Phase 1 advanced configuration settings.....	1626
Phase 2 configuration	1629
Phase 2 advanced configuration settings.....	1629
FortiClient VPN	1632
Manual Key	1633
Manual key configuration settings.....	1633
Concentrator	1635
IPsec Monitor.....	1635
Auto Key phase 1 parameters	1637
Overview	1637
Defining the tunnel ends	1638
Choosing main mode or aggressive mode.....	1638
Choosing the IKE version	1639
Authenticating the FortiGate unit.....	1639
Authenticating the FortiGate unit with digital certificates	1639
Authenticating the FortiGate unit with a pre-shared key	1640
Authenticating remote peers and clients	1642
Enabling VPN access for specific certificate holders	1642
Enabling VPN access by peer identifier.....	1644
Enabling VPN access with user accounts and pre-shared keys	1645
Defining IKE negotiation parameters	1646
Generating keys to authenticate an exchange	1647
Defining IKE negotiation parameters	1647
Using XAuth authentication	1650
Using the FortiGate unit as an XAuth server.....	1651
Using the FortiGate unit as an XAuth client.....	1651
Phase 2 parameters	1653
Basic phase 2 settings.....	1653
Advanced phase 2 settings	1653
P2 Proposals.....	1653
Replay detection	1653
Perfect forward secrecy (PFS)	1654
Keylife	1654
Auto-negotiate	1654

Autokey Keep Alive	1654
DHCP-IPsec	1654
Quick mode selectors	1655
Configure the phase 2 parameters	1655
Specifying the phase 2 parameters	1656
Defining VPN security policies	1659
Defining policy addresses.....	1659
Defining VPN security policies.....	1660
Defining an IPsec security policy for a policy-based VPN.....	1661
Defining security policies for a route-based VPN	1663
Gateway-to-gateway configurations	1665
Configuration overview	1665
General configuration steps.....	1667
Using auto-ipsec.....	1667
Configuring the two VPN peers	1667
Configuring Phase 1 and Phase 2 for both peers.....	1667
Creating security policies.....	1668
How to work with overlapping subnets	1672
Solution for route-based VPN.....	1673
Solution for policy-based VPN.....	1674
Testing	1676
Hub-and-spoke configurations	1679
Configuration overview	1679
Hub-and-spoke infrastructure requirements	1680
Spoke gateway addressing	1680
Authentication	1681
Configure the hub	1681
Define the hub-spoke VPNs.....	1681
Define the hub-spoke security policies.....	1682
Configuring communication between spokes (policy-based VPN)	1684
Configuring communication between spokes (route-based VPN).....	1684
Configure the spokes	1685
Configuring security policies for hub-to-spoke communication.....	1686
Configuring security policies for spoke-to-spoke communication	1687
Dynamic spokes configuration example.....	1689
Configure the hub (FortiGate_1).....	1689
Configure the spokes.....	1692
Dynamic DNS configuration	1695
Dynamic DNS over VPN concepts	1695
Dynamic DNS (DDNS).....	1695
Dynamic DNS over VPN	1696
Dynamic DNS topology	1697
Assumptions	1698

General configuration steps	1698
Configure the dynamically-addressed VPN peer.....	1699
Configuring branch_2 VPN tunnel settings.....	1699
Configuring branch_2 security policies.....	1701
Configure the fixed-address VPN peer	1704
Configuring branch_1 VPN tunnel settings.....	1704
Configuring branch_1 security policies.....	1705
Testing	1707
FortiClient dialup-client configurations.....	1709
Configuration overview	1709
Peer identification	1710
Automatic configuration of FortiClient dialup clients.....	1710
One button FortiGate - to - FortiClient Phase1 VPN	1711
Using virtual IP addresses	1711
FortiClient dialup-client infrastructure requirements	1713
FortiClient-to-FortiGate VPN configuration steps	1714
Configure the FortiGate unit	1714
Configuring FortiGate unit VPN settings.....	1714
Configuring the FortiGate unit as a VPN policy server	1717
Configuring DHCP services on a FortiGate interface	1717
Configure the FortiClient Endpoint Security application	1718
Configuring FortiClient	1718
Adding XAuth authentication	1718
FortiClient dialup-client configuration example.....	1719
Configuring FortiGate_1	1719
Configuring the FortiClient Endpoint Security application.....	1722
FortiGate dialup-client configurations	1724
Configuration overview	1724
FortiGate dialup-client infrastructure requirements	1726
FortiGate dialup-client configuration steps	1727
Configure the server to accept FortiGate dialup-client connections.....	1727
Configure the FortiGate dialup client	1729
Supporting IKE Mode config clients.....	1732
Automatic configuration overview	1732
IKE Mode Config overview	1732
Configuring IKE Mode Config	1732
Configuring an IKE Mode Config client.....	1733
Example: FortiGate unit as IKE Mode Config server	1735
Example: FortiGate unit as IKE Mode Config client	1736
Internet-browsing configuration	1737
Configuration overview	1737
Creating an Internet browsing security policy	1738

Routing all remote traffic through the VPN tunnel.....	1739
Configuring a FortiGate remote peer to support Internet browsing	1739
Configuring a FortiClient application to support Internet browsing.....	1740
Redundant VPN configurations.....	1741
Configuration overview	1741
General configuration steps.....	1742
Configure the VPN peers - route-based VPN.....	1742
Redundant route-based VPN configuration example.....	1745
Configuring FortiGate_1.....	1745
Configuring FortiGate_2.....	1752
Partially-redundant route-based VPN example.....	1758
Configuring FortiGate_1.....	1759
Configuring FortiGate_2.....	1762
Creating a backup IPsec interface.....	1765
Transparent mode VPNs.....	1766
Configuration overview	1766
Transparent VPN infrastructure requirements	1769
Configure the VPN peers	1770
IPv6 IPsec VPNs	1772
Overview of IPv6 IPsec support.....	1772
Certificates.....	1773
Configuring IPv6 IPsec VPNs	1773
Phase 1 configuration	1773
Phase 2 configuration	1773
Security policies.....	1774
Routing.....	1774
Site-to-site IPv6 over IPv6 VPN example	1774
Configure FortiGate A interfaces	1775
Configure FortiGate A IPsec settings	1775
Configure FortiGate A security policies	1776
Configure FortiGate A routing.....	1777
Configure FortiGate B.....	1777
Site-to-site IPv4 over IPv6 VPN example	1778
Configure FortiGate A interfaces	1779
Configure FortiGate A IPsec settings	1779
Configure FortiGate A security policies	1779
Configure FortiGate A routing.....	1780
Configure FortiGate B.....	1780
Site-to-site IPv6 over IPv4 VPN example	1782
Configure FortiGate A interfaces	1782
Configure FortiGate A IPsec settings	1782
Configure FortiGate A security policies	1783
Configure FortiGate A routing.....	1783

Configure FortiGate B	1784
L2TP and IPsec (Microsoft VPN)	1786
Overview	1786
Layer 2 Tunneling Protocol (L2TP).....	1786
Assumptions	1787
Configuring the FortiGate unit	1787
Configuring L2TP users and firewall user group.....	1787
Configuring L2TP	1788
Configuring IPsec.....	1789
Configuring security policies.....	1791
Configuring the Windows PC	1793
Troubleshooting.....	1794
Quick checks	1794
Mac OS X and L2TP	1794
Setting up logging.....	1794
Using the FortiGate unit debug commands.....	1795
GRE over IPsec (Cisco VPN).....	1798
Overview	1798
Configuring the FortiGate unit	1799
Enabling overlapping subnets.....	1799
Configuring the IPsec VPN	1799
Configuring the GRE tunnel	1801
Configuring security policies.....	1802
Configuring routing	1804
Configuring the Cisco router.....	1805
Troubleshooting.....	1805
Quick checks	1805
Setting up logging.....	1806
Protecting OSPF with IPsec	1808
Overview	1808
OSPF over IPsec configuration.....	1809
Configuring the IPsec VPN	1809
Configuring static routing	1810
Configuring OSPF	1810
Creating a redundant configuration.....	1814
Adding the second IPsec tunnel.....	1814
Adding the OSPF interface	1815
Hardware offloading and acceleration	1816
Overview	1816
IPsec session offloading requirements.....	1816
Packet offloading requirements.....	1817
IPsec encryption offloading	1817
HMAC check offloading.....	1817

IPsec offloading configuration examples.....	1817
Accelerated route-based VPN configuration	1818
Accelerated policy-based VPN configuration.....	1820
Monitoring and troubleshooting	1822
Monitoring VPN connections	1822
Monitoring connections to remote peers.....	1822
Monitoring dialup IPsec connections	1822
Testing VPN connections	1823
LAN interface connection	1823
Dialup connection	1824
Troubleshooting VPN connections	1824
Logging VPN events	1825
VPN troubleshooting tips.....	1826
The VPN proposal is not connecting	1826
Attempting hardware offloading beyond SHA1	1826
Check Phase 1 proposal settings	1826
Check your routing	1826
Try enabling XAuth.....	1826
General troubleshooting tips	1826
A word about NAT devices	1827
IPv6 for FortiOS 5.0	1828
IPv6 packet structure.....	1828
Jumbograms and jumbo payloads	1829
Fragmentation and reassembly	1829
Benefits of IPv6.....	1829
IPv6 Features	1830
IPv6 policies.....	1830
IPv6 policy routing	1831
IPv6 security policies	1831
IPv6 explicit web proxy.....	1832
VIP64.....	1833
VIP46.....	1836
IPv6 Network Address Translation	1838
NAT64 and DNS64 (DNS proxy)	1838
NAT66	1841
NAT64 and NAT66 session failover	1842
NAT46	1843
ICMPv6	1843
ICMPv6 Types and Codes	1844
IPv6 in dynamic routing	1847
Dual stack routing.....	1847
IPv6 tunnelling	1848
Tunnel configuration	1848

Tunnelling IPv6 through IPsec VPN	1849
SIP over IPv6	1849
New Fortinet FortiGate IPv6 MIB fields	1850
New OIDs	1851
EXAMPLE SNMP get/walk output	1852
IPv6 Per-IP traffic shaper	1852
DHCPv6	1852
DHCPv6 relay	1852
IPv6 forwarding—Policies, IPS, Application Control, flow-based antivirus, web filtering, and DLP	1853
FortiGate interfaces can get IPv6 addresses from an IPv6 DHCP server	1853
IPv6 Configuration	1855
IPv6 address groups	1855
IPv6 firewall addresses	1856
Scenario: Mail Server	1856
Scenario: First Floor Network	1856
ICMPv6	1857
IPv6 IPsec VPN	1858
Overview of IPv6 IPsec support	1858
Configuring IPv6 IPsec VPNs	1859
Site-to-site IPv6 over IPv6 VPN example	1860
Site-to-site IPv4 over IPv6 VPN example	1864
Site-to-site IPv6 over IPv4 VPN example	1868
BGP and IPv6	1871
RIPng — RIP and IPv6	1872
Network layout and assumptions	1872
Configuring the FortiGate units system information	1873
Configuring RIPng on FortiGate units	1875
Configuring other network devices	1876
Testing the configuration	1877
Debugging IPv6 on RIPng	1877
IPv6 IPS	1878
Blocking IPv6 packets by extension headers	1878
IPv6 Denial of Service policies	1878
Configure hosts in an SNMP v1/2c community to send queries or receive traps	1878
IPv6 PIM sparse mode multicast routing	1879
Chapter 12: Load Balancing for FortiOS 5.0	1880
Before you begin	1880
How this chapter is organized	1880
Configuring load balancing	1881
Load balancing overview	1881

Load balancing, UTM, authentication, and other FortiOS features	1882
Configuring load balancing virtual servers.....	1882
Load balancing methods	1885
Session persistence	1886
Real servers	1886
Health check monitoring	1888
Monitoring load balancing	1890
Load balancing get command	1891
Load balancing diagnose commands.....	1891
Logging Diagnostics	1892
Real server diagnostics.....	1893
Basic load balancing configuration example.....	1893
HTTP and HTTPS load balancing, multiplexing, and persistence	1897
HTTP and HTTPS multiplexing	1898
HTTP and HTTPS persistence	1898
HTTP host-based load balancing	1901
SSL/TLS load balancing	1902
SSL offloading	1903
IP, TCP, and UDP load balancing.....	1910
Load balancing configuration examples	1911
Example: HTTP load balancing to three real web servers.....	1911
Web-based manager configuration	1912
CLI configuration.....	1915
Example: Basic IP load balancing configuration	1917
Example: Adding a server load balance port forwarding virtual IP.....	1917
Example: Weighted load balancing configuration	1919
Web-based manager configuration	1919
CLI configuration.....	1922
Example: HTTP and HTTPS persistence configuration.....	1922
CLI configuration: adding persistence for a specific domain	1926
Chapter 13: Logging and Reporting	1928
.....	1929
Logging and reporting overview	1930
What is logging?	1930
How the FortiGate unit records log messages	1931
FortiOS features available for logging	1931
Traffic	1931
Other Traffic	1932
Event.....	1932
Traffic Shaping.....	1933
Data Leak Prevention.....	1933
NAC Quarantine	1933

Media Access Control (MAC) Address	1933
Application control	1934
Antivirus	1934
Web Filter.....	1935
IPS (attack).....	1935
Packet logs	1935
Email filter	1935
Archives (DLP).....	1936
Network scan.....	1936
Log messages.....	1936
Explanation of a debug log message	1938
Viewing log messages and archives.....	1939
Log files and types.....	1942
Log database and datasets	1943
How to view datasets	1944
How to create datasets (advanced).....	1944
Notifications about network activity	1945
How to configure email notifications	1946
Log devices.....	1946
FortiGate unit's system memory and hard disk.....	1947
FortiAnalyzer unit	1948
Syslog server	1948
WebTrends server.....	1948
How to choose a log device for your network topology.....	1949
How to create a backup solution for logging.....	1949
Reports	1950
What are FortiOS reports?.....	1951
The parts of a FortiOS report.....	1951
What you can do with the default FortiOS report	1951
How to modify the default FortiOS report.....	1952
How to create a FortiOS report.....	1952
Best Practices: Log management	1953
Logging and reporting for small networks	1955
Modifying default log device settings	1955
Modifying the FortiGate unit's system memory default settings	1955
Modifying the FortiGate unit's hard disk default settings	1956
Testing sending logs to the log device	1956
Configuring the backup solution.....	1957
Configuring logging to a FortiCloud server.....	1957
Configuring uploading logs to the FortiAnalyzer unit.....	1958
Testing uploading logs to a FortiAnalyzer unit.....	1958
Modifying the default FortiOS report	1959
Logging and reporting for large networks	1960
Modifying default log device settings	1960

Modifying multiple FortiGate units' system memory default settings.....	1960
Modifying multiple FortiGate units' hard disk default log settings	1961
Testing the modified log settings.....	1962
Configuring the backup solution.....	1962
Configuring logging to multiple FortiAnalyzer units	1963
Configuring logging to the FortiCloud server.....	1964
Modifying the default FortiOS report	1964
Creating datasets.....	1965
Creating charts for the datasets	1965
Uploading the corporate images	1966
Adding a new report cover and page	1966
Advanced logging.....	1967
Configuring logging to multiple Syslog servers	1967
Using Automatic Discovery to connect to a FortiAnalyzer unit	1968
Activating a FortiCloud account for logging purposes	1969
Viewing log storage space.....	1969
Customizing and filtering log messages.....	1970
Viewing logs from the CLI.....	1971
Configuring NAC quarantine logging.....	1971
Logging local-in policies	1972
Tracking specific search phrases in reports.....	1974
Creating a dataset containing attack name instead of attack ID.....	1976
Reverting modified report settings to default settings.....	1976
Customizing FortiOS reports with CLI	1976
Configuring a style	1976
Configuring a theme	1980
Configuring charts	1982
Adding a chart.....	1984
Troubleshooting and logging	1985
Using log messages to help in troubleshooting issues	1985
Using IPS packet logging in diagnostics	1985
Using HA log messages to determine system status	1986
Connection issues between FortiGate unit and logging devices	1986
Unable to connect to a supported log device	1986
FortiGate unit has stopped logging	1986
Log database issues.....	1986
SQL statement syntax errors	1987
Connection problems	1987
SQL database errors.....	1987
Logging daemon (Miglogd).....	1988
Appendix: FortiGate report charts	1989
Traffic charts.....	1989

Web filter charts.....	1990
IPS (or attack) charts	1991
Antivirus charts	1992
Email filter charts	1992
VPN charts.....	1993
Chapter 14: Managing Devices for FortiOS 5.0	1995
Managing “bring your own device”	1996
Device monitoring.....	1996
Device Groups	1997
Creating a custom device group.....	1998
Controlling access with a MAC Address Access Control List.....	1999
Device policies.....	1999
Creating device policies.....	2001
Endpoint Protection	2003
Endpoint Protection overview.....	2003
User experience.....	2003
FortiGate endpoint registration limits	2004
Configuration overview	2005
Changing the FortiClient installer download location	2005
Creating a FortiClient profile.....	2006
Enabling Endpoint Protection in security policies	2008
Configuring endpoint registration over a VPN.....	2009
Endpoint registration on an IPsec VPN.....	2009
Endpoint registration on the SSL VPN.....	2009
Synchronizing endpoint registrations	2009
Monitoring endpoints.....	2010
Modifying the Endpoint Protection replacement messages.....	2010
Vulnerability Scan.....	2011
Configuring vulnerability scans.....	2011
Running a vulnerability scan and viewing scan results	2013
Requirements for authenticated scanning and ports scanned.....	2013
Microsoft Windows hosts - domain scanning	2014
Microsoft Windows hosts - local (non-domain) scanning	2015
Windows firewall settings	2015
Unix hosts	2015
Chapter 15: Unified Threat Management for FortiOS 5.0.....	2018
Security Profiles overview	2019
Traffic inspection	2019
IPS signatures.....	2019
Suspicious traffic attributes.....	2020

Application control.....	2020
Content inspection and filtering.....	2020
AntiVirus.....	2021
FortiGuard Web Filtering.....	2021
Email filter	2021
DLP	2022
Security Profiles components.....	2022
AntiVirus.....	2022
Intrusion Protection System (IPS).....	2022
Web filtering.....	2022
Email filtering.....	2023
Data Leak Prevention (DLP).....	2023
Application Control	2023
ICAP	2023
Security Profiles/lists/sensors.....	2023
Client Reputation.....	2025
Summary of the Client Reputation features.....	2025
Applying client reputation monitoring to your network.....	2026
Viewing client reputation results.....	2026
Changing the client reputation reporting window and database size	2027
Client reputation data update and maintenance intervals	2027
Setting the client reputation profile/definition.....	2028
Expanding client reputation to include more types of behavior	2029
Client reputation execute commands.....	2031
Client reputation diagnose commands.....	2031
AntiVirus	2032
Antivirus concepts	2032
How antivirus scanning works	2032
Antivirus scanning order	2033
Antivirus databases.....	2036
Antivirus techniques.....	2036
FortiGuard Antivirus	2036
Enable antivirus scanning	2037
Antivirus Profiles	2037
Changing the default antivirus database	2038
Configuring the scan buffer size	2038
Configuring archive scan depth.....	2039
Configuring a maximum allowed file size	2039
Configuring client comforting	2040
Grayware scanning	2041
Windows file sharing (CIFS) flow-based antivirus scanning	2041
Advanced Persistent Threat (APT) protection.....	2043
Botnet and phishing protection	2043

FortiGuard Sandbox (in the cloud sandboxing, zero day threat analysis and submission).....	2043
Testing your antivirus configuration.....	2046
Antivirus examples.....	2046
Configuring simple antivirus protection	2046
Protecting your network against malicious email attachments	2048
Email filter	2049
Email filter concepts	2049
Email filter techniques.....	2049
Order of spam filtering	2051
Enable email filtering.....	2052
Configure email traffic types to inspect.....	2052
Configure the spam action	2052
Configure the tag location	2053
Configure the tag format.....	2053
Configure FortiGuard email filters.....	2054
Configure local email filters	2055
Enabling IP address and email address black/white list checking.....	2055
Enabling HELO DNS lookup	2057
Enabling return email DNS checking	2057
Enabling banned word checking	2058
How content is evaluated	2058
Email filter examples.....	2059
Configuring simple antispam protection.....	2059
Blocking email from a user	2060
Intrusion protection.....	2062
IPS concepts	2062
Anomaly-based defense	2062
Signature-based defense.....	2062
Enable IPS scanning.....	2064
General configuration steps.....	2064
Creating an IPS sensor	2064
Creating an IPS filter	2064
Updating predefined IPS signatures.....	2068
Viewing and searching predefined IPS signatures	2068
IPS processing in an HA cluster	2068
Active-passive.....	2068
Active-active	2069
Configure IPS options.....	2069
Hardware Acceleration	2069
Extended IPS Database.....	2069
Configuring the IPS engine algorithm	2070
Configuring the IPS engine-count.....	2070

Configuring fail-open	2070
Configuring the session count accuracy	2070
Configuring the IPS buffer size	2071
Configuring protocol decoders	2071
Configuring security processing modules	2071
IPS signature rate count threshold	2072
Enable IPS packet logging.....	2072
IPS examples.....	2073
Configuring basic IPS protection.....	2073
Using IPS to protect your web server.....	2074
Create and test a packet logging IPS sensor	2076
Configuring a Fortinet Security Processing module	2077
IPS Sensor	2078
Custom Application & IPS Signatures	2080
Creating a custom IPS signature.....	2080
Custom signature syntax.....	2080
Custom signature keywords.....	2081
Information keywords	2081
Session keywords.....	2082
Content keywords.....	2082
IP header keywords	2086
TCP header keywords	2087
UDP header keywords.....	2089
ICMP keywords.....	2089
Other keywords.....	2090
Creating a custom signature to block access to example.com	2091
Creating a custom signature to block the SMTP “vrfy” command.....	2093
Web filter	2095
Web filter concepts.....	2095
Different ways of controlling access.....	2097
Order of web filtering	2097
Inspections Modes	2097
Proxy.....	2097
Flow-based	2097
DNS.....	2098
FortiGuard Web Filtering Service.....	2098
FortiGuard Web Filter and your FortiGate unit	2098
Enabling FortiGuard Web Filter.....	2100
General configuration steps.....	2100
Configuring FortiGuard Web Filter settings	2100
To configure the FortiGuard Web Filter categories	2101
Configuring FortiGuard Web Filter usage quotas	2101
Overriding FortiGuard website categorization.....	2102
The different methods of override.....	2103

Using Alternate Categories	2103
Using Alternate Profiles	2104
SafeSearch	2108
YouTube Education Filter	2109
Enabling YouTube Education Filter in CLI	2109
Deep Scanning Restrictions	2109
Enable HTTPS URL Scan Only	2109
Categories Exempt from Deep Scanning	2109
Web Site Filter	2110
Web Site Filter actions	2111
Status	2112
Configuring a Web Site Filter	2113
Configuring a URL filter list	2113
Web content filter	2113
General configuration steps	2114
Creating a web filter content list	2114
How content is evaluated	2114
Enabling the web content filter and setting the content threshold	2115
Advanced web filter configurations	2116
Allow websites when a rating error occurs	2116
ActiveX filter	2116
Block HTTP redirects by rating	2116
Block Invalid URLs	2116
Cookie filter	2117
Provide Details for Blocked HTTP 4xx and 5xx Errors	2117
HTTP POST action	2117
Java applet filter	2117
Rate Images by URL	2117
Rate URLs by Domain and IP Address	2118
Web resume download block	2118
Working with the Interface	2118
Profile page	2118
New Web Filter Profile page	2119
Profile	2119
URL Filter	2123
Web filtering example	2127
School district	2127
Data leak prevention	2130
Data leak prevention concepts	2130
DLP sensor	2130
DLP filter	2130
Fingerprint	2131
File filter	2131
File size	2131

Regular expression	2131
Watermark.....	2131
Using the FortiExplorer Watermark tool	2132
Installation of the watermark utility on Linux	2133
Enable data leak prevention	2134
General configuration steps.....	2134
Creating a DLP sensor.....	2135
Adding filters to a DLP sensor	2135
DLP document fingerprinting.....	2139
Fingerprinted Documents	2139
File filter	2140
General configuration steps.....	2141
Creating a file filter list	2141
Creating a file pattern	2142
Creating a file type.....	2142
Preconfigured sensors.....	2143
DLP archiving.....	2144
DLP examples.....	2145
Blocking content with credit card numbers.....	2145
Blocking emails larger than 15 MB and logging emails from 5 MB to 15 MB.....	2146
Selective blocking based on a finger print.....	2147
Create policies and attach DLP sensors.....	2150
Application control	2152
Application control concepts.....	2152
Application considerations	2153
Automatically allowing basic applications	2153
IM applications.....	2154
Skype	2154
Application traffic shaping	2154
Direction of traffic shaping.....	2154
Shaper re-use	2155
Application control monitor	2155
Application Control monitor.....	2156
Enable application control	2156
General configuration steps.....	2156
Creating an application sensor	2156
Adding applications to an application sensor.....	2157
Viewing and searching the application list.....	2160
Creating a New Custom Application Signature	2160
Enabling application traffic shaping.....	2161
Application control examples	2161
Blocking all instant messaging	2161

Allowing only software updates	2162
ICAP	2164
The Protocol	2164
Offloading using ICAP	2165
Configuration Settings	2165
Servers	2165
Profiles	2166
Example ICAP sequence	2166
Example Scenerio	2167
Other Security Profiles considerations	2169
Profile Groups	2169
Creating a new group	2170
Security Profiles and Virtual domains (VDOMs).....	2172
Conserve mode	2172
The AV proxy.....	2172
Entering and exiting conserve mode	2172
Conserve mode effects.....	2172
Configuring the av-failopen command	2173
SSL content scanning and inspection	2173
Setting up certificates to avoid client warnings	2174
SSL content scanning and inspection settings	2175
Exeptions	2178
Monitoring Security Profiles activity	2178
Configuring packet logging options.....	2180
Using wildcards and Perl regular expressions.....	2181
Monitor interface reference.....	2183
AV Monitor	2184
Intrusion Monitor.....	2184
Web Monitor	2185
Email Monitor.....	2186
Archive & Data Leak Monitor	2186
Application Monitor.....	2187
FortiGuard Quota	2187
Endpoint Monitor	2188
Chapter 16: SSL VPN for FortiOS 5.0.....	2189
Introduction to SSL VPN	2190
SSL VPN modes of operation.....	2191
Web-only mode	2191
Tunnel mode	2191
Port forwarding mode	2192
Application support.....	2193
SSL VPN and IPv6	2193

Traveling and security	2193
Host check	2193
Cache cleaning	2194
Basic Configuration.....	2195
User accounts and groups	2195
Authentication	2196
MAC host check	2196
IP addresses for users	2196
Authentication of remote users.....	2197
Configuring SSL VPN web portals.....	2199
SSL connection configuration.....	2200
Portal configuration.....	2200
Personal bookmarks	2203
Custom login screen	2203
Tunnel mode and split tunneling.....	2203
The Connection tool widget.....	2203
Configuring security policies	2204
Firewall addresses	2204
Create an SSL VPN security policy.....	2204
Create a tunnel mode security policy	2206
Split tunnel Internet browsing policy	2208
Enabling a connection to an IPsec VPN	2209
Additional configuration options.....	2210
Routing in tunnel mode.....	2211
Changing the port number for web portal connections	2211
SSL offloading	2211
Customizing the web portal login page	2212
Host check.....	2212
Creating a custom host check list	2213
Windows OS check.....	2213
Configuring cache cleaning	2214
Configuring virtual desktop.....	2214
Configuring client OS Check	2215
Adding WINS and DNS services for clients.....	2216
Setting the idle timeout setting	2216
SSL VPN logs.....	2216
Monitoring active SSL VPN sessions.....	2217
Troubleshooting.....	2217
The SSL VPN client.....	2219
FortiClient	2219
Tunnel mode client configuration	2220
Setup examples	2221
Secure internet browsing.....	2221
Creating an SSL VPN IP pool and SSL VPN web portal.....	2221

Creating the SSL VPN user and user group	2221
Creating a static route for the remote SSL VPN user	2222
Creating security policies.....	2222
Results	2223
Split Tunnel.....	2223
Creating a firewall address for the head office server	2224
Results	2226
Multiple user groups with different access permissions example.....	2226
General configuration steps.....	2227
Creating the firewall addresses	2227
Creating the web portals.....	2228
Creating the user accounts and user groups	2229
Creating the security policies.....	2229
Create the static route to tunnel mode clients.....	2231

Chapter 17: Traffic Shaping for FortiOS 5.0..... 2233

The purpose of traffic shaping..... 2234

Quality of Service.....	2234
Traffic policing	2235
Bandwidth guarantee, limit, and priority interactions	2236
FortiGate traffic	2236
Through traffic.....	2237
Important considerations.....	2241

Traffic shaping methods..... 2243

Traffic shaping options	2243
Shared policy shaping	2244
Per policy	2244
All policies.....	2244
Maximum and guaranteed bandwidth.....	2244
Traffic priority.....	2244
VLAN, VDOM and virtual interfaces.....	2245
Shared traffic shaper configuration settings.....	2245
Per-IP shaping	2247
Per-IP traffic shaping configuration settings	2247
Adding Per-IP traffic shapers to a security policy	2248
Application control shaping	2248
Example	2248
Enabling in the security policy	2249
Reverse direction traffic shaping	2249
Setting the reverse direction only	2250
Application control shaper	2250
Type of Service priority	2250
TOS in FortiOS.....	2251

Differentiated Services.....	2251
DSCP examples.....	2253
Tos and DSCP mapping.....	2257
Traffic Shaper Monitor.....	2258
Examples.....	2259
QoS using priority from security policies.....	2259
Sample configuration.....	2260
QoS using priority from ToS or differentiated services.....	2261
Sample configuration.....	2262
Example setup for VoIP.....	2263
Creating the traffic shapers.....	2263
Creating security policies.....	2265
Troubleshooting traffic shaping.....	2266
Interface diagnosis.....	2266
Shaper diagnose commands.....	2266
TOS command.....	2266
Shared shaper.....	2267
Per-IP shaper.....	2267
Packet loss with statistics on shapers.....	2267
Packet lost with the debug flow.....	2268
Session list details with dual traffic shaper.....	2268
Additional Information.....	2269

Chapter 18: Troubleshooting..... 2270

Life of a Packet.....	2271
Stateful inspection.....	2271
Connections over connectionless.....	2272
What is a session?.....	2272
Differences between connections and sessions.....	2272
Flow inspection.....	2273
Proxy inspection.....	2274
Comparison of inspection layers.....	2274
FortiOS functions and security layers.....	2275
Packet flow.....	2275
Packet inspection (Ingress).....	2276
Interface.....	2277
DoS sensor.....	2277
IP integrity header checking.....	2277
IPsec.....	2277
Destination NAT (DNAT).....	2277
Routing.....	2277
Policy lookup.....	2277
Session tracking.....	2278

User authentication.....	2278
Management traffic.....	2278
SSL VPN traffic.....	2278
ICAP traffic.....	2278
Session helpers.....	2278
Flow-based inspection engine.....	2279
Proxy-based inspection engine.....	2279
IPsec.....	2279
Source NAT (SNAT).....	2279
Routing.....	2279
Egress.....	2279
Example 1: client/server connection.....	2279
Example 2: Routing table update.....	2281
Example 3: Dialup IPsec VPN with application control.....	2282
Verifying FortiGate admin access security.....	2285
Install the FortiGate unit in a physically secure location.....	2285
Add new administrator accounts.....	2285
Change the admin account name and limit access to this account.....	2286
Only allow administrative access to the external interface when needed.....	2286
When enabling remote access, configure Trusted Hosts and Two-factor Authentication.....	2287
Configuring Trusted Hosts.....	2287
Configuring Two-factor Authentication.....	2287
Change the default administrative port to a non-standard port.....	2288
Enable Password Policy.....	2288
Maintain short login timeouts.....	2288
Modify administrator account Lockout Duration and Threshold values.....	2288
Administrator account Lockout Duration.....	2289
Administrator account Lockout Threshold.....	2289
Disable auto installation via USB.....	2289
Auditing and Logging.....	2289
Troubleshooting resources.....	2290
Technical Documentation.....	2290
Fortinet Video Library.....	2290
Release Notes.....	2290
Knowledge Base.....	2290
Fortinet Technical Discussion Forums.....	2290
Fortinet Training Services Online Campus.....	2291
Fortinet Customer Support.....	2291
Troubleshooting tools.....	2292
FortiOS diagnostics.....	2292
Check date and time.....	2292

Resource usage	2293
Proxy operation.....	2295
Hardware NIC	2298
Traffic trace	2300
Session table	2300
Firewall session setup rate	2304
Finding object dependencies.....	2305
Flow trace	2306
Packet sniffing and packet capture	2309
FA2 and NP2 based interfaces	2313
Debug command	2314
The execute tac report command.....	2316
Other commands	2316
FortiOS ports	2317
FortiAnalyzer/FortiManager ports	2319
FortiGuard troubleshooting.....	2319
Troubleshooting process for FortiGuard updates	2319
FortiGuard server settings	2320
FortiGuard URL rating.....	2320
.....	2320
.....	2320
Troubleshooting methodologies	2321
Establish a baseline	2321
Define the problem	2322
Gathering Facts	2323
Create a troubleshooting plan	2323
Providing Supporting Elements	2324
Obtain any required additional equipment	2324
Ensure you have administrator level access to required equipment.....	2324
Contact Fortinet customer support for assistance.....	2324
Technical Support Organization Overview	2325
Fortinet Global Customer Services Organization	2325
Creating an account	2326
Registering a device	2326
Reporting problems	2327
Logging online tickets.....	2327
Following up on online tickets	2328
Telephoning a technical support center	2329
Assisting technical support.....	2329
Support priority levels.....	2329
Priority 1	2329
Priority 2.....	2329
Priority 3.....	2330

Priority 4.....	2330
Return material authorization process.....	2330
Chapter 19: Virtual Domains	2331
Virtual Domains	2332
Benefits of Virtual Domains	2332
Improving Transparent mode configuration	2332
Easier administration	2332
Continued security	2333
Savings in physical space and power.....	2333
More flexible MSSP configurations	2334
Enabling and accessing Virtual Domains.....	2334
Enabling Virtual Domains	2334
Viewing the VDOM list	2337
Global and per-VDOM settings.....	2338
Resource settings	2347
Virtual Domain Licensing	2351
Logging in to VDOMs.....	2352
Configuring Virtual Domains	2354
Creating a Virtual Domain	2354
Disabling a Virtual Domain	2355
Deleting a VDOM	2356
Removing references to a VDOM	2356
Administrators in Virtual Domains.....	2357
Virtual Domains in NAT/Route mode.....	2361
Virtual domains in NAT/Route mode	2361
Changing the management virtual domain.....	2361
Configuring interfaces in a NAT/Route VDOM.....	2362
Configuring VDOM routing.....	2365
Configuring security policies for NAT/Route VDOMs	2367
Configuring security profiles for NAT/Route VDOMs.....	2368
Configuring VPNs for a VDOM.....	2368
Example NAT/Route VDOM configuration.....	2368
Network topology and assumptions	2369
General configuration steps.....	2370
Creating the VDOMs	2370
Configuring the FortiGate interfaces.....	2371
Configuring the vdomA VDOM	2373
Configuring the vdomB VDOM	2376
Testing the configuration	2379
Virtual Domains in Transparent mode.....	2380
Transparent operation mode	2380
Broadcast domains.....	2380
Forwarding domains	2380

Spanning Tree Protocol	2381
Differences between NAT/Route and Transparent mode	2382
Operation mode differences in VDOMs	2382
Configuring VDOMs in Transparent mode	2383
Switching to Transparent mode	2383
Adding VLAN subinterfaces	2384
Creating security policies	2384
Example of VDOMs in Transparent mode	2384
Network topology and assumptions	2385
General configuration steps	2386
Configuring common items	2386
Creating virtual domains	2387
Configuring the Company_A VDOM	2387
Configuring the Company_B VDOM	2392
Configuring the VLAN switch and router	2397
Testing the configuration	2399
Inter-VDOM routing	2400
Benefits of inter-VDOM routing	2400
Continued support for secure firewall policies	2400
Configuration flexibility	2400
Getting started with VDOM links	2401
Viewing VDOM links	2401
Creating VDOM links	2403
Deleting VDOM links	2405
NAT to Transparent VDOM links	2405
Inter-VDOM configurations	2406
Standalone VDOM configuration	2407
Independent VDOMs configuration	2408
Management VDOM configuration	2409
Meshed VDOM configuration	2410
Dynamic routing over inter-VDOM links	2410
HA virtual clusters and VDOM links	2411
Example of inter-VDOM routing	2413
Network topology and assumptions	2413
General configuration steps	2414
Creating the VDOMs	2414
Configuring the physical interfaces	2415
Configuring the VDOM links	2417
Configuring the firewall and Security Profile settings	2419
Testing the configuration	2438
Troubleshooting Virtual Domains	2440
VDOM admin having problems gaining access	2440
Confirm the admin's VDOM	2440
Confirm the VDOM's interfaces	2440

Confirm the VDOMs admin access.....	2440
FortiGate unit running very slowly	2440
Too many VDOMs.....	2441
One or more VDOMs are consuming all the resources	2441
Too many Security Features in use	2441
General VDOM tips and troubleshooting.....	2441
Perform a sniffer trace	2441
Debugging the packet flow	2443
Chapter 20: Virtual FortiGate Units for FortiOS 5.0.....	2445
FortiGate VM Overview	2446
FortiGate VM models and licensing.....	2446
FortiGate VM evaluation license	2446
Registering FortiGate VM with Customer Service & Support.....	2447
Downloading the FortiGate VM deployment package.....	2447
Deployment package contents.....	2448
Citrix XenServer	2448
OpenXEN	2448
Microsoft Hyper-V.....	2449
KVM	2449
VMware ESX/ESXi.....	2449
Deploying the FortiGate VM appliance.....	2449
Deployment example: VMware	2451
Open the FortiGate VM OVF file with the vSphere client.....	2451
Configure FortiGate VM hardware settings	2455
Transparent mode configuration.....	2455
Power on your FortiGate VM	2456
Deployment example: MS Hyper-V.....	2457
Create the FortiGate VM virtual machine.....	2457
Configure FortiGate VM hardware settings	2462
FortiGate VM virtual processors	2463
FortiGate VM network adapters.....	2463
FortiGate VM virtual hard disk	2464
Start the FortiGate VM.....	2469
Deployment example: KVM	2470
Create the FortiGate VM virtual machine.....	2470
Configure FortiGate VM hardware settings	2472
Start the FortiGate VM.....	2472
Deployment example: OpenXen.....	2473
Create the FortiGate VM virtual machine (VMM)	2473
Deployment example: Citrix XenServer.....	2477
Create the FortiGate VM virtual machine (XenCenter).....	2477

Configure virtual hardware.....	2479
Configuring number of CPUs and memory size	2479
Configuring disk storage.....	2481
FortiGate VM Initial Configuration	2482
Set FortiGate VM port1 IP address.....	2482
Connect to the FortiGate VM Web-based Manager.....	2484
Upload the FortiGate VM license file	2484
Validate the FortiGate VM license with FortiManager	2485
Configure your FortiGate VM.....	2487

Chapter 21: VoIP Solutions: SIP for FortiOS 5.0..... 2488

FortiGate VoIP solutions: SIP	2489
SIP overview	2489
Common SIP VoIP configurations	2490
Peer to peer configuration	2490
SIP proxy server configuration.....	2491
SIP redirect server configuration	2491
SIP registrar configuration	2492
SIP with a FortiGate unit.....	2493
SIP messages and media protocols.....	2495
SIP request messages	2497
SIP response messages	2498
SIP message start line	2500
SIP headers.....	2500
The SIP message body and SDP session profiles.....	2502
Example SIP messages	2504
The SIP session helper	2505
SIP session helper configuration overview	2506
Configuration example: SIP session helper in Transparent Mode.....	2508
SIP session helper diagnose commands.....	2511
The SIP ALG	2512
SIP ALG configuration overview	2514
Conflicts between the SIP ALG and the session helper	2517
Stateful SIP tracking, call termination, and session inactivity timeout	2518
SIP and RTP/RTCP	2520
How the SIP ALG creates RTP pinholes.....	2520
Configuration example: SIP in Transparent Mode.....	2522
RTP enable/disable (RTP bypass)	2525
Opening and closing SIP register, contact, via and record-route pinholes.	2526
Accepting SIP register responses.....	2527
How the SIP ALG performs NAT	2527
Source address translation	2528
Destination address translation	2528
Call Re-invite messages	2529

How the SIP ALG translates IP addresses in SIP headers	2529
How the SIP ALG translates IP addresses in the SIP body	2531
SIP NAT scenario: source address translation (source NAT)	2532
SIP NAT scenario: destination address translation (destination NAT)	2534
SIP NAT configuration example: source address translation (source NAT)	2536
SIP NAT configuration example: destination address translation (destination NAT)	2539
Additional SIP NAT scenarios	2542
NAT with IP address conservation	2544
Controlling how the SIP ALG NATs SIP contact header line addresses	2545
Controlling NAT for addresses in SDP lines	2546
Translating SIP session destination ports	2546
Translating SIP sessions to multiple destination ports	2548
Adding the original IP address and port to the SIP message header after NAT ..	2549
Enhancing SIP pinhole security	2549
Hosted NAT traversal	2552
Configuration example: Hosted NAT traversal for calls between SIP Phone A and SIP Phone B	2553
Hosted NAT traversal for calls between SIP Phone A and SIP Phone C	2557
Restricting the RTP source IP	2557
SIP over IPv6	2558
Deep SIP message inspection	2558
Actions taken when a malformed message line is found	2559
Logging and statistics	2560
Deep SIP message inspection best practices	2560
Configuring deep SIP message inspection	2560
Blocking SIP request messages	2563
SIP rate limiting	2565
Limiting the number of SIP dialogs accepted by a security policy	2566
SIP logging and DLP archiving	2567
Inspecting SIP over SSL/TLS (secure SIP)	2567
Adding the SIP server and client certificates	2568
Adding SIP over SSL/TLS support to a VoIP profile	2569
SIP and HA: session failover and geographic redundancy	2569
SIP geographic redundancy	2570
Support for RFC 2543-compliant branch parameters	2571
SIP and IPS	2572
SIP debugging	2572
SIP debug log format	2572
SIP-proxy filter per VDOM	2573
SIP-proxy filter command	2574
SIP debug log filtering	2574
SIP debug setting	2575

Display SIP rate-limit data	2575
-----------------------------------	------

Chapter 22: WAN Optimization, Web Cache, Explicit Proxy, and WCCP for FortiOS 5.0.. 2577

Before you begin.....	2577
FortiGate models that support WAN optimization.....	2578
How this chapter is organized	2578
Example network topologies.....	2580
WAN optimization topologies	2580
Basic WAN optimization topologies	2581
Out-of-path topology	2581
Topology for multiple networks	2583
WAN optimization with web caching	2583
WAN optimization and web caching with FortiClient peers.....	2584
Explicit Web proxy topologies	2585
Explicit FTP proxy topologies	2586
Web caching topologies	2587
WCCP topologies	2588
Configuring WAN optimization.....	2590
Client/server architecture.....	2590
WAN optimization peers	2592
Manual (peer-to-peer) and active-passive WAN optimization	2592
Manual (peer to peer) configurations	2592
Active-passive configurations.....	2594
WAN optimization profiles	2595
Processing non-HTTP sessions accepted by a WAN optimization profile with HTTP optimization	2597
Processing unknown HTTP sessions	2597
Protocol optimization.....	2598
Protocol optimization and MAPI	2598
Byte caching	2598
Dynamic data chunking for byte caching	2599
WAN optimization transparent mode	2599
FortiClient WAN optimization.....	2599
Operating modes and VDOMs.....	2600
WAN optimization tunnels	2600
Tunnel sharing.....	2601
WAN optimization and user and device identity policies, load balancing and traffic shaping	2601
Traffic shaping	2602
WAN optimization and HA	2602
WAN optimization, web caching and memory usage.....	2602

Monitoring WAN optimization performance	2603
Traffic Summary.....	2603
Bandwidth Optimization	2604
WAN optimization configuration summary	2604
client-side configuration summary.....	2605
server-side configuration summary	2607
Best practices	2609
Peers and authentication groups.....	2610
Basic WAN optimization peer requirements	2610
Accepting any peers	2610
How FortiGate units process tunnel requests for peer authentication	2611
Configuring peers	2611
Configuring authentication groups	2612
Secure tunneling	2615
Monitoring WAN optimization peer performance	2615
Configuration examples.....	2616
Example: Basic manual (peer-to-peer) WAN optimization configuration	2616
Network topology and assumptions	2616
General configuration steps.....	2617
Configuring basic peer-to-peer WAN optimization - web-based manager .	2617
Configuring basic peer-to-peer WAN optimization - CLI.....	2620
Testing and troubleshooting the configuration	2622
Example: Active-passive WAN optimization.....	2625
Network topology and assumptions.....	2625
General configuration steps.....	2626
Configuring basic active-passive WAN optimization - web-based manager	2626
Configuring basic active-passive WAN optimization - CLI	2630
Testing and troubleshooting the configuration.....	2632
Example: Adding secure tunneling to an active-passive WAN optimization configuration.....	2634
Network topology and assumptions	2634
General configuration steps.....	2635
Configuring WAN optimization with secure tunneling - web-based manager.....	2635
Configuring WAN optimization with secure tunneling - CLI	2639
Web caching and SSL offloading.....	2643
Turning on web caching for HTTP and HTTPS traffic	2644
Turning on web caching and SSL offloading for HTTPS traffic.....	2644
Full mode SSL server configuration.....	2645
Half mode SSL server configuration	2646
Changing the ports on which to look for HTTP and HTTPS traffic to cache.....	2647
Web caching and HA	2647

Web caching and memory usage	2648
Changing web cache settings	2648
Forwarding URLs to forwarding servers and exempting web sites from web caching 2651	
Forwarding URLs and URL patterns to forwarding servers.....	2651
Exempting web sites from web caching.....	2651
Monitoring Web caching performance	2652
Example: Web caching of HTTP and HTTPS Internet content for users on an internal network.....	2652
Example: reverse proxy web caching and SSL offloading for an Internet web server using a static one-to-one virtual IP	2655
Network topology and assumptions.....	2655
General configuration steps.....	2657
Configuration steps - web-based manager.....	2657
Configuration steps - CLI.....	2659
FortiClient WAN optimization.....	2661
FortiClient WAN optimization over SSL VPN configuration example	2661
The FortiGate explicit web proxy	2665
Explicit web proxy configuration overview	2667
General configuration steps.....	2667
Proxy auto-config (PAC) configuration	2671
Unknown HTTP version	2671
Authentication realm	2672
Other explicit web proxy options.....	2672
Restricting the IP address of the explicit web proxy	2672
Restricting the outgoing source IP address of the explicit web proxy	2672
IPv6 Explicit web proxy	2673
Restricting the IP address of the explicit IPv6 web proxy	2674
Restricting the outgoing source IP address of the IPv6 explicit web proxy	2674
Proxy chaining (web proxy forwarding servers)	2674
Adding a web proxy forwarding server.....	2675
Web proxy forwarding server monitoring and health checking	2675
Adding proxy chaining to an explicit web proxy security policy.....	2676
Explicit web proxy authentication.....	2677
IP-Based authentication	2677
Per session authentication.....	2678
Security profiles, client reputation, device identification, and the explicit web proxy 2680	
Web Proxy firewall services and service groups	2681
Example: users on an internal network browsing the Internet through the explicit web proxy with web caching, RADIUS authentication, web filtering and virus scanning 2681	
General configuration steps.....	2682
Configuring the explicit web proxy - web-based manager.....	2682

Configuring the explicit web proxy - CLI	2684
Testing and troubleshooting the configuration	2685
Explicit proxy sessions and user limits	2686
The FortiGate explicit FTP proxy	2689
How to use the explicit FTP proxy to connect to an FTP server	2690
Explicit FTP proxy configuration overview.....	2692
General configuration steps.....	2692
Restricting the IP address of the explicit FTP proxy	2696
Restricting the outgoing source IP address of the explicit FTP proxy	2697
Security profiles, client reputation, device identification, and the explicit FTP proxy	2697
Explicit FTP proxy sessions and protocol options	2697
Explicit FTP proxy sessions and antivirus	2697
Example: users on an internal network connecting to FTP servers on the Internet	
through the explicit FTP with RADIUS authentication and virus scanning	2698
General configuration steps.....	2698
Configuring the explicit FTP proxy - web-based manager	2698
Configuring the explicit FTP proxy - CLI.....	2700
Testing and troubleshooting the configuration	2702
Explicit FTP proxy sessions and user limits	2702
FortiGate WCCP	2703
WCCP service groups, service numbers, service IDs and well known services	2704
Example WCCP server and client configuration for caching HTTP sessions	
(service ID = 0).....	2704
Example WCCP server and client configuration for caching HTTPS sessions	2705
Example WCCP server and client configuration for caching HTTP and HTTPS	
sessions.....	2706
Other WCCP service group options	2706
WCCP configuration overview.....	2707
Example: caching HTTP sessions on port 80 using WCCP	2708
Configuring the WCCP server (WCCP_srv)	2708
Configuring the WCCP client (WCCP_client).....	2710
Example: caching HTTP sessions on port 80 and HTTPS sessions on port 443 using	
WCCP.....	2711
Configuring the WCCP server (WCCP_srv)	2711
Configuring the WCCP client (WCCP_client).....	2712
WCCP packet flow.....	2713
Configuring the forward and return methods and adding authentication	2713
WCCP Messages.....	2714
Troubleshooting WCCP	2714
Real time debugging	2714
Application debugging	2714

Storage	2716
Formatting the hard disk.....	2716
Configuring WAN optimization and Web cache storage	2717
Changing the amount of space allocated for WAN optimization and Web cache storage.....	2717
Adjusting the relative amount of disk space available for byte caching and web caching	2717
Diagnose commands	2719
get test {wa_cs wa_dbd wad wad_diskd wccpd} <test_level>	2719
Examples	2719
diagnose wad	2722
Example: diagnose wad tunnel list	2722
Example: diagnose wad webcache list.....	2724
diagnose wacs.....	2726
diagnose wadbd	2726
diagnose debug application {wa_cs wa_dbd wad wad_diskd wccpd} [<debug_level>]	2726
Index	2728

Change Log

Date	Change Description
June 15, 2015	New section: “Exception to policy order (VIPs)” on page 956. Also added some info about best practices with VIPs to “Virtual IP Addresses (VIPs)” on page 915 and deny policies and VIPs to “Deny Policies” on page 969.
January 24, 2015	New version collecting Misc updates and fixes.
March 28, 2014	New chapters: “Certifications and Compliances” on page 768, “Hardware Acceleration” on page 1067, “IPv6 for FortiOS 5.0” on page 1828, and “Virtual FortiGate Units for FortiOS 5.0” on page 2445. Updated for FortiOS 5.0 Patch 6. Chapters re-organized to be in alphabetical order by title
2013-09-27	New chapter: “Troubleshooting” on page 2270. Updated for FortiOS 5.0 Patch 4
2013-07-11	New chapter: “FortiOS Carrier” on page 631. Updated for FortiOS 5.0 Patch 3
2013-05-27	New FortiOS 5.0 release.

Introduction

This FortiOS™ Handbook is the definitive guide to configuring and operating FortiOS 5.0. It contains concept and feature descriptions, as well as configuration examples worked out in detail for the web-based manager and the CLI. This document also contains operating and troubleshooting information.

This handbook contains the following chapters:

- [Chapter 1, What's New for FortiOS 5.0](#) describes the new features in FortiOS 5.0.
- [Chapter 10, Install and System Administration for FortiOS 5.0](#) describes a number of administrative tasks to configure and setup the FortiGate unit for the first time. It also describes the best practices and sample configuration tips to secure your network and the FortiGate unit itself.
- [Chapter 7, Firewall for FortiOS 5.0](#) describes the concepts and techniques needed to configure the FortiGate firewall on your FortiGate unit.
- [Chapter 13, Logging and Reporting](#) describes how to begin choosing a log device for your logging requirements, the types of log files, how to configure your chosen log device, including detailed explanations of each log type of log message.
- [Chapter 15, Unified Threat Management for FortiOS 5.0](#) describes the Unified Threat Management (UTM) features available on your FortiGate unit, including antivirus, intrusion prevention system (IPS), anomaly protection (DoS), one-armed IPS (sniffer policies), web filtering, email filtering, data leak prevention (DLP), and application control. The chapter includes step-by-step instructions showing how to configure each feature. Example scenarios are included, with suggested configurations.
- [Chapter 3, Authentication for FortiOS 5.0](#) defines authentication and describes the FortiOS options for configuring authentication for FortiOS.
- [Chapter 14, Managing Devices for FortiOS 5.0](#) describes how to control network access for different types of personal mobile devices and apply client reputation.
- [Chapter 11, IPsec VPN for FortiOS 5.0](#) provides a general introduction to IPsec VPN technology, explains the features available with IPsec VPN and gives guidelines to decide what features you need to use, and how the FortiGate unit is configured to implement the features.
- [Chapter 16, SSL VPN for FortiOS 5.0](#) provides a general introduction to SSL VPN technology, explains the features available with SSL VPN and gives guidelines to decide what features you need to use, and how the FortiGate unit is configured to implement the features.
- [Chapter 2, Advanced Routing for FortiOS 5.0](#) provides detailed information about FortiGate dynamic routing including common dynamic routing features, troubleshooting, and each of the protocols including RIP, BGP, and OSPF.
- [Chapter 19, Virtual Domains](#) describes FortiGate Virtual Domains (VDOMs) and is intended for administrators who need guidance on solutions to suit different network needs and information on basic and advanced configuration of VDOMs. Virtual Domains (VDOMs)

multiply the capabilities of your FortiGate unit by using virtualization to partition your resources.

- [Chapter 9, High Availability for FortiOS 5.0](#) describes FortiGate HA, the FortiGate Clustering Protocol (FGCP), FortiGate support of VRRP, and FortiGate standalone TCP session synchronization.
- [Chapter 17, Traffic Shaping for FortiOS 5.0](#) describes how to configure FortiOS traffic shaping.
- [Chapter 4, FortiOS Carrier](#) describes FortiOS Carrier Multimedia messaging service (MMS) protection and GPRS Tunneling Protocol (GTP) protection.
- [Chapter 6, Deploying Wireless Networks for FortiOS 5.0](#) describes how to configure wireless networks with FortiWiFi, FortiGate, and FortiAP units.
- [Chapter 21, VoIP Solutions: SIP for FortiOS 5.0](#) describes FortiOS SIP support.
- [Chapter 22, WAN Optimization, Web Cache, Explicit Proxy, and WCCP for FortiOS 5.0](#) describes how FortiGate WAN optimization, web caching, and web proxy work and also describes how to configure these features.
- [Chapter 12, Load Balancing for FortiOS 5.0](#) describes firewall HTTP, HTTPS, SSL or generic TCP/UDP or IP server load balancing.

Chapter 1 What's New for FortiOS 5.0

This FortiOS Handbook chapter contains the following sections:

- [New Features in FortiOS 5.0 Patch 7](#) highlights the vulnerability fix for the Heartbleed vulnerability that was added to FortiOS 5.0 in Patch 7.
- [New features in FortiOS 5.0 Patch 6](#) highlights some of the changes in FortiOS 5.0 Patch 6.
- [New features in FortiOS 5.0 Patch 5](#) highlights some of the changes in FortiOS 5.0 Patch 5.
- [New features in FortiOS 5.0 Patch 4](#) highlights some of the changes in FortiOS 5.0 Patch 4.
- [New features in FortiOS 5.0 Patch 3](#) highlights some of the changes in FortiOS 5.0 Patch 3.
- [New features in FortiOS 5.0 Patch 2](#) highlights some of the changes in FortiOS 5.0 Patch 2.
- [Security Features](#) describes new Security features.
- [Authentication: users and devices](#) describes what's new for FortiOS user authentication and device management.
- [FortiOS and BYOD](#) outlines how to configure FortiOS device identification and BYOD protection features.
- [Client Reputation](#) introduces the new client reputation feature.
- [Wireless](#) describes new wireless features.
- [IPv6](#) describes new IPv6 features and how to configure many of them.
- [Logging and reporting](#) summarizes new FortiOS 5.0 logging and reporting features.
- [Firewall](#) describes the firewall features new to FortiOS 5.0.
- [WAN optimization and Web Caching](#) provides an overview and some examples that show how you need to change your FortiOS 4.3 WAN optimization configuration to work with FortiOS 5.0 WAN optimization, which is now policy-based.
- [Usability enhancements](#) describes some enhancements that make the web-based manager easier to use and more effective.
- [SSL VPN](#) describes some new SSL VPN features
- [Other new features](#) lists other new features in FortiOS 5.0.

New Features in FortiOS 5.0 Patch 7

This chapter provides information on the vulnerability fix added in FortiOS 5.0 Patch 7.

OpenSSL Vulnerability (Heartbleed) Fixed

An information disclosure vulnerability, known as Heartbleed, has been discovered in OpenSSL version 1.0.1 up to 1.0.1f. This vulnerability may allow an attacker to access sensitive information from memory by sending crafted TLS heartbeat requests.

This vulnerability has been fixed in OpenSSL 1.0.1g. FortiOS 5.0 Patch 7 has been upgraded to use this version of OpenSSL.

For more information about the Heartbleed vulnerability, go to <https://heartbleed.com/>.

New features in FortiOS 5.0 Patch 6

This chapter provides a brief introduction to the following features that were added to Patch 6 of FortiOS 5.0. See the release notes for a complete list of new features/resolved issues in this release.

- [Endpoint Control Daemon Improvement](#)
- [IPS Hardware Acceleration](#)
- [802.11g Protection Mode](#)
- [Miglogd Child Processes](#)
- [IPv6 in CRL/SCEP](#)
- [Extended IPS Database for D-series Desktop Models](#)
- [Logging Options for 3000 and 5000 Series Models](#)
- [Wireless Controller on FortiGate-30D](#)

Endpoint Control Daemon Improvement

Endpoint Control has been improved in several ways:

- For FortiGate models 3000 and higher, the maximum limit of FortiClient registration licence has been increased from 8000 to 16000.
- The intervals between KeepAlive messages are now configurable on the FortiGate unit, so that the interval value can be adjusted to get a trade-off between accuracy and the FortiGate workload. This value can be configured in the CLI:

```
set forticlient-keepalive-interval <interval> (interval measured in seconds).
```
- The intervals between two system update messages can be configured through the CLI:

```
set forticlient-sys-update-interval <interval> (interval measured in minutes).
```
- KeepAlive timestamps are now be stored use a in-memory avl tree structure, in order to decrease the number of CMDB savings.
- A mixed TCP/UDP mechanism now handles the registration sync.
- FortiClient registration information is stored using a hard disk instead of flash disk.

IPS Hardware Acceleration

New CLI commands have been added for configuring IPS hardware acceleration, replacing the previous `set hardware-accel-mode` command, to provide finer control over the settings. There are now two settings that can be chosen, one for the network processor and one for the content processor.

Network processor acceleration can be disabled or set to enable basic acceleration. Content processor acceleration can be disabled or enabled for either basic or advanced acceleration.

Syntax

```
config ips global
  set np-accel-mode {none | basic}
  set cp-accel-mode {none | basic | advanced}
end
```

802.11g Protection Mode

802.11g Protection Mode can now be enabled on managed FortiAP to avoid interference from 802.11b signals. Protection Mode can be set for RTS and CTS protection, or just for CTS.

Syntax

```
config wireless-controller wtp-profile
  edit <name>
    config radio-1
      set protection-mode {ctsonly | disable | rtscts}
    end
  end
end
```

Miglogd Child Processes

The number of miglogd child processes can now be configured directly through CLI to a value between 0-15, in order to keep up with logging requirements. The default number of child processes is 8.

Log messages are not lost if the number of child processes is decreased.

Syntax

```
config system global
  set miglogd-children <integer>
end
```



Increasing the number of child processes may affect a FortiGate unit's performance.

IPv6 in CRL/SCEP

IPv6 is now supported for all certificate revocation list (CRL) and Simple Certificate Enrollment Protocol (SCEP) features.

Extended IPS Database for D-series Desktop Models

The extended IPS database has been added for FortiGate D-series Desktop models. The extended database is disabled by default, but can be enabled in the CLI.

Syntax

```
config ips global
  set database extended
end
```



Enabling the extended IPS database may affect the performance of a FortiGate unit.

Logging Options for 3000 and 5000 Series Models

To increase stability of the 3000 series and 5000 series FortiGate models, the following changes have been made:

- Disk logging is disabled by default, but can be enabled through the CLI.
- When disk logging is disabled, it will not appear disk logging as an option on these models.

The default logging method for these models is now memory logging, which has been reintroduced as an option.

Wireless Controller on FortiGate-30D

The Wireless Controller has been added to the FortiGate-30D, which is now able to manage a maximum of 2 FortiAP units.

By default, this feature is only available through the CLI. To enable the Wireless Controller menu in the web-based manager support, enter the following command in the CLI:

```
config system global
  set gui-wireless-controller enable
end
```

When enabled, the default WiFi menu will be replaced by the full Wireless Controller menu in the web-based manager.

You can also use the command `set gui-ap-profile enable` to enable FortiAP and WIDS profiles.

New features in FortiOS 5.0 Patch 5

This chapter provides a brief introduction to the following features that were added to Patch 5 of FortiOS 5.0. See the release notes for a complete list of new features in this release.

- Improvements to Endpoint Control
- FortiAP LAN port support
- Automatically allowing basic applications
- Pre-authorizing a FortiAP unit
- Preventing IP fragmentation of packets in CAPWAP tunnels
- Limiting access for unauthenticated users
- LDAP browser to import users into a user group
- Dedicated management CPU
- Improvements to the Traffic History and Threat History widgets
- Assigning an IP address to a dynamic IPsec VPN interface
- SSL VPN History widget
- Port Block Allocation (PBA) for CGN to reduce logs
- Neighbor cache table for IPv6
- Improved HA diagnose commands
- Secure disk erasing
- Anonymize user names in logs
- VLAN interface traffic statistics
- Preserving the Class of Service bit
- Front panel illustration
- USB entropy token support
- Station locate for FortiWiFi units
- Switch Controller added models 200D, 240D, 600C, 800C, and 1000C
- Diagnose command for 5000 series FortiGate units
- New platforms for FortiGate-VM
- Supported RFCs

Improvements to Endpoint Control

There have been several improvements made to Endpoint Control.

New menu options

Endpoint Control now has its own menu, which can be found at *User & Device > Endpoint Protection*. This menu contains options for creating FortiClient profiles.

Default profile

A default FortiClient profile has been added that enables AntiVirus, Web Filtering, and VPN for Windows and Mac. All other features are disabled.

The profile creation screen has also been simplified to allow for easier configuration.

Figure 1: The default FortiClient profile

Profile Name

Comments 0/255

FortiClient Configuration Deployment

Windows and Mac

- AntiVirus Protection
- Web Category Filtering
- Disable Web Category Filtering when protected by this FortiGate
- VPN
 - Client VPN Provisioning
- Application Firewall
- Endpoint Vulnerability Scan on Client
- Upload Logs to FortiAnalyzer/FortiManager
- Use FortiManager for client software/signature update
- Dashboard Banner

iOS

- Web Category Filtering
- Client VPN Provisioning
- Distribute Configuration Profile (.mobileconfig file)

Android

- Web Category Filtering
- Client VPN Provisioning

FortiClient Monitor

The FortiClient Monitor displays a variety of information about FortiClient users, including current status, device type, and FortiClient version. It can be found by going to *User & Device > Monitor > FortiClient*.

FortiAP LAN port support

New functions are now available for FortiAP models that have LAN ports (currently the 11C, 14C, and 28C). The LAN port(s) can now be bridged to either an SSID or to the FortiAP unit's WAN port (bridging to the WAN port is the default setting).

LAN port bridging can be done with FortiAP units in either Bridge or Tunnel mode.

Bridging with the FortiAP's SSID(s)

Bridging the LAN port with the FortiAP's SSID(s) allows combines traffic from both sources to provide a single broadcast domain for the wired and wireless users.

This configuration has the following features:

- The IP addresses for LAN clients come from the DHCP server that is serving the wireless clients.
- Traffic from LAN clients is bridged to the VLAN used by the SSID to send traffic to the controller.
- Wireless and LAN clients are on the same network and can communicate locally, via the FortiAP.

Bridging with the WAN port

Bridging the LAN port with the WAN port allows the FortiAP unit to be used as a hub which is also an access point.

This configuration has the following features:

- The IP addresses for LAN clients come from the WAN directly and will typically be in the same range as the AP itself.
- All LAN client traffic is bridged directly to the WAN interface.
- Communication between wireless and LAN clients can only occur if a policy on the FortiGate unit allows it.

Configuring bridging

A FortiAP LAN port can be configured to bridge with an SSID from either the web-based manager or the CLI.

Using the web-based manager


1. Go to *WiFi Controller > WiFi Network > Custom AP Profiles*.



On FortiGate models 100D, 200D, 240D, 600C, 800C, and 1000C, go to *WiFi & Switch Controller > WiFi Network > Custom AP Profiles*.

2. Create a new custom profile or edit the default profile for your FortiAP model.
3. Under *LAN Port*, change *Mode* to *Bridge to* and select the appropriate option.
4. Select *OK*.

Figure 2: Configuring bridging using the web-based manager



Bridging can also be set up configured on a specific FortiAP unit, rather than through the use of an AP profile by going to *WiFi Controller > Managed Devices > Managed FortiAPs*.



On FortiGate models 100D, 200D, 240D, 600C, 800C, and 1000C, go to *WiFi & Switch Controller > Managed Devices > Managed FortiAPs*.

Using the CLI

In the example below, two ports on a FortiAP-28C are configured, with port 1 bridged to the WAN port and port 2 bridged to the SSID(s):

```
config wireless-controller wtp-profile
  edit FAP28C-default
    config lan
      set port1-mode bridge-to-wan
      set port2-mode bridge-to-ssid
    end
  end
end
```

Bridging can also be set up configured on a specific FortiAP unit, rather than through the use of an AP profile:

```
config wireless-controller wtp
  edit FAP28C0123456789
    config lan
      set port1-mode bridge-to-wan
      set port2-mode bridge-to-ssid
    end
  end
end
```

Restrictions

- While the FortiAP-14C has four physical LAN ports, these ports must share the same configuration.
- Any host connected to a LAN port will be taken as authenticated.
- The use of dynamic VLANs for the host behind LAN port is not supported.
- RADIUS MAC authentication for the host behind LAN port is not supported.

Automatically allowing basic applications

Application control profiles can now be configured from the CLI to allow basic, commonly used applications to go through without having to separately configure the profile each application. This is useful when you wish to control the traffic to an entire category of applications without affecting the traffic for basic applications that are required on a daily basis.

For example, an application sensor that blocks the Category "Network.Service" would normally also block DNS service, causing Internet service issues. Using the new command, DNS can now be allowed, eliminating this issue while still blocking other applications within the category.

Basic traffic can also be allowed for ICMP, generic HTTP web browsing, and generic SSL communication.

Syntax

```
config application list
  edit appcontrol
    set options allow-dns allow-icmp allow-http allow-ssl
  end
```



DNS is set to be allowed by default for all application control profiles, while the other settings must be enabled to take effect.

Pre-authorizing a FortiAP unit

Users can now pre-authorize a FortiAP unit by before connecting the unit to the FortiGate unit. To pre-authorize a FortiAP unit, do the following:

1. Go to *WiFi Controller > Managed Access Points > Managed FortiAPs* and select *Create New*.




On FortiGate models 100D, 200D, 240D, 600C, 800C, and 1000C, go to *WiFi & Switch Controller > Managed Devices > Managed FortiAPs*

2. Enter the serial number of the FortiAP unit.
3. Configure the *Wireless Settings* as required.
4. Select *OK*.

The new FortiAP now appear on the Managed FortiAPs list as authorized but off-line. The FortiAP unit can now connect to the FortiGate unit.

Figure 3: Pre-authorizing a FortiAP unit

Serial Number	<input type="text" value="FAP11C3X13000412"/>
Name	<input type="text"/>
Comments	<input type="text" value="Write a comment..."/> 0/35
State	Authorized
Wireless Settings	
<input checked="" type="checkbox"/> Enable WiFi Radio	
SSID	<input checked="" type="radio"/> Automatically Inherit all SSIDs <input type="radio"/> Select SSIDs
Auto TX Power Control	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
TX Power	
<input type="button" value="OK"/> <input type="button" value="Cancel"/>	



If the FortiAP unit will be connecting directly to one of the FortiGate unit's ports, the port will still need to have its *Addressing mode* set to *Dedicate to FortiAP*.

Preventing IP fragmentation of packets in CAPWAP tunnels

A common problem with controller-based WiFi networks is reduced performance due to IP fragmentation of the packets in the CAPWAP tunnel.

Fragmentation can occur because of CAPWAP tunnel overhead increasing packet size. If the original wireless client packets are close to the maximum transmission unit (MTU) size for the network (usually 1500 bytes for Ethernet networks unless jumbo frames are used) the resulting CAPWAP packets may be larger than the MTU, causing the packets to be fragmented. Fragmenting packets can result in data loss, jitter, and decreased throughput.

The FortiOS/FortiAP solution to this problem is to cause wireless clients to send smaller packets to FortiAP devices, resulting in 1500-byte CAPWAP packets and no fragmentation. The following options configure CAPWAP IP fragmentation control:

```
config wireless-controller wtp-profile
  edit new-wtp
    set ip-fragment-preventing {tcp-mss-adjust | icmp-unreachable}
    set tun-mtu-uplink {0 | 576 | 1500}
    set tun-mtu-downlink {0 | 576 | 1500}
  end
end
```

By default, `tcp-mss-adjust` is enabled, `icmp-unreachable` is disabled, and `tun-mtu-uplink` and `tun-mtu-downlink` are set to 0.

To set `tun-mtu-uplink` and `tun-mtu-downlink`, use the default TCP MTU value of 1500. This default configuration prevents packet fragmentation because the FortiAP unit limits the size of TCP packets received from wireless clients so the packets don't have to be fragmented before CAPWAP encapsulation.

The `tcp-mss-adjust` option causes the FortiAP unit to limit the maximum segment size (MSS) of TCP packets sent by wireless clients. The FortiAP does this by adding a reduced MSS value to the SYN packets sent by the FortiAP unit when negotiating with a wireless client to establish a session. This results in the wireless client sending packets that are smaller than the `tun-mtu-uplink` setting, so that when the CAPWAP headers are added, the CAPWAP packets have an MTU that matches the `tun-mtu-uplink` size.

The `icmp-unreachable` option affects all traffic (UDP and TCP) between wireless clients and the FortiAP unit. This option causes the FortiAP unit to drop packets that have the "Don't Fragment" bit set in their IP header and that are large enough to cause fragmentation and then send an ICMP packet -- type 3 "ICMP Destination unreachable" with code 4 "Fragmentation Needed and Don't Fragment was Set" back to the wireless controller. This should cause the wireless client to send smaller TCP and UDP packets.

Limiting access for unauthenticated users

When configuring User Identity policies, if you select the option *Skip this policy for unauthenticated user* the policy will only apply to users who have already authenticated with the FortiGate unit. This feature is intended for networks with two kinds of users:

- Single sign-on users who have authenticated when their devices connected to their network
- Other users who do not authenticate with the network so are “unauthenticated”

Sessions from authenticated users can match this policy and sessions from unauthenticated users will skip this policy and potentially be matched with policies further down the policy list. Typically, you would arrange a policy with *Skip this policy for unauthenticated user* at the top of a policy list.

You can also use the following CLI command to enable skipping policies for unauthenticated users:

```
config firewall policy
  edit <id>
    set identity-based enable
    set fall-through-unauthenticated enable
  next
```

Use case - allowing limited access for unauthenticated users

Consider an office with open use PCs in common areas. Staff and customers do not have to log in to these PCs and can use them for limited access to the Internet. From their desks, employees of this office log into PCs which are logged into the office network. The FortiGate unit on the office network uses single sign-on to get user credentials from the network authentication server.

The open use PCs have limited access to the Internet. Employee PCs can access internal resources and have unlimited access to the Internet.

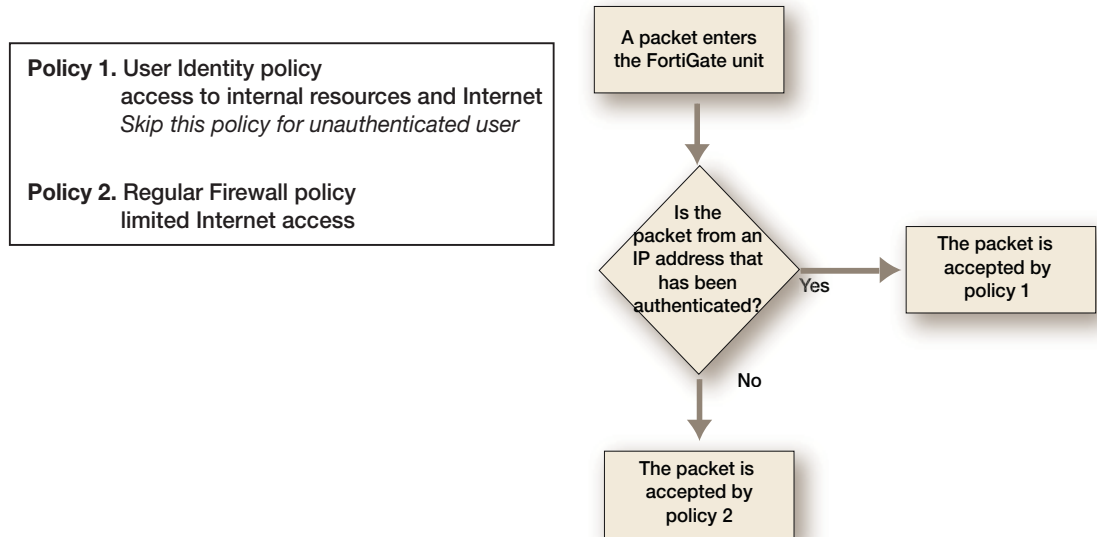
To support these different levels of access you can add a user identity policy to the top of the policy list that allows authenticated users to access internal resources and to have unlimited access to the Internet. In this policy, select *Skip this policy for unauthenticated user*.

Add a normal firewall policy below this policy that allows limited access to the Internet.

Sessions from authenticated PCs will be accepted by the User Identity policy. Sessions from unauthenticated PCs will skip the User Identity policy and be accepted by the normal firewall policy.

Figure 4 shows how the FortiGate unit handles packets received from authenticated and unauthenticated users.

Figure 4: Packet flow for authenticated and unauthenticated users



Use case - multiple levels of authentication

As a variation of the above use case, Policy 2 could be a User Identity policy and *Skip this policy for unauthenticated user* would not be selected. Sessions from unauthenticated users that are accepted by Policy2 would now require users to authenticate before traffic can connect through the FortiGate unit. The result is different levels of authentication: Single sign on for some users and firewall authentication for others.

LDAP browser to import users into a user group

You can use the new LDAP browser to add LDAP users to a user group.

Figure 5: The LDAP browser

Name

Type (RSSO) Firewall Fortinet Single Sign-On (FSSO) Guest RADIUS Single Sign-On

Members

Remote groups

Remote Server	Group Name
FAC_LDAP	Remote_Access

Dedicated management CPU

FortiGate units in the 2U or High-End categories (models 1000 and above) can now be configured to have a dedicated management CPU. This reserves one CPU core, CPU 0, for running management tasks such as the web GUI, as well as the CLI and related daemons. By having a dedicated management CPU, access to the management GUI and CLI is guaranteed even under when the unit is under a heavy traffic load.



Using a dedicated management CPU may have an impact on the overall performance of the FortiGate unit.

The dedicated management CPU is enabled using the CLI:

```
configure system npu
  set dedicated-management-cpu enable
end
```

Improvements to the Traffic History and Threat History widgets

Several changes have been made to the *Traffic History* and *Threat History* widgets:

- The *Show Sessions* and *Show All Incidents* (formerly *Show Threats*) options has been improved to show more information about individual sessions or threats.
- The drilldown page for the *Threat History* widget has been improved by adding new columns and adjusting field formats.
- The *Threat History* widget has replaced the *Reputation Score* monitor used for Client Reputation, which has been removed.

Assigning an IP address to a dynamic IPsec VPN interface

An IP addresses can now be assigned to a dynamic IPsec VPN interface to be used for traffic egressing over the IPsec interface, to avoid traffic being blocked due to an inappropriate address. An IP address is assigned by going to *System > Network > Interfaces* and editing the interface for the IPsec VPN.

Figure 6: Assigning an IP address to a dynamic IPsec VPN interface

Name	fc_vpn
Type	Tunnel Interface
Interface	wan1
<hr/>	
Addressing mode	Manual
IP	<input type="text" value="10.10.20.1"/>
Remote IP	<input type="text" value="1.1.1.1"/>
IPv6 Address	<input type="text" value="::/0"/>

SSL VPN History widget

Login history can now be added to the SSL VPN Portal, which shows a user their past logins. The number of logins shown can be anywhere between 1 and 255 (the default is 5).

The login can be set by going to *VPN > SSL > Portal*, selecting *Include Login History*, and setting an appropriate *Number of history entries*.

It can also be set using the CLI:

```
config vpn ssl web portal
  edit <portal>
    config widget
      edit <ID>
        set type history
        set display-limit <1-255>
      end
    end
  end
end
```

Port Block Allocation (PBA) for CGN to reduce logs

Port Block Allocation (PBA), a Carrier Grade NAT (CGN) feature, can reduce the number of log messages generated by NAT operations.

PBA can be configured using by going to *Firewall Objects > Virtual IPs > IP Pools*. It can also be configured using the CLI.

```
config firewall ippool
  edit ippool
    set type port-block-allocation
    set block-size <integer>
    set num-blocks-per-user <integer>
  end
end
```

You configure PBA by creating a private IP address range and assigning multiple port ranges (or blocks) to that IP address range. When a connection is received from the IP range, the source port is translated to a ports in the first range. A log message is written when this happens.

As more connections are received from this IP address range they are assigned to other ports in the first port block. Eventually all of the ports in the block will be used. When a new connection is received, another block of ports is started and a log message is written.

So instead of writing a log message for every NAT event, log messages are only written when a new block of ports is started and again when its used up.

Neighbor cache table for IPv6

A table has now been added to configure IPv6 neighbor cache entries and to save the entries when the FortiGate unit reboots, using the command `config system ipv6-neighbor-cache`.

In the following example, a neighbor cache entry is configured to use the DMZ interface:

```
config system ipv6-neighbor-cache
  edit 1
    set interface dmz
    set ipv6 6666::11
    set mac 00:09:0f:01:02:03
  end
end
```

Improved HA diagnose commands

The new command `diagnose sys ha dump-by` has replaced the command `diagnose sys ha dump`. The new command has the following syntax:

```
diagnose sys ha dump-by {all-xdb | all-vcluster| rcache | all-group |
  memory | debug-zone | vdom | kernel | device | stat| sesync}
```

Each option displays different types of information about the cluster.

The following new HA diagnose commands have also been added:

```
diagnose sys ha sesync-stats
diagnose sys ha extfile-sig
```

Secure disk erasing

All data on the FortiGate boot device and any hard disks installed in a FortiGate unit can now be securely and permanently erased using the `execute erase-disk` command. This command performs a low-level format and also overwrites every block on the device with random data three times.

Anonymize user names in logs

Log messages can now be configured to replace user names with the word **anonymous**, so that user names are not visible in log messages. This feature can be enabled from the CLI using the following command:

```
config log setting
  set user-anonymize enable
end
```

VLAN interface traffic statistics

A VLAN accounting table has been added to the NP4 driver to poll accounting data from the FortiGate unit in order to monitor traffic statistics from VLAN interfaces. The polling interval is set to 1 second.

Preserving the Class of Service bit

FortiGate units can now preserve the value of the Class of Service (CoS) bit, also called Priority Code Point (PCP), when a packet traverses a VLAN network.

Front panel illustration

An illustrated version of the FortiGate unit's front panel has been added above the list of interfaces, found at *System > Network > Interfaces*. As with the panel found in the *Unit Operation* widget, interfaces appear green when connected and further details are shown when the mouse pointer hovers over a specific port.

USB entropy token support

Use of a USB entropy token during the boot process is now enabled by default when using a FortiGate in Federal Information Processing Standards-Common Criteria (FIPS-CC) mode. If a FortiGate unit in this mode does not have a USB entropy token inserted, it is unable to complete the boot process and will display the following message: `Please insert entropy token to continue boot process.`

Entropy token use can be disabled from the CLI. It can also be enabled on a FortiGate unit in normal mode (by default, entropy tokens are disabled in normal mode).

Syntax

```
config system fips
    set entropy-token {enable | disable}
end
```



The entropy token must be present during boot process when a FortiGate unit is switched to FIPS-CC mode.

Station locate for FortiWiFi units

Station locate allows a FortiWiFi unit to detect all wireless clients whether they are associated or not. A record is kept of MAC address, statistical time interval and RSSI data.

Station locate is enabled using the following CLI command:

```
config wireless-controller wtp-profile
    edit "FAP220B-default"
        config radio-1
            set station-locate enable
            set station-locate-interval 1
        next
        config radio-2
            set station-locate enable
            set station-locate-interval 1
        end
    end
end
```

Switch Controller added models 200D, 240D, 600C, 800C, and 1000C

The Switch Controller, used to managed FortiSwitch units with a FortiGate unit, has been added to the following models: 200D, 240D, 600C, 800C, and 1000C.

Because of this feature, there have been several web-based manager menu changes to these units:

- *WiFi Controller* has changed to *WiFi & Switch Controller*.
- *Managed Access Points* has changed to *Managed Devices* and now contains the *Managed FortiSwitch* option.
- The *Switch Network* menu has been added, which contains the *Virtual Switch* option.

Diagnose command for 5000 series FortiGate units

A new diagnose command, `diagnose test application ipmc_sensord`, is available to view chassis IPMC status from a 5000 series blade installed in a chassis. The command can display:

- Power supply detection
- IPMC sensor status detection
- Comlog enable/disable/info/read/clear
- Smc time set/get
- AMC info
- Microswitch status detection
- HACO info

Because of this change, the following obsolete commands have been removed:

- `get system chassis`
- `get system blades`
- `get chassis status`
- `diag hardware fruinfo`
- `exec bladekvm`

New platforms for FortiGate-VM

FortiGate-VM is now supported for Microsoft Hyper-V and Kernel-based Virtual Machine (KVM).

Supported RFCs

The following RFCs are supported by the new features for FortiOS 5 Patch 5:

Table 1: Supported RFCs

Number	Title
2766	Network Address Translation - Protocol Translation (NAT-PT)

Table 1: Supported RFCs

Number	Title
4787	Network Address Translation (NAT) Behavioral Requirements for Unicast UDP
6691	TCP Options and Maximum Segment Size (MSS)

New features in FortiOS 5.0 Patch 4

This chapter provides a brief introduction to the following features that were added to Patch 4 of FortiOS 5.0. See the release notes for a complete list of new features in this release.

- FortiSandbox
- Wireless Health Dashboard
- IPsec VPN
- Managing FortiAP units
- Dynamic VLANs for SSIDs
- NAT46 & NAT64
- Enhancements to Tables
- FortiAnalyzer and FortiManager log encryption
- FortiToken Mobile
- Load balancing for explicit web proxy forwarding server groups
- Server load balancing enhancements
- Additional filters for IPS and Application Control
- Blocking IPv6 packets by extension headers
- Distinguishing between HTTP GET and POST in DLP
- RADIUS Accounting
- H3C Compatibility
- Web filter administrative overrides
- Configurable idle timeout for console admin login sessions
- TCP reset
- Log Volume Monitor
- Invalid Packet log
- Server limits
- PoE Power Management display
- Other new features

FortiSandbox

The new FortiSandbox unit can be used with a FortiGate unit for sandboxing suspicious files. Sandboxing can also be done using Cloud Sandbox, which was previously known as FortiGuard Analytics. For more information about this feature, see [“FortiSandbox” on page 138](#).

Wireless Health Dashboard

The Wireless Health Dashboard provides an easy method for determining the health of your network’s wireless infrastructure. The dashboard is used to display a variety of widgets, which show information such as AP status, client count over time and login failures.

The dashboard can be found by going to *WiFi Controller > Monitor > Wireless Health*.

Figure 7: The Wireless Health Dashboard



IPsec VPN

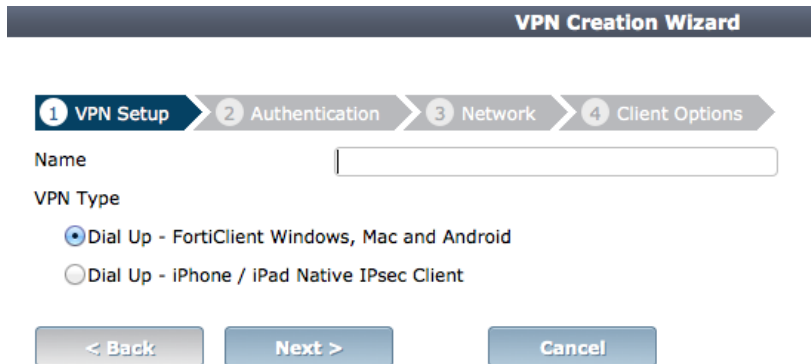
There have been several changes made to how IPsec VPN is configured.

Dial-up IPsec VPN Creation Wizard

A new wizard can be used to create Dial-up IPsec VPNs for FortiClient and Native iOS IPsec clients. The FortiClient configuration can be used for all platforms supported by FortiClient. Find the wizard by going to *VPN > IPsec > Auto Key (IKE)* and selecting *Create VPN Wizard*.

See <http://docs.fortinet.com/sysadmin.html> for more information about using the VPN Creation Wizard.

Figure 8: The VPN Creation Wizard



Show or Hide policy-based IPsec VPN

Policy-based IPsec VPN options have been added to the Feature Select options, which controls what features can be viewed and configured through the web-based manager. For more information on this feature, see “[Feature Select](#)” on page 243.

By default, policy-based IPsec VPN is hidden from the web-based manager and interface-based VPNs are easier to configure.

Managing FortiAP units

There have been several changes to how FortiAP units are managed by a FortiGate unit.

Units remain online when their WiFi Controller goes offline

FortiAP units can now remain online when their connection to the FortiGate unit's WiFi Controller goes offline. During such an outage, WiFi clients already associated with a bridge mode FortiAP unit continue to remain connected to their SSID and can communicate with other WiFi clients. Access to other network resources; however, is suspended until the FortiGate unit is back online.

The FortiAP unit can also continue to authenticate users if the SSID meets these conditions:

- Traffic Mode is Local bridge with FortiAP's Interface
In this mode, the FortiAP unit does not send traffic back to the wireless controller.
- Security Mode is either WPA/WPA2-Personal or Open.
These modes do not require the user database. In WPA/WPA2-Personal authentication, all clients use the same pre-shared key which is known to the FortiAP unit.
- Allow new client association when controller connection is down is enabled.

This field is available only if the other conditions have been met.

Assigning the same profile to multiple FortiAP units

The same profile can now be applied to multiple managed FortiAP units at the same time. To do this, do the following:

1. Go to *WiFi Controller > Managed Access Points > Managed FortiAPs* to view the AP list.
2. Select all FortiAP units you wish to apply the profile to.
3. Right click on one of the selected FortiAPs and select *Assign Profile*.
4. Choose the profile you wish to apply.

Dynamic VLANs for SSIDs

Dynamic VLANs can now be used to divide a single SSID into several VLANs. In Patch 4, Dynamic VLANs are supported for both tunnel and bridge mode SSIDs.

VLAN assignment is based on the credentials supplied by the user. Dynamic VLANs allow individual users to be assigned different VLANs resulting in different levels of access even though all users are connecting to the same SSID.

The task of assigning users to a specific VLAN is handled by a RADIUS authentication server. When a client attempts to associate to a FortiAP registered with a controller, the FortiAP passes the credentials of the user to the RADIUS server for validation. Once the authentication is successful, the RADIUS server passes certain Internet Engineering Task Force (IETF) attributes to the user. These RADIUS attributes include the VLAN ID that should be assigned to the wireless client.

Dynamic VLANs are configured by doing the following:

1. Go to *User & Device > Authentication > RADIUS Servers* and create a new RADIUS server.
2. Go to *WiFi Controller > WiFi Network > SSID* and create a new SSID.
3. Enable Dynamic VLAN in the CLI, using the following command:

```
config wireless-controller vap
  edit <name>
    set dynamic-vlan enable
  end
```

4. Go to *WiFi Controller > WiFi Network > Custom AP* and create a new radio 1 and radio 2 that use the new SSID.
5. Go to *System > Network > Interfaces* and create two or more VLAN interfaces that have DHCP server enabled.
6. Go to *Policy > Policy > Policy* and create policies that allow outbound traffic from the new VLANs.
7. Configure a policy on the RADIUS server for each VLAN.

When users scan for available SSIDs, they can connect to the new SSID and be assigned to one of the VLANs based on their credentials.

See <http://docs.fortinet.com/supplement.html> for some Dynamic VLAN examples.

NAT46 & NAT64

Policies and Virtual IPs for NAT46 and NAT64 can now be configured from the web-based manager. For these options to appear in the web-based manager, this feature must be enabled using Feature Select. For more information, see “[Feature Select](#)” on page 243.

To configure NAT64 policies go to *Policy > Policy > NAT64 Policy*.

To configure NAT46 policies go to *Policy > Policy > NAT46 Policy*.

Enhancements to Tables

Several enhancements have been made to the tables in the web-based manager to improve access to information.

Policy Table

The following enhancements have been made to the Policy Table:

- The Security Features column has been divided so that each feature has an individual menu, allowing the profile used to be easily visible.
- A pulldown menu will appear when a specific element is selected that displays the other options for the field, as well as showing the *Create New* option.
- The right-click menu has been simplified to show only the options related to the location that was selected. The menu will also change depending on whether a specific element name was selected or if the cell background was selected.

Member Display

The member columns for tables under *Firewall Objects* and *User & Device* have been improved to display members in a grid. When there is a high number of members in a single group, some members will be displayed in the grid, with the hidden members viewed via a dropdown menu.

The number of sub-columns displaying members, the width of sub-columns and the number of lines used to display members are all customizable by right-clicking on the header and selecting *Members Column Option*.

Fortinet Top Bar

In order to ensure that the Fortinet Top Bar appears in all browsers, port 8011 must be allowed in the firewall policy being used.

FortiAnalyzer and FortiManager log encryption

Logs sent to a FortiAnalyzer or FortiManager unit from a FortiGate unit can now be encrypted. Encryption is enabled by going to *Log & Report > Log Config > Log Settings*.

FortiToken Mobile

There have been several changes made to FortiToken Mobile:

- A QR code image will be attached to FortiToken activation emails.
- Softtoken polling requests have been extended to 5 minutes.
- The range for two factor FortiToken mobile expiry is now 1-168 hours.
- Notifications will be sent when a local user is created.

Load balancing for explicit web proxy forwarding server groups

Explicit web proxy traffic can now be load balanced among multiple forwarding servers in a forwarding server group.

To configure load balancing, add multiple forwarding servers to a forwarding server group and turn on load balancing for the server group. Then add the forwarding server group to a security policy.

The following example adds three forwarding servers to a forwarding server group. Start by creating the forwarding servers:

```
config web-proxy forward-server
  edit fwd-srv-1
    set ip 10.10.10.10
    set port 8080
  next
  edit fwd-srv-2
    set ip 10.10.10.20
    set port 8080
  next
  edit fwd-srv-3
    set ip 10.10.10.30
    set port 8080
end
```

Then add the forwarding servers to a group:

```
config web-proxy forward-server-group
  edit fwd-srv-grp
    set affinity enable
    set ldb-method weighted
    set group-down-option block
    config server-list
      edit fwd-srv-1
        set weight 10
      next
      edit fwd-srv-2
        set weight 10
      next
      edit fwd-srv-3
        set weight 10
    end
  end
```

Then add the forwarding server group to a web-proxy security policy:

```
config firewall policy
  edit 0
    set srcintf web-proxy
    ...
    set webproxy-forward-server fwd-srv-grp
    ...
  end
```

Server load balancing enhancements

Server load balancing has been enhanced to alert administrators when a server fails and to improve handling of HTTP redirects.

SNMP traps

FortiGate units can now send SNMP traps when the FortiGate unit determines that one of the servers in a server load balance group has gone down. The OID for the trap is *.fgTrapServerLoadBalanceRealServerDown*.

You can use the following CLI command to enable this trap:

```
config system snmp community
  edit 0
    set events load-balance-real-server-down enable
  end
```

HTTP redirects

Server load balancing now also supports checking HTTP redirects and setting the maximum number of redirects when 300-level return codes are received.

You can set `http-max-redirects` in the range 0 to 5. The default value is 0 which means do not check redirects, just assume they are available. This is how redirects functioned previously. When you set this option to 1 or more, the FortiGate unit will check up to 5 redirect URLs, until it finds one that is active. Traffic is then re-directed to the first active URL that is found.

Use the following command to check up to 3 redirects:

```
config firewall ldb-monitor
  edit 0
    set set type http
    set port 80
    set http-get "/index.php"
    set http-max-redirects 3
  end
```

Additional filters for IPS and Application Control

New filters have been added for IPS and Application Control that will be shown when the *Advanced filter* option is selected on the sensor creation page. The new filters for IPS are Application and Protocol and the new filters for Application Control are Vendor and Protocol.

Blocking IPv6 packets by extension headers

FortiOS can now block IPv6 packets based on the extension headers, using the CLI syntax `config firewall ipv6-eh-filter`.

The following commands are now available:

```
set hop-opt {disable | enable}: Block packets with Hop-by-Hop Options header.
set dest-opt {disable | enable}: Block packets with Destination Options header.
set hdopt-type <integar>: Block specific Hop-by-Hop and/or Destination Option types (maximum 7 types, each between 0 and 255).
set routing {disable | enable}: Block packets with Routing header.
set routing-type <integar>: Block specific Routing header types (maximum 7 types, each between 0 and 255).
set fragment {disable | enable}: Block packets with Fragment header.
set auth {disable | enable}: Block packets with Authentication header.
set no-next {disable | enable}: Block packets with No Next header.
```

Distinguishing between HTTP GET and POST in DLP

Data Leak Prevention (DLP) can now distinguish between HTTP GET and POST protocols, allowing the protocols to be selected independently.

HTTP POST protocol can be examined by both message and file filters, while HTTP GET can only be used for file filters.

RADIUS Accounting

Accounting servers can now be configured in a RADIUS setting. For each RADIUS server, four more accounting servers can be created.

The following example adds an accounting server to a RADIUS server:

```
config user radius
  edit rad159
    set server 172.16.62.159
    set secret asdfasdf
    config accounting-server
      edit 1
        set status enable
        set server 175.18.5.36
        set secret asdfasdf
      end
    end
  end
```

H3C Compatibility

FortiOS now has H3C compatibility, allowing 802.1x authentication to be supported with two RADIUS attributes.

The following example enables H3C compatibility:

```
config user radius
  edit rad-jason
    set h3c-compatibility enable
  end
```

Web filter administrative overrides

Administrative web filter overrides can now be configured by going to *Security Profiles > Web Filter > Web Overrides*. Overrides allow specific users to use an alternate web filter profile, in order to access sites that would normally be blocked.

Configurable idle timeout for console admin login sessions

An idle timeout has been added for FortiGate console sessions (admin sessions connecting to a FortiGate console port or USB port). By default the console timeout is set to 0 and console sessions will never timeout. You can enable a timeout in the range of 15-300 seconds from the CLI. Use the following command to set the timeout to 25 seconds:

```
config sys global
  set admin-console-timeout 25
end
```

Use the following command to disable the timeout.

```
config sys global
  unset admin-console-timeout
end
```

TCP reset

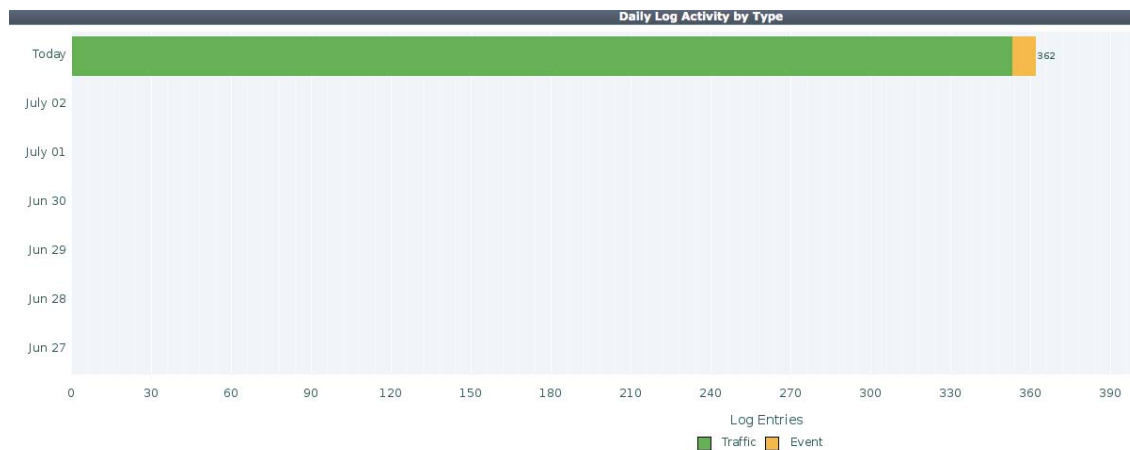
Security policies can now be configured to send a TCP reset when a specific application session times out. The following command is used to enable this function:

```
config firewall policy/policy6
  edit 0
    set timeout-send-rst enable
  end
```

Log Volume Monitor

The Log Monitor has been renamed the Log Volume Monitor and has had several visual enhancements made. It can be found at *Log & Report > Monitor > Logging Volume Monitor*.

Figure 9: The Log Volume Monitor



Invalid Packet log

The Invalid Packet log has been merged with the Local Traffic and Forward Traffic logs. Denied traffic will now appear in either of these logs.

Server limits

The maximum number of virtual and real servers has been increased. The new maximum values vary by FortiGate model and are as follows:

- Desktop FGT: virtual servers, 128 globally; real servers, 4 per entry.
- 1U FGT: virtual servers, 512 globally; real servers, 8 per entry
- 2U FGT: virtual servers, 2048 globally; real servers, 32 per entry.

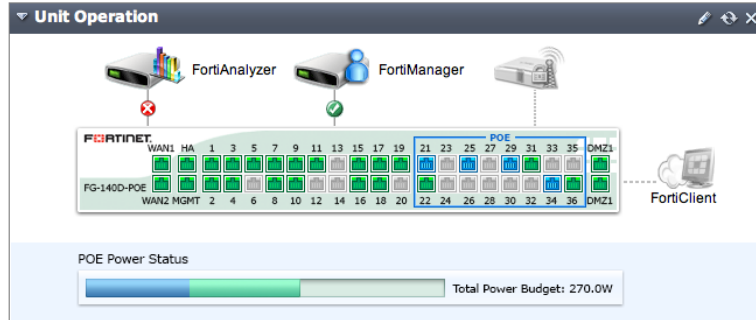
PoE Power Management display

This feature is only available on FortiGate 140D-POE models.

In the *Unit Operation* widget, the *POE Power Status* bar has been added to show the current power output for Power over Ethernet (POE). The bar displays both used power and reserved power.

The port display has also been changed so that the ports that are supplying power appear in blue.

Figure 10:The Unit Operation widget with the POE Power Status bar



Other new features

- Valgrind is now supported in urlfilter daemon, using the following commands:


```
diagnose debug urlfilter valgrind enable
diagnose debug urlfilter valgrind status
diagnose debug urlfilter valgrind memcheck
diagnose debug urlfilter valgrind log
```
- The default post-login banner now displays the time of the last successful and last failed administrator login.

New features in FortiOS 5.0 Patch 3

This chapter provides a brief introduction to the following features that were added to Patch 3 of FortiOS 5.0. See the release notes for a complete list of new features in this release.

- Security Features
- Exempting IP addresses from IPS
- DLP Watermarking Client
- Predefined Device Groups
- Client Reputation Configuration
- Feature Select
- Changes to Endpoint Control
- Managing FortiAP units
- FortiGuard Subscription Services
- Adding Explicit Web Proxy services
- SSO Authentication failover for the Explicit Web Proxy
- User Creation Wizard
- FortiClient Registration
- DSS and ECDSA Certificates for FortiGate SSL-related features
- LDAP Servers
- User Monitor
- Web Filter Profiles
- CAPWAP Administrative Access
- IPS Algorithms
- NAC-Quarantine Traffic Logs
- New System Report Charts
- Memory Logging
- URL-based Web Proxy Forwarding
- Changes to Routing
- RADIUS Support for Dynamic VLANs
- Dedicated Management Port
- URL Filtering
- URL Source Tracking
- IPv6 Denial of Service Policies
- Support for NAT46, VIP64 and VIP46
- Packet Capture Filters
- Configure hosts in an SNMP v1/2c community to send queries or receive traps
- IP in IP tunneling support (RFC 1853)
- GTP-u acceleration on FortiGate units with SP3 processors

Security Features

Features previously known as UTM Security Features are now known as Security Features. For more information about new Security Features in FortiOS 5.0, see [“Security Features” on page 138](#).

Exempting IP addresses from IPS

IPS filters can be configured so that signatures are not applied to traffic from specific IP addresses. For more information about IPS exemptions, see [“Exempting IP addresses from IPS” on page 144](#).

DLP Watermarking Client

The DLP watermarking client is now available for Windows as part of FortiExplorer. For more information about DLP watermarking, see [“DLP watermarking” on page 153](#).

Predefined Device Groups

FortiOS now has Predefined Device Groups for Blackberry Playbook, Router/NAT Device and Windows Tablet. For more information about Predefined Device Groups, see [“Device Groups” on page 175](#).

Client Reputation Configuration

Client Reputation configuration can be found at *Security Profiles > Client Reputation*. For more information about Client Reputation, see [“Client Reputation” on page 181](#).

Feature Select

Feature Select replaces the *Display Options on GUI* feature to control which features can be configured and viewed on the web-based manager. For more information on this feature, see [“Feature Select” on page 243](#).

Changes to Endpoint Control

There have been several changes to Endpoint Control.

Endpoint control for Android

FortiOS now supports endpoint control for Android mobile devices. Endpoint profiles that include Android devices can be configured at *User & Device > Endpoint Protection > FortiClient Profiles*.

Figure 11:Endpoint control for Android

Android

Web Category Filtering

Disable Web Category Filtering when protected by this FortiGate

Client VPN Provisioning

VPN Name	<input type="text"/>
Type	<input checked="" type="radio"/> IPsec VPN <input type="radio"/> SSL-VPN
Remote Gateway	<input type="text"/>
Authentication Method	<input type="text" value="Preshared Key"/>
Preshared Key	<input type="text" value="....."/>

Assigning endpoint profiles to specific users and user groups

Endpoint profiles can now be assigned to specific users and user groups that have been defined on the FortiGate. This can be done from the CLI, using the following command:

```
config endpoint-control profile
  edit <profile-name>
    set users <user-name>
    set <user-groups> <user-group-name>
  end
```

Endpoint profile portal pages

Custom endpoint profile portal pages can be configured. There are five different portals that can be used, depending on the operating system of the endpoint device. The five portals are: Android, Mac, iOS, Windows and other. To access the portal pages, go to *System > Config > Replacement Messages* and select *Extended View*.

Managing FortiAP units

There have been several changes to how FortiAP units are managed by a FortiGate unit.

Firmware Auto-detection

A FortiGate unit now auto-detects what the best firmware version is for the FortiAP units that it manages. If the FortiAP unit is not running the recommended firmware version you can download and install in from the FortiGate web-based manager.

Wireless Device Locating Service

A FortiAP unit can be configured to report all wireless devices that it locates, even if the device does not connect, or is unable to connect, to the FortiAP.

Locating service is enabled from the CLI, using the following command:

```
config wireless-controller wtp-profile
  edit "FAP220B-default"
    set ap-country JP
    config radio-1
      set station-locate enable
    end
    config radio-2
      set station-locate enable
    end
  end
end
```

After this configuration is complete, the list of devices can be found using the following command:

```
diagnose wireless-controller wlac -c sta-locate
```

The command displays a list of currently located wireless devices. The list includes the MAC address of each device as well as wireless-related information about the device.

More Wireless Controller MIB Support

More fields related to wireless controller functionality has been added to the FortiGate MIB. Additions include the following:

- Asynchronous notifications from SNMP agent, including fgTrapWcApUp and fgTrapWcApDown.
- Objects defined for controller level information, for example: controller name and location, WTP capacity and count and station capacity and count.
- A set of objects that display a WLAN interface, for example: assigned SSID and security method.
- An object identifier for a list of tables pertaining to WTPs.
- A set of objects that display a custom WTP profile, for example: profile name, platform type, DTLS policy and country code.
- A set of objects that display a radio in a custom WTP profile, for example: radio mode, band and channel settings, power level and VAP configurations.
- A set of objects that display the configuration of a WTP, for example: WTP ID, name and assigned custom profile. If no custom profile is assigned, then a list of objects that supplement the automatic profile are defined.
- A set of objects that display wireless session information, for example: IP/MAC address, connection state, up time, profile name, WTP HW/SW information, WTP session statistics, CPU load and memory capacity and usage.
- A set of objects that display wireless session radio information, for example: radio mode, operating country code, operating channel, operating power level and client count.
- A set of objects that display a virtual access point (a WLAN allocated on a WTP radio), for example: client count and RX/TX byte counts.

A set of objects that display a wireless station, for example: WTP and radio it connects to, IP/MAC address, VCI/host information, signal/noise level, TX/RX bandwidth, channel, security type and on-line status.

Normal or Remote WTP mode parameter

A new WTP mode parameter has been added in which FortiAP units are classified as either normal or remote. A FortiAP unit in normal mode uses SSID in tunneled mode while remote WTP mode uses only local bridge SSIDs.

This new mode has changed the maximum number of FortiAP units which can be managed by a FortiGate unit, with one value for the maximum number of normal FortiAPs and another for the maximum number of remote FortiAPs. For more information, see the [Maximum Values Table for FortiOS 5.0](#).

FortiGuard Subscription Services

The FortiGuard Subscription Services have been reorganized into three categories: Next Generation Firewall, ATP Services and Other Services.

Figure 12:The FortiGuard Subscription Services

FortiGuard Subscription Services		
Next Generation Firewall		
IPS & Application Control	Valid License (Expires 2014-02-25)	✔
IPS Definitions	4.00345 (Updated 2013-05-23 via Manual Update) [Update]	
IPS Engine	2.00153 (Updated 2013-05-31 via Manual Update)	
<hr/>		
ATP Services		
AntiVirus	Valid License (Expires 2014-02-25)	✔
AV Definitions	1.00000 (Updated 2012-10-17 via Manual Update) [Update]	
AV Engine	5.00146 (Updated 2013-05-21 via Manual Update)	
Web Filtering	Valid License (Expires 2014-02-24)	✔
<hr/>		
Other Services		
Vulnerability Scan	Valid License (Expires 2014-02-25)	✔
		(2013-06-12)
VCM Plugins	1.00316 (Updated 2013-06-12 via Manual Update) [Update]	
Email Filtering	Valid License (Expires 2014-02-24)	✔
Messaging Services	Registered (Expires 2014-02-08)	✔
SMS Messages	4 Allowed (0 Used)	
<hr/>		

Adding Explicit Web Proxy services

Explicit proxy services can now be added and edited by going to *Firewall Objects > Service > Services > Create New > Custom Service* and selecting *Explicit Proxy*. Explicit web proxy services are used in security policies that control access to the explicit web proxy.

Figure 13:Editing an Explicit Web Proxy Service

The screenshot shows the 'Edit Service' configuration window. The 'Name' field is set to 'explicit_proxy'. The 'Service Type' is 'Explicit Proxy'. The 'Protocol' is 'TCP' with a destination port of 80. The 'Destination Port' is set to 'Low' 80 and 'High' is empty. The 'Source Port' is set to 'Low' and 'High' are empty. The 'Category' is 'General' and 'Protocol Type' is 'ALL'. There are 'OK' and 'Cancel' buttons at the bottom.

SSO Authentication failover for the Explicit Web Proxy

SSO authentication failover for the Explicit Web Proxy is now available, allowing two authentication methods to be configured. If the Single Sing-On Method fails, the FortiGate unit will use the *Default Authentication Method*.

To configure failover, the explicit web proxy must first be enabled.

Figure 14:Configuring SSO authentication failover for the explicit web proxy

The screenshot shows the policy configuration window. The 'Policy Type' is 'Firewall'. The 'Policy Subtype' is 'User Identity'. The 'Incoming Interface' is 'web-proxy', 'Source Address' is 'all', 'Outgoing Interface' is 'any', and 'Destination Address' is 'all'. The 'Service' is 'webproxy'. The 'Web Proxy Forwarding Server' checkbox is unchecked. Below this is the 'Configure Authentication Rules' table:

User/Group	Schedule	Security	Traffic Shaping	Logging
ANY	always	-	X	X

Below the table are the 'Explicit Proxy Authentication Options':

- Enable IP Based Authentication
- Single Sign-On Method: Fortinet Single Sign-On (FSSO)
- Default Authentication Method: Basic
- Skip this policy for unauthenticated user
- Disclaimer
- Customize Authentication Messages

The 'Comments' field is empty. There are 'OK' and 'Cancel' buttons at the bottom.

User Creation Wizard

The User Creation Wizard is used to create new users through a four step process. The four steps will vary depending on which type of user is being created (for example, when creating an LDAP user, step 3 requires choosing an LDAP filter).

Figure 15:The User Creation Wizard

The User Creation Wizard can be found at *User & Device > User > User Definition*.

FortiClient Registration

The FortiClient registration process has changed so that the initial confirmation message sent from FortiClient will be ignored by the FortiGate unit and applied only at the end of the registration process, to avoid the registration being rejected.

DSS and ECDSA Certificates for FortiGate SSL-related features

FortiOS now supports DSS and ECDSA certificates for the following features: HTTPS/SSL deep scanning, HTTPS/SSL server load balancing, HTTPS/SSL offloading and HTTPS over the explicit web proxy.

LDAP Servers

A Distinguished Name field and query button have been added to the LDAP Server creation page.

User Monitor

FSSO Logons are now shown in the user monitor, found at *User & Device > Monitor > Firewall*. In order for FSSO Logons to appear, *Show all FSSO Logons* must be enabled.

Web Filter Profiles

URL filters, used for website filtering, are now created as part of a Web Filter Profile. This can be done by going to *Security Profiles > Web Filter > Profiles* and selecting *Enable Web Site Filter*. A filter can then be enabled for all URLs you wish to block.

CAPWAP Administrative Access

CAPWAP Administrative Access can now be configured for all interfaces except Virtual Access Points (VAPs). CAPWAP must be used on any interface used to managed a FortiAP unit.

IPS Algorithms

There is a new algorithms for IPS, “super” mode, that improves performance for FortiGate units with more than 4GB of memory. Improvements have also been made for “low” mode, which is more efficient for FortiGate units with low memory.

The algorithm used for IPS can be changed from the CLI, using the following command:

```
config ips global
  set algorithm {engine-pick | high | low | super}
```

NAC-Quarantine Traffic Logs

Antivirus and DLP NAC-quarantine traffic logs now show whether the IP, user or interface has been banned. In the case of a virus, the name of the virus and the file in which the virus was found are also included.

New System Report Charts

The following charts have been added to the daily FortiGate System Report, based on data collected by event logs:

- VPN Usage
- Client Reputation Summary

Memory Logging

Memory logging is available on all FortiGate models. Logging can be enabled by going to *Log & Report > Log Config > Log Settings* and enabling *Disk*.

Logging to flash is also available for the FortiGate-60D and FortiWiFi-60D.

URL-based Web Proxy Forwarding

In order to configure URL-based web proxy forwarding, *WAN Opt. & Cache* must be enabled using Feature Setting. For more information, see [“Feature Select” on page 243](#).

FortiOS now supports URL-based web proxy forwarded, which is required for explicit proxy installations using Threat Management Gateway or Blue Coat.

URL-based web proxy forwarding can be configured by going to *WAN Opt. & Cache > Cache > URL Match List*.

Changes to Routing

The following changes have occurred for FortiOS Routing:

- The OSPF summary address limit has decreased to 25 from 10.
- More routing community lists can be configured (limits vary by FortiGate model).

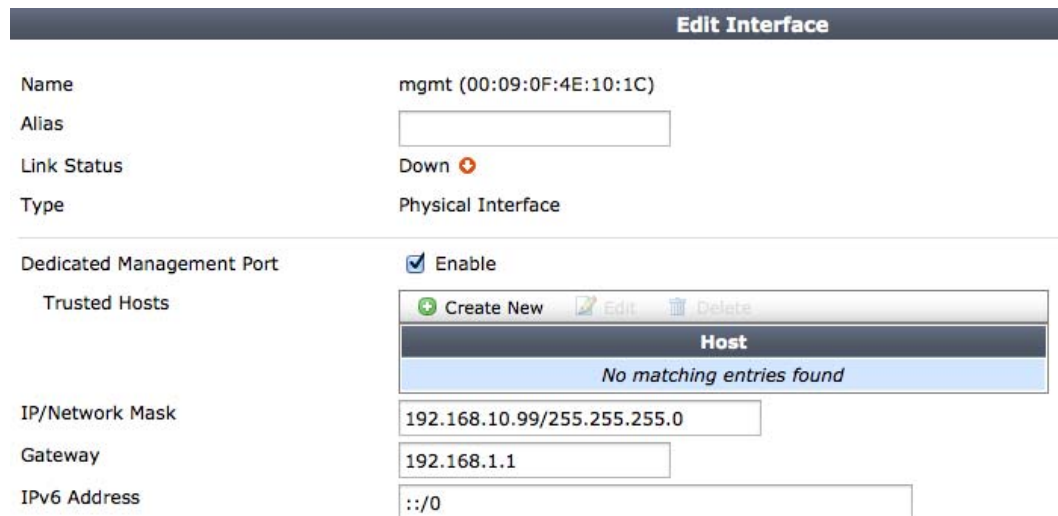
RADIUS Support for Dynamic VLANs

RADIUS authentication can be used to dynamically route authenticated user traffic to a VLAN. The name of this VLAN is then added to the user’s RADIUS record.

Dedicated Management Port

The Management (MGMT) port can now be set as Dedicated to Management, in which case no firewall traffic will be allowed through this port.

Figure 16:MGMT port set to Dedicated to Management



URL Filtering

URL filtering has changed to allow certificate-based URL filtering for HTTPS traffic, which is used when deep-scan is disabled.

Certificate-based filtering extracts the hostname from the TLS handshake. If a valid hostname is found, it is used for the local or FortiGuard category query. If no hostname is found, HTTPS server CN web filtering will be used instead.

URL Source Tracking

URL source tracking has been added to transparent proxy and SSL deep-inspection proxy. There are three current URL source values: HTTP host head, subject CN in the certificate and server name field in the TLS handshake.

IPv6 Denial of Service Policies

Denial of Service (DoS) policies can now be configured by going to *Policy > Policy > IPv6 Dos Policy*.

Support for NAT46, VIP64 and VIP46

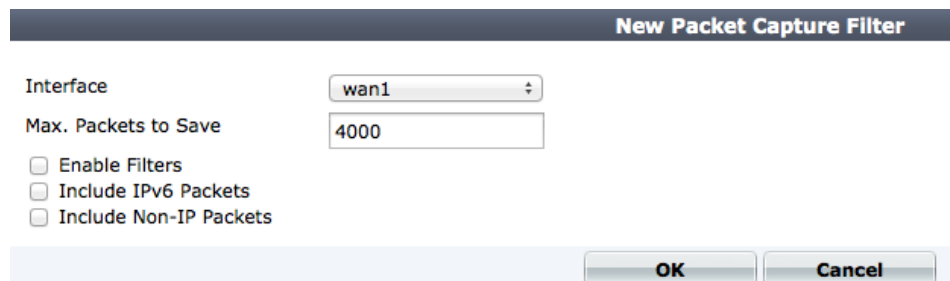
Security policies are now supported for NAT46, VIP64 and VIP46. They can be configured from the CLI, using the following commands:

```
config firewall policy46
config firewall vip64
config firewall vip46
```

Packet Capture Filters

Packet capture filters can be configured by going to *System > Network > Packet Capture*.

Figure 17:Configuring a packet capture filter



The screenshot shows a configuration window titled "New Packet Capture Filter". It includes the following elements:

- Interface:** A dropdown menu currently set to "wan1".
- Max. Packets to Save:** A text input field containing the value "4000".
- Options:** Three unchecked checkboxes: "Enable Filters", "Include IPv6 Packets", and "Include Non-IP Packets".
- Buttons:** "OK" and "Cancel" buttons at the bottom right.

Configure hosts in an SNMP v1/2c community to send queries or receive traps

When you add a host to an SNMP v1/2c community you can now decide whether the FortiGate unit will accept queries from the host or whether the FortiGate unit will send traps to the host. You can also configure the host for both traps and queries.

Use the following command to add two hosts to an SNMP community:

- An IPv4 host that can send queries to the FortiGate unit
- An IPv6 host that the FortiGate unit will send traps to

```
config system snmp community
  config hosts
    edit 1
      set interface port1
      set ip 172.20.120.1
      set host-type query
    end
  config hosts6
    edit 1
      set interface port6
      set ip 2001:db8:0:2::30
      set host-type trap
    end
```

You can add up to 16 IPv4 hosts and up to 16 IPv6 hosts.

IP in IP tunneling support (RFC 1853)

FortiOS 5.0 MR3 supports [RFC 1853 IP in IP Tunneling](#) to provide for differential forwarding to packets. This tunneling mechanism is supported as an interface type and the FortiGate unit does not inspect the traffic in an IP in IP tunnel

To configure IP in IP tunneling:

```
config system ipip-tunnel
  edit tun0
    set interface <interface_name>
    set local-gw <local-gw-ip>
    set remote-gw <remote-gw-ip>
  end
end
```

GTP-u acceleration on FortiGate units with SP3 processors

FortiGate units with SP3 processors can offload GTP-u packet processing to their SP3 processor. The SP3 processor supports:

- GTP-u sanity packet check
- GTP-u rate limiting per gtp profile
- Encapsulated IP traffic filtering

New features in FortiOS 5.0 Patch 2

This chapter provides a brief introduction to the following features that were added to Patch 2 of FortiOS 5.0. See the release notes for a complete list of new features in this release.

- [Endpoint Profile Changes](#)
- [Client Reputation Changes](#)
- [Changes to logging in security policies](#)
- [Configuring the FortiGate unit to be an NTP Server](#)
- [Customizing and viewing the local FortiGate UTM Security Analysis Report](#)
- [Wireless changes: Custom mesh downlink SSIDs and new identifier for local bridge SSIDs](#)
- [SSL-VPN Realm Support \(multiple custom SSL VPN logins\)](#)
- [Automatically add devices found by device identification to the vulnerability scanner configuration](#)
- [The SIP ALG can receive SIP traffic on multiple TCP and UDP ports](#)
- [IPv6 PIM sparse mode multicast routing](#)
- [Wireless RADIUS-Based MAC Authentication](#)

Endpoint Profile Changes

A number of changes have been made to enhance Endpoint Profile functionality.

Client Reputation Changes

A number of changes have been made to enhance Client Reputation functionality. For information about how to configure and use Client Reputation, see [“Client Reputation” on page 181](#).

Changes to logging in security policies

Instead of enabling or disabling traffic logging in security policies three Logging Options are now available:

- *No Log*, do not record log messages about traffic accepted by this security policy
- *Log UTM Events*, record traffic log messages when a UTM event occurs (such as when a virus is found by antivirus, a web page is blocked by web filtering, or the application responsible for a session is identified by application control).
- *Log all Sessions*, record traffic log messages for all sessions. For all sessions, a single traffic log message is recorded when the session ends. If you select this option, you can choose to record a traffic log message when a session starts as well. You can also choose to capture packets.

Enabling logging in a security policy can affect FortiGate performance because of the extra system resources required to record log messages. The performance hit can be reduced by selecting *Log UTM Events*, since fewer log messages will be recorded.

You can also enter the following command to write a log message when a session starts:

```
config firewall policy
  edit <policy-index>
    set logtraffic-start
  end
```

Configuring the FortiGate unit to be an NTP Server

When you configure system time from the System Information dashboard widget, you can configure the FortiGate unit to be an NTP server. As part of the NTP server configuration, you can select one or more interfaces on which to listen for NTP requests.

Figure 18:System time configuration: NTP server

The screenshot displays the 'Time Settings' configuration page. At the top, the 'System Time' is shown as 'Fri Mar 15 13:41:55 2013' with a 'Refresh' button. Below this, the 'Time Zone' is set to '(GMT-8:00)Pacific Time(US&Canada)'. There are two main sections: 'Set Time' and 'Synchronize with NTP Server'. The 'Set Time' section has spinners for Hour (13), Minute (41), Second (55), Year (2013), Month (Mar), and Day (15). The 'Synchronize with NTP Server' section is selected, with 'Use FortiGuard Servers' chosen over 'Specify'. The 'Sync Interval' is set to 60 minutes. The 'Enable NTP Server' checkbox is checked. The 'Listen on Interfaces' field contains 'port1' and 'port2'. At the bottom, there are 'OK' and 'Cancel' buttons.

Customizing and viewing the local FortiGate UTM Security Analysis Report

In order for your FortiGate unit to create a Security Analysis Report, disk logging must be enabled. To enable disk logging, go to *Log & Report > Log Config > Log Settings* and under *Logging and Archiving* select *Disk* and *Enable Local Reports*.

You can go to *Log & Report > Report > Local* to view Local Reports created by the FortiGate unit. Local reports are saved as PDF files that you can view and download at any time.

By default, Local reports are produced every day. You customize how the FortiGate unit to produce reports daily, weekly, or on demand and you can set the day and time when the report is generated. You can control how many users appear in the Top Users by bandwidth summary part of the report. Each user gets a separate summary page. You can also configure the FortiGate unit to email the report to multiple email recipients.

Figure 19:Customizing the FortiGate report

Report Options

Generate Report:

Time:

Top Users By Bandwidth:

Email Generated Reports

Historical Reports

Report Name	Data Range	Size
Schedule-default-2013-03-18-000059	Mar 17, 03:00 AM - Mar 18, 02:59 AM	247.57 KB
Schedule-default-2013-03-17-000059	Mar 16, 03:00 AM - Mar 17, 02:59 AM	247.57 KB
Schedule-default-2013-03-16-000101	Mar 15, 03:00 AM - Mar 16, 02:59 AM	247.57 KB
On-Demand-default-2013-03-15-133320	Mar 14, 04:00 PM - Mar 15, 03:59 PM	247.58 KB
Schedule-default-2013-03-15-000101	Mar 14, 03:00 AM - Mar 15, 02:59 AM	380.11 KB
Schedule-default-2013-03-14-000101	Mar 13, 03:00 AM - Mar 14, 02:59 AM	247.58 KB
Schedule-default-2013-03-13-000100	Completed: Mar 12, 08:01 PM	310.77 KB
Schedule-default-2013-03-12-000100	Completed: Mar 11, 08:01 PM	310.77 KB
Schedule-default-2013-03-10-230100	Completed: Mar 10, 08:01 PM	310.76 KB
Schedule-default-2013-03-10-000100	Completed: Mar 09, 07:01 PM	310.76 KB

80 reports hidden (show all)

You can also select *Run Now* to run a report at any time. The report is created using current data.

You can select *Customize* to change the report layout. You can customize a report to add headings and text, divide a report into sections, add images, and add, remove, and rearrange the individual report charts.

Wireless changes: Custom mesh downlink SSIDs and new identifier for local bridge SSIDs

You can go to *WiFi Controller > WiFi Network > SSID* and select *Create New* to add additional custom mesh downlink SSIDs.

Figure 20:SSID list showing a local-bridge SSID and two mesh downlink SSIDs

SSID	Administrative Status	Traffic Mode	Security Mode	Data Encryption
my-local-bridge	On	Local bridge	WPA/WPA2-Personal	AES
fortinet.mesh.root	On	Mesh Downlink	WPA/WPA2-Personal	AES
mesh-dl-2	On	Mesh Downlink	WPA/WPA2-Personal	AES

Configure a mesh downlink SSID by selecting *Create New*, setting the *Traffic Mode* to *Mesh Downlink* and entering an SSID.

Figure 21:Configuring a custom mesh downlink

SSL-VPN Realm Support (multiple custom SSL VPN logins)

In order to create a custom login page using the web-based manager, this feature must be enabled using Feature Select. For more information, see [“Feature Select” on page 243](#).

You configure a custom SSL VPN login by going to *VPN > SSL > Custom Login* and selecting *Create New*. Users access different portals depending on the URL they enter. The first option in the custom login page is to enter the path of the custom URL. This path is appended to the address of the FortiGate unit interface that SSL VPN users connect to. The actual path for the custom login page appears beside the URL path field. You can also limit the number of users that can access the custom login at any one time. Finally you can use HTML code to customize the appearance of the login page.

Figure 22:Custom SSL VPN login

After adding the custom login, you must associate it with the users that will access the custom login. Do this by going to *Policy > Policy > Policy* and creating an SSL VPN policy. Add an *Authentication Rule* to the policy and select the users and user groups who should access the custom login page. Select *Custom Login* and select the custom login page that you created.

Figure 23: Associating a custom SSL VPN login with a user group

New SSL VPN Authentication Rule	
Group(s)	My-Portal-Users
User(s)	Click to add...
Schedule	always
SSL-VPN Portal	full-access
<input checked="" type="checkbox"/> Custom Login	our-portal
Action	ACCEPT

Automatically add devices found by device identification to the vulnerability scanner configuration

When you go to *System > Network > Interfaces* to configure an interface to detect and identify devices you can also select *Add New Devices to Vulnerability Scan List*. As devices are found, they are added to the Asset Definitions list of the Vulnerability scanner. You can choose to run a scan of all of the devices on the list or of selected items, including devices found using FortiGate device identification.

The SIP ALG can receive SIP traffic on multiple TCP and UDP ports

You also configure the SIP ALG to listen in two different TCP ports and two different UDP ports for SIP sessions. For example, if you receive SIP TCP traffic on port 5060 and 5064 and UDP traffic on ports 5061 and 5065, you can enter the following command to receive the SIP traffic on all of these ports:

```
config system settings
    set sip-tcp-port 5060 5064
    set sip-udp-port 5061 5065
end
```


IPv6 PIM sparse mode multicast routing

FortiOS supports PIM sparse mode multicast routing for IPv6 multicast (multicast6) traffic and is compliant with [RFC 4601: Protocol Independent Multicast - Sparse Mode \(PIM-SM\)](#). You can use the following command to configure IPv6 PIM sparse multicast routing.

```
config router multicast6
  set multicast-routing {enable | disable}
  config interface
    edit <interface-name>
      set hello-interval <1-65535 seconds>
      set hello-holdtime <1-65535 seconds>
    end
  config pim-sm-global
    config rp-address
      edit <index>
        set ipv6-address <ipv6-address>
      end
    end
```

The following diagnose commands for IPv6 PIM sparse mode are also available:

```
diagnose ipv6 multicast status
diagnose ipv6 multicast vif
diagnose ipv6 multicast mroute
```

Wireless RADIUS-Based MAC Authentication

Wireless clients can be supplementally authenticated by MAC address. A RADIUS server stores the allowed MAC address for each client and the wireless controller checks the MAC address independently of other authentication methods.

MAC-based authentication must be configured in the CLI. In the following example, MAC-based authentication is added to an existing access point *vap1* to use a RADIUS server at 192.168.1.95:

```
config wireless-controller vap
  edit vap1
    set radius-mac-auth enable
    set radius-mac-auth-server 192.168.1.95
  end
```

Security Features

Features previously known as UTM Security Features are now known as Security Features.

In order to create new profiles for the Security Features, Multiple Security Profiles must be enabled using Feature Select. For more information, see [“Feature Select” on page 243](#).

New Security Features in FortiOS 5.0 include:

- FortiSandbox
- Botnet and phishing protection
- Windows file sharing (CIFS) flow-based antivirus scanning
- Advanced Application Control and IPS sensor creation
- Custom Application Control signatures and IPS signatures
- Flow-based inspection improvements
- Configuring SSL inspection for flow-based and proxy protection
- Explicit web Proxy Extensions – SSL inspection, IPS, Application Control, and flow-based antivirus, web filtering and DLP
- Replacement messages for flow-based web filtering of HTTPS traffic
- DNS web filtering
- FortiGuard Web Filter quotas can be set based on traffic volume
- Customizing the authentication replacement message for a FortiGuard web filter category
- YouTube Education Filter implemented in Web Filtering Profiles
- IPS hardware acceleration
- New SIP ALG features
- DLP watermarking
- SSH inspection
- Optimizing SSL encryption/decryption performance

FortiSandbox

The new FortiSandbox unit is used for automated sample tracking, or sandboxing, for files from a FortiGate unit. This allows suspicious files to be sent to be inspected without risking network security. If the file exhibits risky behavior, or is found to contain a virus, a new virus signature is created and added to the FortiGuard antivirus signature database.

Cloud Sandbox, formerly known as FortiGuard Analytics, can also be used for sandboxing if you have an active FortiCloud subscription.

Configuration

FortiSandbox is configured by going to *System > Config > FortiSandbox*. After enabling FortiSandbox, select either: *FortiSandbox Appliance* or *Cloud Sandbox (FortiCloud)*.

Figure 24:FortiSandbox Configuration

Sandbox Settings

Enable Sandbox Inspection

FortiSandbox Appliance

IP Address

Notifier Email

Cloud Sandbox (FortiCloud)

Sending files to FortiSandbox

An anti-virus profile can be set up to send files to FortiSandbox. To do this, edit the profile being used and enable *Send Files to FortiSandbox for Inspection*.

Figure 25:An anti-virus profile using FortiSandbox

Name

Comments 21/255

Inspection Mode Proxy Flow-based

Send Files to FortiGuard Sandbox for Inspection

Suspicious Files Only

Suspicious + Clean Files

Tracking submitted files

The *Advanced Threat Protection Statistics* widget shows the number of files that have been submitted to FortiSandbox and the inspection results.

Figure 26:The Advanced Threat Protection Statistics widget

Advanced Threat Protection Statistics	
FortiGate Statistics	
Number of Files Scanned	91648
Detected Malware	6
Detected Zero-Day Malware Variants	0
Suspicious Files	0
Clean	91642
FortiGuard Sandbox Statistics (Last 7 Days)	
# of Files Submitted to FortiGuard Sandbox	0
Detected Malware	0
Clean	0

Botnet and phishing protection

In a proxy or flow-based antivirus profile, you can configure the FortiGate unit to detect and block botnet server connection attempts. This feature also blocks attempted access to phishing URLs. The antivirus database is constantly updated with the addresses of known command and control (C&C) sites that Botnet clients attempt to connect to as well as phishing URLs.

To enable Botnet and phishing protection in either a proxy or flow-based antivirus profile, select *Block Connections to Botnet Servers*.

Figure 27: Adding Botnet and phishing protection to a flow-based antivirus profile

Edit AntiVirus Profile
default

Name:

Comments: 21/255

Inspection Mode: Proxy Flow-based

Block Connections to Botnet Servers

Protocol	Virus Scan and Removal
Web	
HTTP	<input checked="" type="checkbox"/>
Email	
SMTP	<input checked="" type="checkbox"/>
POP3	<input checked="" type="checkbox"/>
IMAP	<input checked="" type="checkbox"/>
MAPI	<input type="checkbox"/>
File Transfer	
FTP	<input checked="" type="checkbox"/>
SMB	<input type="checkbox"/>
IM	
ICQ, Yahoo, MSN Messenger	<input checked="" type="checkbox"/>

Windows file sharing (CIFS) flow-based antivirus scanning

FortiOS 5.0 now supports virus scanning of Windows file sharing traffic. This includes CIFS, SMB and SAMBA traffic. This feature is applied by enabling SMB scanning in an antivirus profile and then adding this profile to a security policy that accepts CIFS traffic. CIFS virus scanning is available only through flow-based antivirus scanning.

FortiOS 5.0 flow-based virus scanning can detect the same number of viruses in CIFS/SMB/SAMBA traffic as it can for all supported content protocols.

Figure 28: Configuring CIFS/SMB/SAMBA virus scanning

New AntiVirus Profile

Name:

Comments: 48/255

Inspection Mode: Proxy Flow-based

Block Connections to Botnet Servers

Protocol	Virus Scan and Removal
Web	
HTTP	<input type="checkbox"/>
Email	
SMTP	<input type="checkbox"/>
POP3	<input type="checkbox"/>
IMAP	<input type="checkbox"/>
MAPI	<input type="checkbox"/>
File Transfer	
FTP	<input type="checkbox"/>
SMB	<input checked="" type="checkbox"/>
IM	
ICQ, Yahoo, MSN Messenger	<input type="checkbox"/>

Use the following command to enable CIFS/SMB/SAMBA virus scanning in an antivirus profile:

```
config antivirus profile
  edit smb-profile
    config smb
      set options scan
      set avdb flow-based
    end
```

Then add this antivirus profile to a security policy that accepts the traffic to be virus scanned. In the security policy the service can be set to ANY, SAMBA, or SMB.

```
config firewall policy
  edit 0
    set service ANY
    ...
    set utm-status enable
    set av-profile smb-profile
  end
```

Note the following about CIFS/SMB/SAMBA virus scanning:

- Some newer version of SAMBA clients and SMB2 can spread one file across multiple sessions, preventing some viruses from being detected.
- Enabling CIFS/SMB/SAMBA virus scanning can affect FortiGate performance.
- SMB2 is a new version of SMB that was first partially implemented in Windows Vista. Currently SMB2 is supported by Windows Vista or later, partly supported by Samba 3.5 and fully support by Samba 3.6.
- The latest version of SMB2.2 will be introduced with Windows 8.
- Most clients still use SMB as default setting.

Advanced Application Control and IPS sensor creation

In FortiOS 5.0, it is much easier to sort through Fortinet’s thousands of application definitions and IPS signatures to find the ones that you want to add to Application Control and IPS sensors. The creation pages for both of these features include filters for severity, category, popularity, technology and risk.

Figure 29:Application list filtering

New Application Filter

Sensor Type Filter Based Specify Applications

[\[Filter Options\]](#)

Category

 Botnet
 Game
 Media
 Proxy
 Storage.Backup

eMail
 General.Interest
 Network.Service
 Remote.Access
 Update

File.Sharing
 IM
 P2P
 Social.Networking
 VoIP

Popularity

 ★★★★★
 ★★★★☆
 ★★★☆☆
 ★★☆☆☆
 ★☆☆☆☆
 ☆☆☆☆☆

Technology

 Browser-Based
 Client-Server
 Network-Protocol
 Peer-to-Peer

Risk

 Botnet
 Excessive-Bandwidth
 None

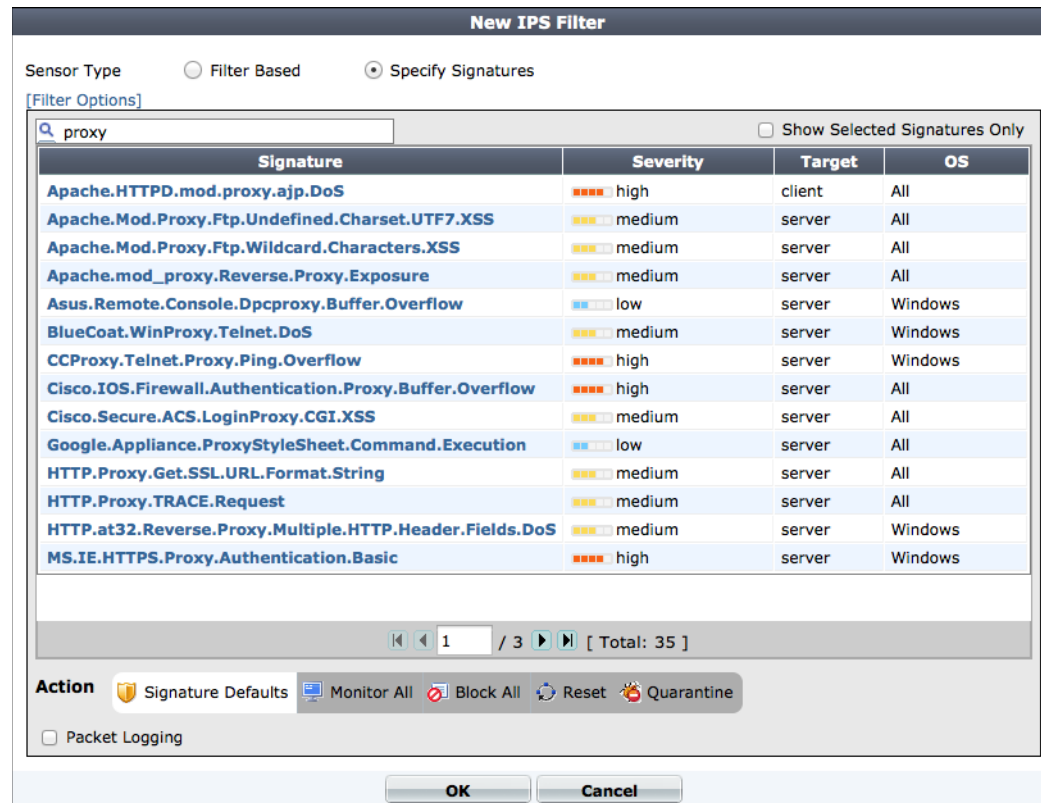
Application Name	Category	Technology	Popularity	Risk
012mail	eMail	Browser-Based	★★★★☆	
Ozz0	File.Sharing	Browser-Based	★★★★☆	Excessive-Bandwidth
1und1.Mail	eMail	Browser-Based	★★★★☆	Excessive-Bandwidth
2Shared.Browse.Upload.File	File.Sharing	Browser-Based	★★★★☆	Excessive-Bandwidth
2Shared.Search.Download.File	File.Sharing	Browser-Based	★★★★☆	Excessive-Bandwidth
2ch	Social.Networking	Browser-Based	★★★★☆	
2ch_Post	Social.Networking	Browser-Based	★★★★☆	
3PC	Network.Service	Network-Protocol	★★★★☆	
4shared	File.Sharing	Browser-Based	★★★★☆	Excessive-Bandwidth
6cn	Media	Browser-Based	★★★★☆	Excessive-Bandwidth
9PFS	Network.Service	Network-Protocol	★★★★☆	
9PTV	P2P	Peer-to-Peer	★★★★☆	Excessive-Bandwidth
24im	IM	Client-Server	★★★★☆	Excessive-Bandwidth
51.Com	Social.Networking	Browser-Based	★★★★☆	
51.Com_BBS	Social.Networking	Browser-Based	★★★★☆	Excessive-Bandwidth

1 / 153 [Total: 2284]

Action

You can also search through the application or signature list by name.

Figure 30:IPS signatures search example



Custom Application Control signatures and IPS signatures

The application control and IPS signatures provide coverage for most applications and network vulnerabilities. You can extend the coverage by adding custom application signatures and custom IPS signatures.

You add custom application signatures by going to *Security Policies > Application Control > Application List* and selecting *Create New*.

You add custom IPS signatures by going to *Security Policies > Intrusion Protection > IPS Signatures* and selecting *Create New*.

Custom application signatures and custom IPS signatures use the same syntax. See the [UTM Guide](#) for a description the signature syntax.

Figure 31:Example custom application signature

Use the following command to add a custom application control signature.

```
config application custom
  edit New-custom-sig
    set signature F-SBID( --attack_id 8640; --name "Block.WMP.Get";
      --default_action drop_session; --protocol tcp; --service
      HTTP; --flow from_client; --pattern "Pragma: xPlayStrm=1";
    )
  end
```

Use the following command to add a custom IPS signature.

```
config ips custom
  edit New-custom-sig
    set signature F-SBID( --attack_id 8640; --name "Block.WMP.Get";
      --default_action drop_session; --protocol tcp; --service
      HTTP; --flow from_client; --pattern "Pragma: xPlayStrm=1";
    )
  end
```

Exempting IP addresses from IPS

IPS filters can be configured so that signatures are not applied to traffic from specific IP addresses. To exempt an IP address, a filter must be created with the *Sensor Type* set to *Specify Signatures*. *Exempt IP* can then be enabled and the necessary exemptions configured using Source and Destination IPs.

Figure 32:Exempting IP addresses from IPS

The screenshot shows the 'Action' configuration page for IPS. At the top, there are tabs for 'Signature Defaults', 'Monitor All', 'Block All', 'Reset', and 'Quarantine'. Below these, there are checkboxes for 'Packet Logging' (unchecked) and 'Exempt IP' (checked). A table below the checkboxes has columns for 'Source IP/Netmask' and 'Destination IP/Netmask', both containing '0.0.0.0/0.0.0.0'. Above the table are buttons for 'Create New', 'Edit', and 'Delete'. At the bottom right, there are 'OK' and 'Cancel' buttons.

Source IP/Netmask	Destination IP/Netmask
0.0.0.0/0.0.0.0	0.0.0.0/0.0.0.0

Flow-based inspection improvements

If your FortiGate unit supports flow-based scanning, you can choose to select flow-based instead of proxy scanning. Flow-based scanning uses the FortiGate IPS engine to examine network traffic as it passes through the FortiGate unit without using a proxy to buffer and potentially change packets.

In FortiOS 5.0, flow-based inspection has been extended to email filtering. The following sections contain information about additional flow-based scanning improvements.

Configuring SSL inspection for flow-based and proxy protection

FortiOS 5.0 fully supports flow-based inspection of SSL sessions. This means that HTTPS, IMAPS, POP3S, SMTPS and FTPS traffic can now be decrypted and inspected by IPS and application control and flow-based antivirus, web filtering and email filtering.

FortiOS 5.0 continues to fully support proxy inspection of SSL sessions. In FortiOS 5.0, configuring proxy SSL inspection has changed as described below.

To enable proxy or flow-based inspection of SSL sessions, you must add an SSL/SSH Inspection profile to a security policy. You can configure SSL/SSH inspection profiles to inspect HTTPS, SMTPS, POP3S, IMAPS and FTPS traffic, as well as SSH traffic. You can configure the profile to control which SSL protocols to inspect, the ports to inspect for each protocol and the certificate to use with SSL sessions.

To apply proxy virus scanning and web filtering to HTTPS, IMAPS, POP3S, SMTPS and FTPS sessions

1. Go to *Policy > Policy > SSL/SSH Inspection* and create or edit an SSL/SSH inspection profile.
2. Under *SSL Inspection Options* select the CA certificate to use for SSL sessions. You can import a new certificate or use one already imported into the FortiGate unit.
3. Under enable the SSL protocols that you want to inspect and set the ports to inspect for each protocol.
4. Configure other settings as required and select *Apply* to save your changes.
5. Go to *Policy > Policy > Policy* and create a new or edit a policy that accepts the SSL traffic to be inspected.
6. Under *Security Profiles*, turn on *AntiVirus* and *Web Filter* and select profiles for them.
7. Turn on *SSL/SSH Inspection* and select the SSL/SSH inspection profile that you configured.
8. Select OK.

To apply flow-based virus scanning and web filtering and application control to HTTPS, and POP3S sessions

This example describes adding factory default antivirus, web filtering, application control and SSL/SSH profiles to a security policy that accepts HTTPS and POP3S traffic to apply flow-based virus scanning, web filtering and application control to the HTTPS and POP3S traffic accepted by the security policy.

1. Go to *Policy > Policy > Policy* and create or edit a policy that accepts the HTTPS and POP3S traffic to be inspected.
2. Under *Security Profiles*, turn on *AntiVirus* and select the *AV-flow* profile.
3. Turn on *Web Filter* and select the *web-filter-flow* profile.
4. Turn on *Application Control* and select the *default* profile.
5. Turn on *SSL/SSH Inspection* and select the *default* profile.
6. Select *OK*.

Explicit web Proxy Extensions – SSL inspection, IPS, Application Control, and flow-based antivirus, web filtering and DLP

FortiOS 5.0 fully supports SSL inspection of explicit web proxy traffic. This means that HTTPS traffic accepted by the explicit web proxy can now be subject to deep inspection for antivirus, web filtering and DLP.

FortiOS 5.0 also fully supports flow-based inspection of explicit web proxy traffic. This includes full support for IPS and application control, as well as flow-based virus scanning and web filtering for HTTP, HTTPS and FTP over HTTP traffic.

SSL content inspection and flow-based inspection are added to explicit web proxy sessions by enabling Security Profiles in a security policy that accepts web-proxy traffic and then selecting profiles that implement flow-based inspection for the features you need.

The explicit FTP proxy and the IPv6 explicit web proxy do not support SSL inspection or IPS, application control, and flow-based antivirus, web filtering and DLP.

Replacement messages for flow-based web filtering of HTTPS traffic

FortiOS 5.0 now supports replacement messages for flow-based HTTPS web filtering. Flow-based HTTP and HTTPS web filtering send the same replacement message as proxy web filtering. For FortiGuard web filtering, the replacement message is the *FortiGuard Block Page* and for URL web filtering, the replacement message is the *URL Block Page*. To edit replacement messages, go to *System > Config > Replacement Messages*.

DNS web filtering

A DNS request is typically the first part of any new session to a new website. DNS web filtering takes advantage of this by including the web site category in DNS responses. When a FortiGate unit resolves a URL, it receives a rating in addition to the IP address of the website.

DNS Web filtering uses the same categories as FortiGuard Web Filtering and requires you to configure your FortiGate unit to use FortiGuard DNS as its DNS Server. DNS web filtering is lightweight in terms of resource usage because it doesn't involve any actual content inspection.

DNS web filtering includes reduced functionality compared to proxy and flow-based web filtering. DNS web filtering does not support:

- Quotas
- Setting web filter categories to Warning or Authenticate (Allow, Monitor and Block are supported)
- Safe Search
- URL only scanning for HTTPS
- Advanced filtering options such as web content filtering, web resume download blocking, blocking invalid URLs, HTTP post action options, Java applet filtering, ActiveX filtering, cookie filtering, image rating, allowing websites when a rating error occurs and blocking HTTP redirects by rating

To configure your FortiGate unit to use DNS web filtering, start by going to *System > Network > DNS* and under *DNS Settings*, make sure *Use FortiGuard Servers* is selected and select *Apply*.

Go to *Security Profiles > Web Filter > Profiles* and edit a web filtering profile or create a new one. Set *Inspection Mode* to DNS. Then you can set *DNS action* to *Block* or *Redirect*. If you select *Redirect*, every time a web page is blocked by DNS web filtering the URL is re-directed to a web page on the FortiGuard network that displays a block message. If you select *Block*, the page is blocked and the user's web browsers display an error message or the connection attempt will time out.

Set the FortiGuard web filtering categories as required. You can configure DNS web filtering to block, allow and monitor web pages in each FortiGuard category. Select *Apply* to save the profile.

Figure 33:DNS web filtering profile

New Web Filter Profile

Name:

Comments: 0/255

Inspection Mode: Proxy Flow-based DNS

DNS Action: Block Redirect

FortiGuard Categories

Show: ▼

- Local Categories
- Potentially Liable
- Adult/Mature Content
- Bandwidth Consuming
- Security Risk
- General Interest - Personal
- General Interest - Business
- Unrated

Enable Web Site Filter

Rate URLs by Domain and IP Address

Go to *Policy > Policy > Policy* and create or edit a security policy, enable web filtering and select the web filtering profile that you configured for DNS web filtering.

All traffic HTTP accepted by the policy will be inspected by DNS web filtering.

FortiGuard Web Filter quotas can be set based on traffic volume

In FortiOS 5.0, FortiGuard web filter quotas can now set based on the amount of traffic as well as time.

You can add traffic quotas to a web filter profile from the CLI. The following command shows how to add a quota of 20 GB for bandwidth consuming web sites. These command assumes you have already set up the profile to monitor, warn or require authentication for bandwidth consuming web sites (category g04).

```
config webfilter profile
  edit default
    config ftgd-wf
      config quota
        edit 0
          set category g04
          set type traffic
          set unit GB
          set value 20
        end
      end
    end
  end
```

Customizing the authentication replacement message for a FortiGuard web filter category

FortiOS 5.0 allows you to customize the replacement message that appears for a specific FortiGuard Web Filtering category. You do this by editing a Web Filter profile, right clicking on a FortiGuard Web Filtering category, selecting *Authenticate* and selecting a user group. Then right-click on the category again and select *Customize*.

A blank customize replacement message window appears and you can create the custom replacement message. Select *Save* and close the replacement message editor. The selected category has an *Authenticate* icon next to it. You can select this icon to edit the replacement message.

Saving the message creates a custom replacement message group. If you go to *System > Config > Replacement Messages Group* and open the replacement message group called *web-filter-default* you can find a *Custom Messages* category that contains the new replacement message.

YouTube Education Filter implemented in Web Filtering Profiles

You can add your organization's YouTube education filter to a web filtering profile. The Educational filter will be implemented on all YouTube sessions accepted by the security policy that the web filter profile is added to. This makes it easier to allow your users to only access educational YouTube content while blocking content not considered educational.

To add a YouTube education filter

- 1 Go to *Security Profiles > Web Filter > Profiles* and edit a web filter profile.
- 2 Select *Enable Safe Search > YouTube Education Filter* and enter the YouTube education filter code.
- 3 Select *Apply* to save the changes to the web filter profile.
- 4 Go to *Policy > Policy > Policy* and edit the security policy that allows users to access the Internet.
- 5 Select *Security Profiles*.

- 6 Select *Enable Web Filter* and select the web filter profile that includes the YouTube education filter.
- 7 Select OK to save the security policy.

Figure 34: Adding a YouTube education filter code to a web filter profile

Use the following CLI command to add a YouTube education code to a web filter profile:

```
config webfilter profile
  edit youtube-EDU
    config web
      set safe-search youtube-edu
      set youtube-edu-filter-id "ABCD1234567890abcdef"
    end
  end
```

IPS hardware acceleration

FortiGate units with CPx and NPx processes can accelerate IPS performance by offloading pattern matching to the CPx or NPx processor. If your FortiGate hardware supports this feature the following CLI command will be available:

```
config ips global
  set hardware-accel-mode {engine-pick | none | CP-only | NP-only |
  NP+CP}
end
```

Where:

- `engine-pick`, let the IPS engine pick the best mode.
- `none`, hardware acceleration disabled.
- `CP-Only`, accelerate with content processors only.
- `NP-only`, accelerate with network processor only.
- `NP+CP`, accelerate with both network and content processors.

New SIP ALG features

FortiOS 5.0 includes the following new SIP ALG features:

- [Inspecting SIP over SSL/TLS \(secure SIP\)](#)
- [Opening and closing SIP via and record-route pinholes](#)
- [Adding the original IP address and port to the SIP header after NAT](#)

Inspecting SIP over SSL/TLS (secure SIP)

Some SIP phones and SIP servers can communicate using SSL or TLS to encrypt the SIP signalling traffic. To allow SIP over SSL/TLS calls to pass through the FortiGate unit, the encrypted signalling traffic has to be unencrypted and inspected. To do this, the FortiGate SIP ALG intercepts, unencrypts and inspects the SIP packets. The packets are then re-encrypted and forwarded to their destination.

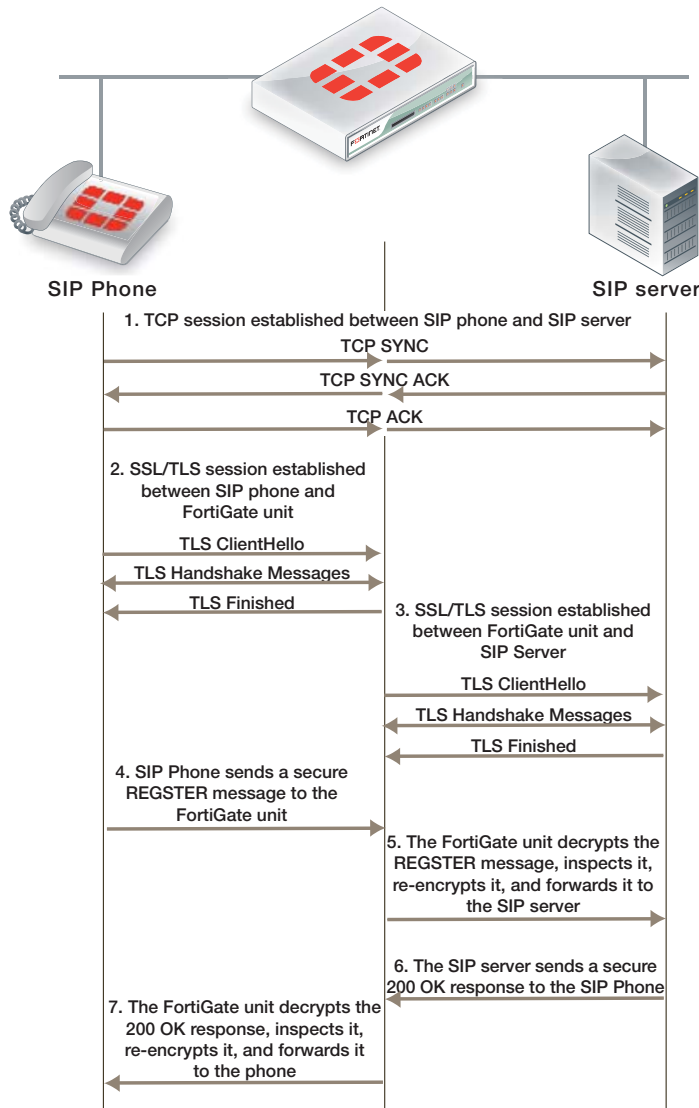
Normally SIP over SSL/TLS uses port 5061. You can use the following command to change the port that the FortiGate listens on for SIP over SSL/TLS sessions to port 5066:

```
config system settings
  set sip-ssl-port 5066
end
```

The SIP ALG supports full mode SSL/TLS only. Traffic between SIP phones and the FortiGate unit and between the FortiGate unit and the SIP server is always encrypted.

You enable SSL/TLS SIP communication by enabling SSL mode in a VoIP profile. You also need to install the SIP server and client certificates on your FortiGate unit and add them to the SSL configuration in the VoIP profile.

Figure 35:SIP over SSL/TLS between a SIP phone and a SIP server



Other than enabling SSL mode and making sure the security policies accept the encrypted traffic, the FortiGate configuration for SSL/TLS SIP is the same as any SIP configuration.

SIP over SSL/TLS is supported for all supported SIP configurations.

Adding the SIP server and client certificates

A VoIP profile that supports SSL/TLS SIP requires one certification for the SIP server and one certificate that is used by all of the clients. Use the following steps to add these certificates to the FortiGate unit. Before you start, make sure the client and server certificate files and their key files are accessible from the management computer.

1. Go to *System > Certificates > Local Certificates* and select *Import*.
2. Set *Type* to *Certificate*.
3. Browse to the *Certificate file* and the *Key file* and select *OK*.
4. Enter a password for the certificate and select *OK*.

The certificate and key are uploaded to the FortiGate unit and added to the *Local Certificates* List.

5. Repeat to upload the other certificate.

The certificates are added to the list of Local Certificates as the filenames you uploaded. You can add comments to make it clear where the certificate is from and how it is intended to be used.

Adding SIP over SSL/TLS support to a VoIP profile

Use the following commands to add SIP over SSL/TLS support to the default VoIP profile. The following command enables SSL mode and adds the client and server certificates and passwords (the same ones you entered when you imported the certificates):

```
config voip profile
  edit default
    config sip
      set ssl-mode full
      set ssl-client-certificate "Client_cert"
      set ssl-server-certificate "Server_cert"
      set ssl-auth-client "check-server"
      set ssl-auth-server "check-server-group"
    end
  end
```

Other SSL mode options are also available:

<code>ssl-send-empty-frags</code> {disable enable}	Enable to send empty fragments to avoid CBC IV attacks. Compatible with SSL 3.0 and TLS 1.0 only. Default is enable.
<code>ssl-client-renegotiation</code> {allow deny secure}	Control how the ALG responds when a client attempts to renegotiate the SSL session. You can allow renegotiation or block sessions when the client attempts to renegotiate. You can also select <code>secure</code> to reject an SSL connection that does not support RFC 5746 secure renegotiation indication. Default is allow.

<code>ssl-algorithm {high low medium}</code>	Select the relative strength of the algorithms that can be selected. You can select <code>high</code> , the default, to allow only AES or 3DES, <code>medium</code> , to allow AES, 3DES, or RC4 or <code>low</code> , to allow AES, 3DES, RC4, or DES.
<code>ssl-pfs {allow deny require}</code>	Select whether to allow, deny, or require perfect forward secrecy (PFS). Default is <code>allow</code> .
<code>ssl-min-version {ssl-3.0 tls-1.0 tls-1.1}</code>	Select the minimum level of SSL support to allow. The default is <code>ssl-3.0</code> .
<code>ssl-max-version {ssl-3.0 tls-1.0 tls-1.1}</code>	Select the maximum level of SSL support to allow. The default is <code>tls-1.1</code> .

Opening and closing SIP via and record-route pinholes

If `open-via-pinhole` is disabled (the default setting), the FortiGate unit does not open pinholes for Via messages. You can enable `open-via-pinhole` so that the FortiGate unit opens pinholes for Via messages. In previous versions of FortiOS, this option was `reg-diff-port`.

If `open-record-route-pinhole` is enabled (the default setting), the FortiGate unit opens pinholes for Record-Route messages. You can disable `open-record-route-pinhole` so that the FortiGate unit does not open pinholes for Record-Route messages.

Usually you would want to open these pinholes. Keeping them closed may prevent SIP from functioning properly through the FortiGate unit. However, they can be disabled for interconnect scenarios (where all SIP traffic is between proxies and traveling over a single session). In some cases, these settings can also be disabled in access scenarios if it is known that all users will be registering regularly so that their contact information can be learned from the register request.

You may also want to prevent pinholes from being opened to avoid creating a pinhole for every register or non-register request. Each pinhole uses additional system memory, which can affect system performance if there are hundreds or thousands of users, and requires refreshing that can take a relatively long amount of time if there are thousands of active calls.

Adding the original IP address and port to the SIP header after NAT

In some cases, your SIP configuration may require that the original IP address and port from the SIP contact request is kept after NAT. For example, the original SIP contact request could include the following:

```
Contact: <sip:0150302438@172.20.120.110:5060>;
```

After the packet goes through the FortiGate unit and NAT is performed, the contact request could look like the following (the IP address translated to a different IP address and the port to a different port):

```
Contact: <sip:0150302438@10.10.10.21:33608>;
```

You can enable `register-contact-trace` in a VoIP profile to have the SIP ALG add the original IP address and port in the following format:

```
Contact: <sip:0150302438@<nated-ip>:<nated-port>;o=<original-ip>:<original-port>>;
```

So the contact line after NAT could look like the following:

```
Contact: <sip:0150302438@10.10.10.21:33608;o=172.20.120.110:5060>;
```


Enter the following command to enable keeping the original IP address and port:

```
config voip profile
  edit Profile_name
    config sip
      set register-contract-trace enable
  end
```

DLP watermarking

DLP watermarking involves using DLP to filter files that pass through the FortiGate unit and contain a corporate identifier (a text string) and a sensitivity level (Critical, Private and Warning) hidden in a watermark. Watermarked files have this information applied in a way that is not visible to the user. DLP watermarking requires a FortiGate unit with a hard disk or flash disk.

You must use the Fortinet watermarking client to apply a watermark to a file. Files should be watermarked before they are distributed. Then, with DLP watermarking enabled, your FortiGate unit can track and optionally block watermarked files that pass through it.

To configure DLP to filter for watermarked files, go to *Security Profiles > Data Leak Prevention > Sensors* and create or edit a DLP sensor. Add a filter to the sensor and set *Filter* to *Files*. Then select *Watermark Sensitivity* and select *Critical*, *Private* or *Warning*. Then enter the *Corporate Identifier*. The corporate identifier is a case-sensitive text string that must exactly match the corporate identifier text string added by the watermarking client.

Select the services in which to look for watermarked files. Usually you would choose all of the email protocols active on your network and HTTP. Then set the action for the watermark filter. When DLP finds a file with a watermark that matches the filter, the action selected in the filter is performed. Actions include writing a log message, blocking the file or quarantining the user, IP address or interface.

Figure 36: DLP filter configuration using a watermark

The screenshot shows the 'New Filter' configuration window in the FortiGate GUI. The 'Filter' section is set to 'Files'. Under 'Files', the 'Watermark Sensitivity' is set to 'Critical' and the 'Corporate Identifier' is 'Do not distribute!'. The 'Examine the following Services' section has checkboxes for SMTP, IMAP, POP3, HTTP, NNTP, and others. The 'Action' is set to 'Block'.

Files can have multiple watermarks in them. The FortiGate unit only has to find one match in a file for DLP watermarking to match it and it will ignore watermarks that don't match. Files without watermarks are ignored by DLP watermarking.

Fortinet watermarking utility

Watermarking uses a digital pattern to mark a file as being proprietary to a specific company. Fortinet has a utility that will apply a digital watermark to any file except a .txt file. The utility adds a small (around 100 bytes) pattern to the file that is recognized by the DLP Watermark filter. This pattern is invisible to the end user.

Currently, FortiGate DLP only works with Fortinet's watermarking client. When watermarking a file, it should be verified that the pattern matches up to a DLP category. Before planning to use watermarking software, it is always best to verify that the software will work with your OS and file types. At the time of writing this document the utility was only available with the current version of FortiExplorer for Windows or through using the CLI for Linux.

Installation of the watermarking client on Linux

Add the watermark file to a location on the system that is in the \$PATH

To see what the path is use the command

```
~$ echo $PATH
```

Example results:

```
/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin:/usr/ga
mes
```

for example you could move or copy the file to the ./bin directory.

Permissions on the watermark file

Check the existing permissions:

The command in Linux for listing file along with the permissions is:

```
ls -l
```

Run the check to see if the permission status. The results may be something along these lines:

```
-rw-r--r-- 1 root root 2053868 Jan 10 11:44
    fortinet-watermark-linux.out
```

You will see that in this case it has no executable permissions

To change the permissions on the watermark file:

It will be assume for this command that the utility is in the bin directory and that you have ownership level access.

```
/bin# chmod o+x /bin/ fortinet-watermark-linux.out
```

To verify the change:

```
/bin# ls -l wa*
-rw-r--r-x 1 root root 2053868 Jan 10 11:44
    fortinet-watermark-linux.out
```

You can see how the x for executable has been added to the permissions for the others group.

Syntax of the watermarking client on Linux

The tool is executed in a Linux environment by passing in files or directories of files to insert a watermark.

USAGE:

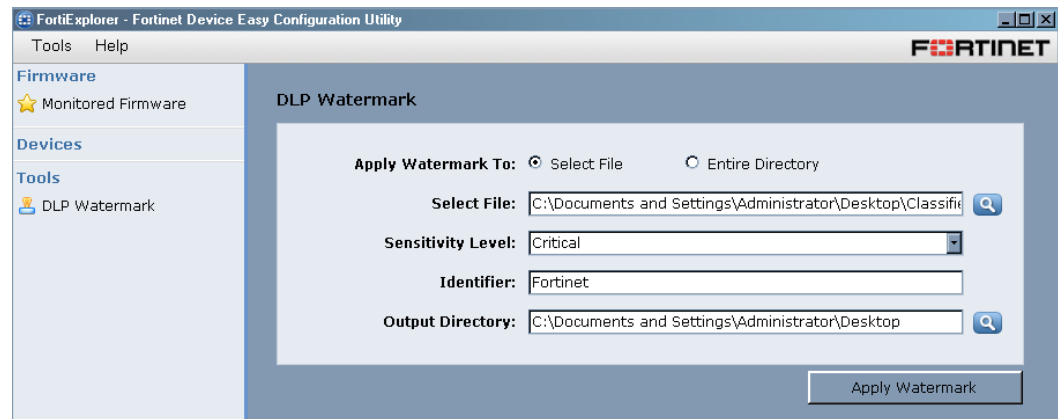
```
fortinet-watermark-linux.out <options> -f <file name> -i <identifier>
    -l <sensitivity level>
fortinet-watermark-linux.out <options> -d <directory> -i <identifier>
    -l <sensitivity level>
```

Options:

- h print help
- v verbose information
- I inplace watermarking (don't copy file)
- o output directory
- e encode <to non-readable>
- a add additional watermark (by default replaces watermarks existing watermarks)
- D delete all watermarks

Using the watermarking client with Windows

The watermarking client is now part of FortiExplorer in Windows, appearing in the Tools menu. Using the client, you can apply a watermark to any files you wish to track and possibly block using DLP.

Figure 37:The Fortinet watermarking client in Windows**Using the watermarking client with Linux**

If you are in your home directory and you want to watermark a file in the Documents directory, you could plan out the command like this:

- ```
watermark [because that is the executable to be used]
-v [so that you can get as much feedback as possible]
-I [because you don't want a new file you just want to watermark the existing one]
-f [because you only want to change the one file not the entire directory]
filename.pdf [the name of the file]
-i 123456 [to set the identifier to 123456 - this is a required setting]
-l Private [to set the sensitivity level to "Private"]
```

Now at the command prompt enter all of these components in order:

```
~/Documents$ fortinet-watermark-linux.out -v -I -f filename.pdf -i
12345 -l Private
Creating watermark. Pattern:
=====identifier=12345
sensitivity=Private=====
Watermarking file: 'filename.pdf'
Inserted watermark size 148
<command prompt>:~/Documents$
```

## SSH inspection

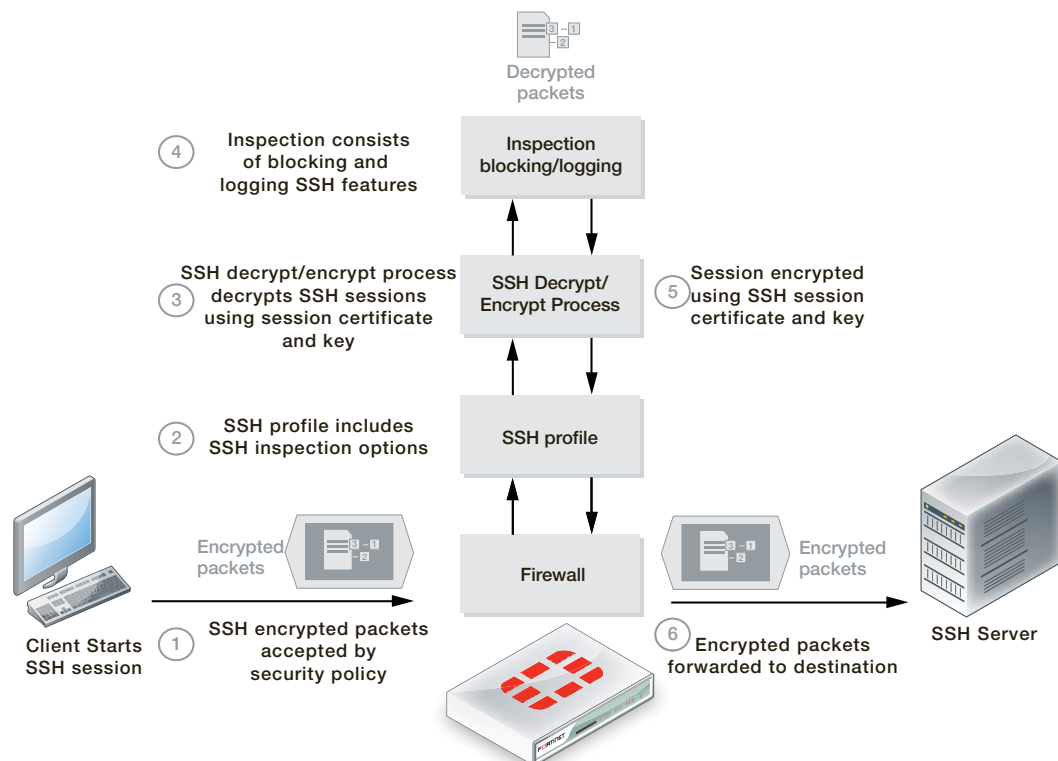
FortiOS 5.0 includes a new SSH proxy, available on selected models, that supports decrypting and inspecting SSH sessions to block or log the following:

- SSH remote execution
- Port forwarding
- SSH remote shell
- x11 server forwarding

Similar to SSL content inspection, the SSH proxy intercepts the SSH key exchange between the SSH client and server when an SSH session is being set up. All traffic that is part of the SSH session is decrypted by the SSH proxy and SSH inspection features are applied according to the SSH profile in the security policy that accepts the SSH traffic. After inspection, the session is re-encrypted and forwarded to the recipient.

SSH inspection is enabled by default in a SSL/SSH inspection profile.

**Figure 38:**SSH inspection



To configure SSH inspection, go to *Policy > Policy > SSL/SSH Inspection* and create or edit an SSL/SSH Inspection profile and ensure that *SSH Inspection Options* are enabled. Then configure the port or ports to look for SSH traffic on. The default port is 22 but you can add more. You can then block or log Exec commands, port-forwarding, SSH shells, and X11 Filters.

Then add this profile to security policy that accepts SSH traffic. Enable the required Security Profiles in the policy.

**Figure 39:**Configuring SSL Inspection Options

**SSH Inspection Options**

Enable SSH Deep Scan

| Protocol     | Inspection Port(s)                                                                                 |
|--------------|----------------------------------------------------------------------------------------------------|
| SSH          | <input type="radio"/> Any <input checked="" type="radio"/> Specify <input type="text" value="22"/> |
| Exec         | <input type="checkbox"/> Block <input type="checkbox"/> Log                                        |
| Port-Forward | <input type="checkbox"/> Block <input type="checkbox"/> Log                                        |
| SSH-Shell    | <input type="checkbox"/> Block <input type="checkbox"/> Log                                        |
| X11-Filter   | <input type="checkbox"/> Block <input type="checkbox"/> Log                                        |

From the CLI:

```
config firewall profile-protocol-options
edit new-profile
config ssh
set port <number> <number> ... (default is port 22)
set inspect-all {enable | disable}
set options {allow-invalid-server-cert | ssl-ca-list }
set oversize-limit <size>
set block {exe | port-forward | ssh-shell | x11-filter}
end
```

## Optimizing SSL encryption/decryption performance

By default, FortiGate units handle SSL decryption/encryption using the SSL functionality built into their FortiASIC processors. In situations where the FortiGate unit processes large amounts of SSL traffic and has more than 4 CPUs, you may be able to optimize SSL encryption/decryption performance by changing how SSL processing is distributed to the CPUs. You can also use the following command to specify the number of CPUs to use for SSL processing (in the command, CPU is called an SSL worker):

```
config system global
set optimize-ssl {enable | disable}
set ssl-worker-count <worker-count>
end
```

The `<worker-count>` is the number of CPUs. The range depends on the number of CPUs in the FortiGate model (this feature only works for FortiGate units with 4 or more CPUs).

You can use the following command to display information about each CPU running in your FortiGate unit:

```
get hardware cpu
```

The command output numbers the CPUs starting at 0. For example, a FortiGate-5001B contains 8 CPUs and the command output for this model contains information about all 8 CPUs numbered 0 to 7. Here is the first few output lines for CPU 7:

```
...
processor : 7
vendor_id : GenuineIntel
cpu family : 6
model : 14
model name : Intel(R) Xeon(R) CPU C5528 @ 2.13GHz
stepping : 4
cpu MHz : 2128.072
...
```

If your FortiGate unit includes multiple CPUs and you want to improve SSL performance, use the following command to begin distributing SSL decryption/encryption to 4 CPUs:

```
config system global
 set optimize-ssl enable
 set ssl-worker-count 4
end
```

Monitor FortiGate performance and if SSL performance improves without affecting other performance, you can either maintain this configuration or add another CPU to the configuration (if one is available). Continue in this manner until you achieve optimum performance for your FortiGate unit.

Continue monitoring performance in case you have to change this setting due changes in your network traffic patterns.

# Authentication: users and devices

The FortiGate authentication umbrella has been expanded from just user authentication (or user identity) to encompass device identification and client reputation. As well, endpoint control and the vulnerability scanner have become part of FortiOS 5.0 user and device detection, identification and authentication.

Endpoint control and vulnerability scanner changes are described in this chapter. For information about device identification, see [“FortiOS and BYOD” on page 173](#). For information about client reputation, see [“Client Reputation” on page 181](#).

New authentication features described in this chapter include:

- [User authentication menu changes](#)
- [User identity policy changes](#)
- [Authentication-based routing](#)
- [Secondary and tertiary RADIUS, LDAP, and TACAS+ servers](#)
- [FortiToken two-factor authentication and FortiToken Mobile](#)
- [SSO using a FortiAuthenticator unit](#)
- [SSO with Windows AD or Novell](#)
- [Citrix Agent support for Single Sign On](#)
- [Configuring guest access](#)
- [Vulnerability Scanning](#)

## User authentication menu changes

The user authentication part of FortiOS is seeing major changes for FortiOS 5.0. To begin, the *User* section of the web-based manager has been renamed *User & Device*. All previously available user authentication features are still available but the menu structure has changed to include device identification, the vulnerability scanner, endpoint control, and client reputation:

- Go to *User & Device > User* to configure users, user groups, and guest users ([“Configuring guest access” on page 167](#))
- Go to *User & Device > Authentication* to configure single sign-on and add RADIUS, LDAP, and TACACS+ servers
- Go to *User & Device > Two-factor Authentication* to configure support for two-factor authentication using FortiToken ([“FortiToken two-factor authentication and FortiToken Mobile” on page 162](#))
- Go to *User & Device > Vulnerability Scan* to configure and operate the FortiGate vulnerability scanner ([“Vulnerability Scanning” on page 170](#))
- Go to *User & Device > Monitor* to view the firewall and banned user authentication lists

## User identity policy changes

The steps for adding user identity-based policies have changed. To add a user identity based policy go to *Policy > Policy > Policy* and create a new security policy. Select the *Firewall* policy type and the *User Identity* subtype. Select the incoming and outgoing interfaces and source addresses. Configure other features such as NAT and so on.

Then select *Create New* to add user authentication rules to the policy. User authentication rules include the destination addresses, user groups and or individual users, schedule, service, action, logging, and UTM security profiles.

You select the destination address separately for each authentication rule. This means that you can apply different features to different user groups depending on the destination address.

**Figure 40:**Adding a user authentication rule

**New Authentication Rule**

Destination Address: all

Group(s): FSSO\_Guest\_Users

User(s): jsmith

Schedule: always

Service: Click to add...

Action: ACCEPT

**Logging Options**

- No Log
- Log Security Events
- Log all Sessions

**Security Profiles**

- AntiVirus: default
- Web Filter: default
- Application Control: default
- IPS: default

## Authentication-based routing

FortiOS 5.0 supports authentication-based routing by creating an identity-based route that associates a user group with one or more routes. This identity-based route is then added to a security policy and all traffic from users authenticated by this user group is routed to the gateway. This feature is configured from the CLI and can be useful for MSSPs who need to route users from different organizations to different Internet gateways.

Enter the following command to add an identity-based route that routes all traffic from users in the company1-user-group and the company2-user-group user groups out the wan1 interface to a next-hop router with IP address 172.20.120.2:

```
config firewall identity-based-route
 edit new-id-route
 config rule
 edit 1
 set gateway 172.20.120.2
 set device wan1
 set groups company1-user-group company2-user-group
 end
 end
 end
```

Enter the following command to add the identity-based route to a security policy:



```

config firewall policy
 edit 1
 ...
 set identity-based enable
 set identity-based-route new-id-route
 ...
 end

```

## Secondary and tertiary RADIUS, LDAP, and TACAS+ servers

You can now add secondary and tertiary servers to RADIUS, LDAP, and TACAS+ remote authentication server configurations. When you add a secondary server, the FortiGate unit will contact the secondary server only if the primary server is unreachable. The FortiGate unit will only contact the tertiary server if the both the primary and secondary servers are unreachable.

Enter the following command to add up to three servers to a RADIUS server configuration. Specify a domain name or IP address for each server as well as the server secret. In the following example, the RADIUS servers are at IP addresses 172.20.120.10, 172.20.120.20, and 172.20.120.30:

```

config user radius
 edit new-radius-server
 set server 172.20.120.10
 set secret 1st-secret
 set secondary-server 172.20.120.20
 set secondary-secret 2nd-secret
 set tertiary-server 172.20.120.30
 set tertiary-secret 3rd-secret
 end

```

Enter the following command to add up to three servers to an LDAP server configuration. Specify a domain name or IP address for each server. Other than the domain name or password, the secondary and tertiary servers must use the same port and LDAP settings such as the cnid and username. In the following example, the LDAP servers are at IP addresses 192.168.10.10, 192.168.10.20, and 192.168.10.30:

```

config user ldap
 edit "test-ldap"
 set server "192.168.10.10"
 set cnid "exAccountName"
 set dn "dc=americas,dc=example,dc=net"
 set port 3268
 set type regular
 set username "CN=example,OU=Service
 Accounts,OU=Admins,DC=example,DC=csplc,DC=net"
 set password ENC AAAEAOZh5R5/oqYeUVkO2OOkh9QV6DAVZoAjbv0sonh
 set member-attr "ASCCGKraftFortinetVPNInternalUsers"
 set secondary-server "192.168.10.20"
 set tertiary-server "192.168.10.30"
 end

```

Enter the following command to add up to three servers to an TACAS+ server configuration. Specify a domain name or IP address and key for each server. In the following example, the TACAS+ servers are at IP addresses 10.10.10.10, 10.10.10.20, and 10.10.10.30:

```


config user tacacs+
 edit "test-tacacs"
 set server "10.10.10.10"
 set key ENC
 2OG/F6wocz2/CpE3eHIJs/Qt8gZsXgeNkQCuTxPWPeBk6BXDu8luM
 set secondary-server "10.10.10.20"
 set secondary-key ENC 2OG/F6wocz2/CpE3eHIJs/Qt8gZ
 set tertiary-server "10.10.10.30"
 set tertiary-key ENC 2OG/F6wocz2/CpE3eHIJs/Qt8gZ
 next
end

```

## FortiToken two-factor authentication and FortiToken Mobile

The web-based manager provides improved management of FortiToken devices. The status of each current FortiToken device is listed under *User & Device > Two-factor Authentication > Fortitokens*. You can also resynchronize FortiToken devices that have gone out of sync. You can also enter new FortiToken devices individually or by importing a list of FortiToken serial numbers in a text file.

**Figure 41:**View status and synchronize a FortiToken



### Configuring FortiToken mobile soft token support

FortiOS 5.0 adds support for FortiToken Mobile, a Fortinet application that enables you to generate One Time Passwords (OTPs) on a mobile device for FortiGate two factor authentication. The user's mobile device and the FortiGate unit must be connected to the Internet to activate FortiToken mobile. Once activated, users can generate OTPs on their mobile device without having network access.

FortiToken Mobile is available for iOS and Android devices from their respective Application stores. No cellular network is required for activation.



The latest FortiToken Mobile documentation is available from the [FortiToken](#) page of the [Fortinet Technical Documentation](#) website.

To use FortiToken Mobile, a user needs to install the application on their mobile device and then activate a token. After the token is activated, the user can begin to generate OTPs on their mobile device.

Two free trial tokens are included with every registered FortiGate unit. Additional tokens can be purchased from your reseller or from Fortinet.

You can generate the two free tokens from the FortiGate CLI by entering the following command:

```
execute fortitoken-mobile import 0000-0000-0000-0000-0000
```

This command adds two FortiToken mobile entries to the FortiGate configuration. To view them, go to *User & Device > Two-factor Authentication > Fortitokens* and open the *MobileToken* list and edit one of the entries.

**Figure 42:** Example free FortiToken Mobile token

To assign a token to a user, go to *User & Device > User > User Definition* and either add a new user or select the user to assign the token to. Configure the user as required and select *Enable Two-factor Authentication*. Select the token to associate with the user.

Select *OK* to assign the token to the user. If you have added the user’s email address or configured SMS settings and configured your FortiGate unit to send email or send SMS messages, the FortiGate unit sends the user an activation code. (To configure your FortiGate unit to send email or SMS messages go to *System > Config > Messaging Servers*.)

**Figure 43:** Assigning a FortiToken Mobile token to a user

If for some reason you cannot send the activation code to the user through email or an SMS message, or would rather send the activation code by other means specific to your operation, you can view the activation code from the CLI. For example, for a token with serial number of FTKMOB28E0CA6018 you can enter the following commands:

```
config user fortitoken
 edit FTKMOB28E0CA6018
 get
 serial-number : FTKMOB28E0CA6018
 activation-code : 8F41F304
 activation-expire : 604800
 comments :
 license : FTMTRIAL00001088
 status : active
```

Send the activation-code (in this example, 8F41F304) to the user. Following the instructions in the *FortiToken Mobile User Guide*, the user can activate their FortiToken Mobile application using this activation code.

The user's mobile device, as well as the FortiGate unit, must be connected to the Internet to complete the activation. When the activation is complete, the *Status* of the FortiToken Mobile token changes to *Provisioned* on the FortiGate unit.

## SSO using a FortiAuthenticator unit

If you use a FortiAuthenticator unit in your network as a single sign-on agent:

- Users can authenticate through a web portal on the FortiAuthenticator unit.
- Users with FortiClient Endpoint Security installed can be automatically authenticated by the FortiAuthenticator unit through the FortiClient SSO Mobility Agent.

The FortiAuthenticator unit can integrate with external network authentication systems such as Windows Active Directory, Novell e-Directory, RADIUS and LDAP to gather user logon information and send it to the FortiGate unit.

### User's view of FortiAuthenticator SSO authentication

There are two different ways users can authenticate through a FortiAuthenticator unit.

#### Users without FortiClient Endpoint Security - SSO widget

To log onto the network, the user accesses the organization's web page with a web browser. Embedded on that page is a simple logon widget. The SSO widget sets a cookie on the user's browser. When the user browses to a page containing the login widget, the FortiAuthenticator unit recognizes the user and updates its database if the user's IP address has changed. The user will not need to re-authenticate until the login expires, which can take up to 30 days.

#### Users with FortiClient Endpoint Security - FortiClient SSO Mobility Agent

All authentication is performed transparently with no request for credentials. IP address changes, such as those due to WiFi roaming, are automatically sent to the FortiAuthenticator unit. When the user logs off or otherwise disconnects from the network, the FortiAuthenticator unit is aware of this and deauthenticates the user.

The FortiClient SSO Mobility Agent, a feature of FortiClient Endpoint Security v5.0, must be configured to communicate with the appropriate FortiAuthenticator unit. After that, the agent automatically provides user name and IP address information to the FortiAuthenticator unit for transparent authentication.

## Administrator's view of FortiAuthenticator SSO authentication

You can configure either or both of these authentication types on your network.

### SSO widget

You need to configure the Single Sign-On portal on the FortiAuthenticator unit. Go to *SSO & Dynamic Policies > SSO > Login Portal* to do this. Copy the *Embeddable login widget* code for use on your organization's home page. Identity-based security policies on the FortiGate unit determine which users or groups of users can access which network resources.

### FortiClient SSO Mobility Agent

Your users must be running FortiClient Endpoint Security v5.0 to make use of this type of authentication.

On the FortiAuthenticator unit, you need to enable *FortiClient Service* when you define the unit's secret key. Go to *SSO & Dynamic Policies > SSO > Options*. You need to provide your users the FortiAuthenticator IP address and secret key so that they can configure the FortiClient SSO Mobility Agent on their computers.

## SSO with Windows AD or Novell

The FortiGate unit can authenticate users transparently based on their Windows Active Directory (AD) or Novell eDirectory privileges. This means that users who have logged on to the network are not asked again for their credentials to access network resources through the FortiGate unit, hence the term "Single Sign-On".

FSSO Collector agent and DC agent have been tested on Windows Server 2003, 2008 and 2012.

On a Microsoft Windows or Novell network, users authenticate with the Microsoft AD or Novell eDirectory at logon. It would be inconvenient if users then had to enter another username and password for network access through the FortiGate unit. FSSO agents installed on the network provide user information, such as IP address and user group memberships, to the FortiGate unit. Security policies on the FortiGate unit allow network access based on the user groups to which the user belongs.

There are several mechanisms for passing user authentication information to the FortiGate unit:

- FSSO Collector agent software installed on a Windows AD network monitors user logons and sends the required information to the FortiGate unit. The FSSO software can obtain this information by polling the AD domain controllers or by using an FSSO agent on each AD domain controller that monitors user logons in real time. New in FortiOS 5.0, a FortiGate unit can obtain group information directly from AD using Lightweight Directory Access Protocol (LDAP).
- On a Windows AD network, the FSSO software can also serve NT LAN Manager (NTLM) requests coming from client browsers (forwarded by the FortiGate unit) with only one or more Controller agents installed.
- FSSO eDirectory agent software installed on a Novell network monitors user logons and sends the required information to the FortiGate unit. The agent can obtain information from the Novell eDirectory using either the Novell API or LDAP.
- A FortiAuthenticator server can act as a replacement for the Collector agent in polling mode in a Windows AD network. FortiAuthenticator can also be configured with internal or external LDAP and RADIUS servers. For more information, see the [FortiAuthenticator Administration Guide](#).

## Citrix Agent support for Single Sign On

FortiOS 5.0 supports single sign on authentication for Citrix environments by installing a Citrix FSSO polling agent on the Citrix server, installing an FSSO collector on the network, and then configuring the FortiGate unit to get user credentials from the Citrix FSSO polling agent.

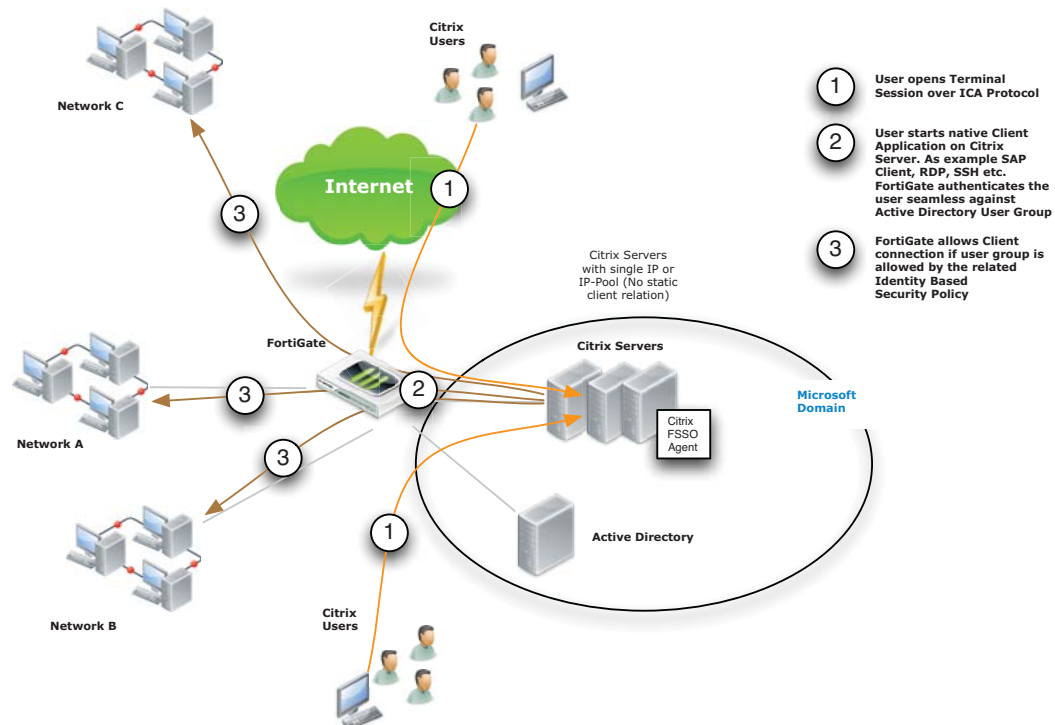
TSAgent running on the Citrix server XenAPP server version 6 has been tested with FSSO build0129 running Windows Server 2008.

Configuration steps include:

- Install the Fortinet Citrix FSSO agent on the Citrix server.
- Install the Fortinet FSSO collector on a server on the network.
- Add the Citrix FSSO agent to the FortiGate Single-sign-On configuration.
- Add Citrix FSSO groups and users to an FSSO user group.
- Add an FSSO identity-based security policy that includes the Citrix FSSO user groups.

After this configuration is complete, Citrix users credentials are made available to the FortiGate unit by the Citrix FSSO agent when a Citrix Terminal Session is started. When the user starts a client application (such as a web browser, SAP client, and so on), the user's session through the FortiGate unit is automatically authenticated and allowed by the FSSO identity-based security policy.

**Figure 44:** Example Citrix single sign on with FSSO network topology



### Installing Citrix/Terminal Service Support Agent (TS Agent)

Install the Citrix/terminal service support agent on the Citrix terminal server, or other terminal, in the same way as you would install the FSSO agent on any platform.

1. Log into the server with an account that has administrator privileges and a password that does not expire.
2. Start the FSSO agent installer.

3. Following the installation wizard prompts.  
During the installation process:
  - the *Host IP Address* is The local IP of the Citrix or Terminal service server.
  - The *Fortinet SSO collector IP and port* is the IP and port of the FSSO collector.
4. Make sure that Launch DC Agent Install Wizard is selected and then select Finish to end the installation.

## Installing the FSSO collector

There are no special requirements for installing the FSSO collector.

## To enable single sign-on using polling mode

1. Go to *User & Device > Authentication > Single Sign-On* and select *Create New* to add a single sign-on server.
2. Select *Fortinet Single-Sign-On Agent*.
3. Enter the Name and the IP address or Name and password for the Citrix server.
4. Select *OK* to save the configuration.

## Verifying the configuration

You can use the following diagnose commands to verify the configuration:

```
diagnose debug authd fsso list
diagnose debug application authd -1
diagnose firewall auth list
```

## Configuring guest access

You can create many guest accounts at once using randomly-generated User IDs and passwords. This reduces administrator workload for large events.

### User's view of guest access

1. The user receives an email, SMS message or printout from a FortiGate administrator listing a User ID and password.
2. The user logs onto the network with the provided credentials.
3. After the expiry time, the credentials are no longer valid.

### Administrator's view of guest access

1. Create one or more guest user groups.  
All members of the group have the same characteristics: type of User ID, type of password, information fields used, type and time of expiry.
2. Create guest accounts using Guest Management.
3. Use captive portal authentication and select the appropriate guest group.

## Creating guest management administrators

The guest management administrator can be a regular FortiGate administrator. Optionally, you can create administrator accounts that can perform only guest management. This type of administrator is also limited to specific guest user groups.

### To create a guest management administrator

1. Go to *System > Admin > Administrators* and create a regular administrator account.  
For detailed information see the System Administration chapter.
2. Select *Restrict to Provision Guest Accounts*.
3. In *Guest Groups*, add the guest groups that this administrator manages.

## Creating guest user groups

The guest group configuration determines the fields that are provided when you create a guest user account.

### To create a guest user group

1. Go to *User & Device > User > User Group* and select *Create New*.
2. Configure the guest user group:

|                            |                                                                                                                                                                                                                       |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Name</b>                | Enter a name for the group.                                                                                                                                                                                           |
| <b>Type</b>                | Guest                                                                                                                                                                                                                 |
| <b>User ID</b>             | Select one of: <ul style="list-style-type: none"> <li>• Email — User's email address</li> <li>• Specify — Administrator assigns user ID</li> <li>• Auto-Generate — FortiGate unit creates a random user ID</li> </ul> |
| <b>Password</b>            | Select one of: <ul style="list-style-type: none"> <li>• Specify — Administrator assigns user ID</li> <li>• Auto-Generate — FortiGate unit creates a random password</li> <li>• Disable — no password</li> </ul>       |
| <b>Enable Name</b>         | If enabled, user must provide a name.                                                                                                                                                                                 |
| <b>Enable Sponsor</b>      | If enabled, user form has Sponsor field. Select <i>Required</i> or <i>Optional</i> .                                                                                                                                  |
| <b>Enable Company</b>      | If enabled, user form has Company field. Select <i>Required</i> or <i>Optional</i> .                                                                                                                                  |
| <b>Enable Email</b>        | If enabled, user is notified by email.                                                                                                                                                                                |
| <b>Enable Phone Number</b> | If enabled, user is notified by SMS. Select whether FortiGuard Messaging Service or a another SMS provider is used. You can add SMS providers in <i>System &gt; Config &gt; Messaging Servers</i> .                   |
| <b>Expire Type</b>         | Choose one of: <ul style="list-style-type: none"> <li>Immediately — expiry time is counted from creation of account</li> <li>After first login — expiry time is counted from user's first login</li> </ul>            |



|                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default Expire Time</b>                 | Set the expire time. The administrator can change this for individual users.                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Enable Batch Guest Account Creation</b> | <p>Create multiple accounts automatically. When this is enabled:</p> <ul style="list-style-type: none"> <li>• <i>User ID</i> and <i>Password</i> are set to <i>Auto-Generate</i>.</li> <li>• The user accounts have only <i>User ID</i>, <i>Password</i>, and <i>Expiration</i> fields. Only the <i>Expiration</i> field is editable. If the expiry time is a duration, such as “8 hours”, this is the time after first login.</li> <li>• You can print the account information. Users do not receive email or SMS notification.</li> </ul> |

**Figure 45:** Adding a Guest user group

## Creating guest user accounts

Guest user accounts are not the same as local user accounts created in *User & Device > User > User Definition*. Guest accounts are not permanent; they expire after a defined time period. You create guest accounts in *User & Device > User > Guest Management*.

### To create a guest user account

1. Go to *User & Device > User > Guest Management*.
2. In *Guest Groups*, select the guest group to manage.
3. Select *Create New* and fill in the fields in the *New User* form.

Fields marked *Optional* can be left blank. The guest group configuration determines the fields that are available.

4. Select *OK*.
5. Select to print the temporary user account information or to email it to the email address of the account.
6. Select *Return*.

## Guest Management Account List

Go to *User & Device > User > Guest Management* to create, view, edit or delete guest user accounts.

|                     |                                                                                                                                                                |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Create New</b>   | Creates a new guest user account.                                                                                                                              |
| <b>Edit</b>         | Edit the selected guest user account.                                                                                                                          |
| <b>Delete</b>       | Delete the selected guest user account.                                                                                                                        |
| <b>Purge</b>        | Remove all accounts from the list.                                                                                                                             |
| <b>Print</b>        | Print all of the user accounts in the group. You can print one or 3 accounts per page.                                                                         |
| <b>Send</b>         | Send the user account information to the guest. Depending on the group settings and user information, the information can be sent to the user by email or SMS. |
| <b>Refresh</b>      | Update the list.                                                                                                                                               |
| <b>Guest Groups</b> | Select the guest group to list. New accounts are added to this group.                                                                                          |
| <b>User ID</b>      | The user ID. Depending on the guest group settings, this can be the user's email address, an ID that the administrator specified, or a randomly-generated ID.  |
| <b>Expires</b>      | Indicates a duration such as "3 hours". A duration on its own is relative to the present time. Or, the duration is listed as "after first login."              |

## Batch guest account creation

You can use the guest user group auto generate options and guest user management options to quickly create any number of guest user accounts in just a few steps. Use the following steps to create 50 users with randomly generated usernames and passwords.

1. Go to *User & Device > User > User Groups* and select *Create New* to add a new user group.
2. Give the group a name, and set *Type* to *Guest*.
3. Select *Enable Batch Guest Account Creation* and select *OK*.
4. Go to *User & Device > User > Guest Management* and in the *Guest Groups* field select the guest user group that you just created.
5. Select *Create New > Multiple Users*.
6. Set the number of accounts to 50, set the expiry date, and select *OK*.

You can edit the individual accounts after they are created to send them to the user. You can also change the user name, password, and expiration time.

## Vulnerability Scanning

The network vulnerability scanner helps you to protect your network assets (servers and workstations) by scanning them for security weaknesses. Configuration and operation of the vulnerability scanner has been simplified and the feature has been moved to *User & Device > Vulnerability Scan*.

This section describes how to configure a single FortiGate unit for network scanning and how to view the results of the scan.

## Running and configuring scans and viewing scan results

You can configure regular network scans on a daily, weekly, or monthly basis.

### To run a vulnerability scan

1. Go to *User & Device > Vulnerability Scan > Scan Definition* and select *Start Scan*.  
The vulnerability starts a scan using the current scanner settings. When the scan is running you can pause or stop it at any time. You can also watch the progress of the scan.
2. When the scan is complete go to *User & Device > Vulnerability Scan > Vulnerability Result* to view the results of the scan.

### To run a vulnerability scan of one device or selected devices

1. Go to *User & Device > Vulnerability Scan > Scan Definition*.
2. Select the devices to scan from the *Asset Definitions* list.  
You can shift-click to select more than one device.
3. Select *Start*.  
The vulnerability starts a scan of the selected devices using the current scanner settings. When the scan is running you can pause or stop it at any time. You can also watch the progress of the scan.
4. When the scan is complete go to *User & Device > Vulnerability Scan > Vulnerability Result* to view the results of the scan. Select any log entry to view log details.

**Figure 46:**Example vulnerability scan results

| #  | Date/Time   | Dst            | Vulnerability                         | Severity |
|----|-------------|----------------|---------------------------------------|----------|
| 3  | 15:01:08    | 172.20.120.14  | AFP.File.Sharing.Guest.Access.Enabled | Info     |
| 4  | 15:01:08    | 172.20.120.14  |                                       | unknown  |
| 5  | 15:01:00    | 172.20.120.14  |                                       |          |
| 6  | 15:01:00    | 172.20.120.14  |                                       |          |
| 7  | 15:01:00    | 172.20.120.14  |                                       |          |
| 8  | 15:00:57    | 172.20.120.14  |                                       |          |
| 9  | 15:00:57    |                |                                       |          |
| 10 | 03-18 10:45 | 172.20.120.220 |                                       |          |
| 11 | 03-18 10:45 | 172.20.120.100 |                                       |          |
| 12 | 03-18 10:45 | 172.20.120.83  |                                       |          |
| 13 | 03-18 10:45 | 172.20.120.51  |                                       |          |
| 14 | 03-18 10:45 | 172.20.120.40  |                                       |          |

|                        |                                                                                         |                      |                                     |
|------------------------|-----------------------------------------------------------------------------------------|----------------------|-------------------------------------|
| <b>Vuln ID</b>         | 33353                                                                                   | <b>Dst</b>           | 172.20.120.14                       |
| <b>Virtual Domain</b>  | root                                                                                    | <b>Severity</b>      | Info                                |
| <b>Level</b>           | notice                                                                                  | <b>Timestamp</b>     | Tue Mar 19 15:01:08 2013            |
| <b>Protocol</b>        | tcp                                                                                     | <b>Vuln Category</b> | Remote Access                       |
| <b>Vulnerability</b>   | AFP.File.Sharing.Guest.Access.Enabled                                                   | <b>Log ID</b>        | 4098                                |
| <b>Sub Type</b>        | vulnerability                                                                           | <b>Date/Time</b>     | 15:01:08 (Tue Mar 19 15:01:08 2013) |
| <b>Reference</b>       | <a href="http://www.fortinet.com/ids/VID33353">http://www.fortinet.com/ids/VID33353</a> | <b>Action</b>        | vuln-detection                      |
| <b>NIST Vuln Score</b> | 5.0                                                                                     | <b>Dst Port</b>      | 548                                 |

### To configure scanning

1. Go to *User & Device > Vulnerability Scan > Scan Definition*.

2. Beside *Schedule* select *Change* to set the scan schedule and mode:

|                                |                                                                                                                                                                          |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Recurrence</b>              | Select <i>Daily</i> , <i>Weekly</i> , or <i>Monthly</i> and configure the details for the option you have selected.                                                      |
| <b>Suspend Scan between</b>    | Set a time during which the scan should be paused if its running.                                                                                                        |
| <b>Vulnerability Scan Mode</b> | <b>Quick</b> — check only the most commonly used ports<br><b>Standard</b> — check the ports used by most known applications<br><b>Full</b> — check all TCP and UDP ports |

3. Select *Apply* to save the schedule and scan type.
4. Select *Create New* under *Asset Definitions* to select the devices on the network to scan.
5. Enter the following information and select *OK*:

|                                                |                                                                                                                                               |
|------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Name</b>                                    | Enter a name for this asset.                                                                                                                  |
| <b>Type</b>                                    | Select <i>IP Address</i> to add a single IP address.<br>Select <i>Range</i> to add a range of IP addresses to scan.                           |
| <b>IP Address</b>                              | Enter the IP address of the asset. ( <i>Type is IP Address.</i> )                                                                             |
| <b>Range</b>                                   | Enter the start and end of the IP address range. ( <i>Type is Range.</i> )                                                                    |
| <b>Enable Scheduled Vulnerability Scanning</b> | Select to allow this asset to be scanned according to the schedule. Otherwise the asset is not scanned during a scheduled vulnerability scan. |
| <b>Windows Authentication</b>                  | Select to use authentication on a Windows operating system. Enter the username and password in the fields provided.                           |
| <b>Unix Authentication</b>                     | Select to use authentication on a Unix operating system. Enter the username and password in the fields provided.                              |

6. Select *Apply* to save the configuration.

# FortiOS and BYOD

FortiOS can control network access for different types of personal mobile devices that your employees bring onto your premises. This is done by:

- Identifying and monitoring the types of devices connecting to your networks, wireless or wired
- Using MAC address based access control to allow or deny individual devices
- Creating policies based on device type
- Enforcing endpoint control on devices that can run FortiClient Endpoint Control software

This section describes:

- [Device monitoring](#)
- [Controlling access with a MAC Address Access Control List](#)
- [Device policies](#)
- [Device policy portal options](#)
- [Creating the WiFi SSID](#)
- [Configuring Internet access for guests with mobile devices](#)

## Device monitoring

The FortiGate unit can monitor your networks and gather information about the devices operating on those networks. Collected information includes:

- Whether the device is currently online
- MAC address
- IP address
- Operating system
- Hostname
- User
- How long ago the device was detected and on which FortiGate interface
- Whether FortiClient is installed on the device

You can go to *User & Device > Device > Device Definitions* to view this information.

**Figure 47:**The Device List

| <span>+</span> Create New <span>✎</span> Edit <span>🗑</span> Delete <span>🔄</span> Refresh <span style="float: right;">Total Devices Tracked: 18</span> |                   |                 |      |                |                   |
|---------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|-----------------|------|----------------|-------------------|
| Online                                                                                                                                                  | Device            | OS              | User | IP Address     | FortiClient State |
| ✓                                                                                                                                                       | 00:09:0f:15:04:86 |                 |      |                |                   |
| ✓                                                                                                                                                       | 00:09:0f:9b:24:e1 | Fortinet OS     |      | 172.20.111.100 |                   |
| ✓                                                                                                                                                       | 18:03:73:89:1b:25 |                 |      | 172.20.120.222 |                   |
| ✓                                                                                                                                                       | 24:b6:fd:28:25:26 |                 |      | 172.20.120.220 |                   |
| ✓                                                                                                                                                       | c4:2c:03:0d:3a:38 |                 |      | 172.20.120.51  |                   |
| ✓                                                                                                                                                       | c4:2c:03:21:a9:8e |                 |      | 172.20.120.83  |                   |
| ✓                                                                                                                                                       | c4:2c:03:21:af:04 | iOS / 5.x, 6.0+ |      | 172.20.120.14  |                   |
| ✓                                                                                                                                                       | f0:4d:a2:f1:bf:a3 |                 |      | 172.20.120.26  |                   |
| ✓                                                                                                                                                       | f0:4d:a2:f1:d3:4a |                 |      | 172.20.120.36  |                   |
| ✓                                                                                                                                                       | f0:4d:a2:f1:d6:60 |                 |      | 172.20.120.46  |                   |
|                                                                                                                                                         | 00:0c:29:0e:64:85 |                 |      | 172.20.120.220 |                   |
|                                                                                                                                                         | 00:0c:29:92:7f:4a |                 |      | 172.20.120.52  |                   |
|                                                                                                                                                         | 00:0c:29:df:22:b0 |                 |      | 172.20.120.225 |                   |
|                                                                                                                                                         | 18:03:73:59:b3:3c |                 |      | 172.20.120.224 |                   |
|                                                                                                                                                         | 18:03:73:b6:f9:e9 |                 |      | 172.20.120.100 |                   |
|                                                                                                                                                         | a8:20:66:06:ac:7d |                 |      | 172.20.120.48  |                   |
|                                                                                                                                                         | a8:20:66:14:fa:da | iOS / 5.x, 6.0+ |      | 172.20.120.221 |                   |
|                                                                                                                                                         | b8:ca:3a:c7:e1:ff |                 |      | 172.20.120.223 |                   |

Device monitoring is enabled separately on each interface.

**To configure device monitoring**

1. Go to *System > Network > Interfaces* and edit a FortiGate interface to use for device monitoring.
2. Under *Device Management* select *Detect and Identify Devices*.
3. If you plan to use the Vulnerability scanner to scan discovered devices for vulnerabilities, select *Add New Devices to Vulnerability Scan List*.
4. Select *OK*.
5. Repeat for all interfaces to use for device monitoring.

**To edit device information**

1. Go to *User & Device > Device > Device Definitions* and double-click the entry to edit it.
2. Enter an *Alias* to identify the device.  
This step is compulsory. The alias replaces the MAC address in the device list.
3. If the device can have more than one MAC address, add them to the device.
4. Optionally add the device to a custom device group.
5. Change other information as needed.
6. Select *OK*.

**To add a device manually**

1. Go to *User & Device > Device > Device Definitions* and select *Create New*.
2. Enter the following information.
  - Alias (required)
  - MAC address
  - Device Type
3. Optionally, add additional MAC addresses, select a *Custom Group* and enter *Comments*.
4. Select *OK*.

## Device Groups

Device Groups are used in device policies to specify which devices match the policy. FortiOS automatically adds detected devices of well-known device types to predefined device groups. You can also create custom device groups so that you can create a different policy for specific, known devices.

**Table 2:** Predefined Device Groups

| Group                | Devices                                                                                               |
|----------------------|-------------------------------------------------------------------------------------------------------|
| Android Phone        | Android-based phones.                                                                                 |
| Android Tablet       | Android-based Tablets.                                                                                |
| BlackBerry Phone     | BlackBerry-based phones.                                                                              |
| BlackBerry Playbook  | BlackBerry-based tablets.                                                                             |
| Collected Emails     | All devices from which FortiOS has collected a user email address.                                    |
| Fortinet Device      | FortiGate, FortiManager, FortiAnalyzer, FortiMail, etc.                                               |
| Gaming Console       | All Gaming consoles listed in the Device Visibility database. This includes Xbox, PS2, PS3, Wii, PSP. |
| iPad                 | IOS-based tablets.                                                                                    |
| iPhone               | IOS-based phones.                                                                                     |
| IP Phone             | IP phones.                                                                                            |
| Linux PC             | Linux-based PCs.                                                                                      |
| Mac                  | Apple Macintosh computers.                                                                            |
| Media Streaming      | Media streaming devices such as Apple TV.                                                             |
| Router/NAT Device    | Routers and other gateway devices.                                                                    |
| Windows Phone        | Windows OS based phones.                                                                              |
| Windows PC           | Windows-based PCs.                                                                                    |
| Windows Tablet       | Windows-based tablets.                                                                                |
| Other Network Device | All other network devices not categorized under any other group.                                      |
| All                  | All devices.                                                                                          |

## Creating a custom device group

The predefined device groups are automatically populated. When you create a custom device group, you choose the members. Adding a device that the FortiGate unit has already detected is easiest. But you can also add a device that has not yet been detected if you know its MAC address.

### To create the custom device group

1. Go to *User & Device > Device > Device Groups* and select *Create New*.

2. Enter a name for the group.
3. Add devices to the group.
4. Select *OK*.

## Controlling access with a MAC Address Access Control List

A MAC Address Access Control List is best used to handle exceptions. If you want to limit network access to a larger group, such as your employees, it is better to create a custom device group and specify that group in your device-based security policies.

A MAC Address Access Control List functions as either a list of blocked devices or a list of allowed devices. This is determined by the *Unknown MAC Address* entry.

- By default, unknown MAC addresses are allowed: *Action* is *Assign IP*. You add an entry for each MAC address that you want to block and set its *Action* to *Block*.
- If you want to restrict access to a limited set of devices, you set the *Unknown MAC Address* entry to *Block* and add an entry for each allowed MAC address with *Action* set to *Assign IP*.

### To create a MAC Address Access Control List

1. In the SSID or other interface configuration, select *Enable DHCP Server*.
2. Enter the required *Address Range* and *Netmask*.
3. Expand *MAC Address Access Control List*.
4. Select *Create New* and enter the device's *MAC Address*.
5. Select *Assign IP* to allow the device or *Block* to block the device and then select *OK*.
6. Repeat Steps 4 and 5 for each additional MAC address entry.

## Device policies

Policies based on device identity enable you to implement policies according to device type. For example:

- Gaming consoles cannot connect to the company network or the Internet.
- Personal tablet and phone devices can connect to the Internet but not to company servers.
- Company-issued laptop computers can connect to the Internet and company servers. Web filtering and antivirus are applied.
- Employee laptop computers can connect to the Internet, but web filtering is applied. They can also connect to company networks, but only if FortiClient Endpoint Security is installed to protect against viruses.

Figure 48 shows these policies implemented for WiFi to the company network.



**Figure 48:**Device policies for WiFi access to the company network

**Edit Policy**

Policy Type:  Firewall  VPN

Policy Subtype:  Address  User Identity  Device Identity

Incoming Interface: wifi (SSID: fortinet)

Source Address: all

Outgoing Interface: internal

Enable NAT

Use Destination Interface Address  Fixed Port

Use Dynamic IP Pool

**Configure Authentication Rules**

| Destination Address | Device                                                                                       | Endpoint Compliance | Service | Schedule | UTM Security | Traffic Shaping | Logging | Action                               |
|---------------------|----------------------------------------------------------------------------------------------|---------------------|---------|----------|--------------|-----------------|---------|--------------------------------------|
| all                 | Gaming Console                                                                               | -                   | ALL     | always   | -            | ⊗               | ⊗       | ⊗ DENY                               |
| all                 | Android Phone<br>Android Tablet<br>BlackBerry Phone<br>BlackBerry PlayBook<br>iPad<br>iPhone | -                   | ALL     | always   | -            | ⊗               | ⊗       | ⊗ DENY                               |
| all                 | company laptop                                                                               | ⊗                   | ALL     | always   | 🛡️           | ⊗               | ⊗       | ✅ ACCEPT                             |
| all                 | employee laptop                                                                              | ✅                   | ALL     | always   | -            | ⊗               | ⊗       | ✅ ACCEPT                             |
| all                 | employee laptop                                                                              | -                   | ALL     | always   | -            | ⊗               | ⊗       | 🚪 Captive Portal - Enforce FortiClic |

Customize Authentication Messages

Comments:  0/255

**OK** **Cancel**

Device-based security policies are similar to policies based on user identity:

- The policy enables traffic to flow from one network interface to another.
- NAT can be enabled.
- Authentication rules can allow or deny specific devices or device groups.
- UTM protection can be applied.

### To create a device identity policy

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. In *Policy Subtype*, select *Device Identity*.
3. Choose *Incoming Interface*, *Source Address*, and *Outgoing Interface* as you would for any security policy.
4. Select *Enable NAT* if appropriate.  
You are now ready to create authentication rules.

### To create an authentication rule

1. Select *Create New*.
2. Enter *Destination*, *Schedule*, and *Service* as you would for any security policy.
3. In *Device*, select the devices or device groups to which this policy applies.  
You can select multiple devices or groups.
4. Select *Compliant with Endpoint Profile* if you want to enforce use of FortiClient Endpoint Security by the client devices. This is available here only if Action is ACCEPT. See [“Adding endpoint control”](#) next.
5. Select one of the following for Action:
  - ACCEPT
  - DENY

6. Configure *UTM Security Profiles* as you would for any security policy.
7. Select *OK*.
8. Select *OK* again to complete creation of the security policy.

### Adding endpoint control

Optionally, you can require that user's devices have FortiClient Endpoint Security software installed. The software provides FortiOS more detailed information about the applications being used. FortiOS pushes its endpoint profile to the FortiClient software, configuring network protection such as antivirus, application control, and web category filtering. Devices without an up-to-date installation of FortiClient software are restricted to a captive portal that provides links from which the user can download a FortiClient installer.

If you have already created an ACCEPT rule for particular device groups, you simply edit this rule and enable *Compliant with Endpoint Profile*. Then you add a second rule that sends the same devices to the Enforce FortiClient Compliance captive portal. Devices lacking the required FortiClient software arrive at this policy because they do not match the preceding policy.

**Figure 49:**Endpoint compliance rule and captive portal rule

| Configure Authentication Rules |                 |                     |         |          |              |                 |         |                                          |
|--------------------------------|-----------------|---------------------|---------|----------|--------------|-----------------|---------|------------------------------------------|
| Destination Address            | Device          | Endpoint Compliance | Service | Schedule | UTM Security | Traffic Shaping | Logging | Action                                   |
| all                            | employee laptop | ✓                   | ALL     | always   | -            | ✗               | ✗       | ✓ ACCEPT                                 |
| all                            | employee laptop | -                   | ALL     | always   | -            | ✗               | ✗       | ✗ Captive Portal - Enforce FortiClient C |

## Device policy portal options

The following portal options are available when configuring a device policy:

- Attempt to detect all Unknown device types before implicit deny
- Redirect all non-compliant/unregistered FortiClient compatible devices to a captive portal  
Custom portals are available for Windows, Mac OS, iPhone/iPad and Android devices. These portals acts as a quarantine for devices that are not protected by FortiClient Endpoint Security. The portal provides links to obtain the FortiClient software. The user can retry connecting after installing the FortiClient software.
- Prompt E-mail Collection Portal for all devices  
This portal is used to collect an email address as a means of identifying the device user. When the email address has been verified, the device is added to the Collected Emails device group.

## Creating the WiFi SSID

In order to create a WiFi SSID using the web-based manager, the WiFi Controller (called WiFi & Switch Controller on FortiGate models 100D, 200D, 240D, 600C, 800C, and 1000C) feature must be enabled using Feature Select. For more information, see [“Feature Select” on page 243](#).

Both guest and employee devices will need an SSID (WiFi network) with open security. This means that no passphrase is required to join the SSID. Device policies will determine who gets access to network resources. By default, open security is not available in the WiFi SSID configuration.

**To configure the SSID**

1. Go to *WiFi Controller > WiFi Network > SSID* and select *Create New*.
2. Enter the following information and select OK:

|                                               |                               |
|-----------------------------------------------|-------------------------------|
| <b>Name</b>                                   | byod-example                  |
| <b>IP/Netmask</b>                             | 10.10.110.1/24                |
| <b>Administrative Access</b>                  | Ping (to assist with testing) |
| <b>Enable DHCP Server</b>                     | Enable                        |
| <b>Address Range</b>                          | 10.10.110.2 - 10.10.110.199   |
| <b>Netmask</b>                                | 255.255.255.0                 |
| <b>Default Gateway</b>                        | Same As Interface IP          |
| <b>SSID</b>                                   | byod-guest                    |
| <b>Security Mode</b>                          | Open                          |
| <b>Block Intra-SSID Traffic</b>               | Select.                       |
| Leave other settings at their default values. |                               |

For detailed information about creating a WiFi SSID, see the *Deploying Wireless Networks* chapter of the FortiOS Handbook.

## Configuring Internet access for guests with mobile devices

Guest devices have access only to the Internet. You need a device policy that allows traffic to flow from the WiFi SSID to the Internet interface. Within that policy, you need an authentication rules to allow access for the various types of devices.

**To create the device policy**

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Enter the following information:

|                           |                 |
|---------------------------|-----------------|
| <b>Policy Type</b>        | Firewall        |
| <b>Policy Subtype</b>     | Device Identity |
| <b>Incoming Interface</b> | byod-example    |
| <b>Source Address</b>     | all             |
| <b>Outgoing Interface</b> | wan1            |
| <b>Enable NAT</b>         | Enable.         |

You are now ready to create the authentication rule.

**To create the authentication rule**

1. In *Configure Authentication Rules*, select *Create New* and enter:

|                                        |                        |
|----------------------------------------|------------------------|
| <b>Destination Address</b>             | all                    |
| <b>Device</b>                          | Device or Device Group |
| <b>Compliant with Endpoint Profile</b> | not selected           |
| <b>Schedule</b>                        | always                 |
| <b>Service</b>                         | ALL                    |
| <b>Action</b>                          | ACCEPT                 |

2. Select *OK*.
3. If asked, confirm that you accept FortiOS will enable device identification on the source interface.  
The rule is now configured.
4. Select *OK* to complete configuration of the security policy.

# Client Reputation

The Security scan types available on FortiGate units are varied and tailored to detect specific attacks. However, sometimes user/client behavior can increase the risk of attack or infection. For example, if one of your network clients receives email viruses on a daily basis while no other clients receive these attachments, extra measures may be required to protect the client, or a discussion with the user about this issue may be worthwhile.

Before you can decide on a course of action, you need to know the problem is occurring. Client reputation can provide this information by tracking client behavior and reporting on activities that you determine are risky or otherwise noteworthy.

Activities you can track include:

- **Bad Connection Attempts:** A typical BOT behavior is to connect to some hosts that do not exist on the Internet. This is because the BOT home needs to constantly change itself to dodge legislative enforcement or to hide from AV vendors. Bad connection attempts are tracked by:
  - Look ups for a DNS name that does not exist.
  - Connection attempts to an IP address that has no route.
  - HTTP 404 errors
- Packets that are blocked by security policies.
- **Intrusion protection:** Attack detected. The effect on reputation increases with severity of attack. A subscription to FortiGuard IPS updates is required.
- **Malware protection:** Malware detected. This requires a subscription to FortiGuard Antivirus updates.
- **Web activity:** Visit to web site in risky categories, including Potentially Liable, Adult/Mature Content, Bandwidth Consuming and Security Risk. A subscription to FortiGuard Web Filtering is required.
- **Application protection:** Client uses software in risky categories, including Botnet, P2P, Proxy, and Games applications. A subscription to FortiGuard IPS updates is required.
- **Geographical locations** that clients are communicating with. Access to the FortiGuard geographic database and a valid Fortinet support contract is required.

You can configure how severely each type of tracked activity will impact the reputation of the client in a sliding scale of Low, Medium, High or Critical. You can also choose to ignore an activity by setting it to Off. When an activity is turned off, it will have no effect on reputation.

You can enable client reputation tracking for your FortiGate unit by going to *Security Profiles > Client Reputation > Threat Level Definition*. Turning on client reputation tracking turns on traffic logging for all security policies, for all DoS policies and for all sniffer policies. While client reputation is enabled, logging cannot be turned off for these policies. Traffic logging must be enabled for data to be added to the client reputation database.



Client reputation only highlights risky activity and does not include tools to stop it. Instead, client reputation is a tool that exposes risky behavior. When you uncover risky behavior that you are concerned about, you can take additional action to stop it. That action could include adding more restrictive security policies to block the activity or increase UTM protection. You can also taking other measures outside your FortiGate unit to stop the activity.

---

To support client reputation your FortiGate unit must be registered, have a valid support contract and be licensed for FortiGuard antivirus, IPS and Web Filtering.

After client reputation is turned on, the FortiGate unit tracks recent behavior using a sliding window and displays current data for this window. The client reputation monitor displays clients and their activities in charts ordered according to how risky the behavior exhibited by the client is.

Client Reputation data is stored in traffic log messages in the newly added client reputation fields (crscore and craction). When you enable client reputation *Log UTM Events* or *Log all Sessions* is enabled in all security policies. *Log UTM Events* records traffic log messages for UTM sessions and *Log all Sessions* records traffic logs for all sessions. When Client Reputation is enabled you cannot select *No Log* in a security policy. Using client reputation data in log messages, you can configure FortiAnalyzer to produce a client reputation report.

Enabling client reputation can affect system performance if you had not been using traffic logging.

This chapter describes:

- [Setting the client reputation profile/definition](#)
- [Applying client reputation monitoring to your network](#)
- [Viewing client reputation results](#)
- [Expanding client reputation to include more types of behavior](#)
- [Client reputation execute commands](#)
- [Client reputation diagnose commands](#)

## Setting the client reputation profile/definition

Configure the client reputation profile by going to *Security Profiles > Client Reputation > Threat Level Definition*. You configure one client reputation profile for all of the activity monitored by the FortiGate unit. The profile sets the risk levels for the types of behavior that client reputation monitors. You can set the risk to off, low, medium, high and critical for the following types of behavior:

- Application Protection
  - Botnet applications
  - P2P applications
  - Proxy applications
  - Games applications
- Intrusion protection (IPS)
  - Critical severity attack detected
  - High severity attack detected
  - Medium severity attack detected
  - Low severity attack detected
  - Informational severity attack detected
- Malware Protection
  - Malware detected
  - Botnet connection detected
- Packet based inspection
  - Blocked by firewall policy
  - Failed connection attempts

- Web Activity
  - All blocked URLs
  - Visit to security risk sites
  - Visit to potentially liable sites
  - Visit to adult/mature content sites
  - Visit to bandwidth consuming sites

**Figure 50: Default client reputation profile**

**Threat Level Definition**

**ON Client Reputation Tracking**

**Application Protection**

- Botnet Applications
- P2P Applications
- Proxy Applications
- Games Applications

**Intrusion Protection**

- Critical Severity Attack Detected
- High Severity Attack Detected
- Medium Severity Attack Detected
- Low Severity Attack Detected
- Informational Severity Attack Detected

**Risk Level Values**

LOW 5 MED 10 HIGH 30 CRIT 50

**Malware Protection**

- Malware Detected
- Botnet Connection Detected

**Packet Based Inspection**

- Blocked by Firewall Policy
- Failed Connection Attempts

**Web Activity**

- All Blocked URLs
- Visit to Security Risk Sites
- Visit to Potentially Liable Sites
- Visit to Adult/Mature Content Sites
- Visit to Bandwidth Consuming Sites

**Apply**

To configure the profile, decide how risky or dangerous each of the types of behavior are to your network and rate them accordingly. The higher you rate a type of behavior, the more visible clients engaging in this behavior will become in the client reputation monitor and the more easily you can detect this behavior.

For example, if you consider malware a high risk for your network, you can set the client reputation profile for malware to high or critical (as it is in the default client reputation profile). Then, whenever any amount of malware is detected, clients that originated the malware will be very visible in the client reputation monitor.

Set the risk to off for types of activity that you do not want client reputation to report on. This does not reduce the performance requirements or the amount of data gathered by client reputation, just the report output.

You can change a profile setting at any time and data that has already been collected will be used.

It is normally not necessary to change the *Risk Level Values* but it can be done if you need to alter the relative importance of the risk settings.

## Applying client reputation monitoring to your network

Client reputation monitoring is applied to network traffic by going to *Security Profiles > Client Reputation > Threat Level Definition* turning on *Client Reputation Tracking* and selecting *Apply*.

You can then either change the client reputation profile used by your FortiGate unit or you can accept the default profile. The client reputation profile indicates how risky you consider different types of client behavior to be. See [“Expanding client reputation to include more types of behavior” on page 184](#) for details.

## Viewing client reputation results

Client reputation results can be viewed in the *Threat History* widget, which is found at *System > Dashboard > Threat History*.

The *Threat History* widget displays threat severity over time. Specific time periods can be selected, at which point drilldown menus are available to view more information about the threats, including information about threat types and the sources of each incident.

**Figure 51:** The *Threat History* widget



## Expanding client reputation to include more types of behavior

You can use the following command to change the client reputation profile from the CLI to include client reputation reporting about more settings:

```
config client-reputation profile
```



In addition to the settings configurable from the web-based manager, you can also set the following options:

- `geolocation` to enable reporting on connections to and from different countries (geographical locations). For example, use the following command to indicate that you consider communication with Aruba to be medium risk:

```
config client-reputation profile
 config geolocation
 edit 0
 set country AW
 set level medium
 end
 end
```

- `url-block-detected` to report on connections blocked by web filtering. Use the following command to enable reporting about blocked URLs and set the risk level to medium:

```
config client-reputation profile
 set url-block-detected medium
end
```

From the CLI you can configure client reputation to report more FortiGuard web filtering categories and more types of applications. For example, to report on social network activity (application control category 23):

```
config client-reputation-profile
 config application
 edit 0
 set category 23
 set level medium
 end
 end
```

To report on the local web filtering category (category 22):

```
config client-reputation-profile
 config web
 edit 0
 set group 22
 set level medium
 end
 end
```

## Client reputation execute commands

The `execute client-reputation` command includes the following options:

- `erase`, deletes all client reputation data.
- `host-count`, lists the clients that started sessions recorded by client reputation
- `host-detail`, for a specified client's IP address, displays the client reputation traffic log messages saved for that client.
- `host-summary`, for a specified client's IP address, displays the client's IP address, total entries, and total score.
- `purge`, deletes all data from the client reputation database.
- `topN`, display the top N clients identified by client reputation.

## Client reputation diagnose commands

The `diagnose client-reputation` command includes the following options

- `convert-timestamp` convert a client reputation database timestamp to date and time
- `test-all` adds log messages from multiple sources to the client reputation database for testing
- `test-app` adds application control log messages to the client reputation database for testing
- `test-ips` adds Intrusion Protection log messages to the client reputation database for testing
- `test-webfilter` adds webfilter log messages to the client reputation database for testing

# Wireless

New wireless features include:

- Wireless IDS
- WiFi performance improvements
- FortiAP web-based manager and CLI
- WiFi guest access provisioning
- FortiAP local bridging (Private Cloud-Managed AP)
- WiFi data channel encryption
- Wireless client load balancing for high-density deployments
- Bridge SSID to FortiGate wired network

## Wireless IDS

FortiGate wireless IDS monitors wireless traffic for a wide range of security threats by detecting and reporting on possible intrusion attempts. When an attack is detected, the FortiGate unit records a log message.

You can create a WIDS profile to enable the following types of intrusion detection among others:

- Unauthorized Device Detection
- Rogue/Interfering AP Detection
- Adhoc Network Detection and Containment
- Wireless Bridge Detection
- Misconfigured AP Detection
- Weak WEP Detection
- Multi Tenancy Protection
- MAC OUI Checking

You can enable wireless IDS by going to *WiFi Controller > WiFi Network > Custom AP Profiles* and editing an access point profile or creating a new one.

Inside the profile, set *WIDS Profile* to the name of a wireless IDS profile to apply wireless IDS protection to the access points that uses the profile. FortiGate units include a *default* wireless IDS profile. You can customize this profile or create additional profiles by going to *WiFi Controller > WiFi Network > WIDS Profiles*.

Figure 52:Configuring a WIDS profile

**Edit Wireless Intrusion Detection System Profile** default

Name:

Comments:

| Intrusion Type                            | Status                              | Threshold                | Interval (sec) |
|-------------------------------------------|-------------------------------------|--------------------------|----------------|
| Asleep Attack                             | <input checked="" type="checkbox"/> |                          |                |
| Association Frame Flooding                | <input checked="" type="checkbox"/> | 30 (1 - 100)             | 10 (5 - 120)   |
| Authentication Frame Flooding             | <input checked="" type="checkbox"/> | 30 (1 - 100)             | 10 (5 - 120)   |
| Broadcasting De-authentication            | <input checked="" type="checkbox"/> |                          |                |
| EAPOL-FAIL Flooding (to AP)               | <input checked="" type="checkbox"/> | 10 (2 - 100)             | 1 (1 - 3600)   |
| EAPOL-LOGOFF Flooding (to AP)             | <input checked="" type="checkbox"/> | 10 (2 - 100)             | 1 (1 - 3600)   |
| EAPOL-START Flooding (to AP)              | <input checked="" type="checkbox"/> | 10 (2 - 100)             | 1 (1 - 3600)   |
| EAPOL-SUCC Flooding (to AP)               | <input checked="" type="checkbox"/> | 10 (2 - 100)             | 1 (1 - 3600)   |
| Invalid MAC OUI                           | <input checked="" type="checkbox"/> |                          |                |
| Long Duration Attack                      | <input checked="" type="checkbox"/> | 8200 (1000 - 32767) usec |                |
| Null SSID Probe Response                  | <input checked="" type="checkbox"/> |                          |                |
| Premature EAPOL-FAIL Flooding (to Client) | <input checked="" type="checkbox"/> | 10 (2 - 100)             | 1 (1 - 3600)   |
| Premature EAPOL-SUCC Flooding (to Client) | <input checked="" type="checkbox"/> | 10 (2 - 100)             | 1 (1 - 3600)   |
| Spoofed De-authentication                 | <input checked="" type="checkbox"/> |                          |                |
| Weak WEP IV (Initialization Vector)       | <input checked="" type="checkbox"/> |                          |                |
| Wireless Bridge                           | <input checked="" type="checkbox"/> |                          |                |

**Apply**

You can also use the `config wireless-controller wids-profile` command to configure Wireless Intrusion Detection (WIDS) profiles.

### Syntax

```
config wireless-controller wids-profile
edit <wids-profile_name>
 set comment <comment_str>
 set asleep-attack {enable | disable}
 set assoc-frame-flood {enable | disable}
 set auth-frame-flood {enable | disable}
 set deauth-broadcast {enable | disable}
 set eapol-fail-flood {enable | disable}
 set eapol-fail-intv <int>
 set eapol-fail-thres <int>
 set eapol-logoff-flood {enable | disable}
 set eapol-logoff-intv <int>
 set eapol-logoff-thres <int>
 set eapol-pre-fail-flood {enable | disable}
 set eapol-pre-fail-intv <int>
 set eapol-pre-fail-thres <int>
 set eapol-pre-succ-flood {enable | disable}
 set eapol-pre-succ-intv <int>
 set eapol-pre-succ-thres <int>
 set eapol-start-flood {enable | disable}
 set eapol-start-intv <int>
```

```

set eapol-start-thres <int>
set eapol-succ-flood {enable | disable}
set eapol-succ-intv <int>
set eapol-succ-thres <int>
set invalid-mac-oui {enable | disable}
set long-duration-attack {enable | disable}
set long-duration-thresh <int>
set null-ssid-probe-resp {enable | disable}
set spoofed-death {enable | disable}
set weak-wep-iv {enable | disable}
set wireless-bridge {enable | disable}
end

```

| Variable                                   | Description                                                                                           | Default     |
|--------------------------------------------|-------------------------------------------------------------------------------------------------------|-------------|
| <wids-profile_name>                        | Enter a name for this WIDS profile.                                                                   | No default. |
| comment <comment_str>                      | Optionally, enter a descriptive comment.                                                              | No default. |
| asleap-attack<br>{enable   disable}        | Enable to detect asleep attack (attempt to crack LEAP security).                                      | disable     |
| assoc-frame-flood<br>{enable   disable}    | Enable to detect association frame flood attack.                                                      | disable     |
| auth-frame-flood<br>{enable   disable}     | Enable to detect authentication frame flood attack.                                                   | disable     |
| death-broadcast<br>{enable   disable}      | Enable to detect deauthentication broadcasts which can disrupt wireless services to multiple clients. | disable     |
| eapol-fail-flood<br>{enable   disable}     | Enable to detect EAP FAIL flood attack.                                                               | disable     |
| eapol-fail-intv <int>                      | Set EAP FAIL detection interval.                                                                      | 1           |
| eapol-fail-thres<br><int>                  | Set EAP FAIL detection threshold.                                                                     | 10          |
| eapol-logoff-flood<br>{enable   disable}   | Enable to detect EAP LOGOFF flood attack.                                                             | disable     |
| eapol-logoff-intv<br><int>                 | Set EAP LOGOFF detection interval.                                                                    | 1           |
| eapol-logoff-thres<br><int>                | Set EAP LOGOFF detection threshold.                                                                   | 10          |
| eapol-pre-fail-flood<br>{enable   disable} | Enable to detect EAP premature FAIL flood attack.                                                     | disable     |
| eapol-pre-fail-intv<br><int>               | Set EAP premature FAIL detection interval.                                                            | 1           |
| eapol-pre-fail-thres<br><int>              | Set EAP premature FAIL detection threshold.                                                           | 10          |
| eapol-pre-succ-flood<br>{enable   disable} | Enable to detect EAP premature SUCC flood attack.                                                     | disable     |
| eapol-pre-succ-intv<br><int>               | Set EAP premature SUCC detection interval.                                                            | 1           |
| eapol-pre-succ-thres<br><int>              | Set EAP premature SUCC detection threshold.                                                           | 10          |

| Variable                                   | Description                                                                                                                                   | Default |
|--------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|---------|
| eapol-start-flood<br>{enable   disable}    | Enable to detect EAP START flood attack.                                                                                                      | disable |
| eapol-start-intv<br><int>                  | Set EAP START detection interval.                                                                                                             | 1       |
| eapol-start-thres<br><int>                 | Set EAP START detection threshold.                                                                                                            | 10      |
| eapol-succ-flood<br>{enable   disable}     | Enable to detect EAP SUCC flood attack.                                                                                                       | disable |
| eapol-succ-intv <int>                      | Set EAP SUCC detection interval.                                                                                                              | 1       |
| eapol-succ-thres<br><int>                  | Set EAP SUCC detection threshold.                                                                                                             | 10      |
| invalid-mac-oui<br>{enable   disable}      | Enable to detect use of spoofed MAC addresses. (The first three bytes should indicate a known manufacturer.)                                  | disable |
| long-duration-attack<br>{enable   disable} | Enable for long duration attack detection based on long-duration-thresh.                                                                      | disable |
| long-duration-thresh<br><int>              | Enter the duration in usec for long-duration attack detection. This is available when long-duration-attack is enable.                         | 8200    |
| null-ssid-probe-resp<br>{enable   disable} | Detect attacks that include an incorrectly formed response packets that include a null SSID. This attack can cause wireless clients to crash. | disable |
| spoofed-deauth<br>{enable   disable}       | Enable to detect spoofed deauthentication packets.                                                                                            | disable |
| weak-wep-iv<br>{enable   disable}          | Enable to detect APs using weak WEP encryption.                                                                                               | disable |
| wireless-bridge<br>{enable   disable}      | Enable to detect wireless bridge operation, which is suspicious if your network doesn't use a wireless bridge.                                | disable |

## WiFi performance improvements

FortiOS 5.0 improves performance for WiFi users connecting to the wifi network on a FortiWiFi unit or remotely on FortiAP units. WiFi traffic is now handled in the FortiOS kernel like other network traffic, rather than in a separate application.

## FortiAP web-based manager and CLI

You can now log into a FortiAP web-based manager to view FortiAP status as well as view and change the FortiAP configuration. Logging into the FortiAP web-based manager is similar to logging into the FortiGate web-based manager.

The FortiAP CLI now includes more configuration commands and a complete set of diagnose commands.

Configuration commands include the following

|                               |                                |
|-------------------------------|--------------------------------|
| <code>cfg -h</code>           | Display help for all commands. |
| <code>cfg -r var</code>       | Remove variables.              |
| <code>cfg -e</code>           | Export variables.              |
| <code>cfg -s</code>           | List variables.                |
| <code>cfg -x</code>           | Reset to factory defaults.     |
| <code>cfg -c</code>           | Commit the change to flash.    |
| <code>cfg -a var=value</code> | Add or change variables.       |

Diagnose commands include:

|                                                                       |                                                        |
|-----------------------------------------------------------------------|--------------------------------------------------------|
| <code>cw_diag help</code>                                             | Display help for all diagnose commands.                |
| <code>cw_diag uptime</code>                                           | Show daemon uptime.                                    |
| <code>cw_diag --tlog &lt;on off&gt;</code>                            | Turn on/off telnet log message.                        |
| <code>cw_diag --clog &lt;on off&gt;</code>                            | Turn on/off console log message.                       |
| <code>cw_diag baudrate [9600   19200   38400   57600   115200]</code> | Set the console baud rate.                             |
| <code>cw_diag plain-ctl [0 1]</code>                                  | Show or change current plain control setting.          |
| <code>cw_diag sniff-cfg ip port</code>                                | Set sniff server ip and port.                          |
| <code>cw_diag sniff [0 1 2]</code>                                    | Enable/disable sniff packet.                           |
| <code>cw_diag stats wl_intf</code>                                    | Show wl_intf status.                                   |
| <code>cw_diag admin-timeout [30]</code>                               | Set shell idle timeout in minutes.                     |
| <code>cw_diag -c wtp-cfg</code>                                       | Show current wtp config parameters in control plane.   |
| <code>cw_diag -c radio-cfg</code>                                     | Show current radio config parameters in control plane. |
| <code>cw_diag -c vap-cfg</code>                                       | Show current vaps in control plane.                    |
| <code>cw_diag -c ap-rogue</code>                                      | Show rogue APs pushed by AC for on-wire scan.          |
| <code>cw_diag -c sta-rogue</code>                                     | Show rogue STAs pushed by AC for on-wire scan.         |
| <code>cw_diag -c arp-req</code>                                       | Show scanned arp requests.                             |
| <code>cw_diag -c ap-scan</code>                                       | Show scanned APs.                                      |
| <code>cw_diag -c sta-scan</code>                                      | Show scanned STAs.                                     |
| <code>cw_diag -c sta-cap</code>                                       | Show scanned STA capabilities.                         |
| <code>cw_diag -c wids</code>                                          | Show scanned WIDS detections.                          |

---

|                                          |                                              |
|------------------------------------------|----------------------------------------------|
| <code>cw_diag -c darrp</code>            | Show darrp radio channel.                    |
| <code>cw_diag -c mesh</code>             | Show mesh status.                            |
| <code>cw_diag -c mesh-veth-acinfo</code> | Show mesh veth ac info, and mesh ether type. |
| <code>cw_diag -c mesh-veth-vap</code>    | Show mesh veth vap.                          |
| <code>cw_diag -c mesh-veth-host</code>   | Show mesh veth host.                         |
| <code>cw_diag -c mesh-ap</code>          | Show mesh ap candidates.                     |
| <code>cw_diag -c scan-clr-all</code>     | Flush all scanned AP/STA/ARPs.               |
| <code>cw_diag -c ap-suppress</code>      | Show suppressed APs.                         |
| <code>cw_diag -c sta-deauth</code>       | Ee-authenticate an STA.                      |

---

## WiFi guest access provisioning

Guest access provisioning allows you to easily add guest accounts to your FortiGate unit. These accounts are mainly used to authenticate guest WiFi users for temporary access to a WiFi network managed by a FortiGate unit.

Guest access configuration begins by going to *User & Device > User Definition > User Group* and adding one or more guest user groups.



**Figure 53:** Adding a Guest user group

**New User Group**

Name:

Type:  Firewall  Fortinet Single Sign-On (FSSO)  Guest  RADIUS Single Sign-On (RSSO)

User ID:

Password:

Enable Name

Enable Sponsor:

Enable Company:

Enable Email

Enable Phone Number:  FortiGuard Messaging Service  Custom - SMS Provider: No SMS providers configured

Expire Type:

Default Expire Time:

Enable Batch Guest Account Creation

Many guest account options are available including:

- Email address or user name to identify the guest account
- Requiring a password or no password to log in
- Require a sponsor or company name
- Sending the user's account information to them using email, FortiGuard Messaging Service, or SMS messages
- Configurable account expiry time, starting immediately or after the first login
- Batch guest account creation using auto-generated user IDs and passwords

Guest users are added, removed and managed from *User & Device > User > Guest Management*. From this page, you select the guest user group to change, then add users to it, edit users that have been added or purge all users. When you add the user, you can customize the account expiration date and time.

To provide guest users with their account information, you can select the account from this Guest Management page and select *Send* to print the guest account credentials or send them to the user as an email or SMS message.

## Adding guest access to a WiFi network

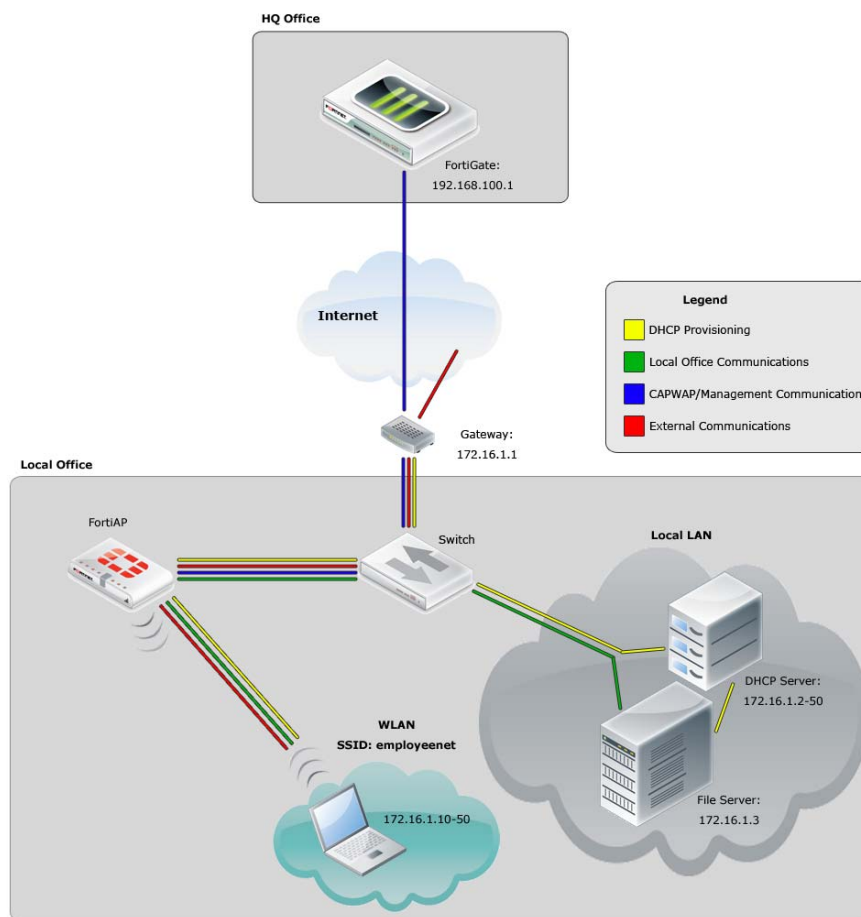
To apply guest access provisioning to a Wifi network, set the *Security Mode* of an SSID to *Captive Portal* and select one or more guest user groups. Guest users can then log into the portal using their guest account.

## FortiAP local bridging (Private Cloud-Managed AP)

A FortiAP unit can provide WiFi access to a LAN, even when the wireless controller is located remotely. This configuration is useful for the following situations:

- Installations where the WiFi controller is remote and most of the traffic is local or uses the local Internet gateway
- Wireless-PCI compliance with remote WiFi controller
- Telecommuting, where the FortiAP unit has the WiFi controller IP address pre-configured and broadcasts the office SSID in the user's home or hotel room. In this case, data is sent in the wireless tunnel across the Internet to the office and should have encryption using DTLS enabled.

**Figure 54:**Remotely-managed FortiAP providing WiFi access to local network



On the remote FortiGate wireless controller, the WiFi SSID is created with the *Traffic Mode* set to *Local Bridge with FortiAP's Interface*. In this mode, no IP addresses are configured. The FortiAP unit's WiFi and Ethernet interfaces behave like a switch. WiFi client devices obtain IP addresses from the same DHCP server as wired devices on the LAN.

There can be only one Bridge mode SSID per FortiAP unit. The Local Bridge feature cannot be used in conjunction with Wireless Mesh features.

**To configure a FortiAP local bridge - web-based manager**

1. Go to *WiFi Controller > WiFi Network > SSID* and select *Create New*.
2. Enter:

|                       |                                       |
|-----------------------|---------------------------------------|
| <b>Interface name</b> | A name for the new WiFi interface.    |
| <b>Traffic Mode</b>   | Local bridge with FortiAP's Interface |

|                                                                        |                                                             |
|------------------------------------------------------------------------|-------------------------------------------------------------|
| <b>SSID</b>                                                            | The SSID visible to users.                                  |
| <b>Security Mode</b><br><b>Data Encryption</b><br><b>Preshared Key</b> | Configure security as you would for a regular WiFi network. |

3. Select OK.
4. Go to *WiFi Controller > Managed Access Points > Managed FortiAP*, select the FortiAP unit for editing.
5. Authorize the FortiAP unit.
6. The FortiAP unit can carry regular SSIDs in addition to the Bridge SSID.

**Figure 55:**SSID configured for Local Bridge operation

### To configure a FortiAP local bridge - CLI

This example creates a WiFi interface “branchbridge” with SSID “LANbridge” using WPA-Personal security, passphrase “Fortinet1”.

```
config wireless-controller vap
 edit "branchbridge"
 set vdom "root"
 set ssid "LANbridge"
 set local-bridging enable
 set security wpa-personal
 set passphrase "Fortinet1"
 end
config wireless-controller wtp
 edit FAP22B3U11005354
 set admin enable
 set vaps "branchbridge"
 end
```

## WiFi data channel encryption

You can enhance the security of communication between a FortiGate wireless controller and a FortiAP unit by applying DTLS encryption to the data channel.

There are data channel encryption settings on both the FortiGate unit and the FortiAP unit. At both ends, you can enable Clear Text, DTLS encryption or both. The settings must agree or the FortiAP unit will not be able to join the WiFi network. By default, both Clear Text and DTLS-encrypted communication are enabled on the FortiAP unit, allowing the FortiGate setting to determine whether data channel encryption is used. If the FortiGate unit also enables both Clear Text and DTLS, Clear Text is used.

Data channel encryption settings are located in the Custom AP profile. If you use Automatic profile, only Clear Text is supported.



Data channel encryption is software-based and can affect performance. Verify that the system meets your performance requirements with encryption enabled.

## Configuring DTLS on the FortiGate unit

To enable DTLS for the FAP320B-default profile, enter:

```
config wireless-controller wtp-profile
 edit FAP320B-default
 set dtls-policy dtls-enabled
 end
```

## Configuring encryption on the FortiAP unit

The FortiAP unit has its own settings for data channel encryption.

### Enabling CAPWAP encryption - FortiAP web-based manager

- 1 On the *System Information* page, in *WTP Configuration > AC Data Channel Security*, select one of:
  - Clear Text
  - DTLS Enabled
  - Clear Text or DTLS Enabled (default)
- 2 Select *Apply*.

### Enabling encryption - FortiAP CLI

You can set the data channel encryption using the AC\_DATA\_CHAN\_SEC variable:

- 0 is Clear Text
- 1 is DTLS Enabled
- 2 is Clear Text or DTLS Enabled (default)

For example, to set security to DTLS and then save the setting, enter

```
cfg -a AC_DATA_CHAN_SEC=1
cfg -c
```

## Wireless client load balancing for high-density deployments

Wireless load balancing allows your wireless network to distribute wireless traffic more efficiently among wireless access points and available frequency bands. FortiGate wireless controllers support the following types of client load balancing:

- Access Point Hand-off - the wireless controller signals a client to switch to another access point.
- Frequency Hand-off - the wireless controller monitors the usage of 2.4GHz and 5GHz bands, and signals clients to switch to the lesser-used frequency.

Load balancing is not applied to roaming clients.

### Access point hand-off

Access point handoff wireless load balancing involves the following:

- If the load on an access point (ap1) exceeds a threshold (for example, 30 clients) then the client with the weakest signal will be signaled by the wireless controller to drop off and join another nearby access point (ap2).
- When one or more access points are overloaded (more than 30 clients) and a new client attempts to join a wireless network, the wireless controller selects the least busy access point that is closest to the new client and this access point is the one that responds to the client and the one that the client joins.

### Frequency hand-off or band-steering

Encouraging clients to use the 5GHz WiFi band if possible enables those clients to benefit from faster interference-free 5GHz communication. The remaining 2.4GHz clients benefit from reduced interference.

The WiFi controller probes clients to determine their WiFi band capability. It also records the RSSI (signal strength) for each client on each band.

If a new client attempts to join the network, the controller looks up that client's MAC address in its wireless device table and determines if it's a dual band device. If it is not a dual band device, then its allowed to join. If it is a dual band device, then its RSSI on 5GHz is used to determine whether the device is close enough to an access point to benefit from movement to 5GHz frequency.

If both conditions of 1) dual band device and 2) RSSI value is strong, then the wireless controller does not reply to the join request of the client. This forces the client to retry a few more times and then timeout and attempt to join the same SSID on 5GHz. Once the controller see this new request on 5GHz, the RSSI is again measured and the client is allowed to join. If the RSSI is below threshold, then the device table is updated and the controller forces the client to timeout again. A client's second attempt to connect on 2.4GHz will be accepted.

### Configuration

From the web-based manager, edit a custom AP profile and select *Frequency Handoff* and *AP Handoff* as required for each radio on the AP.

From the CLI, you configure wireless client load balancing thresholds for each custom AP profile. Enable access point hand-off and frequency hand-off separately for each radio in the custom AP profile.

```
config wireless-controller wtp-profile
 edit new-ap-profile
 set handoff-rssi <rssi_int>
 set handoff-sta-thresh <clients_int>
 config radio-1
 set frequency-handoff {disable | enable}
 set ap-handoff {disable | enable}
 end
 config radio-2
 set frequency-handoff {disable | enable}
 set ap-handoff {disable | enable}
 end
 end
end
```

Where:

- `handoff-rssi` is the RSSI threshold. Clients with a 5 GHz RSSI threshold over this value are load balanced to the 5GHz frequency band. Default is 25. Range is 20 to 30.
- `handoff-sta-thresh` is the access point handoff threshold. If the access point has more clients than this threshold it is considered busy and clients are changed to another access point. Default is 30, range is 5 to 25.
- `frequency-handoff` enable or disable frequency handoff load balancing for this radio. Disabled by default.
- `ap-handoff` enable or disable access point handoff load balancing for this radio. Disabled by default.



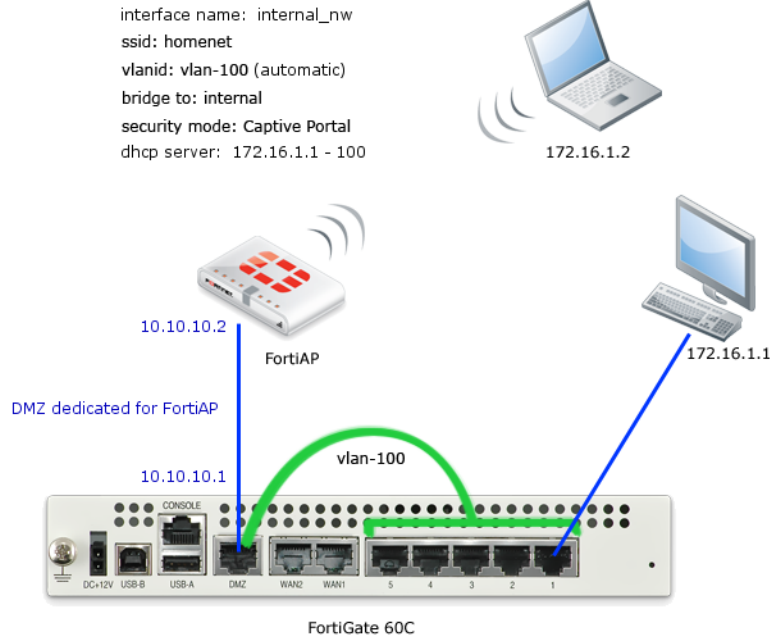
Frequency handoff must be enabled on the 5GHz radio to learn client capability.

---

## Bridge SSID to FortiGate wired network

A WiFi network can be combined with a wired LAN so that WiFi and wired clients are on the same subnet. This is a convenient configuration for users.

Figure 56:A FortiAP unit bridged with the internal network



This configuration cannot be used in conjunction with Wireless Mesh features because it enables the FortiAP Local Bridge option.

To create the bridged WiFi and wired LAN configuration, you need to configure the SSID with the Local Bridge option so that traffic is sent directly over the FortiAP unit's Ethernet interface to the FortiGate unit, instead of being tunneled to the WiFi controller.

Figure 57:SSID configured with Local Bridge option

| New Interface                                                           |                                                                                                |
|-------------------------------------------------------------------------|------------------------------------------------------------------------------------------------|
| Name                                                                    | Local-Bridge                                                                                   |
| Type                                                                    | WiFi SSID                                                                                      |
| Traffic Mode                                                            | Local bridge with FortiAP's Interfac                                                           |
| WiFi Settings                                                           |                                                                                                |
| SSID                                                                    | my-SSID                                                                                        |
| Security Mode                                                           | WPA/WPA2-Personal                                                                              |
| Data Encryption                                                         | <input checked="" type="radio"/> AES <input type="radio"/> TKIP <input type="radio"/> TKIP-AES |
| Pre-shared Key                                                          | ..... (8 - 63 characters)                                                                      |
| Maximum Clients                                                         | <input type="checkbox"/>                                                                       |
| Comments                                                                | Write a comment... 0/255                                                                       |
| <input type="button" value="OK"/> <input type="button" value="Cancel"/> |                                                                                                |

Enter the following command from the CLI:

```
config wireless-controller vap
 edit "homenet_if"
 set vdom "root"
 set ssid "homenet"
 set local-bridging enable
 set security wpa-personal
 set passphrase "Fortinet1"
 end
config wireless-controller wtp
 edit FAP22B3U11005354
 set admin enable
 set vaps "homenet_if"
 end
```



# IPv6

In order to configure IPv6 features using the web-based manager, IPv6 must be enabled using Feature Select. For more information, see “[Feature Select](#)” on page 243.


The following new IPv6 features are available from the FortiOS 5.0 web-based manager:

- IPv6 Policy routing
- IPv6 security policies
- IPv6 Explicit web proxy
- IPv6 NAT – NAT64, DNS64, NAT66
- IPv6 Forwarding Policies - IPS, Application Control, and flow-based antivirus, web filtering and DLP
- New Fortinet FortiGate IPv6 MIB fields
- IPv6 Per-IP traffic shaper
- DHCPv6 relay
- FortiGate interfaces can get IPv6 addresses from an IPv6 DHCP server

## IPv6 Policy routing

IPv6 policy routing functions in the same way as IPv4 policy routing. To add an IPv6 policy route, go to *Router > Static > Policy Routes* and select *Create New > IPv6 Policy Route*.

**Figure 58:** IPv6 policy route



**Figure 59:** Adding an IPv6 Policy route

You can also use the following command to add IPv6 policy routes:

```
config router policy6
 edit 0
 set input-device <interface>
 set src <ipv6_ip>
 set dst <ipv6_ip>
 set protocol <0-255>
 set gateway <ipv6_ip>
 set output-device <interface>
 set tos <bit_pattern>
 set tos-mask <bit_mask>
 end
```

## IPv6 security policies

IPv6 security policies now support all the features supported by IPv4 security policies. The following new features were added in FortiOS 5.0:

- Policy types and subtypes.
- NAT support including using the destination interface IP address, fixed port, and dynamic IP pools.
- All security features (antivirus, web filtering, application control, IPS, email filtering, DLP, VoIP and ICAP).
- All traffic shaping options, including shared traffic shaping, reverse shared traffic shaping and per-IP traffic shaping.
- All user and device authentication options.

## IPv6 Explicit web proxy

With FortiOS 5.0, you can use the explicit web proxy for IPv6 traffic. To do this you need to:

- Enable the Explicit Proxy from the System Information dashboard widget.
- Enable the IPv6 explicit web proxy from the CLI.
- Enable the explicit web proxy for one or more FortiGate interfaces. These interfaces also need IPv6 addresses.
- Add IPv6 web proxy security policies to allow the explicit web proxy to accept IPv6 traffic.

Use the following steps to set up a FortiGate unit to accept IPv6 traffic for the explicit web proxy at the Internal interface and forward IPv6 explicit proxy traffic out the wan1 interface to the Internet.

1. Enter the following CLI command to enable the IPv6 explicit web proxy:

```
config web-proxy explicit
 set status enable
 set ipv6-status enable
end
```

2. Go to *System > Network > Interfaces* and edit the *internal* interface, select *Enable Explicit Web Proxy* and select *OK*.
3. Go to *Policy > Policy > IPv6 Policy* and select *Create New* to add an IPv6 explicit web proxy security policy with the following settings shown in [Figure 60](#).

This IPv6 explicit web proxy policy allows traffic from all IPv6 IP addresses to connect through the explicit web proxy and through the wan1 interface to any IPv6 addresses that are accessible from the wan1 interface.



If you have enabled both the IPv4 and the IPv6 explicit web proxy you can combine IPv4 and IPv6 addresses in a single explicit web proxy policy to allow both IPv4 and IPv6 traffic through the proxy.

---

**Figure 60:**Example IPv6 Explicit Web Proxy security policy

The screenshot shows the 'New Policy' configuration interface. The policy is of type 'Firewall' and subtype 'Address'. The incoming interface is 'web-proxy' and the outgoing interface is 'port2'. Both source and destination IPv6 addresses are set to 'all'. The schedule is 'always' and the service is 'webproxy'. The action is 'ACCEPT'. Logging is enabled for allowed traffic. Under 'UTM Security Profiles', AntiVirus and Web Filter are turned on, while others are off. The UTM Proxy Options are set to 'default'. There are 'OK' and 'Cancel' buttons at the bottom.

## Restricting the IP address of the explicit IPv6 web proxy

You can use the following command to restrict access to the IPv6 explicit web proxy using only one IPv6 IP address. The IPv6 address that you specify must be the IPv6 address of an interface that the explicit HTTP proxy is enabled on. You might want to use this option if the explicit web proxy is enabled on an interface with multiple IPv6 addresses.

For example, to require users to connect to the IPv6 address 2001:db8:0:2::30 to connect to the explicit IPv6 HTTP proxy:

```
config web-proxy explicit
 set incoming-ipv6 2001:db8:0:2::30
end
```

## Restricting the outgoing source IP address of the IPv6 explicit web proxy

You can use the following command to restrict the source address of outgoing web proxy packets to a single IPv6 address. The IP address that you specify must be the IPv6 address of an interface that the explicit HTTP proxy is enabled on. You might want to use this option if the explicit HTTP proxy is enabled on an interface with multiple IPv6 addresses.

For example, to restrict the outgoing packet source address to 2001:db8:0:2::50:

```
config http-proxy explicit
 set outgoing-ip6 2001:db8:0:2::50
end
```

## IPv6 NAT – NAT64, DNS64, NAT66

NAT66, NAT64 and DNS64 are now supported for IPv6. These options provide IPv6 NAT and DNS capabilities IPv6-IPv4 tunnelling or dual stack configurations. These are available only in the CLI.


Fortinet supports all features described in RFC 6146. However, for DNS64 there is no support for handling Domain Name System Security Extensions (DNSSEC). DNSSEC is for securing types of information that are provided by the DNS as used on an IP network or networks. You can find more information about DNS64 in RFC 6147.

### NAT64 and DNS64

NAT64 is used to translate IPv6 addresses to IPv4 addresses so that a client on an IPv6 network can communicate transparently with a server on an IPv4 network. NAT64 is usually implemented in combination with DNS64. DNS64 synthesizes AAAA records from A records and is used to synthesize IPv6 addresses for hosts that only have IPv4 addresses.

With a NAT64 and DNS64 configuration in place on a FortiGate unit, clients on an IPv6 network can transparently connect to addresses on an IPv4 network. NAT64 and DNS64 perform IPv4 to IPv6 transition, allowing clients that have already switched to IPv6 addresses to continue communicating with servers that still use IPv4 addresses.

**Figure 61:** Example NAT64 configuration



### To configure NAT64 to allow a host on the IPv6 network to connect to the Internet server

In this example the Internal IPv6 network address is 2001:db8:1::/48 and the external IPv4 network address is 172.20.120.0/24. NAT64 is configured to allow a user on the internal network to connect to the server at IPv4 address 172.20.120.12. In this configuration, sessions exiting the wan1 interface must have their source address changed to and IPv4 address in the range 172.20.120.200 to 172.20.120.210.

1. Enter the following command to enable NAT64.

```
config system nat64
 set status enable
end
```



Enabling NAT64 with the `config system nat64` command means that all IPv6 traffic received by the current VDOM can be subject to NAT64 if the source and destination address matches an NAT64 security policy.

By default, the setting `always-synthesize-aaaa-record` is not enabled. With this setting disabled, the DNS proxy will attempt to find an AAAA records for queries to domain names and therefore resolve the host names to IPv6 addresses. If the DNS proxy cannot find an AAAA record, it synthesizes one by adding the NAT64 prefix to the A record.

By using the `nat64-prefix` option of the `config system nat64` command to change the default nat64 prefix (the default is the well known prefix `64:ff9b::/96`) and setting `always-synthesize-aaaa-record` to enable, the DNS proxy does not check for AAAA records and always synthesizes AAAA records.

As an alternative to the above entry, there is the optional configuration that would allow the resolution of CNAME queries.

```
config system nat64
 set status enable
 set nat64-prefix 64:ff9b::/96
 set always-synthesize-aaaa-record enable
end
```

2. Enter the following command to add an IPv6 firewall address for the internal network:

```
config firewall address6
 edit internal-net6
 set ip6 2001:db8::/48
 end
```

3. Enter the following command to add an IPv4 firewall address for the external network:

```
config firewall address
 edit external-net4
 set subnet 172.20.120.0/24
 set associated-interface wan1
 end
```

4. Enter the following command to add an IP pool containing the IPv4 address that the should become the source address of the packets exiting the wan1 interface:

```
config firewall ippool
 edit exit-pool4
 set startip 172.20.120.200
 set endip 172.20.120.210
 end
```

5. Enter the following command to add a NAT64 policy that allows connections from the internal IPv6 network to the external IPv4 network:

```
config firewall policy64
 edit 0
 set srcintf internal
 set srcaddr internal-net6
 set dstintf wan1
 set dstaddr external-net4
 set action accept
 set schedule always
 set service ANY
 set logtraffic enable
 set ippool enable
 set poolname exit-pool4
 end
```



The `srcaddr` can be any IPv6 firewall address and the `dstaddr` can be any IPv4 firewall address.

Other NAT64 policy options include `fixedport`, that can be used to prevent NAT64 from changing the destination port. You can also configure traffic shaping for NAT64 policies.

#### How a host on the internal IPv6 network communicates with `example.server.com` that only has IPv4 address on the Internet

1. The host on the internal network does a DNS lookup for `example.server.com` by sending a DNS query for an AAAA record for `example.server.com`.
2. The DNS query is intercepted by the FortiGate DNS proxy.
3. The DNS proxy attempts to resolve the query with a DNS server on the Internet and discovers that there are no AAAA records for `example.server.com`.



The previous step is skipped if `always-synthesize-aaaa-record` is enabled.

4. The DNS proxy performs an A-record query for `example.server.com` and gets back an RRSet containing a single A record with the IPv4 address `172.20.120.12`.
5. The DNS proxy then synthesizes an AAAA record. The IPv6 address in the AAAA record begins with the configured NAT64 prefix in the upper 96 bits and the received IPv4 address in the lower 32 bits. By default, the resulting IPv6 address is `64:ff9b::172.20.120.12`.
6. The host on the internal network receives the synthetic AAAA record and sends a packet to the destination address `64:ff9b::172.20.120.12`.
7. The packet is routed to the FortiGate internal interface where it is accepted by the NAT64 security policy.
8. The FortiGate unit translates the destination address of the packets from IPv6 address `64:ff9b::172.20.120.12` to IPv4 address `172.20.120.12` and translates the source address of the packets to `172.20.120.200` (or another address in the IP pool range) and forwards the packets out the `wan1` interface to the Internet.

## NAT66

NAT66 is used for translating an IPv6 source or destination address to a different IPv6 source or destination address. NAT66 is not as common or as important as IPv4 NAT, as many IPv6 IP addresses do not need NAT66 as much as IPv4 NAT. However, NAT66 can be useful for a number of reasons. For example, you may have changed the IP addresses of some devices on your network but want traffic to still appear to be coming from their old addresses. You can use NAT66 to translate the source addresses of packets from the devices to their old source addresses.

In FortiOS 5.0, NAT66 options can be added to an IPv6 security policy from the CLI.

Configuring NAT66 is very similar to configuring NAT in an IPv4 security policy. For example, use the following command to add an IPv6 security policy that translates the source address of IPv6 packets to the address of the destination interface (similar to IPv4 source NAT):

```
config firewall policy6
 edit 0
 set srcintf internal
 set dstintf wan1
 set srcaddr internal_net
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
 set nat enable
 end
```

It also can be useful to translate one IPv6 source address to another address that is not the same as the address of the exiting interface. You can do this using IP pools. For example, enter the following command to add an IPv6 IP pool containing one IPv6 IP address:

```
configure firewall ippool6
 edit example_6_pool
 set startip 2001:db8::
 set endip 2001:db8::
 end
```

Enter the following command to add an IPv6 firewall address that contains a single IPv6 IP address.

```
configure firewall address6
 edit device_address
 set ip6 2001:db8::132/128
 end
```



Enter the following command to add an IPv6 security policy that accepts packets from a device with IP address 2001:db8::132 and translates the source address to 2001:db8:::

```
config firewall policy6
 edit 0
 set srcintf internal
 set dstintf wan1
 set srcaddr device_address
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
 set nat enable
 set ippool enable
 set poolname example_6_pool
 end
```

## NAT66 destination address translation

NAT66 can also be used to translate destination addresses. This is done in an IPv6 policy by using IPv6 virtual IPs. For example, enter the following command to add an IPv6 virtual IP that maps destination address 2001:db8::dd to 2001:db8::ee

```
configure firewall vip6
 edit example-vip6
 set extip 2001:db8::dd
 set mappedip 2001:db8::ee
 end
```

Enter the following command to add an IPv6 security policy that accepts packets with a destination address 2001:db8::dd and translates that destination address to 2001:db8::ee

```
config firewall policy6
 edit 0
 set srcintf internal
 set dstintf wan1
 set srcaddr all
 set dstaddr example-vip6
 set action accept
 set schedule always
 set service ANY
 end
```

## IPv6 Forwarding Policies - IPS, Application Control, and flow-based antivirus, web filtering and DLP

FortiOS 5.0 fully supports flow-based inspection of IPv6 traffic. This includes full support for IPS, application control, as well as flow-based virus scanning, and web filtering.

To add flow-based inspection to IPv6 traffic go to *Policy > Policy > IPv6 Policy* and select *Create New* to add an IPv6 Security Policy. Configure the policy to accept the traffic to be scanned. Select UTM and select the UTM profiles to apply to the traffic. To apply flow-based

inspection you can select an IPS and an application control profile. You can also select antivirus or web filtering profiles in which flow-based inspection has been selected.

## New Fortinet FortiGate IPv6 MIB fields

The following IPv6 MIB fields have been added to the Fortinet FortiGate MIB. These MIS entries can be used to display IPv6 session and policy statistics.

- IPv6 Session Counters:

- `fgSysSes6Count`
- `fgSysSes6Rate1`
- `fgSysSes6Rate10`
- `fgSysSes6Rate30`
- `fgSysSes6Rate60`

- IPv6 Policy Statistics:

- `fgFwPol6StatsTable`
- `fgFwPol6StatsEntry`
- `FgFwPol6StatsEntry`
- `fgFwPol6ID`
- `fgFwPol6PktCount`
- `fgFwPol6ByteCount`

- IPv6 Session Statistics:

- `fgIp6SessStatsTable`
- `fgIp6SessStatsEntry`
- `FgIp6SessStatsEntry`
- `fgIp6SessNumber`

The `fgSysSesCount` and `fgSysSesRateX` MIBs report statistics for IPv4 plus IPv6 sessions combined. This behavior was not changed.

## New OIDs

The following OIDs have been added:

```
FORTINET-FORTIGATE-MIB:fortinet.fnFortiGateMib.fgSystem.fgSystemInfo
.fgSysSes6Count 1.3.6.1.4.1.12356.101.4.1.15
.fgSysSesRate1 1.3.6.1.4.1.12356.101.4.1.16
.fgSysSesRate10 1.3.6.1.4.1.12356.101.4.1.17
.fgSysSesRate30 1.3.6.1.4.1.12356.101.4.1.18
.fgSysSesRate60 1.3.6.1.4.1.12356.101.4.1.19
```

```
FORTINET-FORTIGATE-MIB:fortinet.fnFortiGateMib.fgFirewall.fgFwPolicies
.fgFwPolTables
.fgFwPol6StatsTable.fgFwPol6StatsEntry.fgFwPol6ID 1.3.6.1.4.1.12356.
101.5.1.2.2.1.1
.fgFwPol6StatsTable.fgFwPol6StatsEntry.fgFwPol6PktCount 1.3.6.1.4.1.
12356.101.5.1.2.2.1.2
.fgFwPol6StatsTable.fgFwPol6StatsEntry.fgFwPol6ByteCount 1.3.6.1.4.1
.12356.101.5.1.2.2.1.3
```

```
FORTINET-FORTIGATE-MIB:fortinet.fnFortiGateMib.fgInetProto.fgInetProto
Tables
.fgIp6SessStatsTable.fgIp6SessStatsEntry.fgIp6SessNumber
1.3.6.1.4.1.12356.101.11.2.3.1.1
```

## EXAMPLE SNMP get/walk output

```
// Session6 stats excerpt from sysinfo:
snmpwalk -v2c -cpublic 192.168.1.111 1.3.6.1.4.1.12356.101.4
FORTINET-FORTIGATE-MIB::fgSysSes6Count.0 = Gauge32: 203
FORTINET-FORTIGATE-MIB::fgSysSes6Rate1.0 = Gauge32: 10 Sessions Per
Second
FORTINET-FORTIGATE-MIB::fgSysSes6Rate10.0 = Gauge32: 2 Sessions Per
Second
FORTINET-FORTIGATE-MIB::fgSysSes6Rate30.0 = Gauge32: 1 Sessions Per
Second
FORTINET-FORTIGATE-MIB::fgSysSes6Rate60.0 = Gauge32: 0 Sessions Per
Second

// FwPolicy6 table:
snmpwalk -v2c -cpublic 192.168.1.111 1.3.6.1.4.1.12356.101.5.1.2.2
FORTINET-FORTIGATE-MIB::fgFwPol6ID.1.3 = INTEGER: 3
FORTINET-FORTIGATE-MIB::fgFwPol6ID.1.4 = INTEGER: 4
FORTINET-FORTIGATE-MIB::fgFwPol6PktCount.1.3 = Counter64: 4329
FORTINET-FORTIGATE-MIB::fgFwPol6PktCount.1.4 = Counter64: 0
FORTINET-FORTIGATE-MIB::fgFwPol6ByteCount.1.3 = Counter64: 317776
FORTINET-FORTIGATE-MIB::fgFwPol6ByteCount.1.4 = Counter64: 0

// IP6SessNumber:
snmpwalk -v2c -cpublic 192.168.1.111 1.3.6.1.4.1.12356.101.11.2.3.1
FORTINET-FORTIGATE-MIB::fgIp6SessNumber.1 = Counter32: 89
```

## IPv6 Per-IP traffic shaper

You can add any Per-IP traffic shaper to an IPv6 security policy using the following command:

```
config firewall policy6
 edit 0
 set per-ip-shaper 'new-perip-shaper'
 end
```

## DHCPv6 relay

You can use the following command to configure a FortiGate interface to relay DHCPv6 queries and responses from one network to a network with a DHCPv6 server and back. The command enables DHCPv6 relay and includes adding the IPv6 address of the DHCP server that the FortiGate unit relays DHCPv6 requests to:

```
config system interface
 edit internal
 config ipv6
 set dhcp6-relay-service enable
 set dhcp6-relay-type regular
 set dhcp6-relay-ip 2001:db8:0:2::30
 end
```

## FortiGate interfaces can get IPv6 addresses from an IPv6 DHCP server

From the CLI you can configure any FortiGate interface to get an IPv6 address from an IPv6 DHCP server. For example, to configure the wan2 interface to get an IPv6 address from an IPv6 DHCP server enter the following command:

```
config system interface
 edit wan2
 config ipv6
 set ip6-mode dhcp
 end
 end
end
```

# Logging and reporting

New logging and reporting features include:

- [Log message reorganization](#)
- [Log Viewer Improvements](#)
- [The FortiGate Security Analysis Report](#)
- [Converting compact log format](#)

FortiCloud, the new logging and reporting feature that replaces FAMS, is described in FortiOS 5.0 [Installation and System Administration](#).

## Log message reorganization

FortiOS 5.0 log messages have been re-organized. The new log message format will be described in the FortiOS 5.0 [Logging and Reporting Guide](#) and [Log Message Reference](#).

## Log Viewer Improvements

In FortiOS 5.0, the Log & Report menu provides access to the following types of log messages:

- Traffic log
  - Forward Traffic - log messages for traffic passing through the FortiGate unit. Includes traffic log messages as well as Security log messages so that you can view messages about Security events (such as a message indicating that a virus was found) in the same location as the traffic log messages that recorded the current traffic at the time. From the forward traffic log viewer, you can also view content logs and quarantined files. Forward traffic log messages also include log messages created when you enable logging for the IPv4 and IPv6 security policy implicit security policies.
  - Local Traffic - Log messages for traffic terminating at the FortiGate unit. All traffic terminating at the FortiGate unit is allowed or denied by a local in policy. To view local in policies go to *Policy > Policy > Local In Policy*. You can enable and disable logging for local in traffic from here as well.
  - Multicast Traffic - Log messages for multicast traffic passing through the FortiGate unit and allowed by multicast traffic policies
  - Invalid Packets - Log messages recorded when the FortiGate unit receives invalid packets.
- Event log - Event log messages organized into the following categories:
  - System
  - Router
  - VPN
  - User
  - WAN Opt. & Cache
  - WiFi

**Figure 62:**Example Forward Traffic log display showing some Security log messages

| #  | Date/Time | Src            | Device | Dst            | Application Name | UTM Action | Sent / R      |
|----|-----------|----------------|--------|----------------|------------------|------------|---------------|
| 1  | Thursday  | 172.20.120.221 | 🇺🇸     | 172.20.120.13  |                  |            | 384 B / 552   |
| 2  | Thursday  | 11.11.11.12    | 🇺🇸     | 17.158.28.36   | SSL              |            | 1.79 KB / 4.1 |
| 3  | Thursday  | 11.11.11.12    | 🇺🇸     | 208.91.113.212 | SSL              |            | 16.85 KB / 1  |
| 4  | Thursday  | 11.11.11.12    | 🇺🇸     | 208.91.113.212 | SSL              |            | 1.66 KB / 1.0 |
| 5  | Thursday  | 11.11.11.12    | 🇺🇸     | 208.91.113.212 | SSL              |            | 2.04 KB / 4.9 |
| 6  | Thursday  | 11.11.11.12    | 🇺🇸     | 199.47.218.147 | HTTP.BROWSER     |            | 594 B / 810   |
| 7  | Thursday  | 11.11.11.12    | 🇺🇸     | 199.47.216.177 | Dropbox          |            | 1.32 KB / 5.0 |
| 8  | Thursday  | 11.11.11.12    | 🇺🇸     | 64.94.18.154   | LogMeIn          |            | 1.20 KB / 3.1 |
| 9  | Thursday  | 11.11.11.12    | 🇺🇸     | 192.168.110.9  | DNS              |            | 59 B / 281 B  |
| 10 | Thursday  | 11.11.11.12    | 🇺🇸     | 17.171.4.21    |                  |            | 380 B / 380   |

Refresh Download Raw Log Log location: Disk

0 / 1293 [ Total: 64607 ]

| Application Control List | client-reputation | Date/Time        | Thursday (Thu Sep 20 13:17:51 2012) |
|--------------------------|-------------------|------------------|-------------------------------------|
| Destination Country      | Reserved          | Dst              | 172.20.120.13                       |
| Dst Interface            | internal          | Dst NAT IP       | 11.11.11.40                         |
| Dst NAT Port             | 3389              | Dst Port         | 3389                                |
| Duration                 | 126               | Level            | notice                              |
| Policy ID                | 2                 | Protocol         | 6                                   |
| Received                 | 552               | Received Packets | 6                                   |

From each of these log message viewers, you can download raw log messages, adjust the column settings to display the ones you are most interested in, filter the messages according to a wide range of criteria and display detailed information for a selected log message. In addition, you can control the source of the log messages that are displayed. For example, setting the log location source can be the disk, if your FortiGate unit has available internal storage.

Other logging improvements include:

- New event log types including event logs for FortiClient usage.
- Traffic log messages for packets blocked by the implicit policies that appear at the bottom of IPv4 and IPv6 policy lists. You can enable logging for these policies by editing them and selecting *Log Violation Traffic*.
- When you enable logging on a security policy and a session begins, the FortiGate unit logs that session as well as logging when that session closes. These traffic logs contain the log fields status=start, when a session has started, and status=close, when the session ends.
- Email filter log messages now contain a new field, the cc field, which provides all the email addresses that were copied (or cc'd) to the original email message.
- There are two new event logs that contain information about the FSSO polling daemon: one event log is recorded when a logged on user adds or replaces an old user; the second event log is recorded when an old user is logged off because of inactivity or when replaced.

From the CLI, all of these and some other options are available from the new `config log setting` command. This command configures traffic logging settings per VDOM.

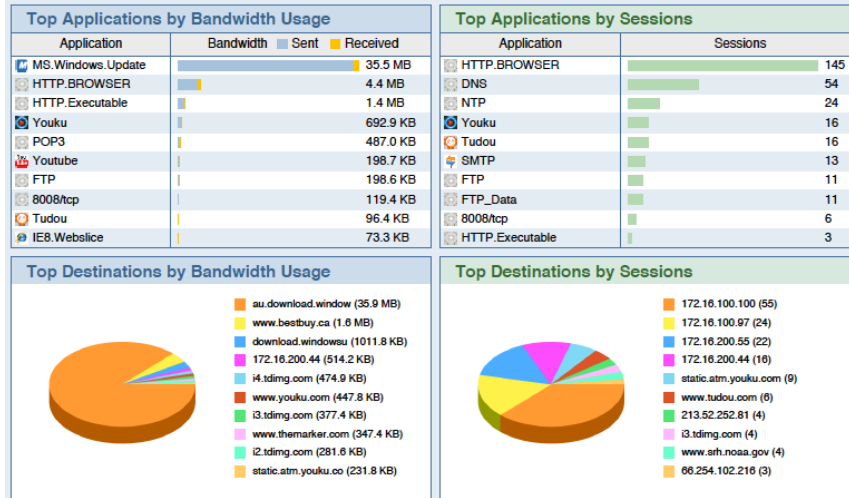
## The FortiGate Security Analysis Report

In order for your FortiGate unit to create a Security Analysis Report, disk logging must be enabled. To enable disk logging, go to *Log & Report > Log Config > Log Settings* and under *Logging and Archiving* select *Disk* and *Enable Local Reports*.

## Viewing the current report

You can go to *Log & Report > Report > Local* and select *Run Now* to view the latest FortiGate Security Analysis Report. This report is constantly updated so you can go here any time to get the current report. The Multi-page report provides wide range of data including bandwidth and application use, users, destinations, streaming usage, email traffic, and so on.

**Figure 63:** Extract from a sample report



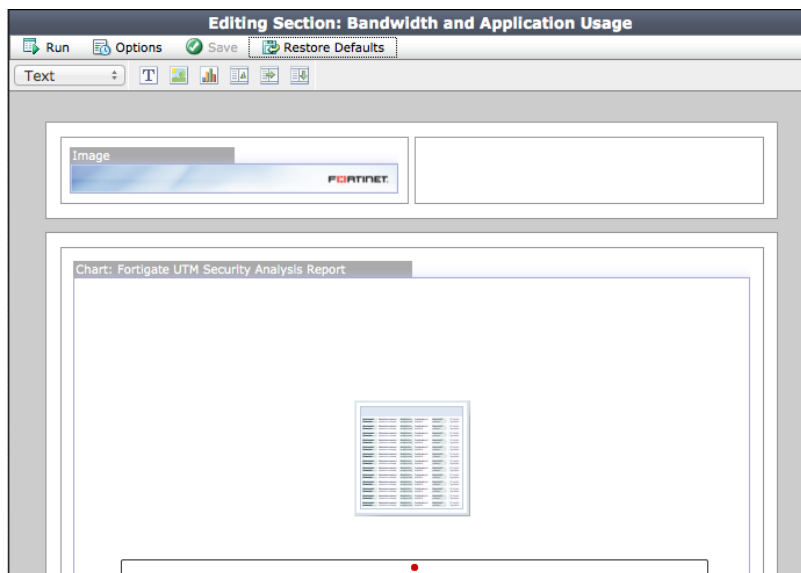
## Viewing the saved (historical) security analysis reports

You can go to *Log & Report > Report > Local* to select saved reports to view.

## Customizing the security analysis report

You can go to *Log & Report > Report > Local* and select *Customize* to customize the look and content of the report. You can also change the report schedule and other options.

**Figure 64:** Customizing report look and content



Select *Options* to change the report schedule to Daily, Weekly or On Demand. You can also change the time of day at which an old report is saved and a new one started.



## Converting compact log format

To convert your compact logs to the new FortiOS 5 format, use the following CLI command:

```
execute log convert-oldlogs
```



This command is available only if you have upgraded from an earlier version of FortiOS and have old compact logs on your system.

---

# Firewall

New firewall features include:

- [Choosing the policy type](#)
- [Reorganized Firewall Services](#)
- [Local in policies](#)
- [Multicast Policies](#)
- [Adding DoS Anomaly protection to a FortiGate interface](#)
- [Changes to security proxy options](#)
- [SSL and SSH inspection](#)

## Choosing the policy type

Creating and editing security policies has been enhanced make creating different types of security policies easier to understand. Now the first step in creating a security policy, after going to *Policy > Policy > Policy* and selecting *Create New* is to select the *Policy Type* (Firewall or VPN) and *Policy Subtype*. The policy subtype selections depend on the policy type:

- If you select *Firewall* you can also select
  - *Address* to create a basic security policy
  - *User Identity* to create a policy that identifies users (user authentication)
  - *Device Identity* to create a policy that identifies devices (device authentication or BYOD)
- If you select *VPN* you can also select
  - *IPsec* to create IPsec VPN policies
  - *SSL-VPN* to create SSL VPN policies

This section describes:

- [Creating a basic security policy](#)
- [Creating a security policy to authenticate users](#)
- [Creating a security policy to authenticate devices for BYOD](#)
- [Creating a policy-based IPsec VPN security policy](#)
- [Creating a route-based IPsec VPN security policy](#)
- [Creating an SSL VPN security policy](#)

## Creating a basic security policy

Use the default policy type settings to create a basic security policy. Select incoming and outgoing interfaces, source and destination addresses, the schedule, services and the action. Select other features such as NAT, logging, Security Features and so on.

**Figure 65:**Creating a basic security policy

New Policy

|                                                                                                                                                                                                                                                 |                                                                                                                    |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| Policy Type                                                                                                                                                                                                                                     | <input checked="" type="radio"/> Firewall <input type="radio"/> VPN                                                |
| Policy Subtype                                                                                                                                                                                                                                  | <input checked="" type="radio"/> Address <input type="radio"/> User Identity <input type="radio"/> Device Identity |
| Incoming Interface                                                                                                                                                                                                                              | port1                                                                                                              |
| Source Address                                                                                                                                                                                                                                  | all                                                                                                                |
| Outgoing Interface                                                                                                                                                                                                                              | port2                                                                                                              |
| Destination Address                                                                                                                                                                                                                             | all                                                                                                                |
| Schedule                                                                                                                                                                                                                                        | always                                                                                                             |
| Service                                                                                                                                                                                                                                         | ALL                                                                                                                |
| Action                                                                                                                                                                                                                                          | ACCEPT                                                                                                             |
| <input checked="" type="checkbox"/> Enable NAT                                                                                                                                                                                                  |                                                                                                                    |
| <input checked="" type="radio"/> Use Destination Interface Address <input type="checkbox"/> Fixed Port<br><input type="radio"/> Use Dynamic IP Pool <input type="text" value="Click to add..."/><br><input type="radio"/> Use Central NAT Table |                                                                                                                    |
| <b>Logging Options</b>                                                                                                                                                                                                                          |                                                                                                                    |
| <input type="radio"/> No Log<br><input checked="" type="radio"/> Log UTM Events<br><input type="radio"/> Log all Sessions                                                                                                                       |                                                                                                                    |

### Creating a security policy to authenticate users

Select the *Firewall* policy type and the *User Identity* subtype. Select incoming and outgoing interfaces, source addresses, and other features. Then select *Create New* to add a user authentication rule to the policy.

User authentication rules include destination addresses, user groups and or individual users, a schedule, services, the action, logging settings, and UTM security profiles.



You select the destination address separately for each authentication rule. This means that you can apply different features to different user groups depending on their destination addresses. You can also now add individual users to authentication rules instead of just user groups.

**Figure 66:**Adding a user authentication rule

New Authentication Rule

|                                                                                                                                |                  |
|--------------------------------------------------------------------------------------------------------------------------------|------------------|
| Destination Address                                                                                                            | all              |
| Group(s)                                                                                                                       | FSSO_Guest_Users |
| User(s)                                                                                                                        | jsmith           |
| Schedule                                                                                                                       | always           |
| Service                                                                                                                        | Click to add...  |
| Action                                                                                                                         | ACCEPT           |
| <b>Logging Options</b>                                                                                                         |                  |
| <input type="radio"/> No Log<br><input type="radio"/> Log Security Events<br><input checked="" type="radio"/> Log all Sessions |                  |
| <b>Security Profiles</b>                                                                                                       |                  |
| <input checked="" type="checkbox"/> AntiVirus                                                                                  | default          |
| <input checked="" type="checkbox"/> Web Filter                                                                                 | default          |
| <input type="checkbox"/> Application Control                                                                                   | default          |
| <input type="checkbox"/> IPS                                                                                                   | default          |

## Creating a security policy to authenticate devices for BYOD

Select the *Firewall* policy type and the *Device Identity* subtype. Select incoming and outgoing interfaces, source addresses, and other features. Then select *Create New* to add device authentication rules to the policy.

Device authentication rules include the destination address, user groups and or individual users, schedule, service, action, logging, and UTM security profiles.

**Figure 67:** Adding a device authentication rule

**New Authentication Rule**

Destination Address: all

Device: BlackBerry Phone

Compliant with Endpoint Profile:

Schedule: always

Service: ALL

Action: ACCEPT

**Logging Options**

No Log

Log Security Events

Log all Sessions

**Security Profiles**

AntiVirus: default

Web Filter: default

Application Control: default

IPS: default

Email Filter: default

DLP Sensor: default

## Creating a policy-based IPsec VPN security policy

Select the *VPN* policy type and the *IPsec* subtype. Select the local and outgoing VPN interfaces, local protected subnet and remote protected subnet addresses, the schedule, services, and the action.

You have two options to configure the *VPN Tunnel* used by the policy.

- You can use a tunnel that has already been added. Select *Use Existing* and select the tunnel to use.
- You can add a new tunnel. Select *Create New* and select either *Site to Site* or *Dialup*. Add a *Name* for the tunnel, the IP address of the remote FortiGate unit (not required for Dialup) and the Preshared Key to be used by the tunnel.

You can also apply Security Profiles and Client Reputation to IPsec VPN traffic.

**Figure 68:**Creating a policy-based IPsec VPN policy

New Policy

|                         |                                                                       |
|-------------------------|-----------------------------------------------------------------------|
| Policy Type             | <input type="radio"/> Firewall <input checked="" type="radio"/> VPN   |
| Policy Subtype          | <input checked="" type="radio"/> IPsec <input type="radio"/> SSL-VPN  |
| Local Interface         | <input type="text" value="port1"/>                                    |
| Local Protected Subnet  | <input type="text" value="all"/> <span style="float: right;">+</span> |
| Outgoing VPN Interface  | <input type="text" value="wan1"/>                                     |
| Remote Protected Subnet | <input type="text" value="all"/> <span style="float: right;">+</span> |
| Schedule                | <input type="text" value="always"/>                                   |
| Service                 | <input type="text" value="ALL"/> <span style="float: right;">+</span> |

**Logging Options**

No Log

Log Security Events

Log all Sessions

---

|                     |                                                                                |
|---------------------|--------------------------------------------------------------------------------|
| VPN Tunnel          | <input checked="" type="radio"/> Create New <input type="radio"/> Use Existing |
|                     | <input checked="" type="radio"/> Site-to-Site <input type="radio"/> Dialup     |
| Name                | <input type="text" value="My-Tunnel"/>                                         |
| Remote FortiGate IP | <input type="text" value="172.20.120.122"/>                                    |
| Preshared Key       | <input type="text" value="*****"/>                                             |
|                     | <input type="checkbox"/> Allow traffic to be initiated from the remote site    |

---

**Security Profiles**

|                                                         |                                      |
|---------------------------------------------------------|--------------------------------------|
| <input checked="" type="checkbox"/> AntiVirus           | <input type="text" value="default"/> |
| <input checked="" type="checkbox"/> Web Filter          | <input type="text" value="default"/> |
| <input checked="" type="checkbox"/> Application Control | <input type="text" value="default"/> |

### Creating a route-based IPsec VPN security policy

You configure route-based IPsec VPN by first adding a Phase 1 and selecting *Enable IPsec Interface Mode*. This adds an IPsec interface with the same name as the Phase 1.

A security policy for this VPN is a standard security policy that allows traffic between the IPsec interface and another FortiGate unit interface. Create a new policy, select the *Firewall* policy type and the *Address* policy subtype. To allow traffic from the internal network to connect to the VPN, set the incoming interface to internal and the outgoing interface to the IPsec interface. Otherwise configure the security policy like any basic security policy.

**Figure 69:**Creating a route-based IPsec VPN policy

New Policy

|                                                 |                                                                                                                    |
|-------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| Policy Type                                     | <input checked="" type="radio"/> Firewall <input type="radio"/> VPN                                                |
| Policy Subtype                                  | <input checked="" type="radio"/> Address <input type="radio"/> User Identity <input type="radio"/> Device Identity |
| Incoming Interface                              | <input type="text" value="port1"/>                                                                                 |
| Source Address                                  | <input type="text" value="all"/> <span style="float: right; color: green;">+</span>                                |
| Outgoing Interface                              | <input type="text" value="My-Phase1"/>                                                                             |
| Destination Address                             | <input type="text" value="all"/> <span style="float: right; color: green;">+</span>                                |
| Schedule                                        | <input type="text" value="always"/>                                                                                |
| Service                                         | <input type="text" value="ALL"/> <span style="float: right; color: green;">+</span>                                |
| Action                                          | <input type="text" value="ACCEPT"/>                                                                                |
| <input type="checkbox"/> Enable NAT             |                                                                                                                    |
| <b>Logging Options</b>                          |                                                                                                                    |
| <input type="radio"/> No Log                    |                                                                                                                    |
| <input checked="" type="radio"/> Log UTM Events |                                                                                                                    |
| <input type="radio"/> Log all Sessions          |                                                                                                                    |

### Creating an SSL VPN security policy

Select the *VPN* policy type and the *SSL-VPN* subtype. Select the incoming interface, the remote address, the local interface, and the address for the local protected subnet. Then select *Create New* to add SSL VPN authentication rules to the policy.

SSL authentication rules include user groups, individual users, schedule, service, SSL VPN portal, action, logging, and UTM security profiles.



FortiOS 5.0 no longer includes SSL VPN users or user groups. Any type of user group can be added to an SSL VPN authentication rule.

**Figure 70:**Creating an SSL VPN policy

New Policy

|                                                                                                                                     |                                                                                                        |
|-------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|
| Policy Type                                                                                                                         | <input type="radio"/> Firewall <input checked="" type="radio"/> VPN                                    |
| Policy Subtype                                                                                                                      | <input type="radio"/> IPsec <input checked="" type="radio"/> SSL-VPN                                   |
| Incoming Interface                                                                                                                  | <input type="text" value="port1"/>                                                                     |
| Remote Address                                                                                                                      | <input type="text" value="all"/> <span style="float: right; color: green;">+</span>                    |
| Local Interface                                                                                                                     | <input type="text" value="port2"/>                                                                     |
| Local Protected Subnet                                                                                                              | <input type="text" value="all"/> <span style="float: right; color: green;">+</span>                    |
| <input type="checkbox"/> SSL Client Certificate Restrictive                                                                         |                                                                                                        |
| Cipher Strength                                                                                                                     | <input type="text" value="cipher_any"/>                                                                |
| <b>Configure SSL-VPN Authentication Rules</b>                                                                                       |                                                                                                        |
| <span style="color: green;">+</span> Create New <span style="color: gray;">✎</span> Edit <span style="color: gray;">🗑</span> Delete |                                                                                                        |
| <b>User/Group</b>                                                                                                                   | <b>Service</b>                                                                                         |
| <b>Schedule</b>                                                                                                                     | <b>UTM Security</b>                                                                                    |
| <b>SSL-VPN Portal</b>                                                                                                               | <b>Logging</b>                                                                                         |
| <b>Action</b>                                                                                                                       |                                                                                                        |
| ANY                                                                                                                                 | ALL                                                                                                    |
| always                                                                                                                              | -                                                                                                      |
| -                                                                                                                                   | ✘                                                                                                      |
| -                                                                                                                                   | 🚫 DENY                                                                                                 |
| <b>Tags</b>                                                                                                                         |                                                                                                        |
| Applied tags                                                                                                                        |                                                                                                        |
| Add tag                                                                                                                             | <input type="text" value=""/> <span style="float: right; color: blue;">+</span>                        |
| Comments                                                                                                                            | <input type="text" value="Write a comment..."/> <span style="float: right; color: gray;">0/1023</span> |
| <input type="button" value="OK"/> <input type="button" value="Cancel"/>                                                             |                                                                                                        |

## Reorganized Firewall Services

To make Firewall services easier to use and more customizable, the former predefined services list and custom services have been merged into one list and the list has been organized into categories. You can go to *Firewall Objects > Service > Services* to see the new service list.

The order that the categories and services in the services list is maintained when you add a service to another configuration object, such as a security policy. You can add categories, change the category that a service is in and rearrange the order of categories in the list.

**Figure 71:** Firewall Services list

| Service Name       | Ports                       | IP/FQDN | Show in Service List                | Ref. |
|--------------------|-----------------------------|---------|-------------------------------------|------|
| <b>General</b>     |                             |         |                                     |      |
| ALL                | ANY                         |         | <input checked="" type="checkbox"/> | 5    |
| ALL_ICMP           | ICMP/ANY                    |         | <input checked="" type="checkbox"/> | 0    |
| ALL_ICMP6          | ICMP6/ANY                   |         | <input checked="" type="checkbox"/> | 0    |
| ALL_TCP            | TCP/1-65535                 | 0.0.0.0 | <input checked="" type="checkbox"/> | 0    |
| ALL_UDP            | UDP/1-65535                 | 0.0.0.0 | <input checked="" type="checkbox"/> | 0    |
| <b>Web Access</b>  |                             |         |                                     |      |
| HTTP               | TCP/80                      | 0.0.0.0 | <input checked="" type="checkbox"/> | 1    |
| HTTPS              | TCP/443                     | 0.0.0.0 | <input checked="" type="checkbox"/> | 2    |
| <b>File Access</b> |                             |         |                                     |      |
| AFS3               | TCP/7000-7009 UDP/7000-7009 | 0.0.0.0 | <input checked="" type="checkbox"/> | 0    |
| FTP                | TCP/21                      | 0.0.0.0 | <input checked="" type="checkbox"/> | 0    |
| FTP_GET            | TCP/21                      | 0.0.0.0 | <input checked="" type="checkbox"/> | 0    |
| FTP_PUT            | TCP/21                      | 0.0.0.0 | <input checked="" type="checkbox"/> | 0    |
| NFS                | TCP/111,2049 UDP/111,2049   | 0.0.0.0 | <input checked="" type="checkbox"/> | 0    |
| SAMBA              | TCP/139                     | 0.0.0.0 | <input checked="" type="checkbox"/> | 1    |
| SMB                | TCP/445                     | 0.0.0.0 | <input checked="" type="checkbox"/> | 1    |
| TFTP               | UDP/69                      | 0.0.0.0 | <input checked="" type="checkbox"/> | 0    |
| <b>Email</b>       |                             |         |                                     |      |
| IMAP               | TCP/143                     | 0.0.0.0 | <input checked="" type="checkbox"/> | 1    |
| IMAPS              | TCP/993                     | 0.0.0.0 | <input checked="" type="checkbox"/> | 1    |
| POP3               | TCP/110                     | 0.0.0.0 | <input checked="" type="checkbox"/> | 1    |
| POP3S              | TCP/995                     | 0.0.0.0 | <input checked="" type="checkbox"/> | 1    |
| SMTP               | TCP/25                      | 0.0.0.0 | <input checked="" type="checkbox"/> | 1    |
| SMTPS              | TCP/465                     | 0.0.0.0 | <input checked="" type="checkbox"/> | 1    |

The single *ANY* service has been replaced with five *ALL* services that match all services or all services of a specific type (ICMP, ICMP6, TCP and UDP). Otherwise, all of the familiar pre-defined services can be found in the new services list.

By default, the list is organized by categories. You can select *Category Settings* to change the order of the categories in the list. You can also choose to organize the list alphabetically by service name.

### Editing and deleting services

You can edit and delete any of the services in the list; however, use caution when deleting any services. Normally you would only delete a custom service that you have created.

You can edit any service to change its name, add a comment, change its icon color, add or remove it from the list, change its category, change its protocol type, add an IP address or fully qualified domain name (FQDN) and change the source and destination ports for the service.

In most cases, you should not need to edit predefined services. Instead you would add custom services. However, in some cases editing pre-defined services can simplify your configuration. For example, if all of your HTTP traffic uses port 8080, you could edit the HTTP service and change its destination port to 8080. Alternatively, if your network uses port 80 and 8080 for HTTP traffic, you could edit the HTTP service and add port 8080.

**Figure 72:**HTTP service customized to include port 80 and port 8080

**Edit Service**

Name: HTTP

Comments: Write a comment... 0/255

Service Type:  Firewall  Explicit Proxy

Color: [Change]

Show in Service List:

Category: Web Access

Protocol Type: TCP/UDP/SCTP

IP/FQDN:

| Protocol | Destination Port |      | Source Port |      |   |
|----------|------------------|------|-------------|------|---|
|          | Low              | High | Low         | High |   |
| TCP      | 80               | -    |             |      | × |
| TCP      | 8080             | -    |             |      | × |

OK Cancel

## Adding an address to a service

Services now include a IP/FQDN field into which you can add an IP address or a fully qualified domain name. Use this field if you want to restrict the network address that the FortiGate unit will accept connections to this service from.

## Adding a new service

To add a new service, go to the service list and select *Create New > Service*. Configure the custom service as shown in [Figure 72](#) and select *OK* to save it.

## Adding a new service category

If you have a group of services that you use often, you can group them into your own service category to make them easy to find in the list. To add a new service category, go to the service list and select *Create New > Category*, add a name for the custom service and select *OK*.

Then select *Category Settings* and change the order in which the services appear in the list. To make your custom category easy to find, move it to the top of the list.

To add services to your category, edit them and set *Category* to the name of your custom category.

## Local in policies

Read-only local in policies show you all the types of traffic that can connect to or terminate at the FortiGate unit. For the FortiGate unit to receive local traffic a policy to receive the traffic must be in the local in policy list. The FortiGate unit needs to be able to receive traffic for a number of reasons. Among them:

- Central management connections from FortiManager
- Networking and routing connections, for example accepting or relaying DHCP requests, accepting routing communication from other routers (for example, OSPF, RIP, VRRP)
- Administrative access to FortiGate interfaces over ICMP, HTTP, HTTPS, and so on.

The local-in policy list includes an action column that shows whether the FortiGate unit accepts or drops sessions identified by the individual local in policies. As you change some



configuration settings those changes are reflected in the local in policies. For example, Administrative Access local in policies change depending on the administrative access settings of your FortiGate interfaces.

From the local in policy page (*Firewall > Policy > Local In Policy*), you can enable or disable logging for local in allowed and denied traffic and for local out traffic.

In addition to the pre-defined local in policies, you can add your own using the following command:

```
config firewall {local-in-policy | local-in-policy6}
 edit 0
 set srcaddr all
 set dstaddr all
 set action {deny | allow}
 set service ALL
 set schedule always
 set auto-asic-offload {disable | enable}
 end
```

## Multicast Policies

A number of popular services use multicast protocols. Examples include the Bonjour service used for finding devices on a network, EIGRP and OSPF. To make it easier to allow multicast traffic through the FortiGate unit, you can now add multicast policies from the web-based manager by going to *Policy > Policy > Multicast Policy* and selecting *Create New*.

Similar to a regular security policy, you configure a multicast policy by selecting incoming and outgoing interfaces, source and destination addresses, enabling NAT, and selecting an action.

**Figure 73:** Adding a multicast policy

The screenshot shows the 'New Multicast Policy' configuration window. The fields are as follows:

- Incoming Interface: port1
- Source Address: all
- Outgoing Interface: port2
- Destination Address: all\_hosts
- Enable NAT:
- DNAT: 0.0.0.0
- Protocol: 0
- Port Range: 9000-9001
- Action: ACCEPT
- Log Allowed Traffic:

Buttons: OK, Cancel

Specific to multicast policies, you can also specify a destination NAT (DNAT) address and select a multicast protocol (options include ANY, ICMP, IGMP, TCP, UDP, OSPF and other). You cannot add or edit these protocols but, if you select *Other*, you can add a protocol number.

The destination address of a multicast policy must be a multicast address firewall object. Multicast addresses are added by going to *Firewall Objects > Address > Addresses* and selecting *Create New > Multicast Address*. The FortiGate default configuration includes some commonly used multicast addresses. [Figure 74](#) shows the configuration of the default Bonjour multicast address.

**Figure 74:**Default Bonjour multicast firewall address

Edit Address

Category  Address  IPv6 Address  Multicast

Address

Name

Color [Change]

Show in Address List

Multicast IP Range

Interface

Comments  0/255

## Adding DoS Anomaly protection to a FortiGate interface

New DoS policies allow you to apply DoS anomalies to all traffic that hits a FortiGate interface. This is the only way to apply DoS anomaly protection.

To add a DoS policy go to *Policy > Policy > DoS Policy* and select Create New to add a new DoS policy. Select the FortiGate interface to add the policy to and select the source and destination addresses and services that will match the packets that you want to apply DoS anomalies to.

Enable one or more DoS anomalies. For each anomaly you can enable logging, set the action to pass or block and change the threshold.

**Figure 75:**DoS policy

New DoS Policy

Incoming Interface

Source Address  +

Destination Address  +

Service  +

**Anomalies**

| Name             | Status                   | Logging                  | Action | Threshold |
|------------------|--------------------------|--------------------------|--------|-----------|
| tcp_syn_flood    | <input type="checkbox"/> | <input type="checkbox"/> | Pass ▾ | 2000      |
| tcp_port_scan    | <input type="checkbox"/> | <input type="checkbox"/> | Pass ▾ | 1000      |
| tcp_src_session  | <input type="checkbox"/> | <input type="checkbox"/> | Pass ▾ | 5000      |
| tcp_dst_session  | <input type="checkbox"/> | <input type="checkbox"/> | Pass ▾ | 5000      |
| udp_flood        | <input type="checkbox"/> | <input type="checkbox"/> | Pass ▾ | 2000      |
| udp_scan         | <input type="checkbox"/> | <input type="checkbox"/> | Pass ▾ | 2000      |
| udp_src_session  | <input type="checkbox"/> | <input type="checkbox"/> | Pass ▾ | 5000      |
| udp_dst_session  | <input type="checkbox"/> | <input type="checkbox"/> | Pass ▾ | 5000      |
| icmp_flood       | <input type="checkbox"/> | <input type="checkbox"/> | Pass ▾ | 250       |
| icmp_sweep       | <input type="checkbox"/> | <input type="checkbox"/> | Pass ▾ | 100       |
| icmp_src_session | <input type="checkbox"/> | <input type="checkbox"/> | Pass ▾ | 300       |
| icmp_dst_session | <input type="checkbox"/> | <input type="checkbox"/> | Pass ▾ | 1000      |
| ip_src_session   | <input type="checkbox"/> | <input type="checkbox"/> | Pass ▾ | 5000      |

Use the following command to add a DoS policy from the CLI that adds syn flood protection for all traffic hitting the wan2 interface:

```
config firewall DoS-policy
 edit 1
 set interface wan2
 set srcaddr all
 set dstaddr all
 set service ALL
 config anomaly
 edit tcp_syn_flood
 set status enable
 set log enable
 set action block
 end
 end
 end
```

DoS sensors no longer available. DoS policies are the most common method for applying DoS anomalies in FortiOS 5.0.

You can also use the following command to apply DoS anomalies to a one-arm sniffer configuration.

```
config firewall sniffer
```

Interface policies are still available in FortiOS 5.0 from the CLI using the following commands:

```
config firewall interface-policy
config firewall interface-policy6
```

You can use interface policies to apply application control, intrusion protection, virus scanning, web filtering, email filtering and data leak protection to traffic received by an interface.

The following commands are also available for adding sniffer interface policies, which are similar to interface policies:

```
config firewall sniff-interface-policy
config firewall sniff-interface-policy6
```

All of these command have similar syntax for applying Security Features to traffic connecting to or sniffed by a FortiGate interface.

## Changes to security proxy options

Security proxy profile options are now configured by feature instead of by protocol. Also, SSL and SSH inspection options have been moved to the new SSL/SSH inspection options profiles.

### Protocol port mapping

For all content protocols, you can configure protocol port mapping to set the ports on which the FortiGate unit looks for the protocols. The FortiGate unit can inspect each kind of traffic on any port or you can specify one or more ports.

**Figure 76:**Configuring Protocol Port Mapping

**SSL Inspection Options**

CA Certificate: Fortinet\_CA\_SSLProxy

Inspect All Ports:

| Enable                              | Protocol | Inspection Port(s) |
|-------------------------------------|----------|--------------------|
| <input checked="" type="checkbox"/> | HTTPS    | 443                |
| <input checked="" type="checkbox"/> | SMTPS    | 465                |
| <input checked="" type="checkbox"/> | POP3S    | 995                |
| <input checked="" type="checkbox"/> | IMAPS    | 993                |
| <input checked="" type="checkbox"/> | FTPS     | 990                |

**SSH Inspection Options**

### Common options, web options and email options

Here you can configure client comforting, whether to block oversized files or email, and whether or not to allow invalid SSL certificates.

Web options include enabling chunked by pass and adding the Fortinet bar (see “Fortinet Top Bar” on page 245).

Email options include allowing fragmented messages and appending a signature to all SMTP email messages.

**Figure 77:**Configuring common options, web options and email options

**Common Options**

Comfort Clients

Block Oversized File/Email

Allow Invalid SSL Certificates

**Web Options**

Enable Chunked Bypass

Add Fortinet Bar

**Email Options**

Allow Fragmented Messages

Append Signature (SMTP)

Email Signature Text: Disclaimer...

**Apply**

## SSL and SSH inspection

To configure how encrypted SSL or SSH traffic is inspected in security policy that accepts the SSL or SSH traffic to be inspected, turn on *SSL/SSH Inspection* and select an SSL/SSH inspection profile.

You can go to *Policy > Policy > SSL/SSH Inspection* to change the default SSL and SSH inspection profile or you can create new profiles.

## SSL inspection options

For SSL traffic, you can select the certificate to use for this traffic and enable inspection of SSL traffic on all ports. You can also select individual SSL protocols and configure the inspection ports and port ranges for them. You can also choose to block or allow invalid SSL certificates.

**Figure 78:**Configuring SSL and SSH Inspection

**Edit Deep Inspection Options**
default ↕

Name

Comments  20/255

**SSL Inspection Options**

CA Certificate

Inspect All Ports

| Enable                              | Protocol | Inspection Port(s)                                    |
|-------------------------------------|----------|-------------------------------------------------------|
| <input checked="" type="checkbox"/> | HTTPS    | <input style="width: 60px;" type="text" value="443"/> |
| <input checked="" type="checkbox"/> | SMTPTS   | <input style="width: 60px;" type="text" value="465"/> |
| <input checked="" type="checkbox"/> | POP3S    | <input style="width: 60px;" type="text" value="995"/> |
| <input checked="" type="checkbox"/> | IMAPS    | <input style="width: 60px;" type="text" value="993"/> |
| <input checked="" type="checkbox"/> | FTPS     | <input style="width: 60px;" type="text" value="990"/> |

**SSH Inspection Options**

Enable SSH Deep Scan

| Protocol     | Inspection Port(s)                                                                                                      |
|--------------|-------------------------------------------------------------------------------------------------------------------------|
| SSH          | <input type="radio"/> Any <input checked="" type="radio"/> Specify <input style="width: 60px;" type="text" value="22"/> |
| Exec         | <input type="checkbox"/> Block <input type="checkbox"/> Log                                                             |
| Port-Forward | <input type="checkbox"/> Block <input type="checkbox"/> Log                                                             |
| SSH-Shell    | <input type="checkbox"/> Block <input type="checkbox"/> Log                                                             |
| X11-Filter   | <input type="checkbox"/> Block <input type="checkbox"/> Log                                                             |

**Common Options**

Allow Invalid SSL Certificates

## SSH inspection options

For SSH traffic, you can enable SSH deep scanning, inspect all ports or specified ports for SSH traffic. You can also block or log all Exec, Port-Forward, SSH-Shell and X-11 SSH activity.

# WAN optimization and Web Caching

In FortiOS 5.0, WAN optimization is enabled in security policies and WAN optimization rules are no longer required. Instead of adding a security policy that accepts traffic to be optimized and then creating WAN optimization rules to apply WAN optimization, in FortiOS 5.0 you create security policies that accept traffic to be optimized and enable WAN optimization in those policies. WAN optimization is applied by WAN optimization profiles which are created separately and added to the required security policies.

Because of this change, you can now apply all Security features to WAN optimization traffic without having to use a configuration that requires two VDOMS (one for applying Security features and one for applying WAN optimization). Instead, you can enable Security features in the security policies that accept WAN optimization traffic.

WAN optimization in policies requires you to add extra security policies with the incoming interface set to the new *wanopt* interface.

In FortiOS 4.3, you could create web caching only WAN optimization rules while in FortiOS 5.0 you cannot create web caching only WAN optimization profiles. Instead, you simply enable web caching in security policies, including WAN optimization policies. You can enable web caching for any WAN optimization policy. You can also enable HTTPS web caching and SSL offloading from the CLI for any security policy.

This chapter describes:

- [Configuring WAN optimization profiles](#)
- [Dynamic data chunking for WAN optimization byte caching](#)
- [Policy-based WAN optimization configuration changes summary](#)
- [Combining web caching for HTTP traffic with WAN optimization](#)
- [Turning on web caching and SSL offloading for HTTPS traffic](#)
- [Changing the ports on which to look for HTTP and HTTPS traffic to cache](#)
- [Web proxy URL debugging](#)
- [FortiOS Web Caching now caches Windows/MS-Office software updates](#)

## Configuring WAN optimization profiles

Use WAN optimization profiles to apply WAN optimization techniques to traffic to be optimized. In a WAN optimization profile, you can select the protocols to be optimized and for each protocol you can enable SSL offloading, secure tunneling, byte caching and set the port the protocol uses. You can also enable transparent mode and select an authentication group. You can edit the default WAN optimization profile or create new ones.

In order to configure WAN optimization profiles using the web-based manager, this feature must be enabled using Feature Select. For more information, see [“Feature Select” on page 243](#).

To configure a WAN optimization profile go to *WAN Opt & Cache > WAN Opt. Profile > Profile* and edit a profile or create a new one.

**Figure 79:**Configuring a WAN optimization profile

**Edit WAN Optimization Profile** default

Name: default

Comments: default WANopt profile 22/255

Transparent Mode

Authentication Group: Auth-Grp

| Protocol                                 | SSL Offloading                      | Secure Tunneling                    | Byte Caching                        | Port    |
|------------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|---------|
| <input checked="" type="checkbox"/> CIFS |                                     | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | 445     |
| <input checked="" type="checkbox"/> FTP  |                                     | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | 21      |
| <input checked="" type="checkbox"/> HTTP | <input checked="" type="checkbox"/> | <input type="checkbox"/>            | <input checked="" type="checkbox"/> | 80      |
| <input checked="" type="checkbox"/> MAPI |                                     | <input type="checkbox"/>            | <input checked="" type="checkbox"/> | 135     |
| <input checked="" type="checkbox"/> TCP  | <input checked="" type="checkbox"/> | <input type="checkbox"/>            | <input type="checkbox"/>            | 1-65535 |

**Apply**

From the CLI you can use the following command to configure a WAN optimization profile to optimize HTTP traffic.

```
config wanopt profile
 edit new-profile
 config http
 set status enable
 end
 end
```

Enter the following command to view WAN optimization profile CLI options:

```
tree wanopt profile
-- [profile] --*name (36)
 |- transparent
 |- comments
 |- auth-group (36)
 |- <http> -- status
 |- secure-tunnel
 |- byte-caching
 |- prefer-chunking
 |- tunnel-sharing
 |- log-traffic
 |- port
 |- ssl
 |- ssl-port
 |- unknown-http-version
 +- tunnel-non-http
 |- <cifs> -- status
 |- secure-tunnel
 |- byte-caching
 |- prefer-chunking
 |- tunnel-sharing
 |- log-traffic
 +- port
 |- <mapi> -- status
 |- secure-tunnel
 |- byte-caching
 |- tunnel-sharing
 |- log-traffic
 +- port
 |- <ftp> -- status
 |- secure-tunnel
 |- byte-caching
 |- prefer-chunking
 |- tunnel-sharing
 |- log-traffic
 +- port
 +- <tcp> -- status
 |- secure-tunnel
 |- byte-caching
 |- byte-caching-opt
 |- tunnel-sharing
 |- log-traffic
 |- port
 |- ssl
 +- ssl-port
```



## Dynamic data chunking for WAN optimization byte caching

Dynamic data chunking helps to detect persistent data chunks in a changed files or in data embedded in traffic using an unknown protocol. For example, Lotus notes uses a private protocol to transfer email attachments in crafted messages. Dynamic data chunking performs byte caching of data in Lotus notes traffic. Dynamic data chunking is available for HTTP, CIFS and FTP.

Use the following command to enable dynamic data chunking for HTTP in the default WAN optimization profile.

```
config wanopt profile
 edit default
 config http
 set prefer-chunking dynamic
 end
```

By default, dynamic data chunking is disabled and `prefer-chunking` is set to `fix`.

## Policy-based WAN optimization configuration changes summary

This section summarizes how basic WAN optimization configurations work in FortiOS 5.0, now that WAN optimization is enabled in security policies.

### On the client side

New features:

- WAN optimization rules are removed and WAN optimization profiles are added. Profiles are configured in the client side.
- New options in firewall policies: `wanopt`, `wanopt-detection`, `wanopt-profile` and `wanopt-peer`. `wanopt-peer` is used only on the client side for manual mode (`wanopt-detection` is off).
- You can add Security features inspection to security policies that accept WAN optimization traffic.

### On the server side

New features:

- New `wanopt` interface which represents the WAN optimization tunnel.
- Add a firewall policy with incoming (source) interface set to the `wanopt` interface to accept WAN optimization tunnel sessions (only required on the server side).
- For active/passive WAN optimization, set the server side to *passive*.
- For manual mode no WAN optimization policy required.
- WAN optimization profiles inherited by the server side.
- You can add Security feature inspection to security policies that accept WAN optimization traffic.

## Client side configuration summary

### WAN optimization profile

```
config wanopt profile
 edit "default"
 set comments "default WANopt profile"
 config http
 set status enable
 set prefer-chunking fix
 end
 config cifs
 set status enable
 set prefer-chunking fix
 end
 config mapi
 set status enable
 end
 config ftp
 set status enable
 set prefer-chunking fix
 end
 config tcp
 set status enable
 set byte-caching-opt mem-disk
 end
 end
end
```

### Local host ID and peer settings

```
config wanopt settings
 set host-id "client"
end
config wanopt peer
 edit "server"
 set ip 10.10.2.82
end
```

### Security policies

Two client side WAN optimization security policy configurations are possible: one for active-passive WAN optimization and one for manual WAN optimization.

**Active/passive mode on the client side**

```
config firewall policy
 edit 2
 set srcintf "internal"
 set dstintf "wan1"
 set srcaddr "all"
 set dstaddr "all"
 set action accept
 set schedule "always"
 set service "ANY"
 set utm-status enable <<< enable UTM
 set av-profile default <<< select an antivirus profile
 set profile-protocol-options default
 set wanopt enable <<< enable WAN optimization
 set wanopt-detection active <<< set the mode to active/passive
 set wanopt-profile "default" <<< select the wanopt profile
 next
end
```

**Manual mode on the client side**

```
config firewall policy
 edit 2
 set srcintf "internal"
 set dstintf "wan1"
 set srcaddr "all"
 set dstaddr "all"
 set action accept
 set schedule "always"
 set service "ANY"
 set utm-status enable <<< enable UTM
 set av-profile default <<< select an antivirus profile
 set profile-protocol-options default
 set wanopt enable <<< enable WAN optimization
 set wanopt-detection off <<< sets the mode to manual
 set wanopt-profile "default" <<< select the wanopt profile
 set wanopt-peer "server" <<< set the only peer to do wanopt with
 (required for manual mode)
 next
end
```

## Server Side configuration summary

### Local host ID and peer settings

```

config wanopt settings
 set host-id "server"
end
config wanopt peer
 edit "client"
 set ip 10.10.2.81
 end

```

### Security policies

Two server side WAN optimization security policy configurations are possible: one for active-passive WAN optimization and one for manual WAN optimization.

#### Active/passive mode on server side

```

config firewall policy
 edit 2 <<< the passive mode policy
 set srcintf "wan1"
 set dstintf "internal"
 set srcaddr "all"
 set dstaddr "all"
 set action accept
 set schedule "always"
 set service "ANY"
 set utm-status enable <<< enable UTM
 set av-profile default <<< select an antivirus profile
 set profile-protocol-options default
 set wanopt enable
 set wanopt-detection passive
 set wanopt-passive-opt transparent
 next
 edit 3 <<< policy that accepts wanopt tunnel connections from the server
 set srcintf "wanopt" <<< wanopt tunnel interface
 set dstintf "internal"
 set srcaddr "all"
 set dstaddr "all"
 set action accept
 set schedule "always"
 set service "ANY"
 next
end

```

**Manual mode on server side**

```

configure firewall policy
 edit 3 <<< wanopt tunnel policy
 set srcintf "wanopt" <<< wanopt tunnel interface
 set dstintf "internal"
 set srcaddr "all"
 set dstaddr "all"
 set action accept
 set schedule "always"
 set service "ANY"
 set utm-status enable <<< enable UTM
 set av-profile default <<< select an antivirus profile
 set profile-protocol-options default
 next
end

```

**Combining web caching for HTTP traffic with WAN optimization**

Web caching can be applied to any HTTP or HTTPS traffic by enabling web caching in a security policy that accepts the traffic. This includes WAN optimization and explicit web proxy traffic. Web caching caches all HTTP traffic accepted by a policy on TCP port 80.

You can add web caching to a WAN optimization security policy to combine web caching with WAN optimization for any WAN optimization security policy. This includes manual, active and passive WAN optimization policies and WAN optimization tunnel policies. You can enable web caching on both the client-side and the server-side FortiGate units or on just one or the other. For optimum performance, you can enable web caching on both the client-side and server-side FortiGate units. In this way, only uncached content is transmitted through the WAN optimization tunnel. All cached content is access locally by clients from the client side FortiGate unit.

**Turning on web caching and SSL offloading for HTTPS traffic**

Web caching can cache the content of HTTPS traffic on TCP port 443. With HTTPS web caching, the FortiGate unit receives the HTTPS traffic on behalf of the client, opens up the encrypted traffic and extracts content to be cached. Then FortiGate unit re-encrypts the traffic and sends it on to its intended recipient. It is very similar to a man-in-the-middle attack. You enable HTTPS web caching from the CLI in a security policy that accepts the traffic to be cached using `webcache-https`:

```

config firewall policy
 edit 0
 .
 .
 .
 set webcache enable
 set webcache-https any
 .
 .
 .
 end

```

The `any` setting causes the FortiGate unit to re-encrypt the traffic with the FortiGate unit's certificate rather than the original certificate. This configuration can cause errors for HTTPS clients because the name on the certificate does not match the name on the web site.

You can stop these errors from happening by configuring HTTPS web caching to use the web server's certificate by setting `webcache-https` to `ssl-server`:

```
config firewall policy
 edit 0
 .
 .
 .
 set webcache enable
 set webcache-https ssl-server
 .
 .
 .
 end
```

The `ssl-server` option causes the FortiGate unit to re-encrypt the traffic with the certificate that you imported into the FortiGate unit. The certificate is added to an SSL server configuration using the following command:

```
config wanopt ssl-server
 edit example_server
 set ip <Web-Server-IP>
 set port 443
 set ssl-mode { full | half}
 set ssl-cert <Web-Server-Cert>
 end
```

Where:

`Web-Server-IP` is the web server's IP address.

`Web-Server-Cert` is the original web server certificate imported into the FortiGate unit.

The SSL server configuration also determines whether the SSL server is operating in half or full mode and the port used for the HTTPS traffic.

Using the SSL server configuration, web caching also supports SSL offloading that uses the FortiGate unit's FortiASIC SSL encryption/decryption engine to accelerate SSL performance.

## Changing the ports on which to look for HTTP and HTTPS traffic to cache

By default, FortiOS assumes HTTP traffic uses TCP port 80 and HTTPS traffic uses port 443 and so web caching is configured for all HTTP traffic accepted by a policy on TCP port 80 and all HTTPS traffic on TCP port 443. If you want to cache HTTP or HTTPS traffic on other ports, you can enable Security features for the security policy and add an SSL/SSH inspection profile that looks for HTTP and HTTPS traffic on other TCP ports.

Setting the HTTP port to *Any* in the an SSL/SSH inspection profile is not compatible with web caching. If you set the HTTP port to any, web caching only caches HTTP traffic on port 80.

## Web proxy URL debugging

You can use the following CLI commands to get debugging information that shows how the web cache is handling specific URLs. You can debug web caching for a single web page (such as docs.fortinet.com/fgt40mr3.html) or for all requests to a URL pattern (such as docs.fortinet.com to debug all connections to any page on docs.fortinet.com). Wildcard characters and regular expressions are not supported.

Normally you would use this feature if the web cache was not caching specific pages and sites. It makes it easier to get debug information just for the pages causing the problem. This feature works for web caching enabled in any security policy including web proxy and WAN optimization security policies.

### Debugging caching of a specific web page

Start by adding the URL to the configuration:

```
config web-proxy debug-url
 edit docs-url
 set url-pattern "docs.fortinet.com/fgt40mr3.html"
 set status enable
 set exact enable
 end
```

Then enter the following commands to enable debugging:

```
diagnose debug application wad 0
diagnose wad debug-url enable
diagnose debug enable
```

The CLI then displays debug information as the wad application processes sessions. However, the `diagnose wad debug-url enable` command isolates and formats the debug output for sessions to and from docs.fortinet.com/fgt40mr3.html.

Example output when a user browses to docs.fortinet.com/fgt40mr3.html with Firefox. You may have to scroll through other debug output to find this, but its should be easy to find because its formatted differently than the other web cache diagnose output.

```
[0x40d977d0] Received request from client: 10.31.101.20:54932

GET /fgt40mr3.html HTTP/1.1
Host: docs.fortinet.com
User-Agent: Mozilla/5.0 (Macintosh; Intel Mac OS X 10.7; rv:5.0.1)
 Gecko/20100101 Firefox/5.0.1
Accept:
 text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8
Accept-Language: en-us,en;q=0.5
Accept-Encoding: gzip, deflate
Accept-Charset: ISO-8859-1,utf-8;q=0.7,*;q=0.7
DNT: 1
Connection: keep-alive
Referer: http://docs.fortinet.com/fgt.html

[0x40d977d0] Connect to server: 208.91.113.43:80

[0x40d977d0] Forward request to server:
```



```

GET /fgt40mr3.html HTTP/1.1
Host: docs.fortinet.com
User-Agent: Mozilla/5.0 (Macintosh; Intel Mac OS X 10.7; rv:5.0.1)
 Gecko/20100101 Firefox/5.0.1
Accept:
 text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8
Accept-Language: en-us,en;q=0.5
Accept-Encoding: gzip, deflate
Accept-Charset: ISO-8859-1,utf-8;q=0.7,*;q=0.7
DNT: 1
Referer: http://docs.fortinet.com/fgt.html
Connection: Keep-Alive

```

[0x40d977d0] Received response from server:

```

HTTP/1.1 200 OK
Server: Microsoft-IIS/5.0
Date: Fri, 10 Feb 2012 17:05:15 GMT
X-Powered-By: ASP.NET
Connection: close
Content-Type: text/html

```

[0x40d977d0] Forward response from sever:

```

HTTP/1.1 200 OK
Server: Microsoft-IIS/5.0
Date: Fri, 10 Feb 2012 17:05:15 GMT
X-Powered-By: ASP.NET
Content-Type: text/html
Connection: close

```

## Debugging caching of multiple web pages

Use the following commands to get URL debugging output when any docs.fortinet.com and www.fortinet.com web page is cached. In this configuration, just the high-level URLs are added to the configuration and `exact` is set to `disable`:

```

config web-proxy debug-url
 edit docs-url
 set url-pattern "docs.fortinet.com"
 set status enable
 set exact disable
 next
 edit docs-url
 set url-pattern "www.fortinet.com"
 set status enable
 set exact disable
 next
end

```

Then enter the following commands to enable debugging:

```
diagnose debug application wad 0
diagnose wad debug-url enable
diagnose debug enable
```

The CLI then displays debug information as the wad application processes sessions, highlighting all connections to docs.fortinet.com and www.fortinet.com.

## FortiOS Web Caching now caches Windows/MS-Office software updates

FortiOS web caching is not always able to cache Windows and MS-Office updates because they are downloaded using HTTP in multipart or chunked mode and typically run through multiple TCP connections. To resolve this issue in FortiOS 5.0, the first request for Windows or MS-Office updates (to download.windowsupdate.com) causes the cache process to download the new update file in the background.

Once the new update file has been downloaded to the cache, it is available to web cache users and all subsequent requests for this update will be downloaded from the cache. Because the update file will not be available in the web cache until it has completely downloaded, the first update request will not be able to get it from the web cache and neither will any updates requested while the file is downloading in the background.

# Usability enhancements

FortiOS 5.0 introduces usability enhancements to make configuration easier and management more effective and efficient.

New usability features include:

- [Feature Select](#)
- [Improved list editing](#)
- [Dynamic comment fields](#)
- [Setup Wizard enhancements](#)
- [Fortinet Top Bar](#)
- [VDOM Mode GUI changes](#)
- [Enhanced Top Sessions dashboard widget](#)
- [Improved CLI syntax for multi-value fields](#)

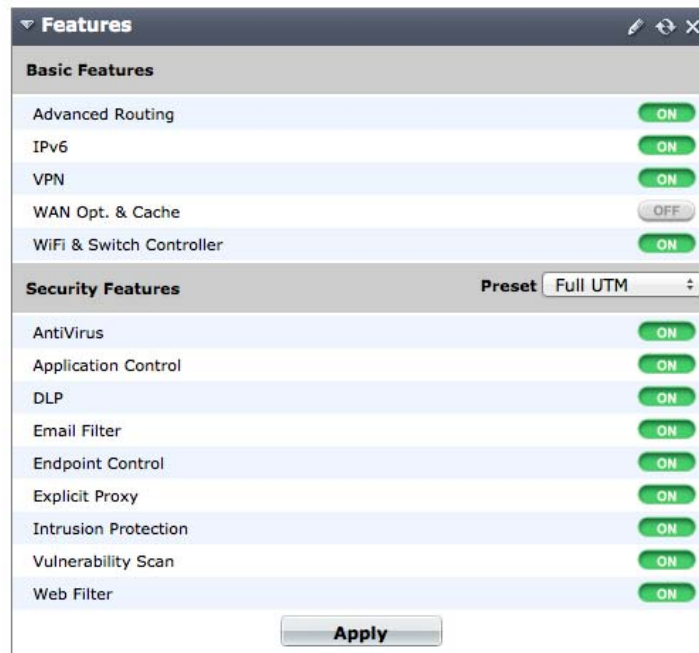
## Feature Select

Feature Select is used to disable features which are not required for network administration. Disabling features also removes all related configuration options from the web-based manager.

This feature replaces the previous GUI display options control.

Feature Select can be managed using the *Features* widget on the *Status* page. They can also be found at *System > Config > Features*, where additional features are also available by selecting *Show More*.

**Figure 80:**The Features widget





If a feature, such as IPv6, has been configured before being removed from the web-based manager, this configuration will still exist as part of the network, even though it is no longer visible using the web-based manager.

## Security Features Presets

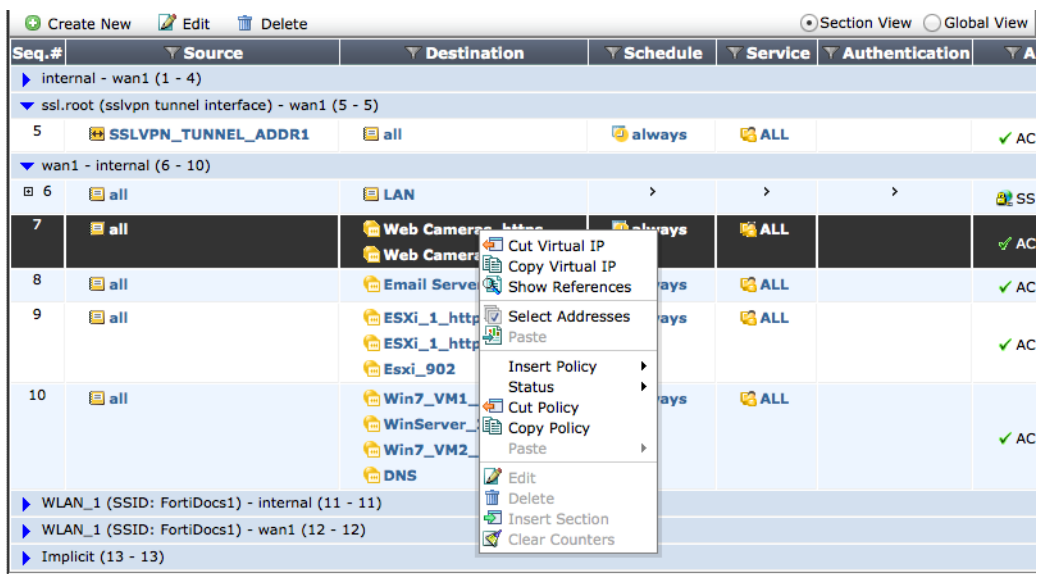
The main Security Features can be turned off individually or the five system presets can be used:

- *UTM* should be chosen for networks that require full protection from FortiOS. UTM is the default setting.
- *WF* should be chosen for networks that require web filtering.
- *ATP* should be chosen for networks that require protection from viruses and other external threats.
- *NGFW* should be chosen for networks that require application control and protection from external attacks.
- *NGFW + ATP* should be chosen for networks that require protection from external threats and attacks.

## Improved list editing

List editing has been enhanced on most lists of configuration items on the FortiOS 5.0 web-based manager. On most list items, you can click on any item to display a list of options. The options available depend on the item and context.

**Figure 81:** Example security policy list address menu



## Dynamic comment fields

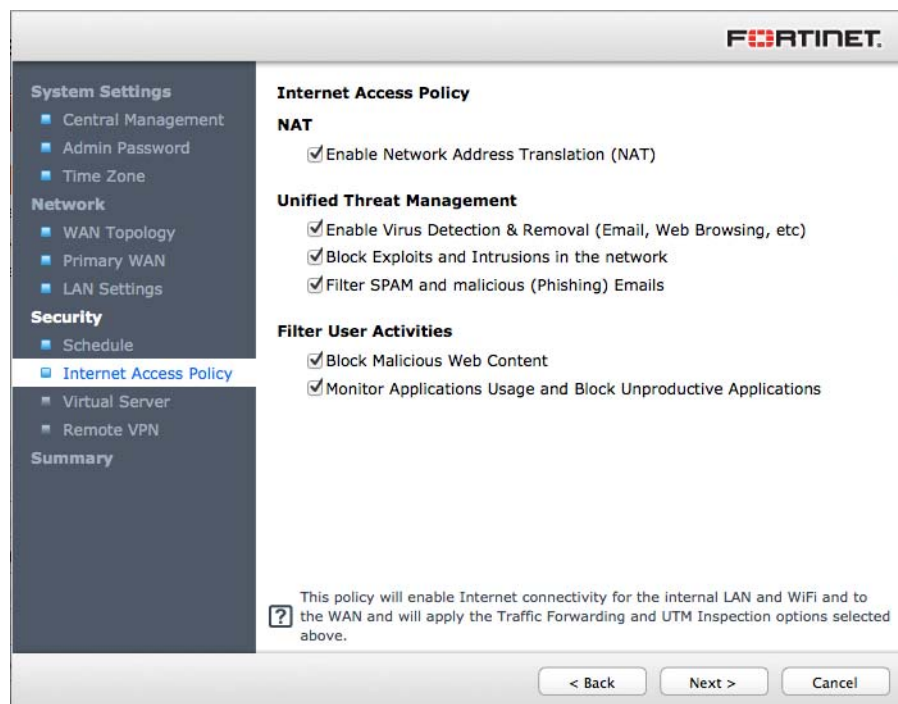
Most comment fields available from the web-based manager and CLI have been extended to a maximum of 255 characters and others to a maximum of 1023 characters. To save system memory, the amount of storage space for comment fields is dynamically allocated based on the size of the comment.

## Setup Wizard enhancements

The Setup Wizard is now available for all FortiGate units. On individual models the wizard can include advanced or model-specific configuration options, such as load balancing, 3G/4G modem, virtual servers, remote VPN and the opportunity to configure all available interfaces.

The Setup Wizard also allows you to enable central management. When this option is selected, much of the Wizard is bypassed because a FortiManager unit supplies the configuration information. The WAN Topology and Primary WAN wizard pages are still presented for configuration because the FortiGate unit must be able to connect to its network before FortiManager can contact it.

**Figure 82:** Setup wizard - Internet access policy setup



## Fortinet Top Bar



You can configure the FortiGate unit to overlay a Fortinet status bar on your user's web pages by going to *Policy > Policy > Proxy Options* and selecting *Add Fortinet Bar* and then by adding this Proxy Options profile to a security policy. Whenever a user accesses a web page through this policy, the Fortinet Top Bar is displayed overlaying the upper right corner of the web page.

The top bar can display a user ID if the user has authenticated with the FortiGate unit (in the example the user ID is bdickie). You can select the user icon to sign out of the FortiGate unit.

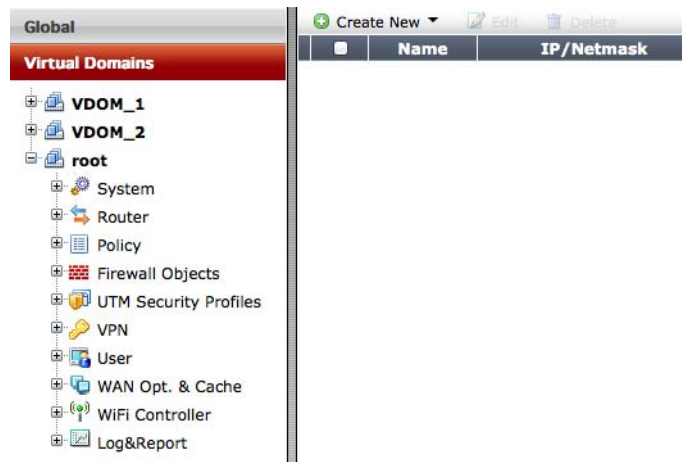
The top bar can also display other information such as:

- Application control violations
- Endpoint control enforcement
- Your web browsing quota
- User ID if the user has authenticated
- SSL VPN status and bookmarks

## VDOM Mode GUI changes

When operating a FortiGate unit with VDOMs enabled, when you log in as a system administrator who can access multiple VDOMs, the VDOMs you can access appear under Virtual Domains in the left web-based manager menu. [Figure 83](#) shows an example VDOM menu for a FortiGate unit with three virtual domains.

**Figure 83:**VDOM menu



## Enhanced Top Sessions dashboard widget

The Top Sessions dashboard widget has been enhanced to allow you to display information about sessions according to their source address, destination address and the application creating the sessions. To demonstrate this new functionality, by default the web-based manager includes three new dashboards.




You may need to reset your dashboard if these dashboard widgets are not appearing. To reset the dashboard, select *Reset Dashboards* from the *Dashboard* menu.

### Top Sources

Top sources displays the top 25 source addresses. For each source address, the widget displays device information, host name, the number of active sessions and the amount of data sent/received by the device. You can hover over the device icon to get more information about

the device. You can also select an entry to see all of the individual sessions from that source address. The session table includes the destination address, security policy, application name and amount of data sent and received by the session.

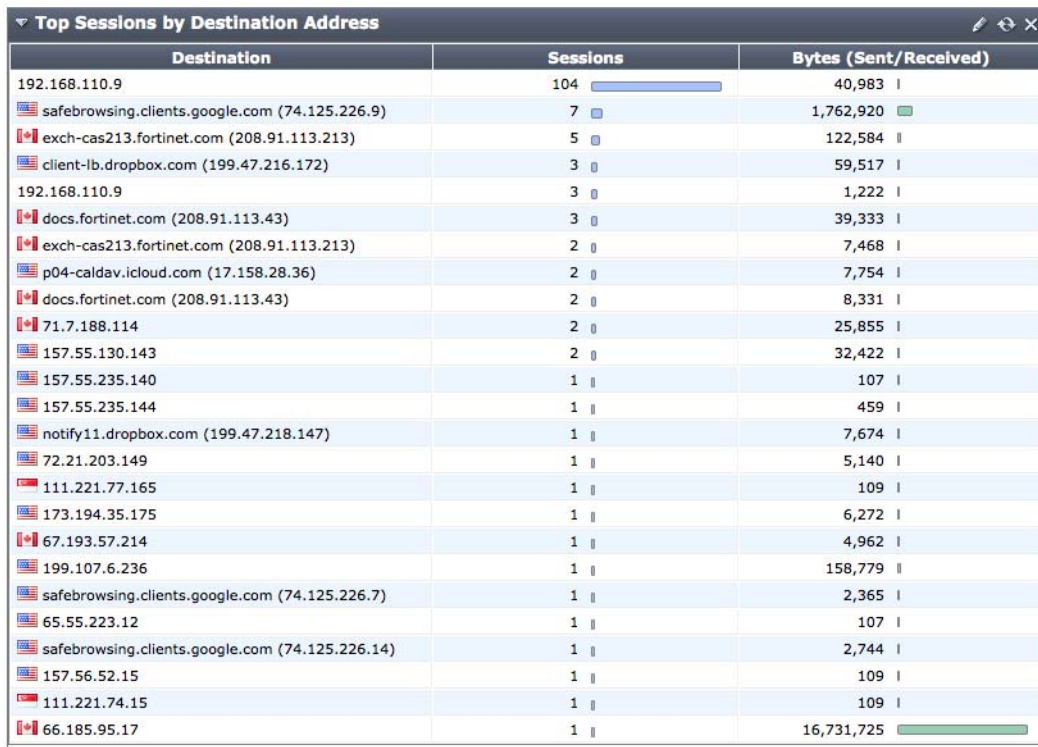
**Figure 84:**Top Sessions by Source Address



### Top Destinations

Top destinations displays the top 25 destination addresses. For each destination address, the widget displays the destination host name and country, the number of sessions going to the destination and the amount of data in bytes sent and received by the destination. You can select an entry to see the individual sessions to that destination address. The session table includes the source address, security policy, application name and amount of data sent received by the session.

**Figure 85:**Top Sessions by Destination Address



## Top Applications

Top applications displays the top 25 applications sending traffic through the FortiGate unit. For this widget to show data, you must enable application control for security policies that allow traffic through the FortiGate unit.

**Figure 86:**Top Sessions by Application



| Application         | Sessions | Bytes (Sent/Received) |
|---------------------|----------|-----------------------|
| DNS                 | 54       | 20,133 I              |
| Skype.Communication | 17       | 2,924 I               |
| Unknown             | 10       | 103,455 I             |
| HTTP.BROWSER        | 8        | 46,488 I              |
| WorldofWarcraft     | 2        | 60,954 I              |
| Youtube             | 2        | 17,504 I              |
| HTTP.Video          | 1        | 33,924,494            |
| Twitter             | 1        | 2,152 I               |
| Dropbox             | 1        | 51,487 I              |

For each application, the widget displays the application name, the number of sessions and the amount of data in bytes sent and received. You can select the application name to get more information about the application. You can select an entry to see the individual sessions for that application. The session table includes the source address, destination address, security policy, application name and amount of data sent received by the session.

**Figure 87:**Details about the sessions for an application



| # | Src               | Dst                | Policy ID | Device            | Host  | Application Name | Bytes (Sent/Received) |
|---|-------------------|--------------------|-----------|-------------------|-------|------------------|-----------------------|
| 1 | 11.11.11.20:49577 | 199.107.6.235:3724 | 1         | 34:15:9e:1c:a3:b4 | wd-mb | WorldofWarcraft  | 30,171                |
| 2 | 11.11.11.20:49578 | 199.107.6.236:3724 | 1         |                   |       |                  | 78,864                |

**Device Details**

Device: 34:15:9e:1c:a3:b4  
 OS: Mac OS X / 10.8  
 Hostname: wd-mb  
 IP Address: 0.0.0.0  
 Last Seen: 1 second ago (internal)

New Sessions per Second: 4 / Total Concurrent Sessions: 264

1 / 1 Total: 2

## Identifying Skype sessions

If Skype is in use on your network, Skype sessions may appear on the *Top Sessions By Application* list with the *Application Name* displayed as *unknown*. You can help the FortiGate unit identify Skype sessions by using the following command to add the public IP address of your network to the FortiGate configuration.

For example, if the IP address of the FortiGate interface connected to the Internet is 172.20.120.14 and if the security policies for connections to the Internet have source NAT enabled, enter the following command to add the public IP address of your network which is the public address used by Skype sessions:

```
config ips global
 set skype-client-public-ipaddr 172.20.120.14
end
```



You can add multiple IP addresses with this command. This can be useful if your network or your Skype sessions have more than one public IP address. For example, you may have multiple Internet connections each with a different IP address. Also, if the external IP address is set using DHCP or PPPoE it may change and you can add multiple IP addresses to help account for this. Use the following command to add multiple public IP addresses (separate the addresses with a comma and no spaces).

```
config ips global
 set skype-client-public-ipaddr 172.20.120.14,10.10.10.20
end
```



You may not have direct knowledge of your network's public IP address. This can happen for a number of reasons, depending on your network configuration. For example, your FortiGate unit may not be connected directly to the Internet. To make sure you are adding the right IP addresses you can use free services such as WhatIsMyIP.com to verify your network's public IP address.

## Customizing the Top Sessions dashboard widget

You can create multiple top sessions dashboard widgets that report sessions by source address, destination address or application. You can customize the widgets with a custom widget name, control the source and destination interfaces of the sessions and determine whether to sort the sessions by bytes or by number of sessions.

**Figure 88:** Customizing the Top Sessions dashboard widget

The screenshot shows a configuration window titled "Dashboard - Custom Top Sessions Display". The settings are as follows:

- Custom Widget Name: wan2 Sessions
- Report By: Destination Address
- Sort By: Sessions
- Source Interface: All
- Destination Interface: wan2
- Top Sessions To Show: 25
- Automatically Refresh:
- Resolve Hostnames:
- IP Version:  IPv4  IPv6  Both

Buttons for "OK" and "Cancel" are visible at the bottom.

## Improved CLI syntax for multi-value fields

Several new subcommands simplify editing of CLI fields that accept multiple values. This eliminates re-typing of lists of options, IP addresses, and so on, saving time and avoiding errors. You can simply add or remove individual items from the list.

**Table 3:** CLI subcommands for multi-value fields

|                                                       |                                          |
|-------------------------------------------------------|------------------------------------------|
| <code>append &lt;field_name&gt; &lt;list&gt;</code>   | Add one or more values to the list.      |
| <code>unselect &lt;field_name&gt; &lt;list&gt;</code> | Remove one or more values from the list. |

---

|                                                     |                                                                                                          |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------|
| <code>select &lt;field_name&gt; &lt;list&gt;</code> | Select one or more values. This is the same as the <code>set</code> subcommand.                          |
| <code>clear &lt;field_name&gt;</code>               | Reset the multi-value field to its default value. This is the same as the <code>unset</code> subcommand. |

---

You can continue to use the `set` and `unset` subcommands on multi-value fields.

The new subcommands support command completion. For example, in the `config system interface` command, if you enter `select ?`, the response shows only the multi-value fields:

```
dhcp-relay-ip dhcp relay ip address
allowaccess Allow management access to the interface
```

### Example

Prior to FortiOS 5.0, to add SSH administrative access to an interface that currently allows HTTPS, FGFM and PING access, you would enter:

```
set allowaccess https fgfm ping ssh
```

In FortiOS 5.0, you can do this:

```
append allowaccess ssh
```

Similarly, to remove `ping` from the list, you would enter:

```
unselect allowaccess ping
```

To reset `allowaccess` to its default, you would enter:

```
clear allowaccess
```

# SSL VPN

New SSL VPN features include:

- [New default SSL VPN portals](#)
- [SSL VPN user groups no longer required](#)
- [SSL VPN policy interface name change](#)
- [Support SSL VPN push configuration of DNS suffix](#)

## New default SSL VPN portals

FortiOS 5.0 includes 3 new default SSL VPN portal configurations:

- full-access is a general use portal that includes tunnel mode and web mode and supports all possible supported applications over the SSL VPN. Split tunneling and including the FortiClient download is not enabled.
- tunnel-access only includes support for tunnel mode (and not web mode). Split tunneling is enabled and remote users are prompted to download FortiClient.
- web-access only includes support for web mode. The connection tool and FortiClient download options are disabled.

## SSL VPN user groups no longer required

The distinction between SSL VPN user groups and firewall user groups has been removed. Any use group can be used for SSL VPN authentication, except FSSO user groups.

## SSL VPN policy interface name change

The former SSL.root interface used in SSL VPN security policies as the source or destination interface for SSL VPN traffic has been renamed to *sslvpn tunnel interface*.

## Support SSL VPN push configuration of DNS suffix

You can now assign one or more DNS suffixes to the FortiGate SSL VPN configuration so that SSL VPN clients do not need to use full-qualified host names to connect to internal resources. If you add one suffix, it is always attached to DNS queries.

If you add more than one suffix the FortiGate unit will attempt a DNS lookup by adding each suffix and use the first one that can be found in the DNS database. Multiple suffixes should be added in the proper search order. You can use up to 253 characters to add one or more DNS suffixes. Separate the suffixes with a space.

For example, if an organization requires DNS suffixes for example.com and example.org and you want DNS queries to try example.com first, you can use the following command to add these suffixes to the SSL VPN configuration:

Use the following command to add a DNS suffix that is used for SSL VPN sessions:

```
config vpn ssl settings
 set dns-suffix "example.com example.org"
end
```

# Other new features

This chapter provides a brief introduction to the following new features:

- New FortiGuard features
- FortiGate Auto-config using DHCP
- FortiGate Session Life Support Protocol (FGSP)
- HA failover supports more features
- New HA mode: Fortinet redundant UTM protocol (FRUP)
- ICAP and the explicit web proxy
- New interface features - DHCP server and authentication
- Replacement Message Improvements
- Acceleration of Inter-VDOM Traffic (by NP4)
- Virtual Hardware Switch
- FortiExplorer for iOS devices
- Inter-VDOM links between NAT mode and Transparent mode VDOMs
- Sniffer modes: one-armed and normal
- Integrated switch fabric (ISF) access control list (ACL) short-cut path
- Generalized TTL Security Mechanism (GTSM) support
- Firewall services

## New FortiGuard features

The following FortiGuard services are available for any FortiGate unit free of charge:

- System time from FortiGuard NTP servers
- Default DNS configuration uses FortiGuard DNS servers

With a valid support contract, the following FortiGuard services are also available:

- Antivirus database updates
- IPS signature updates
- Vulnerability scan signature updates
- FortiGuard Web Filtering lookups
  - DNS-based Web Filtering; the FortiGuard DNS server network returns web filter ratings. DNS web filtering uses less CPU time, system memory and network bandwidth than proxy or flow-based FortiGuard web filtering, resulting in better performance.
  - IP address reputation scores from the FortiGuard DNS server network.
- FortiGuard Email filtering lookups
- Geographic address database for geographic firewall addressing
- BYOD device signature updates
- USB Modem Updates

## FortiGate Auto-config using DHCP

FortiOS 5.0 supports uploading a configuration file from a TFTP server to the FortiGate unit to automatically configure the FortiGate unit with one simple step. Similar to an auto-configuration feature used for VoIP phones, you can store the domain name or IP address of a TFTP server and a configuration file name in your DHCP server configuration.

- DHCP option 66 is used for the TFTP server domain name ([RFC 2132](#))
- DHCP option 67 is used for the configuration file name ([RFC 2132](#))

For example, to use auto-configuration to configure a FortiGate unit, add the TFTP server information and configuration file name to your DHCP server. Make sure the TFTP server is running and includes the configuration file. Then, from the CLI of the FortiGate unit to be auto-configured, enter the following command (assuming the FortiGate internal interface is connected to the same network as the TFTP server).

```
execute restore config dhcp internal
```

The FortiGate unit gets the information it needs from the DHCP server, downloads and installs the configuration file from the TFTP server and restarts running its new configuration.

If the TFTP server is only available on a VLAN network (for example, VLAN id 224), you can use the following command to access the TFTP server on the VLAN network:

```
execute restore config dhcp internal 224
```

## FortiGate Session Life Support Protocol (FGSP)

In a network that already includes load balancing (either with load balancers or routers) for traffic redundancy, two FortiGate units can be integrated into the load balancing configuration using the FortiGate Session Life Support Protocol (FGSP). The external load balancers or routers can distribute IPv4 and IPv6 TCP, UDP, ICMP and expectation, and NAT sessions among the FortiGate units and the FGSP performs **session synchronization** to keep the session tables of both FortiGate units synchronized.

If one of the FortiGate units fails, session failover occurs and active sessions fail over to the unit that is still operating. This failover occurs without any loss of data. As well, the external load balancers or routers detect the failover and re-distribute all sessions to the unit that is still operating.

Load balancing and session failover is done by external routers or load balancers and not by the FGSP. The FortiGate units perform session synchronization which allows session failover to occur without packet loss.


The FGSP also includes **configuration synchronization**, allowing you to make configuration changes at once for both FortiGate units, instead of requiring duplicate configuration changes on each unit.

Settings that identify the FortiGate unit to the network, for example, interface IP addresses and BGP neighbor settings, are not synchronized so that each FortiGate unit maintains its identity on the network. These settings must be configured separately for each FortiGate unit.



In previous versions of FortiOS, the FGSP was called TCP session synchronization or standalone session synchronization. However, the FGSP has been expanded to include configuration synchronization and session synchronization of connectionless sessions, expectation sessions, and NAT sessions.

---

**Figure 89:**Example FGSP HA configuration

## HA failover supports more features

FortiOS 5.0 HA and the FortiGate Clustering Protocol (FGCP) support the following new types of failover:

- IPv6 session failover: if session pickup is enabled, IPv6 sessions are synchronized between cluster members and, after an HA failover, IPv6 sessions will resume with only minimal interruption.
- NAT64 session failover: if session pickup is enabled, NAT64 sessions are synchronized between cluster members and, after an HA failover, NAT64 sessions will resume with only minimal interruption.
- Full support for NAT 66 session failover: if session pickup is enabled, after an HA failover, NAT66 sessions will resume with only minimal interruption.
- SSL VPN authentication failover support: if session pick is enabled, SSL VPN sessions will resume after a failover without requiring SSL VPN users to re-authenticate.
- Device identification and management (BYOD).

## New HA mode: Fortinet redundant UTM protocol (FRUP)

FortiOS 5.0 includes an extension to the FortiGate Clustering Protocol that combines Switching HA and Firewall HA into a single unified design. This feature is initially available on the FortiGate-100D and will be considered for other models in future releases.

A FRUPS setup consists of 2 (and only 2) identical FortiGate-100D units. The setup supports dual redundant HA links between the units for sharing session and configuration data.

To see a FRUP example, please refer to [Cookbook Beta - FortiGate Redundant UTM Protocol](#).

## ICAP and the explicit web proxy

Internet Content Adaptation Protocol (ICAP) profiles can be added to explicit web proxy security policies. ICAP is a light-weight response/request protocol that allows the FortiGate unit to offload explicit web proxy traffic to external servers for different kinds of processing. ICAP is often used for offloading virus scanning and web filtering but has many other applications.

If you enable ICAP in a web proxy security policy, HTTP traffic intercepted by the explicit web proxy is transferred to an ICAP server in the ICAP profile added to the policy. Responses from the ICAP server are returned to the FortiGate unit which forwards them to their destination.

You can offload HTTP responses or HTTP requests (or both) to the same or different ICAP servers. You can also enable streaming media bypass.

### Example ICAP sequence for an ICAP server performing web URL filtering on web proxy HTTP requests


- 1 A user opens a web browser and sends an HTTP request to connect to a web server.
- 2 The FortiGate unit intercepts the HTTP request and forwards it to an ICAP server.
- 3 The ICAP server receives the request and determines if the request is for URL that should be blocked or allowed.
  - If the URL should be blocked, the ICAP server sends a response to the FortiGate unit. The FortiGate unit returns this response to the user's web browser. This response could be a message informing the user that their request was blocked.
  - If the URL should be allowed, the ICAP server sends a request to the FortiGate unit. The FortiGate unit forwards the request to the web server that the user originally attempted to connect to.

When configuring ICAP on the FortiGate unit, you must configure an ICAP profile that contains the ICAP server information; this profile is then applied to a security policy.

### Example ICAP configuration

The following example shows how to configure the FortiGate unit to offload processing to an ICAP server. The ICAP server IP address is 172.16.122.151 and port it is listening on is 453. The ICAP server request path /service-in and its response path is /service-out.

**Figure 90:**Example ICAP network configuration





## Adding ICAP to a web proxy security policy - web-based manager

In order to configure ICAP using the web-based manager, this feature must be enabled using Feature Select. For more information, see [“Feature Select” on page 243](#).

The following is an example of configuring the ICAP feature on the FortiGate unit and applying an ICAP profile to an existing web proxy security policy.

- 1 Go to *Security Profiles > ICAP > Servers* and select *Create New* to add the following ICAP server:

|                   |                 |
|-------------------|-----------------|
| <b>Name</b>       | New ICAP Server |
| <b>IP Type</b>    | IPv4            |
| <b>IP Address</b> | 172.16.122.151  |
| <b>Port</b>       | 453             |

- 2 Go to *Security Profiles > ICAP > Profiles* and select *Create New* to add an ICAP profile names *New ICAP Profile*.
- 3 Select *Enable Request Processing* and configure the following:

|                   |                 |
|-------------------|-----------------|
| <b>Server</b>     | New ICAP Server |
| <b>Path</b>       | /service-in     |
| <b>On Failure</b> | Error           |

- 4 Select *Enable Response Processing* and configure the following:

|                   |                 |
|-------------------|-----------------|
| <b>Server</b>     | New ICAP Server |
| <b>Path</b>       | /service-out    |
| <b>On Failure</b> | Error           |

- 5 Select *Enable Streaming Media Bypass* and select *OK*.
- 6 Go to *Policy > Policy > Policy* and edit the security policy that accepts the traffic to be processed by the ICAP server.
- 7 Under *Security Policies*, select *Enable ICAP* and set *New ICAP Server*.
- 8 Select *OK*.

## Adding ICAP to a web proxy security policy - CLI

The following is an example of configuring the ICAP feature on the FortiGate unit and applying an ICAP profile to an existing web proxy security policy.

- 1 Log in to the CLI.
- 2 Enter the following to configure the ICAP server:

```
config icap server
 edit "New ICAP Server"
 set ip-address 172.16.122.151
 set ip-version 4
 set max-connections 100
 set port 453
 end
```

- 3 Enter the following to configure the ICAP profile to then apply to a security policy:

```
config icap profile
 edit "New ICAP Profile"
 set request enable
 set request-failure error
 set request-path "/service-in"
 set request-server icap_server
 set response enable
 set response-failure error
 set response-path "/service-out"
 set response-server "New ICAP Server"
 set streaming-content-bypass enable
 end
```

- 4 In the config firewall policy command, apply the ICAP profile to a security policy:

```
config firewall policy
 edit 0
 set srcintf web-proxy
 ...
 set utm-status enable
 set icap-profile "New ICAP Profile"
 end
```

## New interface features - DHCP server and authentication

You can add a DHCP server and authentication to any FortiGate interface. This includes physical interfaces, WiFi interfaces (SSIDs), switch interfaces (software and hardware switch interfaces), aggregate interfaces, redundant interfaces, loopback interfaces and VLAN interfaces.

### Adding a DHCP server to an interface

To add a DHCP server to an interface, edit the interface and select *Enable DHCP Server*. Then you can specify the address range, netmask, default gateway and DNS servers provided by the DHCP server. An interface must have a static IP address to add a DHCP server to it.

**Figure 91:** Adding a DHCP server to a FortiGate interface

Enable DHCP Server

Address Range  -

Netmask

Default Gateway  Same As Interface IP  Specify

DNS Server  Same As System DNS  Specify

▼ MAC Address Access Control List

[+ Create New](#) [Edit](#) [Delete](#)

| <input type="checkbox"/> | MAC                   | IP or Action |
|--------------------------|-----------------------|--------------|
| <input type="checkbox"/> | c4:2c:03:21:af:04     | 10.10.10.200 |
| <input type="checkbox"/> | Unknown MAC Addresses | Assign IP    |

## Reserving, assigning and blocking MAC addresses

While adding a DHCP server to an interface you can select *MAC Address Control list* and select *Create New* to configure the DHCP server to:

- Always assign the same IP address to a device according to its MAC address (*Reserved IP address*)
- *Block* access to a device according to its MAC address (MAC address filtering)
- The default action is to *Assign* an IP address to a device.

**Figure 92:**Reserving an IP address for a device with a specific MAC address


## Authentication - Captive Portal

To add authentication to an interface, edit an interface and set *Security Mode* to *Captive Portal*. Then select one or more *User Groups*. Users who attempt to connect through the interface must first use HTTP or HTTPS to connect to a captive portal and enter a user name and password.

**Figure 93:**Reserving an IP address for a device with a specific MAC address

You can customize the captive portal for each interface or select from a saved portal.

**Figure 94:**Customizing the captive portal



## Replacement Message Improvements


Go to *System > Config > Replacement Messages* to edit replacement messages. Editing replacement messages has been enhanced with a new editor that you can use to select and edit replacement messages and view your changes in real time as you make them.

**Figure 95:**Editing replacement messages


| UTM                   |                                                      |
|-----------------------|------------------------------------------------------|
| FortiGuard Block Page | Replacement HTML for FortiGuard Webfilter block page |
| URL Block Page        | Replacement HTML for HTTP url blocked page           |
| Virus Block Page      | Replacement HTML for antivirus block page            |
| Virus Block Message   | Replacement text for antivirus block message         |
| DLP Block Page        | Replacement HTML for DLP block page                  |
| DLP Block Message     | Replacement text for DLP block message               |

Save
Restore Default
Message Format: text/html Message Size: 2584/8192



Powered by FortiGuard



## Web Page Blocked!

You have tried to access a web page which is in violation of your internet usage policy.

URL: www.example.com/  
Category: Personal Websites and Blogs

override

To have the rating of this web page re-evaluated [please click here](#).

```
<!DOCTYPE html PUBLIC "-//W3C//DTD H
<html>
<head>
<meta http-equiv="Content-Type"
<title>
 Web Filter Violation
</title>
<style type="text/css">
html, body {
 margin: 0;
 padding: 0;
 font-family: Verdana, Arial, s
 font-size: 10pt;
}
h1, h2 {
 height: 82px;
 text-indent: -999em;
 margin: 0;
 padding: 0;
 margin: 0;
}
div {
 margin: 0;
 padding: 0;
}
div.header {
 background: url(%IMAGE:logo_v
 height: 82px;
}
div.header h1 {
```

In addition, more replacement messages are available and the replacement message editor includes simple and extended view. The simple view includes the most commonly edited replacement messages while the extended view includes all of them.

## Acceleration of Inter-VDOM Traffic (by NP4)

On high-end FortiGate units that include NP4 processors, you can add inter-VDOM links where the traffic is accelerated by NP4 processors. This means enhanced performance for traffic passing between VDOMs that will improve the overall performance and capacity of many multiple VDOM implementations.

If your FortiGate unit supports accelerated inter-VDOM links, when it is operating in multiple VDOM mode, the interface list includes interfaces with names such as npu0-vlink0, npu1-vlink and so on (see [Figure 96](#)).

**Figure 96:**FortiGate-5001B interface list showing NP4 accelerated inter-VDOM links

	Name	Virtual Domain	IP/Netmask	Access	Administrative Status	Link Status	Type	Ref
<input type="checkbox"/>	<b>mesh.root</b> (SSID: fortinet.mesh.root)	root	0.0.0.0 / 0.0.0.0				WiFi Interface	<a href="#">0</a>
<input type="checkbox"/>	<b>mgmt1</b>	root	172.20.120.177 / 255.255.255.0	HTTP,HTTPS,PING,SSH,TELNET			Physical	<a href="#">0</a>
<input type="checkbox"/>	<b>mgmt2</b>	root	192.168.100.99 / 255.255.255.0	PING,FMG-Access			Physical	<a href="#">1</a>
<input type="checkbox"/>	<b>npu0-vlink</b> (VDOM Link)	root, root	-				VDOM Link	<a href="#">0</a>
<input type="checkbox"/>	npu0-vlink0	root	0.0.0.0 / 0.0.0.0				Pair	<a href="#">0</a>
<input type="checkbox"/>	npu0-vlink1	root	0.0.0.0 / 0.0.0.0				Pair	<a href="#">0</a>
<input type="checkbox"/>	<b>npu1-vlink</b> (VDOM Link)	root, root	-				VDOM Link	<a href="#">0</a>
<input type="checkbox"/>	npu1-vlink0	root	0.0.0.0 / 0.0.0.0				Pair	<a href="#">0</a>
<input type="checkbox"/>	npu1-vlink1	root	0.0.0.0 / 0.0.0.0				Pair	<a href="#">0</a>
<input type="checkbox"/>	<b>port1</b>	root	0.0.0.0 / 0.0.0.0				Physical	<a href="#">0</a>
<input type="checkbox"/>	<b>port2</b>	root	0.0.0.0 / 0.0.0.0				Physical	<a href="#">1</a>
<input type="checkbox"/>	<b>port3</b>	root	0.0.0.0 / 0.0.0.0				Physical	<a href="#">0</a>
<input type="checkbox"/>	<b>port4</b>	root	0.0.0.0 / 0.0.0.0				Physical	<a href="#">0</a>
<input type="checkbox"/>	<b>port5</b>	root	0.0.0.0 / 0.0.0.0				Physical	<a href="#">0</a>
<input type="checkbox"/>	<b>port6</b>	root	0.0.0.0 / 0.0.0.0				Physical	<a href="#">0</a>
<input type="checkbox"/>	<b>port7</b>	root	0.0.0.0 / 0.0.0.0				Physical	<a href="#">0</a>
<input type="checkbox"/>	<b>port8</b>	root	0.0.0.0 / 0.0.0.0				Physical	<a href="#">0</a>

By default, all of these links are associated with the root VDOM. However, you can edit each interface in the link and add it another VDOM, creating an inter-VDOM link between 2 VDOMs

## Virtual Hardware Switch

In previous versions of FortiOS, you can use the software switch feature to group independent interfaces into a single logical switch. In this virtual software switch, all of the interfaces share the same IP address and be connected to the same subnet and traffic would pass between them as if they were switch ports, with no firewall or other FortiGate features applied to the traffic. However, the virtual software switch feature just simulates a switch and, since the FortiGate CPU must process the switch traffic, performance can be affected if the FortiGate unit becomes busy processing a lot of traffic.

In FortiOS 5.0, for FortiGate models that have internal hardware switches, you can use the following command to group interfaces in the hardware switch into virtual hardware switches in which all traffic between the switch ports is processed on the switch itself and the FortiGate CPU is not involved resulting in improved performance.

Recent FortiGate models with internal hardware switches support this feature.

Use the following command to create a virtual hardware switch using ports p1, p2, p3, and p4:

```
config system virtual-switch
 edit virt-sw-1
 set physical-switch sw0
 config port
 edit 1
 set port p1
 set speed <speed>
 set duplex { up | down}
 next
 edit 2
 set port p2
 set speed <speed>
 set duplex { up | down}
 end
 edit 3
 set port p3
 set speed <speed>
 set duplex { up | down}
 next
 edit 4
 set port p4
 set speed <speed>
 set duplex { up | down}
 end
 end
 end
```

## FortiExplorer for iOS devices

You can use FortiExplorer for iOS to manage most FortiGate models running FortiOS 5.0 firmware from an iOS device (iPhone, iPad, or iPod Touch running iOS 5.0 or later). FortiExplorer is a free download from the Apple iOS App Store ([App Store Link](#)).



## Connecting to and logging into a FortiGate unit



Use FortiExplorer for iOS by connecting your iOS device to any FortiGate USB port, using the USB cable that came with your iOS device. (Connect to any FortiGate USB-A port. There is no need to connect to the USB-B port required for the PC or Mac OS versions of FortiExplorer).

Start FortiExplorer on your iOS device and select *Setup* and log into the FortiGate unit using any administrator account user name and password.

After logging in, you can use FortiExplorer to change the firmware running on the FortiGate unit, configure network settings, and change general system settings. You also have one-step access to the configuration of each FortiGate interface.

## Updating firmware and configuring network settings

Usually you would use FortiExplorer to upgrade firmware and configure network and basic settings. Then log into the web-based manager for more advanced configuration. However, you can select *Web* to connect directly to the FortiGate unit's web-based manager from your iOS device.

You can also select *Firmware* to view the available firmware versions for any FortiGate model and download new firmware and install it on your FortiGate unit.



## Inter-VDOM links between NAT mode and Transparent mode VDOMs

FortiOS 5.0 supports inter-VDOM links between NAT and Transparent mode VDOMs. No special configuration is required and you can create an inter-VDOM link between NAT and Transparent mode VDOMs in the same way as creating an inter-VDOM link between two NAT mode VDOMs.



## About inter-VDOM links between NAT and Transparent mode VDOMs

Inter-VDOM links between NAT and Transparent mode VDOMs can be useful for configurations where the NAT based VDOMs that share a common Internet service route, which can be routed through a Transparent VDOM that provides additional functionality, like common Security inspection, WAN optimization, explicit proxying and so on.

Other examples include:

- Performing SSL offloading in the Transparent mode VDOM and providing Internet access through a NAT mode VDOM.
- Applying WAN optimization in a Transparent mode VDOM and other security features in the NAT mode VDOM.
- Using a dedicated Transparent mode VDOM for the explicit web proxy in front of a NAT mode VDOM that applies other security features.
- An ISP configuration with multiple per-tenant NAT mode VDOMs all sharing a single Internet connection but where the ISP only presents a single routed subnet. Each tenant can then be assigned an IP from the subnet for their respective VDOM link interface while using a single physical port to connect to the ISP router.

## Sniffer modes: one-armed and normal

FortiGate units can operate in one-arm sniffer mode or as a regular traffic sniffer. When the FortiGate unit has an interface dedicated to its exclusive use (one-arm sniffer mode), all traffic entering the interface is processed by the sniffer. The traffic is compared to the configured filters and data that doesn't match the filters is discarded. The selected Security profiles process the remaining traffic and log their findings. At the same time, the packets triggering the configured Security features are saved for later examination. After all examination of the traffic is complete, it is discarded. This continues until the configured maximum number of packets are saved, when the sniffer stops.

When the sniffer does not have an interface dedicated to its exclusive use, the traffic is examined by the sniffer, then processed normally by the FortiGate unit. That is, traffic is sniffed and can then leave the FortiGate unit depending on how it is configured. The sniffed interface traffic is examined for traffic matching the sniffer filters and matching packets are saved. Security features can not be used to limit the traffic the sniffer examines when not in one-armed mode. When the configured maximum number of packets is saved, the sniffer stops. The FortiGate unit continues to process network traffic as normal.

### Configuring an interface to operate as a one-arm sniffer

Connect the interface to the network to be analyzed. Go to *System > Network > Interfaces* and edit the interface and select *One-Arm Sniffer* and select *Apply*.

Configure the sniffer by selecting *Enable Filters* to filter traffic by IP address (host), address range, port number, VLANs and protocols. You can also configure the sniffer to *Include IPv6 Packets* and to *Include Non-IP Packets*.

Finally, under *Security Profiles* you can select Security profiles to apply to the sniffer.

Select *Apply* to save your changes.

**Figure 97:**Example one-arm sniffer configuration

Edit Interface	
Name	port5 (00:09:0F:4E:10:23)
Alias	Sniffer
Link Status	Down <span style="color: red;">⬇</span>
Type	Physical Interface
Addressing mode	<input type="radio"/> Manual <input type="radio"/> DHCP <input type="radio"/> PPPoE <input checked="" type="radio"/> One-Arm Sniffer <input type="radio"/> Dedicate to FortiAP/FortiSwitch
<input checked="" type="checkbox"/> Enable Filters	
Host(s)	72.20.120.100-172.20.120.120 <span>?</span>
Port(s)	80, 8080 <span>?</span>
VLAN(s)	23 <span>?</span>
Protocol	<span>?</span>
<input checked="" type="checkbox"/> Include IPv6 Packets	
<input type="checkbox"/> Include Non-IP Packets	
<b>Security Profiles</b>	
<input checked="" type="checkbox"/> Enable AntiVirus	default
<input checked="" type="checkbox"/> Enable Web Filter	dns-wf
<input checked="" type="checkbox"/> Enable Application Control	block-p2p
<input checked="" type="checkbox"/> Enable IPS	all_default
<input checked="" type="checkbox"/> Enable Email Filter	default

## Integrated switch fabric (ISF) access control list (ACL) short-cut path

On FortiGate models that include NP4 and XLR ports and an integrated switch fabric (for example, the FortiGate-3x40 and 3950/1 models), you can create an ISF ACL security policy that allows some traffic (for example, multicast traffic) to bypass security inspection, resulting in reduced CPN and NP4 processor load.

This feature is only available in Transparent mode and only between port pairs.



Traffic accepted and forwarded by an ISF policy is not subject to security inspection. Normally, you should only create ISF policies for traffic that you consider very low risk.

Use the following command to add an ISF ACL shortcut policy:

```
config firewall isf-acl
 config port-pair-1
 edit 1
 set type binary
 set ingressport <port1 | port2>
 set offset
 set length
 set matchpattern <patter in hex>
 set action <bypass|block>
 edit 2
 set type 5-tuple
 set srcaddr: a.b.c.d/32
 set dstaddr 239.A.A.a/32
 set proto UDP
 set port XXX
 set action <bypass|block>
 end
```

## Generalized TTL Security Mechanism (GTSM) support

Generalized TTL Security Mechanism (GTSM), defined in [RFC 5082](#), prevents attacks based on forged protocol packets sent from outside the network.

In IP packets, the TTL (time-to-live) value sets the maximum number of routers the packet can pass through to reach its destination. Each router decrements the TTL value and the packet is discarded if TTL reaches zero before the packet reaches its destination. In IPv6, TTL is called Hop Limit.

Most protocol-related packets pass between adjacent routers, so the TTL value at the destination is within a predictable range. TTL is difficult to spoof, especially the value of 255 which occurs if the sender is directly connected to the destination router.

On the FortiGate firewall, you can define TTL policies that specify the acceptable TTL range for a particular packet source, destination and service. You do this using the new `config firewall ttl-policy` command.

Use the following command to add a TTL policy that sets the TTL range to from 20 to 30:

```
config firewall ttl-policy
 edit 0
 set srcintf port1
 set srcaddr example_net
 set service ALL
 set schedule always
 set ttl 20-30
 end
```

## Firewall services

The CLI command `get firewall service predefined` command has been removed. All predefined services have been moved to `{get | set} firewall service custom`.

# Chapter 2 Advanced Routing for FortiOS

## 5.0

This chapter describes advanced static routing concepts and how to implement dynamic routing on FortiGate units.

This FortiOS Handbook chapter contains the following sections:

[Advanced Static Routing](#) explains universal and static routing concepts, equal cost multipath (ECMP) and load balancing, policy routing, and routing in transparent mode.

[Dynamic Routing Overview](#) provides an overview of dynamic routing, compares static and dynamic routing, and helps you decide which dynamic routing protocol is best for you.

[Routing Information Protocol \(RIP\)](#) describes a distance-vector routing protocol intended for small, relatively homogeneous networks.

[Border Gateway Protocol \(BGP\)](#) describes classless inter-domain routing, and aggregate routes. BGP is the only routing protocol to use TCP for a transport protocol.

[Open Shortest Path First \(OSPF\)](#) provides background on the specific protocol explaining terms used and how the protocol works, as well as providing some troubleshooting information and examples on configuring the protocols in different situations.

[Intermediate System to Intermediate System Protocol \(IS-IS\)](#), which describes the link state protocol, is well-suited to smaller networks and with near universal support on routing hardware. The section also provides troubleshooting information and configuration examples.

# Advanced Static Routing

Advanced static routing includes features and concepts that are used in more complex networks. Dynamic routing is not addressed in this section.

This section includes:

- [Routing concepts](#)
- [Static routing tips](#)
- [Policy routing](#)
- [Transparent mode static routing](#)
- [Static routing example](#)
- [Advanced static example: ECMP failover and load balancing](#)

## Routing concepts

Many routing concepts apply to static routing. However without first understanding these basic concepts, it is difficult to understand the more complex dynamic routing.

This section includes:

- [Routing in VDOMs](#)
- [Default route](#)
- [Adding a static route](#)
- [Routing table](#)
- [Building the routing table](#)
- [Static routing security](#)
- [Multipath routing and determining the best route](#)

## Routing in VDOMs

Routing on FortiGate units is configured per-VDOM. This means if VDOMs are enabled, you must enter a VDOM to do any routing configuration. This allows each VDOM to operate independently, with its own default routes and routing configuration.

In this guide, the procedures assume your FortiGate unit has VDOMs disabled. This is stated in the assumptions for the examples. If you have VDOMs enabled you will need to perform the following steps in addition to the procedure's steps.

### **To route in VDOMs - web-based manager**

Select the VDOM that you want to view or configure at the bottom of the main menu.

### **To route in VDOMs - CLI**

Before following any CLI routing procedures with VDOMs enabled, enter the following commands. For this example, it is assumed you will be working in the root VDOM. Change root to the name of your selected VDOM as needed.

```
config vdom
edit root
```

Following these commands, you can enter any routing CLI commands as normal.

## Default route

The default route is used if either there are no other routes in the routing table or if none of the other routes apply to a destination. Including the gateway in the default route gives all traffic a next-hop address to use when leaving the local network. The gateway address is normally another router on the edge of the local network.

All routers, including FortiGate units, are shipped with default routes in place. This allows customers to set up and become operational more quickly. Beginner administrators can use the default route settings until a more advanced configuration is warranted.

FortiGate units come with a default static route with an IPv4 address of 0.0.0.0, an administration distance of 10, and a gateway IPv4 address.

## Adding a static route

To add or edit a static route, go to *Router > Static > Static Routes* and select *Create New*.

<b>Destination IP / Mask</b>	Enter the destination IP address and netmask. A value of 0.0.0.0/0.0.0.0 is universal.
<b>Device</b>	Select the name of the interface which the static route will connect through.
<b>Gateway</b>	Enter the gateway IP address.
<b>Distance</b>	Enter the distance value, which will affect which routes are selected first by different protocols for route management or load balancing. The default is 10.
<b>Priority</b>	Enter the priority if desired, which will artificially weight the route during route selection. The higher the number, the less likely the route is to be selected over others. The default is 0.

## Routing table

When two computers are directly connected, there is no need for routing because each computer knows exactly where to find the other computer. They communicate directly.

Networking computers allows many computers to communicate with each other. This requires each computer to have an IP address to identify its location to the other computers. This is much like a mailing address - you will not receive your postal mail at home if you do not have an address for people to send mail to. The routing table on a computer is much like an address book used to mail letters to people in that the routing table maintains a list of how to reach computers. Routing tables may also include information about the quality of service (QoS) of the route, and the interface associated with the route if the device has multiple interfaces.

Looking at routing as delivering letters is more simple than reality. In reality, routers lose power or have bad cabling, network equipment is moved without warning, and other such events happen that prevent static routes from reaching their destinations. When any changes such as these happen along a static route, traffic can no longer reach the destination — the route goes down. Dynamic routing can address these changes to ensure traffic still reaches its destination. The process of realizing there is a problem, backtracking and finding a route that is operational is called convergence. If there is fast convergence in a network, users won't even know that re-routing is taking place.

The routing table for any device on the network has a limited size. For this reason, routes that aren't used are replaced by new routes. This method ensures the routing table is always populated with the most current and most used routes—the routes that have the best chance of being reused. Another method used to maintain the routing table's size is if a route in the table and a new route are to the same destination, one of the routes is selected as the best route to that destination and the other route is discarded.

Routing tables are also used in unicast reverse path forwarding (uRPF). In uRPF, the router not only looks up the destination information, but also the source information to ensure that it exists. If there is no source to be found, then that packet is dropped because the router assumes it to be an error or an attack on the network.

The routing table is used to store routes that are learned. The routing table for any device on the network has a limited size. For this reason, routes that aren't used are replaced by new routes. This method ensures the routing table is always populated with the most current and most used routes — the routes that have the best chance of being reused. Another method used to maintain the routing table's size is if a route in the table and a new route are to the same destination, one of the routes is selected as the best route to that destination and the other route is discarded.

Some actions you can perform on the routing table include:

- [Viewing the routing table in the web-based manager](#)
- [Viewing the routing table in the CLI](#)
- [Searching the routing table](#)

## Viewing the routing table in the web-based manager

VDOM

By default, all routes are displayed in the Routing Monitor list. The default static route is defined as 0.0.0.0/0, which matches the destination IP address of “any/all” packets.

To display the routes in the routing table, go to *Router > Monitor > Routing Monitor*.

**Figure 98** shows the Routing Monitor list belonging to a FortiGate unit that has interfaces named “port1”, “port4”, and “lan”. The names of the interfaces on your FortiGate unit may be different.

**Figure 99** shows the Routing Monitor list when IPv6 has been selected. Note that the information available for IPv6 is limited.

**Figure 98:**Routing Monitor list - IPv4

IP Version: <span>IPv4</span> Type: <span>All</span> Network: <input type="text"/>		Gateway: <input type="text"/>		<span>Apply Filter</span>			
Type	Subtype	Network	Distance	Metric	Gateway	Interface	Up Time (d h:m:s)
Connected		10.10.10.0/24	0	0	0.0.0.0	port4	
Connected		172.20.120.0/24	0	0	0.0.0.0	port1	

**Figure 99:**Routing Monitor list - IPv6

IP Version: <span>IPv6</span>		Interface	Network	Gateway
		havlink1	fe80:::10	
		havlink1	ff00::/8	

<b>IP version:</b>	Select IPv4 or IPv6. This is available only when IPv6 is enabled in the web-based manager. The fields displayed in the table depend on which IP version is selected.
<b>Type:</b>	<p>Select one of the following route types to search the routing table and display routes of the selected type only:</p> <p><i>All</i> — all routes recorded in the routing table.</p> <p><i>Connected</i> — all routes associated with direct connections to FortiGate unit interfaces.</p> <p><i>Static</i> — the static routes that have been added to the routing table manually.</p> <p><i>RIP</i> — all routes learned through RIP. For more information see <a href="#">“Routing Information Protocol (RIP)” on page 319</a>.</p> <p><i>RIPNG</i> — all routes learned through RIP version 6 (which enables the sharing of routes through IPv6 networks).</p> <p><i>BGP</i> — all routes learned through BGP. For more information see <a href="#">“Border Gateway Protocol (BGP)” on page 358</a>.</p> <p><i>OSPF</i> — all routes learned through OSPF. For more information see <a href="#">“Open Shortest Path First (OSPF)” on page 396</a>.</p> <p><i>OSPF6</i> — all routes learned through OSPF version 6 (which enables the sharing of routes through IPv6 networks).</p> <p><i>IS-IS</i> — all routes learned through IS-IS. For more information see <a href="#">“Intermediate System to Intermediate System Protocol (IS-IS)” on page 438</a>.</p> <p><i>HA</i> — RIP, OSPF, and BGP routes synchronized between the primary unit and the subordinate units of a high availability (HA) cluster. HA routes are maintained on subordinate units and are visible only if you are viewing the router monitor from a virtual domain that is configured as a subordinate virtual domain in a virtual cluster.</p> <p>Not displayed when IP version IPv6 is selected.</p> <p>For details about HA routing synchronization, see the <a href="#">FortiGate HA User Guide</a>.</p>
<b>Network:</b>	Enter an IP address and netmask (for example, 172.16.14.0/24) to search the routing table and display routes that match the specified network.
<b>Gateway:</b>	Enter an IP address and netmask (for example, 192.168.12.1/32) to search the routing table and display routes that match the specified gateway.
<b>Apply Filter</b>	<p>Select to search the entries in the routing table based on the specified search criteria and display any matching routes.</p> <p>Not displayed when IP version IPv6 is selected.</p>
<b>Type</b>	<p>The type values assigned to FortiGate unit routes (Static, Connected, RIP, OSPF, or BGP).</p> <p>Not displayed when IP version IPv6 is selected.</p>



<b>Subtype</b>	<p>If applicable, the subtype classification assigned to OSPF routes.</p> <p>An empty string implies an intra-area route. The destination is in an area to which the FortiGate unit is connected.</p> <p><i>OSPF inter area</i> — the destination is in the OSPF AS, but the FortiGate unit is not connected to that area.</p> <p><i>External 1</i> — the destination is outside the OSPF AS. This is known as OSPF E1 type. The metric of a redistributed route is calculated by adding the external cost and the OSPF cost together.</p> <p><i>External 2</i> — the destination is outside the OSPF AS. This is known as OSPF E2 type. In this case, the metric of the redistributed route is equivalent to the external cost only, expressed as an OSPF cost.</p> <p><i>OSPF NSSA 1</i> — same as External 1, but the route was received through a not-so-stubby area (NSSA).</p> <p><i>OSPF NSSA 2</i> — same as External 2, but the route was received through a not-so-stubby area.</p> <p>For more information on OSPF subtypes, see <a href="#">“OSPF Background and concepts” on page 396</a>.</p> <p>Not displayed when IP version 6 is selected.</p>
<b>Network</b>	<p>The IP addresses and network masks of destination networks that the FortiGate unit can reach.</p>
<b>Distance</b>	<p>The administrative distance associated with the route. A value of 0 means the route is preferable compared to other routes to the same destination, and the FortiGate unit may routinely use the route to communicate with neighboring routers and access servers.</p> <p>Modifying this distance for dynamic routes is route distribution. See <a href="#">“Redistributing and blocking routes in BGP” on page 389</a></p> <p>Not displayed when IP version 6 is selected.</p>
<b>Metric</b>	<p>The metric associated with the route type. The metric of a route influences how the FortiGate unit dynamically adds it to the routing table. The following are types of metrics and the protocols they are applied to.</p> <p><i>Hop count</i> — routes learned through RIP.</p> <p><i>Relative cost</i> — routes learned through OSPF.</p> <p><i>Multi-Exit Discriminator (MED)</i> — routes learned through BGP. However, several attributes in addition to MED determine the best path to a destination network. For more information on BGP attributes, see <a href="#">“BGP attributes” on page 364</a>. By default, the MED value associated with a BGP route is zero. However, the MED value can be modified dynamically. If the value was changed from the default, the Metric column will display a non-zero value.</p> <p>Not displayed when IP version 6 is selected.</p>
<b>Gateway</b>	<p>The IP addresses of gateways to the destination networks.</p>

<b>Interface</b>	The interface through which packets are forwarded to the gateway of the destination network.
<b>Up Time</b>	The total accumulated amount of time that a route learned through RIP, OSPF, or BGP has been reachable.  Not displayed when IP version IPv6 is selected.

## Viewing the routing table in the CLI

In the CLI, you can easily view the static routing table just as in the web-based manager or you can view the full routing table.

When viewing the list of static routes using the CLI command `get route static`, it is the configured static routes that are displayed. When viewing the routing table using the CLI command `get router info routing-table all`, it is the entire routing table information that is displayed including configured and learned routes of all types. The two are different information in different formats.



If VDOMs are enabled on your FortiGate unit, all routing related CLI commands must be performed within a VDOM and not in the global context.

### To view the routing table

```
get router info routing-table all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
 area
* - candidate default

S* 0.0.0.0/0 [10/0] via 192.168.183.254, port2
S 1.0.0.0/8 [10/0] via 192.168.183.254, port2
S 2.0.0.0/8 [10/0] via 192.168.183.254, port2
C 10.142.0.0/23 is directly connected, port3
B 10.160.0.0/23 [20/0] via 10.142.0.74, port3, 2d18h02m
C 192.168.182.0/23 is directly connected, port2
```

### Examining an entry:

```
B 10.160.0.0/23 [20/0] via 10.142.0.74, port3, 2d18h02m
```

<b>B</b>	BGP. The routing protocol used.
<b>10.160.0.0/23</b>	The destination of this route including netmask.
<b>[20/0]</b>	20 indicates and administrative distance of 20 out of a range of 0 to 255. 0 is an additional metric associated with this route, such as in OSPF
<b>10.142.0.74</b>	The gateway, or next hop.
<b>port3</b>	The interface used by this route.
<b>2d18h02m</b>	How old this route is, in this case almost three days old.

### To view the kernel routing table

```
get router info kernel
```

```
tab=254 vf=0 scope=253 type=1 proto=2 prio=0
 0.0.0.0/0.0.0.0/0->10.11.201.0/24 pref=10.11.201.4 gwy=0.0.0.0
 dev=5(external1)
```

```
tab=254 vf=0 scope=253 type=1 proto=2 prio=0
 0.0.0.0/0.0.0.0/0->172.20.120.0/24 pref=172.20.120.146 gwy=0.0.0.0
 dev=6(internal)
```

The parts of the routing table entry are:

<b>tab</b>	table number. This will be either 254 (unicast) or 255 (multicast).
<b>vf</b>	virtual domain of the firewall. This is the vdom index number. If vdoms are not enabled, this number will be 0.
<b>type</b>	type of routing connection. Valid values include: 0 - unspecified 1 - unicast 2 - local 3 - broadcast 4 - anycast 5 - multicast 6 - blackhole 7 - unreachable 8 - prohibited

<b>proto</b>	type of installation. This indicates where the route came from. Valid values include: 0 - unspecified 2 - kernel 11 - ZebOS routing module 14 - FortiOS 15 - HA 16 - authentication based 17 - HA1
<b>prio</b>	priority of the route. Lower priorities are preferred.
<b>-&gt;10.11.201.0/24 (-&gt;x.x.x.x/mask)</b>	the IP address and subnet mask of the destination
<b>pref</b>	preferred next hop along this route
<b>gwy</b>	gateway - the address of the gateway this route will use
<b>dev</b>	outgoing interface index. This number is associated with the interface for this route, and if VDOMs are enabled the VDOM will be included here as well. If an interface alias is set for this interface it will also be displayed here.

## Searching the routing table

You can apply a filter to search the routing table and display certain routes only. For example, you can display one or more static routes, connected routes, routes learned through RIP, OSPF, or BGP, and routes associated with the network or gateway that you specify.

If you want to search the routing table by route type and further limit the display according to network or gateway, all of the values that you specify as search criteria must match corresponding values in the same routing table entry in order for that entry to be displayed — an implicit AND condition is applied to all of the search parameters you specify.

For example, if the FortiGate unit is connected to network 172.16.14.0/24 and you want to display all directly connected routes to network 172.16.14.0/24, you must select *Connected* from the *Type* list, type 172.16.14.0/24 in the *Network* field, and then select *Apply Filter* to display the associated routing table entry or entries. Any entry that contains the word “Connected” in its *Type* field and the specified value in the *Gateway* field will be displayed.

In this example, you will apply a filter to search for an entry for static route to 10.10.10.10/24

### To search the FortiGate unit routing table in the web-based manager

1. Go to *Router > Monitor > Routing Monitor*.
2. From the *Type* list, select the type of route to display. In our example, select *Static*.
3. If you want to display routes to a specific network, type the IP address and netmask of the network in the *Networks* field. In our example, enter 10.10.10.10/24.
4. If you want to display routes to a specific gateway, type the IP address of the gateway in the *Gateway* field.

## 5. Select *Apply Filter*.



All of the values that you specify as search criteria must match corresponding values in the same routing table entry in order for that entry to be displayed.

### To search the FortiGate unit routing table in the CLI

```
FGT # get router info routing-table details 10.10.10.10
Routing entry for 10.10.10.10/24
 Known via "static", distance 10, metric 0, best
```

If there are multiple routes that match your filter, they will all be listed, with the best match at the top of the list as indicated by the word best.

## Building the routing table

In the factory default configuration, the FortiGate unit routing table contains a single static default route. You can add routing information to the routing table by defining additional static routes.

It is possible that the routing table is faced with several different routes to the same destination—the IP addresses of the next-hop router specified in those routes or the FortiGate interfaces associated with those routes may vary. In this situation, the “best” route is selected from the table.

The FortiGate unit selects the “best” route for a packet by evaluating the information in the routing table. The “best” route to a destination is typically associated with the shortest distance between the FortiGate unit and the closest gateway, also known as a next-hop router. In some cases, the next best route may be selected if the best route is unavailable.

The FortiGate unit installs the best available routes in the unit’s forwarding table, which is a subset of the unit’s routing table. Packets are forwarded according to the information in the forwarding table.

## Static routing security

Securing the information on your company network is a top priority for network administrators. Security is also required as the routing protocols used are internationally known standards that typically provide little or no inherent security by themselves.

The two reasons for securing your network are the sensitive and proprietary information on your network, and also your external bandwidth. Hackers not only can steal your information, but they can also steal your bandwidth. Routing is a good low level way to secure your network, even before UTM features are applied.

Routing provides security to your network in a number of ways including obscuring internal network addresses with NAT and blackhole routing, using RPF to validate traffic sources, and maintaining an access control list (ACL) to limit access to the network.

This section includes:

- [Network Address Translation \(NAT\)](#)
- [Access Control List \(ACL\)](#)
- [Blackhole Route](#)
- [Reverse path lookup](#)

## Network Address Translation (NAT)

Network address translation (NAT) is a method of changing the address traffic appears to originate from. This practice is used to hide the IP address on company's internal networks, and helps prevent malicious attacks that use those specific addresses.

This is accomplished by the router connected to that local network changing all the IP addresses to its externally connected IP address before sending the traffic out to the other networks, such as the Internet. Incoming traffic uses the established sessions to determine which traffic goes to which internal IP address. This also has the benefit of requiring only the router to be very secure against external attacks, instead of the whole internal network as would be the case without NAT. Securing one computer is much cheaper and easier to maintain.

Configuring NAT on your FortiGate unit includes the following steps.

1. Configure your internal network. For example use the 10.11.101.0 subnet.
2. Connect your internal subnet to an interface on your FortiGate unit. For example use `port1`.
3. Connect your external connection, for example an ISP gateway of 172.20.120.2, to another interface on your Fortigate unit, for example `port2`.
4. Configure security policies to allow traffic between `port1` and `port2` on your FortiGate unit, ensuring that the NAT feature is enabled.

The above steps show that traffic from your internal network will originate on the 10.11.101.0 subnet and pass on to the 172.20.120.0 network. The FortiGate unit moves the traffic to the proper subnet. In doing that, the traffic appears to originate from the FortiGate unit interface on that subnet — it does not appear to originate from where it actually came from.

NAT “hides” the internal network from the external network. This provides security through obscurity. If a hacker tries to directly access your network, they will find the Fortigate unit, but will not know about your internal network. The hacker would have to get past the security-hardened FortiGate unit to gain access to your internal network. NAT will not prevent hacking attempts that piggy back on valid connections between the internal network and the outside world. However other UTM security measures can deal with these attempts.

Another security aspect of NAT is that many programs and services have problems with NAT. Consider if someone on the Internet tries to initiate a chat with someone on the internal network. The outsider only can access the FortiGate unit's external interface unless the security policy allows the traffic through to the internal network. If allowed in, the proper internal user would respond to the chat. However if its not allowed, the request to chat will be refused or time-out. This is accomplished in the security policy by allowing or denying different protocols.

## Access Control List (ACL)

An access control list (ACL) is a table of addresses that have permission to send and receive data over a router's interface or interfaces. The router maintains an ACL, and when traffic comes in on a particular interface it is buffered, while the router looks up in the ACL if that traffic is allowed over that port or not. If it is allowed on that incoming interface, then the next step is to check the ACL for the destination interface. If the traffic passes that check as well the buffered traffic is delivered to its accentuation. If either of those steps fail the ACL check, the traffic is dropped and an error message may be sent to the sender. The ACL ensures that traffic follows expected paths, and any unexpected traffic is not delivered. This stops many network attacks. However, to be effective the ACL must be kept up to date — when employees or computers are removed from the internal network their IP addresses must also be removed from the ACL. For more information on the ACL, see the router chapter of the [FortiGate CLI Reference](#).

## Blackhole Route

A blackhole route is a route that drops all traffic sent to it. It is very much like `/dev/null` in Linux programming.

Blackhole routes are used to dispose of packets instead of responding to suspicious inquiries. This provides added security since the originator will not discover any information from the target network.

Blackhole routes can also limit traffic on a subnet. If some subnet addresses are not in use, traffic to those addresses (traffic which may be valid or malicious) can be directed to a blackhole for added security and to reduce traffic on the subnet.

The loopback interface, a virtual interface that does not forward traffic, was added to enable easier configuration of blackhole routing. Similar to a normal interface, this loopback interface has fewer parameters to configure, and all traffic sent to it stops there. Since it cannot have hardware connection or link status problems, it is always available, making it useful for other dynamic routing roles. Once configured, you can use a loopback interface in security policies, routing, and other places that refer to interfaces. You configure this feature only from the CLI. For more information, see the system chapter of the [FortiGate CLI Reference](#).

## Reverse path lookup

Whenever a packet arrives at one of the FortiGate unit's interfaces, the unit determines whether the packet was received on a legitimate interface by doing a reverse lookup using the source IP address in the packet header. This is also called anti-spoofing. If the FortiGate unit cannot communicate with the computer at the source IP address through the interface on which the packet was received, the FortiGate unit drops the packet as it is likely a hacking attempt.

If the destination address can be matched to a local address (and the local configuration permits delivery), the FortiGate unit delivers the packet to the local network. If the packet is destined for another network, the FortiGate unit forwards the packet to a next-hop router according to a policy route and the information stored in the FortiGate forwarding table.

## Multipath routing and determining the best route

Multipath routing occurs when more than one entry to the same destination is present in the routing table. When multipath routing happens, the FortiGate unit may have several possible destinations for an incoming packet, forcing the FortiGate unit to decide which next-hop is the best one.

It should be noted that some IP addresses will be rejected by routing protocols. These are called Martian addresses. They are typically IP addresses that are invalid and not routable because they have been assigned an address by a misconfigured system, or are spoofed addresses.

Two methods to manually resolve multiple routes to the same destination are to lower the administrative distance of one route or to set the priority of both routes. For the FortiGate unit to select a primary (preferred) route, manually lower the administrative distance associated with one of the possible routes. Setting the priority on the routes is a FortiGate unit feature and may not be supported by non-Fortinet routers.

Administrative distance is based on the expected reliability of a given route. It is determined through a combination of the number of hops from the source and the protocol used. A hop is when traffic moves from one router to the next. More hops from the source means more possible points of failure. The administrative distance can be from 1 to 255, with lower numbers being preferred. A distance of 255 is seen as infinite and will not be installed in the routing table.

Here is an example to illustrate how administrative distance works — if there are two possible routes traffic can take between two destinations with administrative distances of 5 (always up) and 31 (sometimes not available), the traffic will use the route with an administrative distance of 5. If for some reasons the preferred route (admin distance of 5) is not available, the other route will be used as a backup.

Different routing protocols have different default administrative distances. These different administrative distances are based on a number of factors of each protocol such as reliability,

speed, and so on. The default administrative distances for any of these routing protocols are configurable.

**Table 4:** Default administrative distances for routing protocols and connections

Routing protocol	Default administrative distance
Direct physical connection	1
Static	10
EBGP	20
OSPF	110
IS-IS	115
RIP	120
IBGP	200

Another method to determine the best route is to manually change the priority of both routes in question. If the next-hop administrative distances of two routes on the FortiGate unit are equal, it may not be clear which route the packet will take. Manually configuring the priority for each of those routes will make it clear which next-hop will be used in the case of a tie. The priority for a route be set in the CLI, or when editing a specific static route, as described in the next section. Lower priority routes are preferred. Priority is a Fortinet value that may or may not be present in other brands of routers.

All entries in the routing table are associated with an administrative distance. If the routing table contains several entries that point to the same destination (the entries may have different gateways or interface associations), the FortiGate unit compares the administrative distances of those entries first, selects the entries having the lowest distances, and installs them as routes in the FortiGate unit forwarding table. As a result, the FortiGate unit forwarding table contains only those routes having the lowest distances to every possible destination. While only static routing uses administrative distance as its routing metric, other routing protocols such as RIP can use metrics that are similar to administrative distance.

## Route priority

After the FortiGate unit selects static routes for the forwarding table based on their administrative distances, the priority field of those routes determines routing preference. Priority is a Fortinet value that may or may not be present in other brands of routers.

You can configure the priority field through the CLI or the web-based manager. Priority values can range from 0 to 4 294 967 295. The route with the lowest value in the priority field is considered the best route. It is also the primary route.

### To change the priority of a route - web-based manager

1. Go to *Router > Static > Static Routes*.
2. Select the route entry, and select *Edit*.
3. Select *Advanced*.
4. Enter the *Priority* value.
5. Select *OK*.



## To change the priority of a route - CLI

The following command changes the priority to 5 for a route to the address 10.10.10.1 on the port1 interface.

```
config router static
 edit 1
 set device port1
 set gateway 10.10.10.10
 set dst 10.10.10.1
 set priority 5
 end
```

If there are other routes set to priority 10, the route set to priority 5 will be preferred. If there are routes set to priorities less than 5, those other routes will be preferred instead.

In summary, because you can use the CLI to specify which sequence numbers or priority field settings to use when defining static routes, you can prioritize routes to the same destination according to their priority field settings. For a static route to be the preferred route, you must create the route using the `config router static` CLI command and specify a low priority for the route. If two routes have the same administrative distance and the same priority, then they are equal cost multipath (ECMP) routes.

Since this means there is more than one route to the same destination, it can be confusing which route or routes to install and use. However, if you have enabled load balancing with ECMP routes, then different sessions will resolve this problem by using different routes to the same address.

## Troubleshooting static routing

When there are problems with your network that you believe to be static routing related, there are a few basic tools available to locate the problem.

These tools include:

- [Ping](#)
- [Traceroute](#)
- [Examine routing table contents](#)

### Ping

Beyond the basic connectivity information, ping can tell you the amount of packet loss (if any), how long it takes the packet to make the round trip, and the variation in that time from packet to packet.

If there is no packet loss detected, your basic network connectivity is OK.

If there is some packet loss detected, you should investigate:

- possible ECMP, split horizon, network loops
- cabling to ensure no loose connections

If there is total packet loss, you should investigate:

- hardware - ensure cabling is correct, and all equipment between the two locations is accounted for
- addresses and routes - ensure all IP addresses and routing information along the route is configured as expected
- firewalls - ensure all firewalls are set to allow PING to pass through

#### **To ping from a Windows PC**

1. Go to a DOS prompt. Typically you go to *Start > Run*, enter `cmd` and select OK.
2. Enter `ping 10.11.101.100` to ping the default internal interface of the FortiGate unit with four packets.

#### **To ping from an Apple computer**

1. Open the Terminal.
2. Enter `ping 10.11.101.100`.
3. If the ping fails, it will stop after a set number of attempts. If it succeeds, it will continue to ping repeatedly. Press `Control+C` to end the attempt and see gathered data.

#### **To ping from a Linux PC**

1. Go to a command line prompt.
2. Enter `"/bin/etc/ping 10.11.101.101"`.

### **Traceroute**

Where ping will only tell you if it reached its destination and came back successfully, traceroute will show each step of its journey to its destination and how long each step takes. If ping finds an outage between two points, traceroute can be used to locate exactly where the problem is.

#### **To use traceroute on a Windows PC**

1. Go to a DOS prompt. Typically you go to *Start > Run*, enter `cmd` and select OK.
2. Enter `tracert fortinet.com` to trace the route from the PC to the Fortinet website.

#### **To use traceroute from an Apple computer**

1. Open the Terminal.
2. Enter `traceroute fortinet.com`.
3. The terminal will list the number of steps made. Upon reaching the destination, it will list three asterisks per line. Press `Control+C` to end the attempt.

#### **To use traceroute on a Linux PC**

1. Go to a command line prompt.
2. Enter `"/bin/etc/traceroute fortinet.com"`.

The Linux traceroute output is very similar to the MS Windows traceroute output.

### **Examine routing table contents**

The first place to look for information is the routing table.

The routing table is where all the currently used routes are stored for both static and dynamic protocols. If a route is in the routing table, it saves the time and resources of a lookup. If a route isn't used for a while and a new route needs to be added, the oldest least used route is bumped if the routing table is full. This ensures the most recently used routes stay in the table. Note that if your FortiGate unit is in Transparent mode, you are unable to perform this step.

If the FortiGate is running in NAT mode, verify that all desired routes are in the routing table: local subnets, default routes, specific static routes, and dynamic routing protocols.

To check the routing table in the web-based manager, use the Routing Monitor — go to *Router > Monitor > Routing Monitor*. In the CLI, use the command `get router info routing-table all`.

## Static routing tips

When your network goes beyond basic static routing, here are some tips to help you plan and manage your static routing.

### Always configure a default route

The first thing configured on a router on your network should be the default route. And where possible the default routes should point to either one or very few gateways. This makes it easier to locate and correct problems in the network. By comparison, if one router uses a second router as its gateway which uses a fourth for its gateway and so on, one failure in that chain will appear as an outage for all the devices downstream. By using one or very few addresses as gateways, if there is an outage on the network it will either be very localized or network-wide — either is easy to troubleshoot.

### Have an updated network plan

A network plan lists different subnets, user groups, and different servers. Essentially it puts all your resources on the network, and shows how the parts of your network are connected. Keeping your plan updated will also help you troubleshoot problems more quickly when they arise.

A network plan helps your static routing by eliminating potential bottlenecks, and helping troubleshoot any routing problems that come up. Also you can use it to plan for the future and act on any changes to your needs or resources more quickly.

### Plan for expansion

No network remains the same size. At some time, all networks grow. If you take future growth into account, there will be less disruption to your existing network when that growth happens. For example allocating a block of addresses for servers can easily prevent having to re-assign IP addresses to multiple servers due to a new server.

With static routing, if you group parts of your network properly you can easily use network masks to address each part of your network separately. This will reduce the amount of administration required both to maintain the routing, and to troubleshoot any problems.

### Configure as much security as possible

Securing your network through static routing methods is a good low level method to defend both your important information and your network bandwidth.

- Implement NAT to obscure your IP address is an excellent first step.
- Implement black hole routing to hide which IP addresses are in use or not on your local network.
- Configure and use access control list (ACL) to help ensure you know only valid users are using the network.

All three features limit access to the people who should be using your network, and obscure your network information from the outside world and potential hackers.

## Policy routing

Policy routing enables you to redirect traffic away from a static route. This can be useful if you want to route certain types of network traffic differently. You can use incoming traffic's protocol, source address or interface, destination address, or port number to determine where to send the traffic. For example, generally network traffic would go to the router of a subnet, but you might want to direct SMTP or POP3 traffic directly to the mail server on that subnet.

If you have configured the FortiGate unit with routing policies and a packet arrives at the FortiGate unit, the FortiGate unit starts at the top of the Policy Route list and attempts to match the packet with a policy. If a match is found and the policy contains enough information to route the packet (a minimum of the IP address of the next-hop router and the FortiGate interface for forwarding packets to it), the FortiGate unit routes the packet using the information in the policy. If no policy route matches the packet, the FortiGate unit routes the packet using the routing table.



Most policy settings are optional, so a matching policy alone might not provide enough information for forwarding the packet. The FortiGate unit may refer to the routing table in an attempt to match the information in the packet header with a route in the routing table. For example, if the outgoing interface is the only item in the policy, the FortiGate unit looks up the IP address of the next-hop router in the routing table. This situation could happen when the interfaces are dynamic (such as DHCP or PPPoE) and you do not want or are unable to specify the IP address of the next-hop router.

Policy route options define which attributes of an incoming packet cause policy routing to occur. If the attributes of a packet match all the specified conditions, the FortiGate unit routes the packet through the specified interface to the specified gateway.

To view policy routes go to *Router > Static > Policy Routes*.

<b>Create New</b>	Add a policy route. See <a href="#">“Adding a policy route” on page 285</a> .
<b>Delete icon</b>	Delete the selected policy route.
<b>Edit icon</b>	Edit the selected policy route.
<b>Move To icon</b>	Move the selected policy route. Enter the new position and select OK. For more information, see <a href="#">“Moving a policy route” on page 287</a> .
<b>#</b>	The ID numbers of configured route policies. These numbers are sequential unless policies have been moved within the table.
<b>Incoming</b>	The interfaces on which packets subjected to route policies are received.
<b>Outgoing</b>	The interfaces through which policy routed packets are routed.
<b>Source</b>	The IP source addresses and network masks that cause policy routing to occur.
<b>Destination</b>	The IP destination addresses and network masks that cause policy routing to occur.

## Adding a policy route

To add a policy route, go to *Router > Static > Policy Route* and select *Create New*.

<b>Protocol</b>	<p>Enter the protocol number to match. The Internet Protocol Number is found in the IP packet header. <a href="#">RFC 5237</a> describes protocol numbers and you can find a list of the assigned protocol numbers <a href="#">here</a>. The range is from 0 to 255. A value of 0 disables the feature.</p> <p>Commonly used <i>Protocol</i> settings include 6 for TCP sessions, 17 for UDP sessions, 1 for ICMP sessions, 47 for GRE sessions, and 92 for multicast sessions.</p>
<b>Incoming Interface</b>	Select the name of the interface through which incoming packets subjected to the policy are received.
<b>Source Address / Mask</b>	To perform policy routing based on IP source address, type the source address and network mask to match. A value of 0.0.0.0/0.0.0.0 disables the feature.
<b>Destination Address / Mask</b>	To perform policy routing based on the IP destination address of the packet, type the destination address and network mask to match. A value of 0.0.0.0/0.0.0.0 disables the feature.
<b>Destination Ports</b>	<p>To perform policy routing based on the port on which the packet is received, type the same port number in the From and To fields. To apply policy routing to a range of ports, type the starting port number in the From field and the ending port number in the To field. A value of 0 disables this feature.</p> <p>The Destination Ports fields are only used for TCP and UDP protocols. The ports are skipped over for all other protocols.</p>
<b>Type of Service</b>	Use a two digit hexadecimal bit pattern to match the service, or use a two digit hexadecimal bit mask to mask out. For more information, see <a href="#">“Type of Service” on page 286</a> .
<b>Outgoing Interface</b>	Select the name of the interface through which packets affected by the policy will be routed.
<b>Gateway Address</b>	Type the IP address of the next-hop router that the FortiGate unit can access through the specified interface.

### Example policy route

Configure the following policy route to send all FTP traffic received at `port1` out the `port10` interface and to a next hop router at IP address `172.20.120.23`. To route FTP traffic set protocol to 6 (for TCP) and set both of the destination ports to 21, the FTP port.

<b>Protocol</b>	6
<b>Incoming interface</b>	port1
<b>Source address / mask</b>	0.0.0.0/0.0.0.0
<b>Destination address / mask</b>	0.0.0.0/0.0.0.0

<b>Destination Ports</b>	From 21 to 21
<b>Type of Service</b>	bit pattern: 00 (hex) bit mask: 00 (hex)
<b>Outgoing interface</b>	port10
<b>Gateway Address</b>	172.20.120.23

## Type of Service

Type of service (TOS) is an 8-bit field in the IP header that enables you to determine how the IP datagram should be delivered, with such qualities as delay, priority, reliability, and minimum cost.

Each quality helps gateways determine the best way to route datagrams. A router maintains a ToS value for each route in its routing table. The lowest priority TOS is 0, the highest is 7 - when bits 3, 4, and 5 are all set to 1. The router tries to match the TOS of the datagram to the TOS on one of the possible routes to the destination. If there is no match, the datagram is sent over a zero TOS route.

Using increased quality may increase the cost of delivery because better performance may consume limited network resources. For more information, see RFC 791 and RFC 1349.

**Table 5:** The role of each bit in the IP header TOS 8-bit field

<b>bits 0, 1, 2</b>	<b>Precedence</b>	Some networks treat high precedence traffic as more important traffic. Precedence should only be used within a network, and can be used differently in each network. Typically you do not care about these bits.
<b>bit 3</b>	<b>Delay</b>	When set to 1, this bit indicates low delay is a priority. This is useful for such services as VoIP where delays degrade the quality of the sound.
<b>bit 4</b>	<b>Throughput</b>	When set to 1, this bit indicates high throughput is a priority. This is useful for services that require lots of bandwidth such as video conferencing.
<b>bit 5</b>	<b>Reliability</b>	When set to 1, this bit indicates high reliability is a priority. This is useful when a service must always be available such as with DNS servers.
<b>bit 6</b>	<b>Cost</b>	When set to 1, this bit indicates low cost is a priority. Generally there is a higher delivery cost associated with enabling bits 3,4, or 5, and bit 6 indicates to use the lowest cost route.
<b>bit 7</b>	<b>Reserved for future use</b>	Not used at this time.

For example, if you want to assign low delay, and high reliability, say for a VoIP application where delays are unacceptable, you would use a bit pattern of `xxx1x1xx` where an 'x' indicates that bit can be any value. Since all bits are not set, this is a good use for the bit mask; if the mask is set to `0x14`, it will match any TOS packets that are set to low delay and high reliability.

## Moving a policy route

A routing policy is added to the bottom of the routing table when it is created. If you prefer to use one policy over another, you may want to move it to a different location in the routing policy table.

The option to use one of two routes happens when both routes are a match, for example 172.20.0.0/255.255.0.0 and 172.20.120.0/255.255.255.0. If both of these routes are in the policy table, both can match a route to 172.20.120.112 but you consider the second one as a better match. In that case the best match route should be positioned before the other route in the policy table.

To change the position of a policy route in the table, go to *Router > Static > Policy Routes* and select *Move To* for the policy route you want to move.

<b>Before/After</b>	Select Before to place the selected Policy Route before the indicated route. Select After to place it following the indicated route.
<b>Policy route ID</b>	Enter the Policy route ID of the route in the Policy route table to move the selected route before or after.

## Transparent mode static routing

FortiOS operating modes allow you to change the configuration of your FortiGate unit depending on the role it needs to fill in your network.

NAT/Route operating mode is the standard mode where all interfaces are accessed individually, and traffic can be routed between ports to travel from one network to another.

In transparent operating mode, all physical interfaces act like one interface. The FortiGate unit essentially becomes a bridge — traffic coming in over any interface is broadcast back out over all the interfaces on the FortiGate unit.

In transparent mode, there is no entry for routing at the main level of the menu on the web-based manager display as there is in NAT/Route mode. Routing is instead accessed through the network menu option.

To view the routing table in transparent mode, go to *System > Network > Routing Table*.

When viewing or creating a static route entry in transparent mode there are only three fields available.

---

<b>Destination IP/Mask</b>	The destination of the traffic being routed. The first entry is attempted first for a match, then the next, and so on until a match is found or the last entry is reached. If no match is found, the traffic will not be routed.  Use 0.0.0.0 to match all traffic destinations. This is the default route.
<b>Gateway</b>	Specifies the next hop for the traffic. Generally the gateway is the address of a router on the edge of your network.
<b>Priority</b>	The priority is used if there is more than one match for a route. This allows multiple routes to be used, with one preferred. If the preferred route is unavailable the other routes can be used instead.  Valid range of priority can be from 0 to 4 294 967 295.  If more than one route matches and they have the same priority it becomes an ECMP situation and traffic is shared among those routes. See <a href="#">“Route priority” on page 280</a> .

---

When configuring routing on a FortiGate unit in transparent mode, remember that all interfaces must be connected to the same subnet. That means all traffic will be coming from and leaving on the same subnet. This is important because it limits your static routing options to only the gateways attached to this subnet. For example, if you only have one router connecting your network to the Internet then all static routing on the FortiGate unit will use that gateway. For this reason static routing on FortiGate units in transparent mode may be a bit different, but it is not as complex as routing in NAT/Route mode.

## Static routing example

This is an example of a typical small network configuration that uses only static routing.

This network is in a dentist office that includes a number of dentists, assistants, and office staff. The size of the office is not expected to grow significantly in the near future, and the network usage is very stable—there are no new applications being added to the network.

The users on the network are:

- admin staff - access to local patient records, and go online for billing purposes
- dentists - access and update local patient records, research online from desk
- assistants - access and update local patient records in exam rooms

The distinction here is mainly that only the admin staff and dentist's office proper need access to the internet—all the other traffic is local and doesn't need to need to leave the local network. Routing is only required for the outbound traffic, and the computers that have valid outbound traffic.



Only configuring routing on computers that need it will act as an additional layer of security by helping prevent malicious traffic from leaving the network.

---

This section includes the following topics:

- [Network layout and assumptions](#)
- [General configuration steps](#)
- [Configure FortiGate unit](#)
- [Configure Admin PC and Dentist PCs](#)
- [Testing network configuration](#)

### Network layout and assumptions


The computers on the network are admin staff computers, dentist office computers, and dental exam room computers. While there are other devices on the local network such as printers, they do not need Internet access or any routing.

This networked office equipment includes 1 admin staff PC, 3 dentist PCs, and 5 exam room PCs. There are also a network printer, and a router on the network as well.



Assumptions about these computers, and network include:

- the FortiGate unit is a model with interfaces labeled port1 and port2
- the FortiGate unit has been installed and is configured in NAT/Route mode
- VDOMs are not enabled
- the computers on the network are running MS Windows software
- any hubs required in the network are not shown in the network diagram
- the network administrator has access to the ISP IP addresses, and is the super\_admin administrator on the FortiGate unit



**Table 6:** Static routing example device names, IP addresses, and level of access

Device Name(s)	IP address	Need external access?
<b>Router</b>	192.168.10.1	YES
<b>Admin</b>	192.168.10.11	YES
<b>Dentist1-3</b>	192.168.10.21-23	YES
<b>Exam1-5</b>	192.168.10.31-35	NO
<b>Printer</b>	192.168.10.41	NO

## General configuration steps

The steps to configuring routing on this network are:

1. [Get your ISP information such as DNS, gateway, etc.](#)
2. [Configure FortiGate unit](#)
3. [Configure Admin PC and Dentist PCs](#)
4. [Testing network configuration](#)

## Get your ISP information such as DNS, gateway, etc.

Your local network connects to the Internet through your Internet Service Provider (ISP). They have IP addresses that you need to configure your network and routing.

The addresses needed for routing are your assigned IP address, DNS servers, and the gateway.

## Configure FortiGate unit

The FortiGate unit will have two interfaces in use—one connected to the internal network and one connected to the external network. Port1 will be the internal interface, and port2 will be the external interface.

To configure the FortiGate unit:

1. [Configure the internal interface \(port1\)](#)
2. [Configure the external interface \(port2\)](#)
3. [Configure networking information](#)
4. [Configure basic security policies](#)
5. [Configure static routing](#)

### Configure the internal interface (port1)

#### To configure the internal interface (port1) - web based manager

1. Go to *System > Network > port1* and select *Edit*.
2. Enter the following:

<b>Addressing Mode</b>	Manual
<b>IP/Netmask</b>	172.100.1.1/255.255.255.0
<b>Administrative Access</b>	HTTPS, PING, TELNET
<b>Description</b>	Internal network

#### To configure the internal interface (port1) - CLI

```
config system interface
 edit port1
 set IP 192.168.10.1 255.255.255.0
 set allowaccess https ping telnet
 set description "internal network"
 end
end
```

### Configure the external interface (port2)

The external interface connects to your ISP's network. You need to know the IP addresses in their network that you should connect to. Use their addresses when you get them, however for this example we will assume the address your ISP gave you is 172.100.20.20 will connect to the gateway at 172.100.20.5 on their network, and their DNS servers are 172.11.22.33 and 172.11.22.34.

#### To configure the internal interface (port2) - web based manager

1. Go to *System > Network > port2* and select *Edit*.

2. Enter the following:

<b>Addressing Mode</b>	Manual
<b>IP/Netmask</b>	172.100.20.20/255.255.255.0
<b>Administrative Access</b>	HTTPS, PING, TELNET
<b>Description</b>	Internal network

### To configure the internal interface (port2) - CLI

```
configure system interface
 edit port2
 set IP 172.100.20.20 255.255.255.0
 set allowaccess https ping telnet
 set description "internal network"
 end
end
```

### Configure networking information

Networking information includes the gateway, and DNS servers. Your FortiGate unit requires a connection to the Internet for antivirus and other periodic updates.

### To configure networking information - web-based manager

1. Go to *System > Network > Options*.
2. Enter the primary and secondary DNS addresses.
3. Select *Apply*.

### To configure networking information - CLI

```
config system global
 set dns_1 172.11.22.33
 set dns_2 172.11.22.34
end
```

### Configure basic security policies

For traffic to flow between the internal and external ports in both directions, two security policies are required as a minimum. More can be used to further limit or direct traffic as needed, but will not be included here.

Before configuring the security policies, a firewall address group is configured for the PCs that are allowed Internet access. This prevents PC without Internet privileges from accessing the Internet.

The security policy assumptions are:

- Only the basic networking services have been listed as allowed for added security. Others can easily be added as the users require them.
- In this example to keep things simple, both incoming and outgoing security policies are the same. In a real network there are applications that are allowed out but not in, and vice versa.
- Endpoint control has been enabled to ensure that all computers on the local network are running FortiClient and those installs are up to date. This feature ensures added security on your local network without the need for the network administrator to continually bother users

to update their software. The FortiGate unit can store an up to date copy of the FortiClient software and offer a URL to it for users to install it if they need to.

**To configure security policies - web-based manager**

1. Go to *Firewall Objects > Address*.
2. Create a new Firewall Address entry for each of:

PC Name	IP Address	Interface
Admin	192.168.10.11	port1
Dentist1	192.168.10.21	port1
Dentist2	192.168.10.22	port1
Dentist3	192.168.10.23	port1

3. Go to *Firewall Objects > Address > Groups*.
4. Select *Create New*.
5. Name the group *Internet\_PCs*.
6. Select *Admin, Dentist1, Dentist2, and Dentist3* to be members of the group.
7. Select *OK*.
8. Go to *Policy > Policy > Policy*.
9. Select *Create New*.
10. Enter the following: *DH - port2(external) -> port1(internal)*

<b>Source Interface/Zone</b>	port2
<b>Source Address</b>	all
<b>Destination Interface/Zone</b>	port1
<b>Destination Address</b>	Internet_PCs
<b>Schedule</b>	always
<b>Service</b>	Multiple. Select <i>DHCP, DNS, FTP, HTTP, HTTPS, NTP, POP3, SMTP, SSH</i> .
<b>Action</b>	ACCEPT
<b>Enable Endpoint Control Check</b>	Enabled Enforce FortiClient AV Up-to-date
<b>Log Allowed Traffic</b>	Enabled

11. Select *OK*.
12. Select *Create New*.
13. Enter the following:

<b>Source Interface/Zone</b>	port1
<b>Source Address</b>	Internet_PCs
<b>Destination Interface/Zone</b>	port2
<b>Destination Address</b>	all
<b>Schedule</b>	always
<b>Service</b>	Multiple. Select <i>DHCP, DNS, FTP, HTTP, HTTPS, NTP, POP3, SMTP, SSH.</i>
<b>Action</b>	ACCEPT
<b>Enable Endpoint Control Check</b>	Enabled Enforce FortiClient AV Up-to-date
<b>Log Allowed Traffic</b>	Enabled

14. Select OK.

#### To configure security policies - CLI

```

config firewall address
 edit "Admin"
 set associated-interface "port1"
 set subnet 192.168.10.11 255.255.255.255
 next
 edit "Dentist1"
 set associated-interface "port1"
 set subnet 192.168.10.21 255.255.255.255
 next
 edit "Dentist2"
 set associated-interface "port1"
 set subnet 192.168.10.22 255.255.255.255
 next
 edit "Dentist3"
 set associated-interface "port1"
 set subnet 192.168.10.23 255.255.255.255
 end
config firewall addrgrp
 edit Internet_PCs
 set member Admin Dentist1 Dentist2 Dentist3
 end
config firewall policy
 edit 1
 set srcintf port1
 set dstintf port2
 set srcaddr Internet_PCs

```

```

set dstaddr all
set action accept
set schedule always
set service "DHCP" "DNS" "FTP" "HTTP" "HTTPS" "NTP" "POP3"
 "SMTP" "SSH"
set logtraffic enable
set endpoint-check enable
set label "Section2"
set endpoint-restrict-check no-av db-outdated
next
edit 2
set srcintf port2
set dstintf port1
set srcaddr all
set dstaddr Internet_PCs
set action accept
set schedule always
set service "DHCP" "DNS" "FTP" "HTTP" "HTTPS" "NTP" "POP3"
 "SMTP" "SSH"
set logtraffic enable
set endpoint-check enable
set label "Section2"
set endpoint-restrict-check no-av db-outdated
end
end

```

## Configure static routing

With the rest of the FortiGate unit configured, static routing is the last step before moving on to the rest of the local network. All traffic on the local network will be routed according to this static routing entry.

### To configure Fortinet unit static routing - web-based manager

1. Go to *Routing > Static*.
2. Select *Edit* for the top route on the page.
3. Enter the following:

<b>Destination IP/Mask</b>	172.100.20.5
<b>Device</b>	port2
<b>Gateway</b>	172.100.20.5
<b>Distance</b>	10

4. Select *OK*.

## To configure Fortinet unit static routing - CLI

```
configure routing static
 edit 1
 set gateway 172.100.20.5
 set distance 10
 set device port2
 set dst 0.0.0.0
 end
end
```

## Configure Admin PC and Dentist PCs

With the router configured, next we need to configure the computers that need Internet access. These computers need routing to be configured on them. As the other computers do not require routing, they are not included here.

The procedure to configure these computers is the same. Repeat the following procedure for the corresponding PCs.



The Windows CLI procedure does not configure the DNS entries. It just adds the static routes.

---

## To configure routing and DNS on Admin and Dentist PCs - Windows GUI

1. On PC, select *Start > Control Panel > Network Connections*.
2. Right click on the network connection to your local network that has a status of Connected, and select *Properties*.
3. Under the *General* tab, from the list select *TCP/IP*, and *Properties*.
4. Under *Gateway*, enter the FortiGate unit address (192.168.10.1).
5. Enter the primary and secondary DNS server addresses from your ISP (172.11.22.33 and 172.11.22.34).
6. Select *OK*.

## To configure routing on Admin and Dentist PCs - Windows CLI

1. On PC, select *Start > Run*, enter "cmd", and select *OK*.
2. At the command prompt, type

```
route ADD 0.0.0.0 MASK 0.0.0.0 172.100.20.5 METRIC 10
route ADD 192.168.10.0 MASK 255.255.255.0 192.168.10.1 METRIC 5
```
3. Confirm these routes have been added. Type:

```
route PRINT
```

If you do not see the two routes you added, try adding them again paying attention to avoid spelling mistakes.
4. Test that you can communicate with other computers on the local network, and with the Internet. If there are no other computers on the local network, connect to the FortiGate unit.

## Configure other PCs on the local network

The PCs on the local network without Internet access (the exam room PCs) can be configured now.

As this step does not require any routing, details have not been included.

## Testing network configuration

There are three tests to run on the network to ensure proper connectivity.

- To test that PCs on the local network can communicate
- Test that Internet\_PCs on the local network can access the Internet.
- Test that non-Internet\_PCs can not access the Internet.

### To test that PCs on the local network can communicate

1. Select any two PCs on the local network, such as Exam4 and Dentist3.
2. On the Exam4 PC, at the command prompt enter `ping 192.168.10.23`.

The output from this command should appear similar to the following.

```
Pinging 192.168.10.23 with 32 bytes of data:
```

```
Reply from 192.168.10.23: bytes=32 time<1m TTL=255
Reply from 192.168.10.23: bytes=32 time<1m TTL=255
Reply from 192.168.10.23: bytes=32 time<1m TTL=255
```

3. At the command prompt enter `exit` to close the window.
4. On the Dentist3 PC, at the command prompt enter `ping 192.168.10.34`.

The output from this command should appear similar to the following.

```
Pinging 192.168.10.34 with 32 bytes of data:
```

```
Reply from 192.168.10.34: bytes=32 time<1m TTL=255
Reply from 192.168.10.34: bytes=32 time<1m TTL=255
Reply from 192.168.10.34: bytes=32 time<1m TTL=255
```

5. At the command prompt enter `exit` to close the window.
6. Repeat these steps for all PCs on the local network.

If the output does not appear similar to above, there is a problem with the network configuration between these two PCs.

### To test that Internet\_PCs can reach the Internet

The easiest way to access the internet is with an Internet browser. However, if that doesn't work its best to do a traceroute to see at what point the problem is. This can help determine if it is a networking problem such as cabling, or if its an access problem such as this PC not having Internet access.

1. Select any PC on the local network that is supposed to have Internet access, such as Admin.
2. On the Admin PC, open an Internet browser and attempt to access a website on the Internet such as <http://www.fortinet.com>.

If this is successful, this PC has Internet access.



3. If step 2 was not successful, at the command prompt on the PC enter `tracert 22.11.22.33`.

The output from this command should appear similar to:

Pinging 22.11.22.33 with 32 bytes of data:

```
Reply from 22.11.22.33: bytes=32 time<1m TTL=255
```

```
Reply from 22.11.22.33: bytes=32 time<1m TTL=255
```

```
Reply from 22.11.22.33: bytes=32 time<1m TTL=255
```

## Advanced static example: ECMP failover and load balancing

Equal Cost Multi-Path (ECMP) load balancing and failover are methods that extend basic static routing. They allow you to use your network bandwidth more effectively and with less down time than if you used basic static routing alone.

The concepts in this section include:

- [Equal-Cost Multi-Path \(ECMP\)](#)
- [Configuring interface status detection for gateway load balancing](#)
- [Configuring spillover or usage-based ECMP](#)
- [Configuring weighted static route load balancing](#)

### Equal-Cost Multi-Path (ECMP)

FortiOS uses equal-cost multi-path (ECMP) to distribute traffic to the same destination such as the Internet or another network. Using ECMP you can add multiple routes to the destination and give each of those routes the same distance and priority.



If multiple routes to the same destination have the same priority but different distances, the route with the lowest distance is used. If multiple routes to the same destination have the same distance but different priorities, the route with the lowest priority is used. Distance takes precedence over priority. If multiple routes to the same destination have different distances and different priorities, the route with the lowest distance is always used even if it has the highest priority.

If more than one ECMP route is available, you can configure how the FortiGate unit selects the route to be used for a communication session. If only one ECMP route is available (for example, because an interface cannot process traffic because interface status detection does not receive a reply from the configured server) then all traffic uses this route.

Previous versions of FortiOS provided source IP-based load balancing for ECMP routes, but now FortiOS includes three configuration options for ECMP route failover and load balancing:

---

**Source based  
(also called source IP  
based)**

The FortiGate unit load balances sessions among ECMP routes based on the source IP address of the sessions to be load balanced. This is the default load balancing method. No configuration changes are required to support source IP load balancing.

---

---

**Weighted Load Balance (also called weight-based)**

The FortiGate unit load balances sessions among ECMP routes based on weights added to ECMP routes. More traffic is directed to routes with higher weights. After selecting weight-based you must add weights to static routes.

---

**Spillover (also called usage-based)**

The FortiGate unit distributes sessions among ECMP routes based on how busy the FortiGate interfaces added to the routes are.

After selecting spill-over you add route Spillover Thresholds to interfaces added to ECMP routes. The FortiGate unit sends all ECMP-routed sessions to the lowest numbered interface until the bandwidth being processed by this interface reaches its spillover threshold. The FortiGate unit then spills additional sessions over to the next lowest numbered interface.

The Spillover Thresholds range is 0-2097000 KBps.

---

You can configure only one of these ECMP route failover and load balancing methods in a single VDOM. If your FortiGate unit is configured for multiple VDOM operation, each VDOM can have its own ECMP route failover and load balancing configuration.

**To configure the ECMP load balancing method from the web-based manager**

1. Go to *Router > Static > Settings*.
2. Set *ECMP Load Balance Method* to *Source IP based*, *Weighted Load Balance*, or *Spillover*.

**To configure the ECMP load balancing method from the CLI**

For example, to set the load balancing method to usage-based, enter the following:

```
config system settings
 set v4-ecmp-mode usage-based
end
```

## ECMP routing of simultaneous sessions to the same destination IP address

When the FortiGate unit selects an ECMP route for a session, a route cache is created that matches the route with the destination IP address of the session. All new sessions to the same destination IP address use the same route until the route is flushed from the cache. Routes are flushed from the cache after a period of time when no new sessions to the destination IP address are received.


The route cache improves FortiGate unit routing performance by reducing how often the FortiGate unit looks up routes in the routing table.

If the FortiGate unit receives a large number of sessions with the same destination IP address, because all of these sessions will be processed by the same route, it may appear that sessions are not distributed according to the ECMP route failover and load balancing configuration.

## Configuring interface status detection for gateway load balancing

Interface status detection is used for ECMP route failover and load balancing. Interface status detection consists of the unit confirming that packets sent from an interface result in a response from a server. You can use up to three different protocols to confirm that an interface can connect to the server. Usually the server is the next-hop router that leads to an external network or the Internet. Interface status detection sends a packet using the configured protocols. If a response is received from the server, the unit assumes the interface can connect to the network. If a response is not received, the unit assumes that the interface cannot connect to the network.

Since it is possible that a response may not be received, even if the server and the network are operating normally, the dead gateway detection configuration controls the time interval between testing the connection to the server and the number of times the test can fail before the unit assumes that the interface cannot connect to the server.

	<p>As long as the unit receives responses for at least one of the protocols that you select, the unit assumes the server is operating and can forward packets. Responding to more than one protocol does not enhance the status of the server or interface.</p>
-----------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

**To configure gateway failover detection for an interface**

1. Go to *Router > Static > Settings*.
2. Under *Dead Gateway Detection*, select *Create New*.
3. Enter the following information:

<b>Interface</b>	Select the interface to test.
<b>Gateway IP</b>	Enter the IP address of the gateway.
<b>Ping Server</b>	Enter the IP address of the server to test.
<b>Detect Protocol</b>	Select one of the following protocols.
<b>ICMP Ping</b>	Use standard ICMP ping to confirm that the server is responding. Ping confirms that the server can respond to an ICMP ping request.
<b>TCP Echo</b>	<p>Use TCP echo to confirm that the server is responding. Select this option if the server is configured to provide TCP echo services. In some cases a server may be configured to reply to TCP echo requests but not to reply to ICMP pings.</p> <p>TCP echo uses TCP packets on port number 7 to send a text string to the server and expect an echo reply back from the server. The echo reply just echoes back the same text to confirm that the server can respond to TCP requests.</p> <p>FortiGate units do not recognize RST (reset) packets from TCP Echo servers as normal TCP echo replies. If the unit receives an RST response to a TCP echo request, the unit assumes the server is unreachable.</p>
<b>UDP Echo</b>	<p>Use UDP echo to detect the server. Select this option if the server is configured to provide UDP echo services. In some cases a server may be configured to reply to UDP echo requests but not to reply ICMP pings.</p> <p>UDP echo uses UDP packets on port number 7 to send a text string to the server and expects an echo reply from the server. The echo reply just echoes back the same text to confirm that the server can respond to UDP requests.</p>
<b>Ping Interval</b>	Enter the interval between pings, in seconds.
<b>Failover Threshold</b>	Enter the number of times the test can fail before the unit assumes that the interface cannot connect to the server.
<b>HA Priority</b>	Set the HA priority, if configuring an HA cluster.

4. Select *OK*.

#### To configure gateway failover detection for an interface - CLI

```
config router gwdetect
 edit port1
 set protocol ping
 set server 10.10.10.1
 set interval 5
 set failtime 5
 end
```

## Configuring spillover or usage-based ECMP

Spill-over or usage-based ECMP routes new sessions to interfaces that have not reached a configured bandwidth limit (called the *Spillover Threshold* or a route-spillover threshold). To configure spill-over or usage-based ECMP routing, you enable spill-over ECMP, add ECMP routes, and add a *Spillover Threshold* to the interfaces used by the ECMP routes. Set the *Spillover Thresholds* to limit the amount of bandwidth processed by each interface. The range is 0 to 2 097 000 Kbps. The threshold counts only outgoing traffic.

With spill-over ECMP routing configured, the FortiGate unit routes new sessions to an interface used by an ECMP route until that interface reaches its *Spillover Threshold*. Then, when the threshold of that interface is reached, new sessions are routed to one of the other interfaces used by the ECMP routes.

#### To add Spillover Thresholds to interfaces - web-based manager

Use the following steps to enable usage based ECMP routing, add Spillover Thresholds to FortiGate interfaces port3 and port4, and then to configure ECMP routes with device set to port3 and port4.

1. Go to *Router > Static > Settings*.
2. Set *ECMP Load Balance Method* to *Spillover*.
3. Go to *Router > Static > Static Routes*.
4. Add ECMP routes for port3 and port4.

<b>Destination IP/Mask</b>	192.168.20.0/24
<b>Device</b>	port3
<b>Gateway</b>	172.20.130.3
<b>Advanced</b>	
<b>Distance</b>	10

<b>Destination IP/Mask</b>	192.168.20.0/24
<b>Device</b>	port4
<b>Gateway</b>	172.20.140.4
<b>Advanced</b>	
<b>Distance</b>	10

5. Go to *System > Network > Settings*.
6. Edit port3 and port4 and add the following spillover-thresholds:

<b>Interface</b>	port3
<b>Spillover Threshold</b>	100

<b>Interface</b>	port4
<b>Spillover Threshold</b>	200

### To add Spillover Thresholds to interfaces - CLI

```

config system settings
 set v4-ecmp-mode usage-based
end
config router static
 edit 1
 set device port3
 set dst 192.168.20.0 255.255.255.0
 set gateway 172.20.130.3
 next
 edit 2
 set device port4
 set dst 192.168.20.0 255.255.255.0
 set gateway 172.20.140.4
 end
config system interface
 edit port3
 set spillover-threshold 100
 next
 edit port4
 set spillover-threshold 200
 end

```

### Detailed description of how spill-over ECMP selects routes

When you add ECMP routes they are added to the routing table in the order displayed by the routing monitor or by the `get router info routing-table static` command. This order is independent of the configured bandwidth limit.

The FortiGate unit selects an ECMP route for a new session by finding the first route in the routing table that sends the session out a FortiGate unit interface that is not processing more traffic than its configured route spill-over limit.



A new session to a destination IP address that already has an entry in the routing cache is routed using the route already added to the cache for that destination address. See [“ECMP routing of simultaneous sessions to the same destination IP address”](#) on page 298.

For example, consider a FortiGate unit with interfaces port3 and port4 both connected to the Internet through different ISPs. ECMP routing is set to usage-based and route spillover for to 100 Kbps for port3 and 200 Kbps for port4. Two ECMP default routes are added, one for port3 and one for port4.

If the route to port3 is higher in the routing table than the route to port4, the FortiGate unit sends all default route sessions out port3 until port3 is processing 100Kbps of data. When port3 reaches its configured bandwidth limit, the FortiGate unit sends all default route sessions out

port4. When the bandwidth usage of port3 falls below 100Kbps, the FortiGate again sends all default route sessions out port3.

New sessions with destination IP addresses that are already in the routing cache; however, use the cached routes. This means that even if port3 is exceeding its bandwidth limit, new sessions can continue to be sent out port3 if their destination addresses are already in the routing cache. As a result, new sessions are sent out port4 only if port3 exceeds its bandwidth limit and if the routing cache does not contain a route for the destination IP address of the new session.

Also, the switch over to port4 does not occur as soon as port3 exceeds its bandwidth limit. Bandwidth usage has to exceed the limit for a period of time before the switch over takes place. If port3 bandwidth usage drops below the bandwidth limit during this time period, sessions are not switched over to port4. This delay reduces route flapping.

FortiGate usage-based ECMP routing is not actually load balancing, since routes are not distributed evenly among FortiGate interfaces. Depending on traffic volumes, most traffic would usually be processed by the first interface with only spillover traffic being processed by other interfaces.

If you are configuring usage-based ECMP in most cases you should add spillover thresholds to all of the interfaces with ECMP routes. The default spillover threshold is 0 which means no bandwidth limiting. If any interface has a spillover threshold of 0, no sessions will be routed to interfaces lower in the list unless the interface goes down or is disconnected. An interface can go down if *Detect interface status for Gateway Load Balancing* does not receive a response from the configured server.

### Determining if an interface has exceeded its Spillover Threshold

You can use the `diagnose netlink dstmac list` CLI command to determine if an interface is exceeding its Spillover Threshold. If the command displays `over_bps=1` the interface is exceeding its threshold. If `over_bps=0` the interface has not exceeded its threshold.

```
dev=Wifi mac=00:00:00:00:00:00 src-vis-os src-vis-host src-vis-user
rx_tcp_mss=0 tx_tcp_mss=0 overspill-threshold=0 bytes=0 over_bps=0
sampler_rate=0
```

## Configuring weighted static route load balancing

Configure weighted load balancing to control how the FortiGate unit distributes sessions among ECMP routes by adding weights for each route. Add higher weights to routes that you want to load balance more sessions to.

With the ECMP load balancing method set to weighted, the FortiGate unit distributes sessions with different destination IPs by generating a random value to determine the route to select. The probability of selecting one route over another is based on the weight value of each route. Routes with higher weights are more likely to be selected.

Large numbers of sessions are evenly distributed among ECMP routes according to the route weight values. If all weights are the same, sessions are distributed evenly. The distribution of a small number of sessions; however, may not be even. For example, its possible that if there are two ECMP routes with the same weight; two sessions to different IP addresses could use the same route. On the other hand, 10,000 sessions with different destination IPs should be load balanced evenly between two routes with equal rates. The distribution could be 5000:5000 or 50001:4999. Also, 10 000 sessions with different destination IP addresses should be load balanced as 3333:6667 if the weights for the two routes are 100 and 200.

Weights only affect how routes are selected for sessions to new destination IP addresses. New sessions to IP addresses already in the routing cache are routed using the route for the session already in the cache. So in practice sessions will not always be distributed according to the routing weight distribution.

**To add weights to static routes from the web-based manager**

1. Go to *Router > Static > Settings*.
2. Set *ECMP Load Balance Method* to *Weighted Load Balance*.
3. Go to *Router > Static > Static Routes*.
4. If needed, add new static routes, for example:

<b>Destination IP/Mask</b>	192.168.20.0/24
<b>Device</b>	port1
<b>Gateway</b>	172.20.110.1
<b>Distance</b>	10

<b>Destination IP/Mask</b>	192.168.20.0/24
<b>Device</b>	port2
<b>Gateway</b>	172.20.120.2
<b>Distance</b>	10

5. Go to *Router > Static > Settings*.
6. Select a number next to an interface name, and choose *Edit* to change it.

For example, set the weight of port1 to 100 and the weight of port2 to 200.

# Dynamic Routing Overview

This section provides an overview of dynamic routing, and how it compares to static routing. For details on various dynamic routing protocols, see the following chapters for detailed information.

The following topics are included in this section:

- [What is dynamic routing?](#)
- [Comparison of dynamic routing protocols](#)
- [Choosing a routing protocol](#)
- [Dynamic routing terminology](#)
- [IPv6 in dynamic routing](#)

## What is dynamic routing?

Dynamic routing uses a dynamic routing protocol to automatically select the best route to put into the routing table. So instead of manually entering static routes in the routing table, dynamic routing automatically receives routing updates, and dynamically decides which routes are best to go into the routing table. It's this intelligent and hands-off approach that makes dynamic routing so useful.

Dynamic routing protocols vary in many ways and this is reflected in the various administrative distances assigned to routes learned from dynamic routing. These variations take into account differences in reliability, speed of convergence, and other similar factors. For more information on these administrative distances, see [“Multipath routing and determining the best route” on page 279](#).

This section includes:

- [Comparing static and dynamic routing](#)
- [Dynamic routing protocols](#)
- [Minimum configuration for dynamic routing](#)

## Comparing static and dynamic routing

A common term used to describe dynamic routing is convergence. Convergence is the ability to work around network problems and outages — for the routing to come together despite obstacles. For example, if the main router between two end points goes down, convergence is the ability to find a way around that failed router and reach the destination. Static routing has zero convergence beyond trying the next route in its limited local routing table — if a network administrator doesn't fix a routing problem manually, it may never be fixed, resulting in a downed network. Dynamic routing solves this problem by involving routers along the route in the decision-making about the optimal route, and using the routing tables of these routers for potential routes around the outage. In general, dynamic routing has better scalability, robustness, and convergence. However, the cost of these added benefits include more



complexity and some overhead: the routing protocol uses some bandwidth for its own administration.

**Table 7:** Comparing static and dynamic routing

Feature	Static Routing	Dynamic Routing
<b>Hardware support</b>	Supported by all routing hardware	May require special, more expensive routers
<b>Router Memory Required</b>	Minimal	Can require considerable memory for larger tables
<b>Complexity</b>	Simple	Complex
<b>Overhead</b>	None	Varying amounts of bandwidth used for routing protocol updates
<b>Scalability</b>	Limited to small networks	Very scalable, better for larger networks
<b>Robustness</b>	None - if a route fails it has to be fixed manually	Robust - traffic routed around failures automatically
<b>Convergence</b>	None	Varies from good to excellent

## Dynamic routing protocols

A dynamic routing protocol is an agreed-on method of routing that the sender, receiver, and all routers along the path (route) support. Typically the routing protocol involves a process running on all computers and routers along that route to enable each router to handle routes in the same way as the others. The routing protocol determines how the routing tables are populated along that route, how the data is formatted for transmission, and what information about a route is included with that route. For example RIP, and BGP use distance vector algorithms, where OSPF uses a shortest path first algorithm. Each routing protocol has different strengths and weaknesses — one protocol may have fast convergence, while another may be very reliable, and a third is very popular for certain businesses like Internet Service Providers (ISPs).

Dynamic routing protocols are different from each other in a number of ways, such as:

- [Classful versus classless routing protocols](#)
- [Interior versus exterior routing protocols](#)
- [Distance vector versus link-state protocols](#)

### Classful versus classless routing protocols

Classful or classless routing refers to how the routing protocol handles the IP addresses. In classful addresses there is the specific address, and the host address of the server that address is connected to. Classless addresses use a combination of IP address and netmask.

Classless Inter-Domain Routing (CIDR) was introduced in 1993 (originally with RFC 1519 and most recently with RFC 4632) to keep routing tables from getting too large. With Classful routing, each IP address requires its own entry in the routing table. With Classless routing, a series of addresses can be combined into one entry potentially saving vast amounts of space in routing tables.

Current routing protocols that support classless routing out of necessity include RIPv2, BGP, IS-IS, and OSPF. Older protocols such as RIPv1 do not support CIDR addresses.

## Interior versus exterior routing protocols

The names *interior* and *exterior* are very descriptive. Interior routing protocols are designed for use within a contained network of limited size, whereas exterior routing protocols are designed to link multiple networks together. They can be used in combination in order to simplify network administration. For example, a network can be built with only border routers of a network running the exterior routing protocol, while all the routers on the network run the interior protocol, which prevents them from connecting outside the network without passing through the border. Exterior routers in such a configuration must have both exterior and interior protocols, to communicate with the interior routers and outside the network.

Nearly all routing protocols are interior routing protocols. Only BGP is commonly used as an exterior routing protocol.

You may see interior gateway protocol (IGP) used to refer to interior routing protocols, and exterior gateway protocol (EGP) used to refer to exterior routing protocols.

## Distance vector versus link-state protocols

Every routing protocol determines the best route between two addresses using a different method. However, there are two main algorithms for determining the best route — Distance vector and Link-state.

### Distance vector protocols

In distance vector protocols, routers are told about remote networks through neighboring routers. The distance part refers to the number of hops to the destination, and in more advanced routing protocols these hops can be weighted by factors such as available bandwidth and delay. The vector part determines which router is the next step along the path for this route. This information is passed along from neighboring routers with routing update packets that keep the routing tables up to date. Using this method, an outage along a route is reported back along to the start of that route, ideally before the outage is encountered.

On distance vector protocols, RFC 1058 which defines RIP v1 states the following:

Distance vector algorithms are based on the exchange of only a small amount of information. Each entity (gateway or host) that participates in the routing protocol is assumed to keep information about all of the destinations within the system. Generally, information about all entities connected to one network is summarized by a single entry, which describes the route to all destinations on that network.

There are four main weaknesses inherent in the distance vector method. Firstly, the routing information is not discovered by the router itself, but is instead reported information that must be relied on to be accurate and up-to-date. The second weakness is that it can take a while for the information to make its way to all the routers who need the information — in other words it can have slow convergence. The third weakness is the amount of overhead involved in passing these updates all the time. The number of updates between routers in a larger network can significantly reduce the available bandwidth. The fourth weakness is that distance vector protocols can end up with routing-loops. Routing loops are when packets are routed for ever around a network, and often occur with slow convergence. The bandwidth required by these infinite loops will slow your network to a halt. There are methods of preventing these loops however, so this weakness is not as serious as it may first appear.

### Link-state protocols

Link-state protocols are also known as shortest path first protocols. Where distance vector uses information passed along that may or may not be current and accurate, in link-state protocols each router passes along only information about networks and devices directly connected to it. This results in a more accurate picture of the network topology around your router, allowing it to make better routing decisions. This information is passed between routers using link-state advertisements (LSAs). To reduce the overhead, LSAs are only sent out when information

changes, compared to distance vector sending updates at regular intervals even if no information has changed. The more accurate network picture in link-state protocols greatly speed up convergence and avoid problems such as routing-loops.

## Minimum configuration for dynamic routing

Dynamic routing protocols do not pay attention to routing updates from other sources, unless you specifically configure them to do so using CLI redistribute commands within each routing protocol.

The minimum configuration for any dynamic routing to function is to have dynamic routing configured on one interface on the FortiGate unit, and one other router configured as well. Some protocols require larger networks to function as designed.

**Table 8:** Minimum configuration based on dynamic protocol

	<b>BGP</b>	<b>RIP</b>	<b>OSPF / IS-IS</b>
<b>Interface</b>	yes	yes	yes
<b>Network</b>	yes	yes	yes
<b>AS</b>	local and neighbor	no	yes
<b>Neighbors</b>	at least one	at least one	at least one
<b>Version</b>	no	yes	no
<b>Router ID</b>	no	no	yes

## Comparison of dynamic routing protocols

Each dynamic routing protocol was designed to meet a specific routing need. Each protocol does some things well, and other things not so well. For this reason, choosing the right dynamic routing protocol for your situation is not an easy task.

### Features of dynamic routing protocols

Each protocol is better suited for some situations over others.

Choosing the best dynamic routing protocol depends on the size of your network, speed of convergence required, the level of network maintenance resources available, what protocols the networks you connect to are using, and so on. For more information on these dynamic routing protocols, see [“Routing Information Protocol \(RIP\)” on page 319](#), [“Border Gateway Protocol \(BGP\)” on page 358](#), [“Open Shortest Path First \(OSPF\)” on page 396](#), and [“Intermediate System to Intermediate System Protocol \(IS-IS\)” on page 438](#).

**Table 9:** Comparing RIP, BGP, and OSPF dynamic routing protocols

<b>Protocol</b>	<b>RIP</b>	<b>BGP</b>	<b>OSPF / IS-IS</b>
<b>Routing algorithm</b>	Distance Vector, basic	Distance Vector, advanced	Link-state
<b>Common uses</b>	Small non-complex networks	Network backbone, ties multinational offices together	Common in large, complex enterprise networks

**Table 9:** Comparing RIP, BGP, and OSPF dynamic routing protocols (Continued)

Protocol	RIP	BGP	OSPF / IS-IS
<b>Strengths</b>	Fast and simple to implement Near universal support Good when no redundant paths	Graceful restart BFD support Only needed on border routers Summarize routes	Fast convergence Robust Little management overhead No hop count limitation Scalable
<b>Weakness</b>	Frequent updates can flood network Slow convergence Maximum 15 hops may limit network configuration	Required full mesh in large networks can cause floods Route flap Load-balance multi-homed networks Not available on low-end routers	Complex No support for unequal cost multipath routing Route summary can require network changes
<b>Authentication</b>	Optional authentication using text string or MD5 password. (RIP v1 has no authentication)		
<b>IPv6 Support</b>	Only in RIPng	Only in BGP4+	Only in OSPF6 / Integrated IS-IS

## Routing protocols

**Routing Information Protocol (RIP)** uses classful routing, as well as incorporating various methods to stop incorrect route information from propagating, such as the poisoned horizon method. However, on larger networks its frequent updates can flood the network and its slow convergence can be a problem.

**Border Gateway Protocol (BGP)** has been the core Internet backbone routing protocol since the mid 1990s, and is the most used interior gateway protocol (IGP). However, some configurations require full mesh connections which flood the network, and there can be route flap and load balancing issues for multihomed networks.

**Open Shortest Path First (OSPF)** is commonly used in large enterprise networks. It is the protocol of choice mainly due to its fast convergence. However, it can be complicated to setup properly.

**Intermediate System to Intermediate System (IS-IS) Protocol** allows routing of ISO's OSI protocol stack Connectionless Network Service (CLNS). IS-IS is an Interior Gateway Protocol (IGP) not intended to be used between Autonomous Systems (ASes). IS-IS is a link state protocol well-suited to smaller networks that is in widespread use and has near universal support on routing hardware.

**Multicast** addressing is used to broadcast from one source to many destinations efficiently. Protocol Independent Multicast (PIM) is the protocol commonly used in enterprises, multimedia content delivery, and stock exchanges.

## Routing algorithm

Each protocol uses a slightly different algorithm for choosing the best route between two addresses on the network. The algorithm is the “intelligent” part of a dynamic protocol because the algorithm is responsible for deciding which route is best and should be added to the local routing table. RIP and BGP use distance vector algorithms, where OSPF and IS-IS use link-state or a shortest path first algorithm.

Vector algorithms are essentially based on the number of hops between the originator and the destination in a route, possibly weighting hops based on how reliable, fast, and error-free they are.

The link-state algorithm used by OSPF and IS-IS is called the Dijkstra algorithm. Link-state treats each interface as a link, and records information about the state of the interface. The Dijkstra algorithm creates trees to find the shortest paths to the routes it needs based on the total cost of the parts of the routes in the tree.

For more information on the routing algorithm used, see [“Distance vector versus link-state protocols” on page 306](#).

## Authentication

If an attacker gains access to your network, they can masquerade as a router on your network to either gain information about your network or disrupt network traffic. If you have a high quality firewall configured, it will help your network security and stop many of this type of threat. However, the main method for protecting your routing information is to use authentication in your routing protocol. Using authentication on your FortiGate unit and other routers prevents access by attackers — all routers must authenticate with passwords, such as MD5 hash passwords, to ensure they are legitimate routers.

When configuring authentication on your network, ensure you configure it the same on all devices on the network. Failure to do so will create errors and outages as those forgotten devices fail to connect to the rest of the network.

For example, to configure an MD5 key of 123 on an OSPF interface called `ospf_test`, enter the following CLI command:

```
config router ospf
 config ospf-interface
 edit ospf_test
 set authentication md5
 set md5-key 123
 end
 end
end
```

## Convergence

Convergence is the ability of a networking protocol to re-route around network outages. Static routing cannot do this. Dynamic routing protocols can all converge, but take various amounts of time to do this. Slow convergence can cause problems such as network loops which degrade network performance.

You may also hear robustness and redundancy used to describe networking protocols. In many ways they are the same thing as convergence. Robustness is the ability to keep working even though there are problems, including configuration problems as well as network outages. Redundancy involves having duplicate parts that can continue to function in the event of some malfunction, error, or outage. It is relatively easy to configure dynamic routing protocols to have backup routers and configurations that will continue to function no matter the network problem short of a total network failure.

## IPv6 Support

IPv4 addressing is in common use everywhere around the world. IPv6 has much larger addresses and it is used by many large companies and government departments. IPv6 is not as common as IPv4 yet, but more companies are adopting it.

If your network uses IPv6, your dynamic routing protocol must support it. None of the dynamic routing protocols originally supported IPv6, but they all have additions, expansions, or new versions that do support IPv6. For more information, see [“RIP and IPv6” on page 320](#), [“BGP and IPv6” on page 359](#), [“OSPFv3 and IPv6” on page 397](#), or [“Integrated IS-IS” on page 444](#).

## When to adopt dynamic routing

Static routing is more than enough to meet your networking needs when you have a small network. However, as your network grows, the question you need to answer is at what point do you adopt dynamic routing in your networking plan and start using it in your network? The main factors in this decision are typically:

- [Budget](#)
- [Current network size and topology](#)
- [Expected network growth](#)
- [Available resources for ongoing maintenance](#)

### Budget

When making any business decision, the budget must always be considered. Static routing does not involve special hardware, fancy software, or expensive training courses.

Dynamic routing can include all of these extra expenses. Any new hardware, such as routers and switches, will need to support your chosen routing protocols. Network management software and routing protocol drivers may be necessary as well to help configure and maintain your more complex network. If the network administrators are not well versed in dynamic routing, either a training course or some hands-on learning time must be budgeted so they can administer the new network with confidence. Together, these factors can impact your budget.

Additionally, people will always account for network starting costs in the budgets, but usually leave out the ongoing cost of network maintenance. Any budget must provide for the hours that will be spent on updating the network routing equipment, and fixing any problems. Without that money in the budget, you may end up back at static routing before you know it.

### Current network size and topology

As stated earlier static routing works well on small networks. At those networks get larger, routing takes longer, routing tables get very large, and general performance isn't what it could be.

Topology is a concern as well. If all your computers are in one building, its much easier to stay with static routing longer. However, connecting a number of locations will be easier with the move to dynamic routing.

If you have a network of 20 computers, you can still likely use static routing. If those computers are in two or three locations, static routing will still be a good choice for connecting them. Also, if you just connect to your ISP and don't worry about any special routing to do that, you are likely safe with just static routing.

If you have a network of 100 computers in one location, you can use static routing but it will be getting slower, more complex, and there won't be much room for expansion. If those 100 computers are spread across three or more locations, dynamic routing is the way to go.

If you have 1000 computers, you definitely need to use dynamic routing no matter how many locations you have.

Hopefully this section has given you an idea of what results you will likely experience from different sized networks using different routing protocols. Your choice of which dynamic routing protocol to use is partly determined by the network size, and topology.

### Expected network growth

You may not be sure if your current network is ready for dynamic routing. However, if you are expecting rapid growth in the near future, it is a good idea to start planning for that growth now so you are ready for the coming expansion.

Static routing is very labor intensive. Each network device's routing table needs to be configured and maintained manually. If there is a large number of new computers being added to the network, they each need to have the static routing table configured and maintained. If devices are being moved around the network frequently, they must also be updated each time.

Instead, consider putting dynamic routing in place before those new computers are installed on the network. The installation issues can be worked out with a smaller and less complex network, and when those new computers or routers are added to the network there will be nowhere near the level of manual configuration required. Depending on the level of growth, this labor savings can be significant. For example, in an emergency you can drop a new router into a network or AS, wait for it to receive the routing updates from its neighbors, and then remove one of the neighbors. While the routes will not be the most effective possible, this method is much less work than static routing in the same situation, with less chance of mistakes.

Also, as your network grows and you add more routers, those new routers can help share the load in most dynamic routing configurations. For example if you have 4 OSPF routers and 20,000 external routes those few routers will be overwhelmed. But in a network with 15 OSPF routers they will better be able to handle that number of routes. Be aware though that adding more routers to your network will increase the amount of updates sent between the routers, which will use up a greater part of your bandwidth and use more bandwidth overall.

### Available resources for ongoing maintenance

As touched on in the budget section, there must be resources dedicated to ongoing network maintenance, upgrades, and troubleshooting. These resources include administrator hours to configure and maintain the network, training for the administrator if needed, extra hardware and software as needed, and possible extra staff to help the administrator in emergencies. Without these resources, you will quickly find the network reverting to static routing out of necessity.

This is because:

- Routing software updates will require time.
- Routing hardware updates will require time.
- Office reorganizations or significant personnel movement will require time from a networking point of view.
- Networking problems that occur, such as failed hardware, require time to locate and fix the problem.

If the resources to accomplish these tasks are not budgeted, they will either not happen or not happen at the required level to continue operation. This will result in both the network administration staff and the network users being very frustrated.

A lack of maintenance budget will also result in increasingly heavy reliance on static routing as the network administrators are forced to use quick fixes for problems that come up. This invariably involves going to static routing, and dropping the more complex and time-consuming dynamic routing.

## Choosing a routing protocol

One of that hardest decisions in routing can be choosing which routing protocol to use on your network. It can be easy to decide when static routing will not meet your needs, but how can you tell which dynamic routing protocol is best for your network and situation?

Here is a brief look at the routing protocols including their strongest and weakest points. The steps to choosing your routing protocol are:

1. [Answer questions about your network](#)
2. [Dynamic routing terminology](#)
3. [Evaluate your chosen protocol](#)
4. [Implement your dynamic routing protocol](#)

### Answer questions about your network

Before you can decide what is best for your situation, you need to examine what the details of your situation are such as what you have for budget, equipment, and users.

The following questions will help you form a clear idea of your routing needs:

*How many computers or devices are on your network?*

It matters if you only have a few computers, or if you have many and if they are all at one location or not as well. All routing protocols can be run on any sized network, however it can be inefficient to run some on very small networks. However, routers and network hardware that support dynamic routing can be more expensive than more generic routers for static routing.

*What applications typically run over the network?*

Finding out what application your users are running will help you determine their needs and the needs of the network regarding bandwidth, quality of service, and other such issues.

*What level of service do the users expect from the network?*

Different network users have different expectations of the network. Its not critical for someone surfing the Internet to have 100% uptime, but it is required for a stock exchange network or a hospital.

*Is there network expansion in your near future?*

You may have a small network now, but if it will be growing quickly, you should plan for the expected size so you don't have to chance technologies again down the road.

*What routing protocols do your networks connect to?*

This is most often how routing protocol decisions are made. You need to be able to communicate easily with your service provider and neighbors, so often people simply use what everyone else is using.

*Is security a major concern?*

Some routing protocols have levels of authentication and other security features built in. Others do not. If security is important to you, be aware of this.

*What is your budget — both initial and maintenance?*

More robust and feature laden routing protocols generally mean more resources are required to keep them working well. Also more secure configurations require still more resources. This includes both set up costs, as well as ongoing maintenance costs. Ignore these costs at the risk of having to drop the adoption of the new routing protocol mid-change.



## Evaluate your chosen protocol

Once you have examined the features of the routing protocols listed above and chosen the one that best meets your needs, you can set up an evaluation or test install of that protocol.

The test install is generally set up in a sandbox configuration so it will not affect critical network traffic. The aim of the test install is to prove that it will work on a larger scale on your network. So be sure that the test install mirrors your larger network well enough for you to discover any problems. If its too simplistic, these problems may not appear.

If your chosen protocol does not meet your goals choose a different protocol and repeat the evaluation process until either a protocol meets your needs, or you change your criteria.

## Implement your dynamic routing protocol

You have examined your needs, selected the best matching dynamic routing protocol, tested it, and now you are ready to implement it with confidence.

This guide will help you configure your FortiGate unit to support your chosen dynamic routing protocol. Refer to the various sections in this guide as needed during your implementation to help ensure a smooth transition. Examples for each protocol have been included to show proper configurations for different types of networks.

## Dynamic routing terminology

Dynamic routing is a complex subject. There are many routers on different networks and all can be configured differently. It become even more complicated when you add to this each routing protocol having slightly different names for similar features, and many configurable features for each protocol.

To better understand dynamic routing, here are some explanations of common dynamic routing terms.

- [Aggregated routes and addresses](#)
- [Autonomous system \(AS\)](#)
- [Area border router \(ABR\)](#)
- [Neighbor routers](#)
- [Route maps](#)
- [Access lists](#)
- [Bi-directional forwarding detection \(BFD\)](#)

For more details on a term as it applies to a dynamic routing protocol, see one of “[Border Gateway Protocol \(BGP\)](#)” on page 358, “[Routing Information Protocol \(RIP\)](#)” on page 319, or “[Open Shortest Path First \(OSPF\)](#)” on page 396.

### Aggregated routes and addresses

Just as an aggregate interface combines multiple interfaces into one virtual interface, an aggregate route combines multiple routes into one. This reduces the amount of space those routes require in the routing tables of the routers along that route. The trade-off is a small amount of processing to aggregate and de-aggregate the routes at either end.

The benefit of this method is that you can combine many addresses into one, potentially reducing the routing table size immensely. The weakness of this method is if there are holes in the address range you are aggregating you need to decide if its better to break it into multiple ranges, or accept the possibility of failed routes to the missing addresses.

For information on aggregated routes in BGP, see “[ATOMIC\\_AGGREGATE](#)” on page 367, and “[Aggregate routes and addresses](#)” on page 370.

### To manually aggregate the range of IP addresses from 192.168.1.100 to 192.168.1.103

1. Convert the addresses to binary

```
192.168.1.100 = 11000000 10101000 00000001 01100100
192.168.1.101 = 11000000 10101000 00000001 01100101
192.168.1.102 = 11000000 10101000 00000001 01100110
192.168.1.103 = 11000000 10101000 00000001 01100111
```

2. Determine the maximum number of matching bits common to the addresses.

There are 30-bits in common, with only the last 2-bits being different.

- 3 Record the common part of the address.

```
11000000 10101000 00000001 0110010X = 192.168.1.100
```

- 4 For the netmask, assume all the bits in the netmask are 1 except those that are different which are 0.

```
11111111 11111111 11111111 11111100 = 255.255.255.252
```

- 5 Combine the common address bits and the netmask.

```
192.168.1.100/255.255.255.252
```

Alternately the IP mask may be written as a single number:

```
192.168.1.100/2
```

- 6 As required, set variables and attributes to declare the routes have been aggregated, and what router did the aggregating.

### Autonomous system (AS)

An Autonomous System (AS) is one or more connected networks that use the same routing protocol, and appear to be a single unit to any externally connected networks. For example an ISP may have a number of customer networks connected to it, but to any networks connected externally to the ISP it appears as one system or AS. An AS may also be referred to as a routing domain.

It should be noted that while OSPF routing takes place within one AS, the only part of OSPF that deals with the AS is the AS border router (ASBR).

There are multiple types of AS defined by how they are connected to other ASes. A multihomed AS is connected to at least two other ASes and has the benefit of redundancy — if one of those ASes goes down, your AS can still reach the Internet through its other connection. A stub AS only has one connection, and can be useful in specific configurations where limited access is desirable.

Each AS has a number assigned to it, known as an ASN. In an internal network, you can assign any ASN you like (a private AS number), but for networks connected to the Internet (public AS) you need to have an officially registered ASN from Internet Assigned Numbers Authority (IANA). ASNs from 1 to 64,511 are designated for public use.



As of January 2010, AS numbers are 4 bytes long instead of the former 2 bytes. RFC 4893 introduced 32-bit ASNs, which FortiGate units support for BGP and OSPF

---

## Do you need your own AS?

The main factors in deciding if you need your own AS or if you should be part of someone else's are:

- exchanging external routing information
- many prefixes should exist in one AS as long as they use the same routing policy
- when you use a different routing protocol than your border gateway peers (for example your ISP uses BGP, and you use OSPF)
- connected to multiple other AS (multi-homed)

You should not create an AS for each prefix on your network. Neither should you be forced into an AS just so someone else can make AS-based policy decisions on your traffic.

There can be only one AS for any prefix on the Internet. This is to prevent routing issues.

## What AS number to use?

In addition to overseeing IP address allocation and Domain Name Systems (DNS), the Internet Assigned Numbers Authority (IANA) assigns public AS numbers. The public AS numbers are from 1 to 64,511. The ASNs 0, 54272–64511, and 65535 are reserved by the IANA. These ASNs should not be used.

ASNs are assigned in blocks by the Internet Assigned Numbers Authority (IANA) to Regional Internet Registries (RIRs) who then assign ASNs to companies within that RIRs geographic area. Usually these companies are ISPs, and to receive an ASN you must complete the application process of the local RIR and be approved before being assigned an ASN. The RIRs names and regions are:

<b>AFRINIC</b>	Serves the African continent
<b>APNIC</b>	Asia-Pacific including China, India, and Japan
<b>ARIN</b>	American registry including Canada and United States
<b>LACNIC</b>	Latin America, including Mexico, Caribbean, Central and South America
<b>RIPE NCC</b>	Europe, the Middle East, former USSR, and parts of Central Asia

AS numbers from 64512 to 65534 are reserved for private use. Private AS numbers can be used for any internal networks with no outside connections to the Internet such as test networks, classroom labs, or other internal-only networks that do not access the outside world. You can also configure border routers to filter out any private ASNs before routing traffic to the outside world. If you must use private ASNs with public networks, this is the only way to configure them. However, it is risky because many other private networks could be using the same ASNs and conflicts will happen. It would be very much like your local 192.168.0.0 network being made public — the resulting problems would be widespread.

In 1996, when RFC 1930 was written only 5,100 ASes had been allocated and a little under 600 ASes were actively routed in the global Internet. Since that time many more public ASNs have been assigned, leaving only a small number. For this reason 32-bit ASNs (four-octet ASNs) were defined to provide more public ASNs. RFC 4893 defines 32-bit ASNs, and FortiGate units support these larger ASNs.

## Area border router (ABR)

Routers within an AS advertise updates internally and only to each other. However, routers on the edge of the AS must communicate both with routers inside their AS and with routers external to their AS, often running a different routing protocol. These routers are called Area Border Routers (ABRs) or edge routers. Often ABRs run multiple routing protocols to be able to

redistribute traffic between different ASes that are running different protocols, such as the edge between an ISP's IS-IS routing network and a large company's OSPF network.

OSPF defines ABRs differently from other routers. In OSPF, an ABR is an OSPF router that connects another AS to the backbone AS, and is a member of all the areas it connects to. An OSPF ABR maintains a LSA database for each area that it is connected to. The concept of the edge router is present, but it's the edge of the backbone instead of the edge of the OSPF supported ASes.

## Neighbor routers

Routing involves routers communicating with each other. To do this, routers need to know information about each other. These routers are called neighbor routers, and are configured in each routing protocol. Each neighbor has custom settings since some routers may have functionality others routers lack. Neighbour routers are sometimes called peers.

Generally neighbor routers must be configured, and discovered by the rest of the network before they can be integrated to the routing calculations. This is a combination of the network administrator configuring the new router with its neighbor router addresses, and the routing network discovering the new router, such as the hello packets in OSPF. That discovery initiates communication between the new router and the rest of the network.

## Route maps

Route maps are a way for the FortiGate unit to evaluate optimum routes for forwarding packets or suppressing the routing of packets to particular destinations. Compared to access lists, route maps support enhanced packet-matching criteria. In addition, route maps can be configured to permit or deny the addition of routes to the FortiGate unit routing table and make changes to routing information dynamically as defined through route-map rules.

Route maps can be used for limiting both received route updates, and sent route updates. This can include the redistribution of routes learned from other types of routing. For example if you don't want to advertise local static routes to external networks, you could use a route map to accomplish this.

The FortiGate unit compares the rules in a route map to the attributes of a route. The rules are examined in ascending order until one or more of the rules in the route map are found to match one or more of the route attributes.

As an administrator, route maps allow you to group a set of addresses together and assign them a meaningful name. Then during your configuration, you can use these route-maps to speed up configuration. The meaningful names ensure fewer mistakes during configuration as well.

The default rule in the route map (which the FortiGate unit applies last) denies all routes. For a route map to take effect, it must be called by a FortiGate unit routing process.

The syntax for route maps are:

```
config router route-map
 edit <route_map_name>
 set comments
 config rule
 edit <route_map_rule_id>
 set action
 set match-*
 set set-*
 ...
 end
```

The `match-*` commands allow you to match various parts of a route. The `set-*` commands allow you to set routing information once a route is matched.

For an example of how route maps can be used to create receiving or sending “groups” in routing, see [“Redistributing and blocking routes in BGP” on page 389](#).

## Access lists

Use this command to add, edit, or delete access lists. Access lists are filters used by FortiGate unit routing processes. For an access list to take effect, it must be called by a FortiGate unit routing process (for example, a process that supports RIP or OSPF). Use `access-list6` for IPv6 routing.

Access lists can be used to filter which updates are passed between routers, or which routes are redistributed to different networks and routing protocols. You can create lists of rules that will match all routes for a specific router or group of routers.

Each rule in an access list consists of a prefix (IP address and netmask), the action to take for this prefix (permit or deny), and whether to match the prefix exactly or to match the prefix and any more specific prefix.



If you are setting a prefix of 128.0.0.0, use the format 128.0.0.0/1. The default route, 0.0.0.0/0 can not be exactly matched with an access-list. A prefix-list must be used for this purpose.

---

The FortiGate unit attempts to match a packet against the rules in an access list starting at the top of the list. If it finds a match for the prefix, it takes the action specified for that prefix. If no match is found the default action is deny.

The syntax for access lists is:

```
config router access-list, access-list6
 edit <access_list_name>
 set comments
 config rule
 edit <access_list_id>
 set action
 set exact-match
 set prefix
 set prefix6
 set wildcard
```

For an example of how access lists can be used to create receiving or sending “groups” in routing, see [“Redistributing and blocking routes in BGP” on page 389](#).

## Bi-directional forwarding detection (BFD)

Bi-directional Forwarding Detection (BFD) is a protocol used to quickly locate hardware failures in the network. Routers running BFD send packets to each other at a negotiated rate. If packets from a BFD-protected router fail to arrive, then that router is declared down. BFD communicates this information to the routing protocol and the routing information is updated.

BFD neighbors establish if BFD is enabled in OSPF or BFP routers that establish as neighbors.

The CLI commands associated with BFD include:

```
config router bgp
 config neighbor
 set bfd

config router ospf
 set bfd
```

Per-VDOM configuration:

```
config system settings
 set bfd
 set bfd-desired-min-tx
 set bfd-required-min-rx
 set bfd-detect-mult
 set bfd-dont-enforce-src-port
```

Per-interface (override) configuration:

```
config system interface
 edit <interface_name>
 set bfd enable
 set bfd-desired-min-tx
 set bfd-detect-mult
 set bfd-required-min-rx
```

For more information about BFD in BGP, see [“Bi-directional forwarding detection \(BFD\)”](#) on page 374.

## IPv6 in dynamic routing

Unless otherwise stated, routing protocols apply to IPv4 addressing. This is the standard address format used. However, IPv6 is becoming more popular and new versions of the dynamic routing protocols have been introduced.

Dynamic routing supports IPv6 on your FortiGate unit. The new versions of these protocols and the corresponding RFCs are:

- **RIP next generation (RIPng)** — RFC 2080 - Routing Information Protocol next generation (RIPng). See [“RIP and IPv6”](#) on page 320.
- **BGP4+** — RFC 2545, and RFC 2858 Multiprotocol Extensions for IPv6 Inter-Domain Routing, and Multiprotocol Extensions for BGP-4 (MP-BGP) respectively. See [“BGP and IPv6”](#) on page 359.
- **OSPFv3** — RFC 2740 Open Shortest Path First version 3 (OSPFv3) for IPv6 support. See [“OSPFv3 and IPv6”](#) on page 397.
- **Integrated IS-IS** — RFC 5308 for IPv6 support. See [“Integrated IS-IS”](#) on page 444.

As with most advanced routing features on your FortiGate unit, IPv6 settings for dynamic routing protocols must be enabled before they will be visible in the GUI. To enable IPv6 configuration in the GUI, enable it in *System > Admin > Settings*. Alternatively, you can directly configure IPv6 for RIP, BGP, or OSPF protocols using CLI commands.

# Routing Information Protocol (RIP)

This section describes the Routing Information Protocol (RIP).

The following topics are included in this section:

- [RIP background and concepts](#)
- [Troubleshooting RIP](#)
- [Simple RIP example](#)
- [RIPng — RIP and IPv6](#)

## RIP background and concepts

This section contains:

- [Background](#)
- [Parts and terminology of RIP](#)
- [How RIP works](#)

### Background

Routing Information Protocol (RIP) is a distance-vector routing protocol intended for small, relatively homogeneous networks. Its widespread use started when an early version of RIP was included with BSD v4.3 Linux as the routed daemon. The routing algorithm used by RIP, the Bellman–Ford algorithm, first saw widespread use as the initial routing algorithm of the ARPANET.

RIP benefits include being well suited to smaller networks, is in widespread use, near universal support on routing hardware, quick to configure, and works well if there are no redundant paths. However, RIP updates are sent out node-by-node so it can be slow to find a path around network outages. RIP also lacks good authentication, can not choose routes based on different quality of service methods, and can create network loops if you are not careful.

The FortiGate implementation of RIP supports RIP version 1 (see RFC 1058), RIP version 2 (see RFC 2453), and the IPv6 version RIPng (see RFC 2080).

### RIP v1

In 1988 RIP version 1, defined in RFC 1058, was released. The RFC even states that RIP v1 is based on Linux routed due to it being a “defacto standard”.

It uses classful addressing and uses broadcasting to send out updates to router neighbors. There is no subnet information included in the routing updates in classful routing, and it does not support CIDR addressing — subnets must all be the same size. Also, route summarization is not possible.

RIP v1 has no router authentication method, so it is vulnerable to attacks through packet sniffing, and spoofing.

## RIP v2

In 1993, RIP version 2 was developed to deal with the limitations of RIP v1. It was not standardized until 1998. This new version supports classless routing, and subnets of various sizes.

Router authentication was added in RIP v2 — it supports MD5. MD5 hashes are an older encryption method, but this is much improved over no security at all.

In RIP v2 the hop count limit remained at 15 to be backwards compatible with RIP v1.

RIP v2 uses multicasting to send the entire routing table to router neighbors, thereby reducing the traffic for devices that are not participating in RIP routing.

Routing tags were added as well, which allow internal routes or redistributed routes to be identified as such.

## RIPng

RIPng, defined in RFC 2080, is an extension of RIP2 designed to support IPv6. However, RIPng varies from RIPv2 in that it is not fully backwards compatible with RIPv1.

- RIPng does not support RIPv1 update authentication, it relies on IPsec
- RIPng does not allow attaching tags to routes as in RIPv2
- RIPng requires specific encoding of the next hop for a set of route entries, unlike RIPv2 that encodes the next-hop into each route entry.

## Parts and terminology of RIP

Before you can understand how RIP functions, you need to understand some of the main concepts and parts of RIP.

This section includes:

- [RIP and IPv6](#)
- [Default information originate option](#)
- [Garbage, timeout, and update timers](#)
- [Authentication and key-chain](#)
- [Access Lists](#)

### RIP and IPv6

RIP Next Generation (RIPng) is a new version of RIP was released that includes support for IPv6.

The FortiGate unit command `config router ripng` is almost the same as `config router rip`, except that IPv6 addresses are used. Also if you are going to use prefix or access lists with RIPng, you must use the `config router access-list6` or `config prefix-list6` versions of those commands.

If you want to troubleshoot RIPng, it is the same as with RIP but specify the different protocol, and use IPv6 addresses. This applies to commands such as `get router info6` when you want to see the routing table, or other related information.

If you want to route IPv4 traffic over an IPv6 network, you can use the command `config system ip6-tunnel` to configure the FortiGate unit to do this. The IPv6 interface is configured under `config system interface`. All subnets between the source and destination addresses must support IPv6. This command is not supported in Transparent mode.



For example, you want to set up a tunnel on the port1 interface starting at 2002:C0A8:3201:: on your local network and tunnel it to address 2002:A0A:A01:: where it will need access to an IPv4 network again. Use the following command:

```
config system ipv6-tunnel
 edit test_tunnel
 set destination 2002:A0A:A01::
 set interface port1
 set source 2002:C0A8:3201::
 end
end
```

The CLI commands associated with RIPng include:

```
config router ripng
config router access-list6
config router prefix-list6
config system ipv6-tunnel
get router info6 *
```

### Default information originate option

This is the second advanced option for RIP in the web-based manager, right after metric. Enabling default-information-originate will generate and advertise a default route into the FortiGate unit's RIP-enabled networks. The generated route may be based on routes learned through a dynamic routing protocol, routes in the routing table, or both. RIP does not create the default route unless you use the always option.

Select *Disable* if you experience any issues or if you wish to advertise your own static routes into RIP updates.

You can enable or disable default-information-originate in *Router > Dynamic > RIP*, under *Advanced Options*, or use the CLI.

The CLI commands associated with default information originate include:

```
config router rip
 set default-information-originate
end
```

### Garbage, timeout, and update timers

RIP uses various timers to regulate its performance including a garbage timer, timeout timer, and update timer. The FortiGate unit default timer settings (30, 180, and 120 seconds respectively) are effective in most configurations — if you change these settings, ensure that the new settings are compatible with local routers and access servers.



The Timeout period should be at least three times longer than the Update period. If the Update timer is smaller than Timeout or Garbage timers, you will experience an error.

You can set the three RIP timers in *Router > Dynamic > RIP*, under *Advanced Options*, or use the CLI.

The CLI commands associated with garbage, timeout, and update timers include:

```
config router rip
 set garbage-timer
 set timeout-timer
 set update-timer
end
```

### **Garbage timer**

The garbage timer is the amount of time (in seconds) that the FortiGate unit will advertise a route as being unreachable before deleting the route from the routing table. If this timer is shorter, it will keep more up-to-date routes in the routing table and remove old ones faster. This will result in a smaller routing table which is useful if you have a very large network, or if your network changes frequently.

### **Update timer**

The update timer determines the interval between routing updates. Generally, this value is set to 30 seconds. There is some randomness added to help prevent network traffic congestion, which could result from all routers simultaneously attempting to update their neighbors. The update timer should be at least three times smaller than the timeout timer, otherwise you will experience an error.

If you are experiencing significant RIP traffic on your network, you can increase this interval to send fewer updates per minute. However, ensure you increase the interval for all the routers on your network or you will experience time outs that will degrade your network speed.

### **Timeout timer**

The timeout timer is the maximum amount of time (in seconds) that a route is considered reachable while no updates are received for the route. This is the maximum time the FortiGate unit will keep a reachable route in the routing table while no updates for that route are received. If the FortiGate unit receives an update for the route before the timeout period expires, the timer is restarted. The timeout period should be at least three times longer than the update period, otherwise you will experience an error.

If you are experiencing problems with routers not responding in time to updates, increase this timer. However, remember that longer timeout intervals result in longer overall update periods — it may be considerable time before the time the FortiGate unit is done waiting for all the timers to expire on unresponsive routes.

## **Authentication and key-chain**

RIP version 2 uses authentication keys to ensure that the routing information exchanged between routers is reliable. RIP version 1 has no authentication. For authentication to work both the sending and receiving routers must be set to use authentication, and must be configured with the same keys.

The sending and receiving routers need to have their system dates and times synchronized to ensure both ends are using the same keys at the proper times. However, you can overlap the key lifetimes to ensure that a key is always available even if there is some difference in the system times.

A key chain is a list of one or more authentication keys including the send and receive lifetimes for each key. Keys are used for authenticating routing packets only during the specified lifetimes. The FortiGate unit migrates from one key to the next according to the scheduled send and receive lifetimes.

Key-chain is a CLI router command. You use this command to manage RIP version 2 authentication keys. You can add, edit or delete keys identified by the specified key number.

This example shows how to configure a key-chain with two keys that are valid sequentially in time. This example creates a key-chain called "rip\_key" that has a password of "fortinet". The accepted and send lifetimes are both set to the same values — a start time of 9:00am February 23, 2010 and an end time of 9:00am March 17, 2010. A second key is configured with a password of "my\_fortigate" that is valid from March 17, 2010 9:01am to April 1 2010 9:00am. This "rip\_key" keychain is then used on the port1 interface in RIP.

```
config router key-chain
 edit "rip_key"
 config key
 edit 1
 set accept-lifetime 09:00:00 23 02 2010 09:00:00 17 03 2010
 set key-string "fortinet"
 set send-lifetime 09:00:00 23 02 2010 09:00:00 17 03 2010
 next
 edit 2
 set accept-lifetime 09:01:00 17 03 2010 09:00:00 1 04 2010
 set key-string "my_fortigate"
 set send-lifetime 09:01:00 17 03 2010 09:00:00 1 04 2010
 next
 end
 end
config router rip
 config interface
 edit port1
 set auth-keychain "rip_key"
 end
 end
end
```

The CLI commands associated with authentication keys include:

```
config router key-chain

config router rip
 config interface
 edit <interface>
 set auth-keychain
 set auth-mode
 set auth-string
 end
 end
end
```

## Access Lists

Access lists are filters used by FortiGate unit RIP and OSPF routing. An access list provides a list of IP addresses and the action to take for them — essentially an access list makes it easy to group addresses that will be treated the same into the same group, independent of their subnets or other matching qualities. You add a rule for each address or subnet that you want to include, specifying the action to take for it. For example if you wanted all traffic from one department to be routed a particular way, even in different buildings, you can add all the addresses to an access list and then handle that list all at once.

Each rule in an access list consists of a prefix (IP address and netmask), the action to take for this prefix (permit or deny), and whether to match the prefix exactly or to match the prefix and any more specific prefix.

The FortiGate unit attempts to match a packet against the rules in an access list starting at the top of the list. If it finds a match for the prefix, it takes the action specified for that prefix. If no match is found the default action is deny.

Access lists greatly speed up configuration and network management. When there is a problem, you can check each list instead of individual addresses. Also its easier to troubleshoot since if all addresses on one list have problems, it eliminates many possible causes right away.

If you are using the RIPng or OSPF+ IPv6 protocols you will need to use access-list6, the IPv6 version of access list. The only difference is that access-list6 uses IPv6 addresses.

For example, if you want to create an access list called `test_list` that only allows an exact match of 10.10.10.10 and 11.11.11.11, enter the command:

```
config router access-list
 edit test_list
 config rule
 edit 1
 set prefix 10.10.10.10 255.255.255.255
 set action allow
 set exact-match enable
 next
 edit 2
 set prefix 11.11.11.11 255.255.255.255
 set action allow
 set exact-match enable
 end
 end
 end
```

Another example is if you want to deny ranges of addresses in IPv6 that start with the IPv6 equivalents of 10.10.10.10 and 11.11.11.11, enter the command access-list6 as follows:

```
config router access-list6
 edit test_list_ip6
 config rule
 edit 1
 set prefix6 2002:A0A:A0A:0:0:0:0:0/48
 set action deny
 next
 edit 2
 set prefix6 2002:B0B:B0B:0:0:0:0:0/48
 set action deny
 end
 end
 end
```

To use an access\_list, you must call it from a routing protocol such as RIP. The following example uses the access\_list from the earlier example called test\_list to match routes coming in

on the port1 interface. When there is a match, it will add 3 to the hop count metric for those routes to artificially increase . Enter the following command:

```
config router rip
 config offset-list
 edit 5
 set access-list test_list
 set direction in
 set interface port1
 set offset 3
 set status enable
 end
```

If you are setting a prefix of 128.0.0.0, use the format 128.0.0.0/1. The default route, 0.0.0.0/0 can not be exactly matched with an access-list. A prefix-list must be used for this purpose

## How RIP works

As one of the original modern dynamic routing protocols, RIP is straightforward. Its routing algorithm is not complex, there are some options to allow fine tuning, and it's relatively simple to configure RIP on FortiGate units.

From RFC 1058:

Distance vector algorithms are based on the exchange of only a small amount of information. Each entity (gateway or host) that participates in the routing protocol is assumed to keep information about all of the destinations within the system. Generally, information about all entities connected to one network is summarized by a single entry, which describes the route to all destinations on that network.

This section includes:

- [RIP versus static routing](#)
- [RIP metric – hop count](#)
- [The Bellman–Ford routing algorithm](#)
- [Passive versus active RIP interfaces](#)
- [RIP packet structure](#)

### RIP versus static routing

RIP was one of the earliest dynamic routing protocols to work with IP addresses. As such, it is not as complex as more recent protocols. However, RIP is a big step forward from simple static routing.

While RIP may be slow in response to network outages, static routing has zero response. The same is true for convergence — static routing has zero convergence. Both RIP and static routing have the limited hop count, so its not a strength or a weakness. Count to infinity can be a problem, but typically can be fixed as it happens or is the result of a network outage that would cause even worse problems on static routing network.

This compares to static routing where each time a packet needs to be routed, the FortiGate unit can only send it to the next hop towards the destination. That next hop then forwards it, and so on until it arrives at its destination. RIP keeps more routing information on each router so your FortiGate unit can send the packet farther towards its destination before it has to be routed again towards its destination. RIP reduces the amount of table lookups and therefore fewer network resources than static routing. Also since RIP is updated on neighboring routes it is aware of new routes or dead routes that static routing would not be aware of.

Overall, RIP is a large step forward when compared to static routing.

### RIP metric — hop count

RIP uses hop count as the metric for choosing the best route. A hop count of 1 represents a network that is connected directly to the FortiGate unit, while a hop count of 16 represents a network that cannot be reached. Each network that a packet travels through to reach its destination usually counts as one hop. When the FortiGate unit compares two routes to the same destination, it adds the route having the lowest hop count to the routing table. As you can see in [“RIP packet structure” on page 329](#), the hop count is part of a RIP v2 packet.

Similarly, when RIP is enabled on an interface, the FortiGate unit sends RIP responses to neighboring routers on a regular basis. The updates provide information about the routes in the FortiGate unit’s routing table, subject to the rules that you specify for advertising those routes. You can specify how often the FortiGate unit sends updates, the period of time a route can be kept in the routing table without being updated, and for routes that are not updated regularly you can specify the period of time that the unit advertises a route as unreachable before it is removed from the routing table.

If hops are weighted higher than one, it becomes very easy to reach the upper limit. This higher weighting will effectively limit the size of your network depending on the numbers used. Merely changing from the default of 1.0 to 1.5 will lower the effective hop count from 15 to 10. This is acceptable for smaller networks, but can be a problem as your network expands over time.

In RIP, you can use the offset command to artificially increase the hop count of a route. Doing this will make this route less preferred, and in turn it will get less traffic. Offsetting routes is useful when you have network connections of different bandwidths, different levels of reliability, or different costs. In each of these situations you still want the redundancy of multiple route access, but you don’t want the bulk of your traffic using these less preferred routes. For an example of RIP offset, see [“Access Lists” on page 323](#).

### The Bellman–Ford routing algorithm

The routing algorithm used by RIP was first used in 1967 as the initial routing algorithm of the ARPANET. The Bellman–Ford algorithm is distributed because it involves a number of nodes (routers) within an Autonomous system, and consists of the following steps:

1. Each node calculates the distances between itself and all other nodes within the AS and stores this information as a table.
2. Each node sends its table to all neighboring nodes.
3. When a node receives distance tables from its neighbors, it calculates the shortest routes to all other nodes and updates its own table to reflect any changes.

To examine how this algorithm functions let’s look at a network with 4 routers — routers 1 through 4. The distance from router1 to router2 is 2 hops, 1 to 3 is 3 hops, and 2 to 3 is 4 hops. Router4 is only connected to routers 2 and 3, each distance being 2 hops.

1. Router1 finds all the distance to the other three routers — router 2 is 2, router 3 is 3. Router1 doesn’t have a route to router 4.
2. Routers 2 through 4 do the same calculations from their point of views.
3. Once router 1 gets an update from router 2 or 3, it will get their route to router 4. At that point it now has a route to router 4 and installs that in its local table.
4. If router1 gets an update from router3 first, it has a hop count of 5 to reach router4. But when router2 sends its update, router1 will go with router2’s shorter 4 hops to reach router4. Future updates don’t change this unless they are shorter than 4 hops, or the routing table route goes down.

**Figure 100:**RIP algorithm example in 4 steps


**Step 1**

Router1 finds the distance to other routers in the network.

It currently has no route to Router4.

Router1 routing table:


- Distance to Router2 = 2 hops.
- Distance to Router3 = 3 hops.



**Step 2**

All routers do the same as Router1, and send out updates containing their routing table.

Note that Router1 and Router4 do not update each other, but rely on Router2 and Router3 to pass along accurate updates.




**Step 3**

Each router looks at the updates it has received, and adds any new or shorter routes to its table.

Router1 updated table:

- Distance to Router2 = 2 hops.
- Distance to Router3 = 3 hops.
- Distance to Router4 = 4 or 5 hops.




**Step 4**

Router1 installs the shortest route to Router4, and the other routes to it are removed from the routing table.

Router1 complete table:

- Distance to Router2 = 2 hops.
- Distance to Router3 = 3 hops.
- Distance to Router4 = 4 hops.



The good part about the Bellman-Ford algorithm in RIP is that the router only uses the information it needs from the update. If there are no newer, better routes than the ones the router already has in its routing table, there is no need to change its routing table. And no change means no additional update, so less traffic. But even when there is update traffic, the RIP packets are very small so it takes many updates to affect overall network bandwidth. For more information about RIP packets, see [“RIP packet structure” on page 329](#).

The main disadvantage of the Bellman–Ford algorithm in RIP is that it doesn't take weightings into consideration. While it is possible to assign different weights to routes in RIP, doing so severely limits the effective network size by reducing the hop count limit. Also other dynamic routing protocols can take route qualities, such as reliability or delay, into consideration to provide not only the physically shortest but also the fastest or more reliable routes as you choose.

Another disadvantage of the Bellman-Ford algorithm is due to the slow updates passed from one RIP router to the next. This results in a slow response to changes in the network topology, which in turn results in more attempts to use routes that are down, which wastes time and network resources.

### Passive versus active RIP interfaces

Normally the FortiGate unit's routing table is kept up to date by periodically asking the neighbors for routes, and sending your routing updates out. This has the downside of generating a lot of extra traffic for large networks. The solution to this problem is passive interfaces.

An standard interface that supports RIP is active by default — it both sends and receives updates by actively communicating with its neighbors. A passive RIP interface does not send out updates — it just listens to the updates of other routers. This is useful in reducing network traffic, and if there are redundant routers in the network that would be sending out essentially the same updates all the time.

The following example shows how to create a passive RIP v2 interface on port1, using MD5 authentication and a key-chain called `passiveRIPv2` that has already been configured. Note that in the CLI, you enable passive by disabling `send-version2-broadcast`.

#### To create a passive RIP interface - web-based manager

1. Go to *Router > Dynamic > RIP*.
2. Under *Interfaces*, select *Create*.
3. Select port1 as the *Interface*.
4. Select 2 as both the *Send Version* and *Receive Version*.
5. Select MD5 for *Authentication*.
6. Select the `passiveRIPv2` *Key-chain*.
7. Select *Passive Interface*.
8. Select *OK* to accept this configuration, and return to the main RIP display page.

#### To create a passive RIP v2 interface on port1 using MD5 authentication- CLI

```
config router rip
 config interface
 edit port1
 set send-version2-broadcast disable
 set auth-keychain "passiveRIPv2"
 set auth-mode md5
 set receive-version 2
 set send-version 2
 end
 end
end
```



## RIP packet structure

It is hard to fully understand a routing protocol without knowing what information is carried in its packets. Knowing what information is exchanged between routers and how will help you better understand the RIP protocol, and better configure your network for it.

This section provides information on the contents of RIP 1 and RIP 2 packets.

### RIP version 1

RIP version 1, or RIP IP packets are 24 bytes in length, with some empty areas left for future expansion.

**Table 10:** RIP IP packets

1-byte command	1-byte version	2-byte zero field	2-byte AFI	2-byte zero field
4-byte IP address	4-byte zero field	4-byte zero field	4-byte metric	

A RIP 1 packet contains the following fields:

**Command** — Indicates whether the packet is a request or a response. The request asks that a router send all or part of its routing table. The response can be an unsolicited regular routing update or a reply to a request. Responses contain routing table entries. Multiple RIP packets are used to convey information from large routing tables.

**Version** — Specifies the RIP version used. This field can signal different potentially incompatible versions.

**Zero field** — This field defaults to zero, and is not used by RFC 1058 RIP.

**Address-family identifier (AFI)** — Specifies the address family used. RIP is designed to carry routing information for several different protocols. Each entry has an address-family identifier to indicate the type of address being specified. The AFI for IP is 2.

**IP Address** — Specifies the IP address for the entry.

**Metric** — This is the number of hops or routers traversed along the route on its trip to the destination. The metric is between 1 and 15 for that number of hops. If the route is unreachable the metric is 16.

### RIP version 2

RIP version 2 has more features than RIP 1, which is reflected in its packets which carry more information. All but one of the empty zero fields in RIP 1 packets are used in RIP 2.

**Table 11:** RIP 2 packets

1-byte command	1-byte version	2-byte unused	2-byte AFI	2-byte route tag	4-byte IP address	4-byte subnet	4-byte next hop	4-byte metric
----------------	----------------	---------------	------------	------------------	-------------------	---------------	-----------------	---------------

A RIP 2 packet contains fields described above in RIP 1, as well as the following:

**Unused** — Has a value set to zero, and is intended for future use

**Route tag** — Provides a method for distinguishing between internal routes learned by RIP and external routes learned from other protocols.

**Subnet mask** — Contains the subnet mask for the entry. If this field is zero, no subnet mask has been specified for the entry.

**Next hop** — Indicates the IP address of the next hop to which packets for the entry should be forwarded.

## Troubleshooting RIP

This section is about troubleshooting RIP. For general troubleshooting information, see the Troubleshooting chapter.

This section includes:

- [Routing Loops](#)
- [Holddowns and Triggers for updates](#)
- [Split horizon and Poison reverse updates](#)
- [Debugging IPv6 on RIPng](#)

### Routing Loops

Normally in routing, a path between two addresses is chosen and traffic is routed along that path from one address to the other. When there is a routing loop, that normal path doubles back on itself creating a loop. When there are loops, the network has problems getting information to its destination and also prevents it from returning to the source to report the inaccessible destination.

A routing loop happens when a normally functioning network has an outage, and one or more routers are offline. When packets encounter this, an alternate route is attempted to maneuver around the outage. During this phase it is possible for a route to be attempted that involves going back a hop, and trying a different hop forward. If that hop forward is blocked by the outage as well, a hop back and possibly the original hop forward may be selected. You can see if this continues, how it can consume not only network bandwidth but also many resources on those routers affected. The worst part is this situation will continue until the network administrator changes the router settings, or the downed routers come back online.

#### Routing loops' effect on the network

In addition to this “traffic jam” of routed packets, every time the routing table for a router changes that router sends an update out to all of the RIP routers connected to it. In a network loop, its possible for a router to change its routes very quickly as it tries and fails along these new routes. This can quickly result in a flood of updates being sent out, which can effectively grind the network to a halt until the problem is fixed.

#### How can you spot a routing loop

Any time network traffic slows down, you will be asking yourself if it is a network loop or not. Often slowdowns are normal, they are not a full stoppage, and normal traffic resumes in a short period of time.

If the slow down is a full halt of traffic or a major slowdown does not return to normal quickly, you need to do serious troubleshooting quickly.

If you aren't running SNMP, dead gateway detection, or you have non-Fortinet routers in your network, you can use networking tools such as ping and traceroute to define the outage on your network and begin to fix it. Ping, traceroute, and other basic troubleshooting tools are largely the same between static and dynamic, and are covered in [“Troubleshooting static routing” on page 281](#).

### Check your logs

If your routers log events to a central location, it can be easy to check the logs for your network for any outages.

On your FortiGate unit, go to *Log & Report*. You will want to look at both event logs and traffic logs. Events to look for will generally fall under CPU and memory usage, interfaces going offline (due to dead gateway detection), and other similar system events.

Once you have found and fixed your network problem, you can go back to the logs and create a report to better see how things developed during the problem. This type of forensics analysis can better help you prepare for next time.

### Use SNMP network monitoring

If your network had no problems one minute and slows to a halt the next, chances are something changed to cause that problem. Most of the time an offline router is the cause, and once you find that router and bring it back online, things will return to normal.

If you can enable a hardware monitoring system such as SNMP or sFlow on your routers, you can be notified of the outage and where it is exactly as soon as it happens.

Ideally you can configure SNMP on all your FortiGate routers and be alerted to all outages as they occur.

### To use SNMP to detect potential routing loops

1. Go to *System > Config > SNMP*.

2. Enable *SMTP Agent* and select *Apply*.

Optionally enter the *Description*, *Location*, and *Contact* information for this device for easier location of the problem report.

3. Under *SNMP v1/v2* or *SNMP v3* as appropriate, select *Create New*.

#### SNMP v3

<b>User Name</b>	Enter the SNMP user ID.
<b>Security Level</b>	Select authentication or privacy as desired. Select the authentication or privacy algorithms to use and enter the required passwords.
<b>Notification Host</b>	Enter the IP addresses of up to 16 hosts to notify.
<b>Enable Query</b>	Select. The <i>Port</i> should be 161. Ensure that your security policies allow ports 161 and 162 (SNMP queries and traps) to pass.

#### SNMP v1/v2

<b>Hosts</b>	Enter the IP addresses of up to 8 hosts to notify. You can also specify the network <i>Interface</i> , or leave it as <i>ANY</i> .
--------------	------------------------------------------------------------------------------------------------------------------------------------

---

<b>Queries</b>	Enable <i>v1</i> and/or <i>v2</i> as needed. The <i>Port</i> should be 161. Ensure that your security policies allow port 161 to pass.
----------------	----------------------------------------------------------------------------------------------------------------------------------------

---

<b>Traps</b>	Enable <i>v1</i> and/or <i>v2</i> as needed. The <i>Port</i> should be 162. Ensure that your security policies allow port 162 to pass.
--------------	----------------------------------------------------------------------------------------------------------------------------------------

---

4. Select the events for which you want notification. For routing loops this should include *CPU usage is high*, *Memory is low*, and possibly *Log disk space is low*. If there are problems the log will be filling up quickly, and the FortiGate unit's resources will be overused.
5. Configure SNMP host (manager) software on your administration computer. This will monitor the SNMP information sent out by the FortiGate unit. Typically you can configure this software to alert you to outages or CPU spikes that may indicate a routing loop.

### Use dead gateway detection and e-mail alerts

Another tool available to you on FortiGate units is the dead gateway detection. This feature allows the FortiGate unit to ping a gateway at regular intervals to ensure it is online and working. When the gateway is not accessible, that interface is marked as down.

#### To detect possible routing loops with dead gateway detection and e-mail alerts

1. To configure dead gateway detection, go to *Router > Static > Settings* and select *Create New*.
2. Enter the *Ping Server* IP address and select the *Interface* that connects to it.
3. Set the *Ping Interval* (how often to send a ping), and *Failover Threshold* (how many lost pings is considered a failure). A smaller interval and smaller number of lost pings will result in faster detection, but will create more traffic on your network.

#### To configure notification of failed gateways

1. Go to *Log & Report > Log Config > Alert E-mail*.
2. Enter your email details.
3. Select the *Configuration changes* event.
4. Select *Apply*.

You might also want to log CPU and Memory usage as a network outage will cause your CPU activity to spike.



If you have VDOMs configured, you will have to enter the basic SMTP server information in the Global section, and the rest of the configuration within the VDOM that includes this interface.

---

After this configuration, when this interface on the FortiGate unit cannot connect to the next router, the FortiGate unit will bring down the interface and alert you with an email about the outage.

### Look at the packet flow

If you want to see what is happening on your network, look at the packets travelling on the network. This is same idea as police pulling over a car and asking the driver where they have been, and what the conditions were like.

The method used in the troubleshooting sections [“Debugging IPv6 on RIPng” on page 334](#) and on debugging the packet flow apply here as well. In this situation, you are looking for routes that have metrics higher than 15 as that indicates they are unreachable.

Ideally if you debug the flow of the packets, and record the routes that are unreachable, you can create an accurate picture of the network outage.

### Action to take on discovering a routing loop

Once you have mapped the problem on your network, and determined it is in fact a routing loop there are a number of steps to take in correcting it.

1. Get any offline routers back online. This may be a simple reboot, or you may have to replace hardware. Often this first step will restore your network to its normal operation, once the routing tables finish being updated.
2. Change your routing configuration on the edges of the outage. Even if step 1 brought your network back online, you should consider making changes to improve your network before the next outage occurs. These changes can include configuring features like holddowns and triggers for updates, split horizon, and poison reverse updates.

### Holddowns and Triggers for updates

One of the potential problems with RIP is the frequent routing table updates that are sent every time there is a change to the routing table. If your network has many RIP routers, these updates can start to slow down your network. Also if you have a particular route that has bad hardware, it might be going up and down frequently, which will generate an overload of routing table updates.

One of the most common solutions to this problem is to use holddown timers and triggers for updates. These slow down the updates that are sent out, and help prevent a potential flood.

#### Holddown Timers

The holddown timer activates when a route is marked down. Until the timer expires, the router does not accept any new information about that route. This is very useful if you have a flapping route because it will prevent your router from sending out updates and being part of the problem in flooding the network. The potential down side is if the route comes back up while the timer has not expired, that route will be unavailable for that period of time. This is only a problem if this is a major route used by the majority of your traffic. Otherwise, this is a minor problem as traffic can be re-routed around the outage.

#### Triggers

Triggered RIP is an alternate update structure that is based around limiting updates to only specific circumstances. The most basic difference is that the routing table will only be updated when a specific request is sent to update, as opposed to every time the routing table changes. Updates are also triggered when a unit is 'powered on', which can include addition of new interfaces or devices to the routing structure, or devices returning to being available after being unreachable.

### Split horizon and Poison reverse updates

Split horizon is best explained with an example. You have three routers linked serially, let's call them A, B, and C. A is only linked to B, C is only linked to B, and B is linked to both A and C. To get to C, A must go through B. If the link to C goes down, it is possible that B will try to use A's route to get to C. This route is A-B-C, so it will loop endlessly between A and B.

This situation is called a split horizon because from B's point of view the horizon stretches out in each direction, but in reality it only is on one side.

Poison reverse is the method used to prevent routes from running into split horizon problems. Poison reverse "poisons" routes away from the destination that use the current router in their

route to the destination. This “poisoned” route is marked as unreachable for routers that cannot use it. In RIP this means that route is marked with a distance of 16.

## Debugging IPv6 on RIPng

The debug commands are very useful to see what is happening on the network at the packet level. There are a few changes to debugging the packet flow when debugging IPv6.

The following CLI commands specify both IPv6 and RIP, so only RIPng packets will be reported. The output from these commands will show you the RIPng traffic on your FortiGate unit including RECV, SEND, and UPDATE actions.

The addresses are in IPv6 format.

```
diagnose debug enable
diagnose ipv6 router rip level info
diagnose ipv6 router rip all enable
```

These three commands will:

- turn on debugging in general
- set the debug level to information, a verbose reporting level
- turn on all rip router settings

Part of the information displayed from the debugging is the metric (hop count). If the metric is 16, then that destination is unreachable since the maximum hop count is 15.

In general, you should see an update announcement, followed by the routing table being sent out, and a received reply in response.

For more information, see [“Testing the IPv6 RIPng information” on page 357](#)

## Simple RIP example

This is an example of a typical medium sized network configuration using RIP routing.

Your company has 3 small local networks, one for each department. These networks are connected by RIP, and then connected to the Internet. Each subnet has more than one route, for redundancy. There are two central routers that are both connected to the internet, and to the other networks. If one of those routers goes down, the whole network can continue to function normally.

The ISP is running RIP, so no importing or exporting routes is required on the side of the network. However, since the internal networks have static networking running those will need to be redistributed through the RIP network.

To keep the example simple, there will be no authentication of router traffic.

With RIP properly configured, if the device fails or temporarily goes offline, the routes will change and traffic will continue to flow. RIP is good for a smaller network due to its lack of complex configurations.

This section includes the following topics:

- [Network layout and assumptions](#)
- [General configuration steps](#)
- [Configuring the FortiGate units system information](#)
- [Configuring other networking devices](#)
- [Testing network configuration](#)

## Network layout and assumptions

### Basic network layout

Your company has 3 departments each with their own network – Sales, R&D, and Accounting. Each network has routers that are not running RIP as well as FortiGate units running RIP.

The R&D network has two RIP routers, and each is connected to both other departments as well as being connected to the Internet through the ISP router. The links to the Internet are indicated in black.

The three internal networks do not run RIP. They use static routing because they are small networks. This means the FortiGate units have to redistribute any static routes they learn so that the internal networks can communicate with each other.


Where possible in this example, the default values will be used or the most general settings. This is intended to provide an easier configuration that will require less troubleshooting.

In this example the routers, networks, interfaces used, and IP addresses are as follows. Note that the Interfaces that connect Router2 and Router3 also connect to the R&D network.

**Table 12:** Rip example network topology

Network	Router	Interface & Alias	IP address
Sales	Router1	port1 (internal)	10.11.101.101
		port2 (router2)	10.11.201.101
		port3 (router3)	10.11.202.101
R&D	Router2	port1 (internal)	10.12.101.102
		port2 (router1)	10.11.201.102
		port3 (router4)	10.14.201.102
		port4 (ISP)	172.20.120.102
	Router3	port1 (internal)	10.12.101.103
		port2 (router1)	10.11.201.103
		port3 (router4)	10.14.202.103
		port4 (ISP)	172.20.120.103
Accounting	Router4	port1 (internal)	10.14.101.104
		port2 (router2)	10.14.201.104
		port3 (router3)	10.14.202.104

**Figure 101:**Network topology for the simple RIP example



## Assumptions

The following assumptions have been made concerning this example.

- All FortiGate units have 5.0 firmware, and are running factory default settings.
- All CLI and web-based manager navigation assumes the unit is running in NAT/Route operating mode, with VDOMs disabled.
- All FortiGate units have interfaces labelled port1 through port4 as required.
- All firewalls have been configured for each FortiGate unit to allow the required traffic to flow across interfaces.
- Only FortiGate units are running RIP on the internal networks.
- Router2 and Router3 are connected through the internal network for R&D.
- Router2 and Router3 each have their own connection to the Internet, indicated in black in Figure 101.

## General configuration steps

This example is very straight forward. The only steps involved are:

- [Configuring the FortiGate units system information](#)
- [Configuring FortiGate unit RIP router information](#)
- [Configuring other networking devices](#)
- [Testing network configuration](#)

## Configuring the FortiGate units system information

Each FortiGate unit needs their hostname, and interfaces configured.

For IP numbering, Router2 and Router3 use the other routers numbering where needed.



Router2 and Router3 have dead gateway detection enabled on the ISP interfaces using Ping. Remember to contact the ISP and confirm their server has ping enabled.

## Configure the hostname, interfaces, and default route

### To configure Router1 system information - web-based manager

1. Go to *System > Dashboard > Status > System Information*.
2. Next to *Host Name* select *Change*, and enter "Router1".
3. Go to *Router > Static > Static Routes*.
4. Edit the default route and enter the following information:

---

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Device</b>	port2 (router2)
<b>Gateway</b>	172.20.120.5/255.255.255.0
<b>Distance</b>	40

---

5. Enter a second default route and enter the following information:

---

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Device</b>	port3 (router3)
<b>Gateway</b>	172.20.120.5/255.255.255.0
<b>Distance</b>	40

---

6. Go to *System > Network > Interfaces*.
7. Edit port1 (internal) interface.
8. Set the following information, and select *OK*.

---

<b>Alias</b>	internal
<b>IP/Netmask</b>	10.11.101.101/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Internal sales network
<b>Administrative Status</b>	Up

---

9. Edit port2 (router2) interface.
10. Set the following information, and select *OK*.

---

<b>Alias</b>	router2
<b>IP/Netmask</b>	10.11.201.101/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING

---

<b>Description</b>	Link to R&D network & internet through Router2
<b>Administrative Status</b>	Up

11. Edit port3 (router3) interface.

12. Set the following information, and select *OK*.

<b>Alias</b>	router3
<b>IP/Netmask</b>	10.11.202.101/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Link to R&D network and internet through Router3
<b>Administrative Status</b>	Up

**To configure Router1 system information - CLI**

```
config system global
 set hostname Router1
end
```

```

config router static
 edit 1
 set device "port2"
 set distance 45
 set gateway 10.11.201.102
 next
 edit 2
 set device "port3"
 set distance 45
 set gateway 10.11.202.103
 end
end

config system interface
 edit port1
 set alias internal
 set ip 10.11.101.101/255.255.255.0
 set allowaccess https ssh ping
 set description "Internal sales network"
 next
 edit port2
 set alias ISP
 set allowaccess https ssh ping
 set ip 10.11.201.101/255.255.255.0
 set description "Link to R&D network & internet through Router2"
 next
 edit port3
 set alias router3
 set ip 10.11.202.101/255.255.255.0
 set allowaccess https ssh ping
 set description "Link to R&D network & internet through Router2"
 end
end

```

### To configure Router2 system information - web-based manager

1. Go to *System > Dashboard > Status > System Information*.
2. Next to *Host Name* select *Change*, and enter "Router2".
3. Go to *Router > Static > Static Routes*.
4. Edit the default route and enter the following information:

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Device</b>	port4 (ISP)
<b>Gateway</b>	172.20.120.5/255.255.255.0
<b>Distance</b>	5

5. Go to *System > Network > Interfaces*.
6. Edit port1 (internal) interface.
7. Set the following information, and select *OK*.

<b>Alias</b>	internal
<b>IP/Netmask</b>	10.12.101.102/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	R&D internal network and Router3
<b>Administrative Status</b>	Up

8. Edit port2 (router1) interface.
9. Set the following information, and select *OK*.

<b>Alias</b>	router1
<b>IP/Netmask</b>	10.12.201.102/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Link to Router1 and the Sales network
<b>Administrative Status</b>	Up

10. Edit port3 (router4) interface.
11. Set the following information, and select *OK*.

<b>Alias</b>	router4
<b>IP/Netmask</b>	10.12.301.102/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Link to Router4 and the accounting network
<b>Administrative Status</b>	Up

12. Edit port4 (ISP) interface.
13. Set the following information, and select *OK*.

<b>Alias</b>	ISP
<b>IP/Netmask</b>	172.20.120.102/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Detect Interface Status for Gateway Load Balancing</b>	enable
<b>Detect Server</b>	172.20.120.5
<b>Detect Protocol</b>	Ping

<b>Detect Interface Status for Gateway Load Balancing</b>	enable
<b>Description</b>	Internet through ISP
<b>Administrative Status</b>	Up

### To configure Router2 system information - CLI

```

config system global
 set hostname Router2
end
config router static
 edit 1
 set device "port4"
 set distance 5
 set gateway 172.20.130.5
 end
end
config system interface
 edit port1
 set alias internal
 set ip 10.11.101.102/255.255.255.0
 set allowaccess https ssh ping
 set description "Internal RnD network and Router3"
 next
 edit port2
 set alias router1
 set allowaccess https ssh ping
 set ip 10.11.201.102/255.255.255.0
 set description "Link to Router1"
 next
 edit port3
 set alias router3
 set ip 10.14.202.102/255.255.255.0
 set allowaccess https ssh ping
 set description "Link to Router4"
 next
 edit port4
 set alias ISP
 set ip 172.20.120.102/255.255.255.0
 set allowaccess https ssh ping
 set description "ISP and internet"
 end
end
end

```

### To configure Router3 system information - web-based manager

1. Go to *System > Dashboard > Status > System Information*.
2. Next to *Host Name* select *Change*, and enter "Router3".
3. Go to *Router > Static > Static Route*.

4. Edit the default route and enter the following information:

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Device</b>	port4 (ISP)
<b>Gateway</b>	172.20.120.5/255.255.255.0
<b>Distance</b>	5

5. Go to *System > Network > Interfaces*.  
6. Edit port1 (internal) interface.  
7. Set the following information, and select *OK*.

<b>Alias</b>	internal
<b>IP/Netmask</b>	10.12.101.103/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	R&D internal network and Router2
<b>Administrative Status</b>	Up

8. Edit port2 (router1) interface.  
9. Set the following information, and select *OK*.

<b>Alias</b>	router1
<b>IP/Netmask</b>	10.13.201.103/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Link to Router1 and Sales network
<b>Administrative Status</b>	Up

10. Edit port3 (router4) interface.  
11. Set the following information, and select *OK*.

<b>Alias</b>	router4
<b>IP/Netmask</b>	10.13.301.103/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Link to Router4 and accounting network
<b>Administrative Status</b>	Up

12. Edit port4 (ISP) interface.

13. Set the following information, and select *OK*.

<b>Alias</b>	ISP
<b>IP/Netmask</b>	172.20.120.103/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Detect Interface Status for Gateway Load Balancing</b>	enable
<b>Detect Server</b>	172.20.120.5
<b>Detect Protocol</b>	Ping
<b>Description</b>	Internet and ISP
<b>Administrative Status</b>	Up

### To configure Router3 system information - CLI

```
config system global
 set hostname Router3
end

config router static
 edit 1
 set device "port4"
 set distance 5
 set gateway 172.20.130.5
 end
end

config system interface
 edit port1
 set alias internal
 set ip 10.12.101.103/255.255.255.0
 set allowaccess https ssh ping
 set description "Internal RnD network and Router2"
 next
 edit port2
 set alias ISP
 set allowaccess https ssh ping
 set ip 10.11.201.103/255.255.255.0
 set description "Link to Router1"
 next
 edit port3
 set alias router3
 set ip 10.14.202.103/255.255.255.0
 set allowaccess https ssh ping
 set description "Link to Router4"
 next
 edit port4
 set alias ISP
 set ip 172.20.120.103/255.255.255.0
 set allowaccess https ssh ping
 set description "ISP and internet"
 end
end
```

### To configure Router4 system information - web-based manager

1. Go to *System > Dashboard > Status > System Information*.
2. Next to *Host Name* select *Change*, and enter "Router4".
3. Go to *Router > Static > Static Routes*.
4. Edit the default route and enter the following information:

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Device</b>	port2 (router2)



<b>Gateway</b>	172.20.120.5/255.255.255.0
<b>Distance</b>	40

5. Enter a second default route and enter the following information:

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Device</b>	port3 (router3)
<b>Gateway</b>	172.20.120.5/255.255.255.0
<b>Distance</b>	40

6. Go to *System > Network > Interfaces*.  
 7. Edit port 1 (internal) interface.  
 8. Set the following information, and select *OK*.

<b>Alias</b>	internal
<b>IP/Netmask</b>	10.14.101.104/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Internal accounting network
<b>Administrative Status</b>	Up

9. Edit port 2 (router2) interface.  
 10. Set the following information, and select *OK*.

<b>Alias</b>	router2
<b>IP/Netmask</b>	10.14.201.104/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Link to R&D network & internet through Router2
<b>Administrative Status</b>	Up

11. Edit port 3 (router3) interface.  
 12. Set the following information, and select *OK*.

<b>Alias</b>	router3
<b>IP/Netmask</b>	10.14.301.104/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Link to R&D network and internet through Router3
<b>Administrative Status</b>	Up

## To configure Router4 system information - CLI

```
config system global
 set hostname Router4
end

config router static
 edit 1
 set device "port2"
 set distance 45
 set gateway 10.14.201.102
 next
 edit 2
 set device "port3"
 set distance 45
 set gateway 10.14.202.103
 end
end

config system interface
 edit port1
 set alias internal
 set ip 10.14.101.104/255.255.255.0
 set allowaccess https ssh ping
 set description "Internal sales network"
 next
 edit port2
 set alias router2
 set allowaccess https ssh ping
 set ip 10.14.201.104/255.255.255.0
 set description "Link to R&D network & internet through Router2"
 next
 edit port3
 set alias router3
 set ip 10.14.202.104/255.255.255.0
 set allowaccess https ssh ping
 set description "Link to R&D network & internet through Router2"
 end
end
```

## Configuring FortiGate unit RIP router information

With the interfaces configured, RIP can now be configured on the FortiGate units.

For each FortiGate unit the following steps will be taken:

- configure RIP version used
- redistribute static networks
- add networks serviced by RIP
- add interfaces that support RIP on the FortiGate unit

Router1 and Router4 are configured the same. Router2 and Router3 are configured the same. These routers will be grouped accordingly for the following procedures – repeat the procedures once for each FortiGate unit.

### Configure RIP settings on Router1 and Router4 - web-based manager

1. Go to *Router > Dynamic > RIP*.
2. Select 2 for *RIP Version*.
3. In *Advanced Options*, under *Redistribute* enable *Static*.
4. Leave the other *Advanced Options* at default values.
5. Enter the following networks, and select *Add* after each:
  - 10.11.0.0/255.255.0.0
  - 10.12.0.0/255.255.0.0
  - 10.14.0.0/255.255.0.0
  - 172.20.120.0/255.255.255.0

6. For interface, select *Create New* and set the following information.

---

<b>Interface</b>	port1 (internal)
<b>Send Version</b>	Both
<b>Receive Version</b>	Both
<b>Authentication</b>	None
<b>Passive Interface</b>	disabled

---

7. For interface, select *Create New* and set the following information.

---

<b>Interface</b>	port2 (router2)
<b>Send Version</b>	Both
<b>Receive Version</b>	Both
<b>Authentication</b>	None
<b>Passive Interface</b>	disabled

---

8. For interface, select *Create New* and set the following information.

---

<b>Interface</b>	port3 (router3)
<b>Send Version</b>	Both
<b>Receive Version</b>	Both
<b>Authentication</b>	None
<b>Passive Interface</b>	disabled

---

## Configure RIP settings on Router1 and Router4 - CLI

```
config router rip
 set version 2
 config interface
 edit "port1"
 set receive-version 1 2
 set send-version 1 2
 next
 edit "port2"
 set receive-version 1 2
 set send-version 1 2
 next
 edit "port3"
 set receive-version 1 2
 set send-version 1 2
 end
 config network
 edit 1
 set prefix 10.11.0.0 255.255.0.0
 next
 edit 2
 set prefix 10.12.0.0 255.255.0.0
 next
 edit 3
 set prefix 10.14.0.0 255.255.0.0
 next
 edit 4
 set prefix 172.20.120.0 255.255.255.0
 end
 config redistribute "static"
 set status enable
 end
end
```

## Configure RIP settings on Router2 and Router3- web-based manager

1. Go to *Router > Dynamic > RIP*.
2. Select 2 for *RIP Version*.
3. In *Advanced Options*, under *Redistribute* enable *Static*.
4. Leave the other *Advanced Options* at default values.
5. Enter the following networks, and select *Add* after each:
  - 10.11.0.0/255.255.0.0
  - 10.12.0.0/255.255.0.0
  - 10.14.0.0/255.255.0.0
  - 172.20.120.0/255.255.255.0

6. For interface, select *Create New* and set the following information.

---

<b>Interface</b>	port1 (internal)
<b>Send Version</b>	Both
<b>Receive Version</b>	Both
<b>Authentication</b>	None
<b>Passive Interface</b>	disabled

---

7. For interface, select *Create New* and set the following information.

---

<b>Interface</b>	port2 (router1)
<b>Send Version</b>	Both
<b>Receive Version</b>	Both
<b>Authentication</b>	None
<b>Passive Interface</b>	disabled

---

8. For interface, select *Create New* and set the following information.

---

<b>Interface</b>	port3 (router4)
<b>Send Version</b>	Both
<b>Receive Version</b>	Both
<b>Authentication</b>	None
<b>Passive Interface</b>	disabled

---

9. For interface, select *Create New* and set the following information.

---

<b>Interface</b>	port4 (ISP)
<b>Send Version</b>	Both
<b>Receive Version</b>	Both
<b>Authentication</b>	None
<b>Passive Interface</b>	disabled

---

## Configure RIP settings on Router2 and Router3- web-based manager

```
config router rip
 set version 2
 config interface
 edit "port1"
 set receive-version 1 2
 set send-version 1 2
 next
 edit "port2"
 set receive-version 1 2
 set send-version 1 2
 next
 edit "port3"
 set receive-version 1 2
 set send-version 1 2
 end
 edit "port4"
 set receive-version 1 2
 set send-version 1 2
 end
 config network
 edit 1
 set prefix 10.11.0.0 255.255.0.0
 next
 edit 2
 set prefix 10.12.0.0 255.255.0.0
 next
 edit 3
 set prefix 10.14.0.0 255.255.0.0
 next
 edit 4
 set prefix 172.20.120.0 255.255.255.0
 end
 config redistribute "static"
 set status enable
 end
end
```

## Configuring other networking devices

In this example there are two groups of other devices on the the network — internal devices, and the ISP.

The first is the internal network devices on the Sales, R&D, and Accounting networks. This includes simple static routers, computers, printers and other network devices. Once the FortiGate units are configured, the internal static routers need to be configured using the internal network IP addresses. Otherwise there should be no configuration required.

The second group of devices is the ISP. This consists of the RIP router the FortiGate routers 2 and 3 connect to. You need to contact your ISP and ensure they have your information for your

network such as the IP addresses of the connecting RIP routers, what version of RIP your network supports, and what authentication (if any) is used.

## Testing network configuration

Once the network has been configured, you need to test that it works as expected.

The two series of tests you need to run are to test the internal networks can communicate with each other, and that the internal networks can reach the internet.

Use ping, traceroute, and other networking tools to run these tests.

If you encounter problems, for troubleshooting help consult [“Troubleshooting RIP” on page 330](#).

## RIPng – RIP and IPv6

RIP next generation, or RIPng, is the version of RIP that supports IPv6.

This is an example of a typical small network configuration using RIPng routing.

Your internal R&D network is working on a project for a large international telecom company that uses IPv6. For this reason, you have to run IPv6 on your internal network and you have decided to use only IPv6 addresses.

Your network has two FortiGate units running the RIPng dynamic routing protocol. Both FortiGate units are connected to the ISP router and the internal network. This configuration provides some redundancy for the R&D internal network enabling it to reach the internet at all times.

This section includes the following topics:

- [Network layout and assumptions](#)
- [Configuring the FortiGate units system information](#)
- [Configuring RIPng on FortiGate units](#)
- [Configuring other networking devices](#)
- [Testing network configuration](#)

### Network layout and assumptions

#### Basic network layout

Your internal R&D network is working on a project for a large international telecom company that uses IPv6. For this reason, you have to run IPv6 on your internal network and you have decided to use only IPv6 addresses.

Your network has two FortiGate units running the RIPng dynamic routing protocol. Both FortiGate units are connected to the ISP router and the internal network. This configuration provides some redundancy for the R&D internal network enabling it to reach the internet at all times.

All internal computers use RIP routing, so no static routing is required. And all internal computers use IPv6 addresses.


Where possible in this example, the default values will be used or the most general settings. This is intended to provide an easier configuration that will require less troubleshooting.

In this example the routers, networks, interfaces used, and IP addresses are as follows.

**Table 13:** Rip example network topology

Network	Router	Interface & Alias	IPv6 address
R&D	Router1	port1 (internal)	2002:A0B:6565:0:0:0:0:0
		port2 (ISP)	2002:AC14:7865:0:0:0:0:0
	Router2	port1 (internal)	2002:A0B:6566:0:0:0:0:0
		port2 (ISP)	2002:AC14:7866:0:0:0:0:0

**Figure 102:**Network topology for the IPV6 RIPng example



### Assumptions

The following assumptions have been made concerning this example.

- All FortiGate units have 5.0 firmware, and are running factory default settings.
- All CLI and web-based manager navigation assumes the unit is running in NAT/Route operating mode, with VDOMs disabled.
- All FortiGate units have interfaces labelled port1 and port2 as required.
- All firewalls have been configured for each FortiGate unit to allow the required traffic to flow across interfaces.
- All network devices are support IPv6 and are running RIPng.

### Configuring the FortiGate units system information

Each FortiGate unit needs IPv6 enabled, a new hostname, and interfaces configured.



### To configure system information on Router1 - web-based manager

1. Go to *System > Dashboard > Status*.
2. For *Host name*, select *Change*.
3. Enter "Router1".
4. Go to *System > Admin > Settings*.
5. In *Display Options on GUI*, enable *IPv6*, and select *Apply*.
6. Go to *System > Network > Interfaces*.
7. Edit port1 (internal) interface.
8. Set the following information, and select *OK*.

---

<b>Alias</b>	internal
<b>IP/Netmask</b>	2002:A0B:6565::/0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Internal RnD network
<b>Administrative Status</b>	Up

---

9. Edit port2 (ISP) interface.
10. Set the following information, and select *OK*.

---

<b>Alias</b>	ISP
<b>IP/Netmask</b>	2002:AC14:7865::/0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	ISP and internet
<b>Administrative Status</b>	Up

---

## To configure system information on Router1 - CLI

```
config system global
 set hostname Router1
 set gui-ipv6 enable
end
config system interface
 edit port1
 set alias internal
 set allowaccess https ping ssh
 set description "Internal RnD network"
 config ipv6
 set ip6-address 2002:a0b:6565::/0
 end
 end
next
 edit port2
 set alias ISP
 set allowaccess https ping ssh
 set description "ISP and internet"
 config ipv6
 set ip6-address 2002:AC14:7865::
 end
 end
end
```

## To configure system information on Router2 - web-based manager

1. Go to *System > Dashboard > Status*.
2. For *Host name*, select *Change*.
3. Enter "Router2".
4. Go to *System > Admin > Settings*.
5. In *Display Options on GUI*, enable *IPv6*, and select *Apply*.
6. Go to *System > Network > Interfaces*.
7. Edit port1 (internal) interface.
8. Set the following information, and select *OK*.

<b>Alias</b>	internal
<b>IP/Netmask</b>	2002:A0B:6566::/0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Internal RnD network
<b>Administrative Status</b>	Up

9. Edit port2 (ISP) interface.
10. Set the following information, and select *OK*.

<b>Alias</b>	ISP
<b>IP/Netmask</b>	2002:AC14:7866::/0

<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	ISP and internet
<b>Administrative Status</b>	Up

### To configure system information on Router2 - CLI

```

config system global
 set hostname Router2
 set gui-ipv6 enable
end
config system interface
 edit port1
 set alias internal
 set allowaccess https ping ssh
 set description "Internal RnD network"
 config ipv6
 set ip6-address 2002:a0b:6566::/0
 end
 next
 edit port2
 set alias ISP
 set allowaccess https ping ssh
 set description "ISP and internet"
 config ipv6
 set ip6-address 2002:AC14:7866::
 end
 end
end

```

## Configuring RIPng on FortiGate units

Now that the interfaces are configured, you can configure RIPng on the FortiGate units.

There are only two networks and two interfaces to include — the internal network, and the ISP network. There is no redistribution, and no authentication. In RIPng there is no specific command to include a subnet in the RIP broadcasts. There is also no information required for the interfaces beyond including their name.

As this is a CLI only configuration, configure the ISP router and the other FortiGate unit as neighbors. This was not part of the previous example as this feature is not offered in the web-based manager. Declaring neighbors in the configuration like this will reduce the discovery traffic when the routers start up.

Since RIPng is not supported in the web-based manager, this section will only be entered in the CLI.

### To configure RIPng on Router1 - CLI

```
config router ripng
 config interface
 edit port1
 next
 edit port2
 end
 config neighbor
 edit 1
 set interface port1
 set ipv6 2002:a0b:6566::/0
 next
 edit 2
 set interface port2
 set ipv6 2002:AC14:7805::/0
 end
```

### To configure RIPng on Router2 - CLI

```
config router ripng
 config interface
 edit port1
 next
 edit port2
 end
 config neighbor
 edit 1
 set interface port1
 set ipv6 2002:a0b:6565::/0
 next
 edit 2
 set interface port2
 set ipv6 2002:AC14:7805::/0
 end
```

## Configuring other network devices

The other devices on the internal network all support IPv6, and are running RIPng where applicable. They only need to know the internal interface network addresses of the FortiGate units.

The ISP routers need to know the FortiGate unit information such as IPv6 addresses.

## Testing the configuration

In addition to normal testing of your network configuration, you must also test the IPv6 part of this example.

For troubleshooting problems with your network, see the Troubleshooting chapter.

For troubleshooting problems with RIP, see [“Troubleshooting RIP” on page 330](#).

## Testing the IPv6 RIPng information

There are some commands to use when checking that your RIPng information is correct on your network. These are useful to check on your RIPng FortiGate units on your network. Comparing the output between devices will help you understand your network better, and also track down any problems.

```
diagnose ipv6 address list
```

View the local scope IPv6 addresses used as next-hops by RIPng on the FortiGate unit.

```
diagnose ipv6 route list
```

View ipv6 addresses that are installed in the routing table.

```
get router info6 routing-table
```

View the routing table. This information is almost the same as the previous command (diagnose ipv6 route list) however it is presented in an easier to read format.

```
get router info6 rip interface external
```

View brief output on the RIP information for the interface listed. The information includes if the interface is up or down, what routing protocol is being used, and whether passive interface or split horizon are enabled.

```
get router info6 neighbor-cache list
```

View the IPv6/MAC address mapping. This also displays the interface index and name associated with the address.

# Border Gateway Protocol (BGP)

This section describes Border Gateway Protocol (BGP).

The following topics are included in this section:

- [BGP background and concepts](#)
- [Troubleshooting BGP](#)
- [Dual-homed BGP example](#)
- [Redistributing and blocking routes in BGP](#)

## BGP background and concepts

The border gateway protocol contains two distinct subsets — internal BGP (iBGP) and external BGP (eBGP). iBGP is intended for use within your own networks. eBGP is used to connect many different networks together, and is the main routing protocol for the Internet backbone. FortiGate units support iBGP, and eBGP only for communities.

The following topics are included in this section:

- [Background](#)
- [Parts and terminology of BGP](#)
- [How BGP works](#)

### Background

BGP was first used in 1989. The current version, BGP-4, was released in 1995 and is defined in RFC 1771. That RFC has since been replaced by the more recent RFC 4271. The main benefits of BGP-4 are classless inter-domain routing, and aggregate routes. BGP is the only routing protocol to use TCP for a transport protocol. Other routing protocols use UDP.

BGP makes routing decisions based on path, network policies and rulesets instead of the hop-count metric as RIP does, or cost-factor metrics as OSPF does.

BGP-4+ supports IPv6. It was introduced in RFC 2858 and RFC 2545.

BGP is the routing protocol used on the Internet. It was designed to replace the old Exterior Gateway Protocol (EGP) which had been around since 1982, and was very limited. In doing so, BGP enabled more networks to take part in the Internet backbone to effectively decentralize it and make the Internet more robust, and less dependent on a single ISP or backbone network.

### Parts and terminology of BGP

In a BGP network, there are some terms that need to be explained before going ahead. Some parts of BGP are not explained here as they are common to other dynamic routing protocols as well. When determining your network topology, note that the number of available or supported routes is not set by the configuration but depends on your FortiGate's available memory. For more information on parts of BGP that are not listed here, see [“Dynamic routing terminology” on page 313](#).

## BGP and IPv6

FortiGate units support IPv6 over BGP using the same `config router bgp` command as IPv4, but different subcommands.

The main CLI keywords have IPv6 equivalents that are identified by the “6” on the end of the keyword, such as with `config network6` or `set allowas-in6`. For more information about IPv6 BGP keywords, see the [FortiGate CLI Reference](#).

IPv6 BGP commands include:

```
config router bgp
 set activate6 {enable | disable}
 set allowas-in6 <max_num_AS_integer>
 set allowas-in-enable6 {enable | disable}
 set as-override6 {enable | disable}
 set attribute-unchanged6 [as-path] [med] [next-hop]
 set capability-default-originate6 {enable | disable}
 set capability-graceful-restart6 {enable | disable}
 set capability-orf6 {both | none | receive | send}
 set default-originate-route-map6 <routemap_str>
 set distribute-list-in6 <access-list-name_str>
 set distribute-list-out6 <access-list-name_str>
 set filter-list-in6 <aspath-list-name_str>
 set filter-list-out6 <aspath-list-name_str>
 set maximum-prefix6 <prefix_integer>
 set maximum-prefix-threshold6 <percentage_integer>
 set maximum-prefix-warning-only6 {enable | disable}
 set next-hop-self6 {enable | disable}
 set prefix-list-in6 <prefix-list-name_str>
 set prefix-list-out6 <prefix-list-name_str>
 set remove-private-as6 {enable | disable}
 set route-map-in6 <routemap-name_str>
 set route-map-out6 <routemap-name_str>
 set route-reflector-client6 {enable | disable}
 set route-server-client6 {enable | disable}
 set send-community6 {both | disable | extended | standard}
 set soft-reconfiguration6 {enable | disable}
 set unsuppress-map6 <route-map-name_str>
 config network6
 config redistribute6
end
```

## Roles of routers in BGP networks

Dynamic routing has a number of different roles routers can fill such as those covered in [“Dynamic routing terminology” on page 313](#). BGP has a number of custom roles that routers can fill. These include:

- Speaker routers
- Peer routers or neighbors
- Route reflectors (RR)

### Speaker routers

Any router configured for BGP is considered a BGP speaker. This means that a speaker router advertises BGP routes to its peers.

Any routers on the network that are not speaker routers, are not treated as BGP routers.

### Peer routers or neighbors

In a BGP network, all neighboring BGP routers or peer routers are routers that are connected to your FortiGate unit. Your FortiGate unit learns about all other routers through these peers.

You need to manually configure BGP peers on your FortiGate unit as neighbors. Otherwise these routers will not be seen as peers, but instead as simply other routers on the network that don't support BGP. You can optionally use MD5 authentication to password protect BGP sessions with those neighbors. (see RFC 2385).

You can configure up to 1000 BGP neighbors on your FortiGate unit. You can clear all or some BGP neighbor connections (sessions) using the `execute router clear bgp` command.

For example, if you have 10 routes in the BGP routing table and you want to clear the specific route to IP address 10.10.10.1, enter the command:

```
execute router clear bgp ip 10.10.10.1
```

To remove all routes for AS number 650001, enter the command:


```
execute router clear bgp as 650001
```

To remove route flap dampening information for the 10.10.0.0/16 subnet, enter the command:

```
execute router clear bgp dampening 10.10.0.0/16
```

In Figure 1, Router A is directly connected to five other routers in a network that contains 12 routers overall. These routers, the ones in the blue circle, are Router A's peers or neighbors.

**Figure 103:**Router A and its 5 peer routers



Router A's peer routers

As a minimum, when configuring BGP neighbors you must enter their IP address, and the AS number (*remote-as*). This is all the information the web-based manager interface allows you to enter for a neighbor.



The BGP commands related to neighbors are quite extensive and include:

```
config router bgp
 config neighbor
 edit <neighbor_address_ipv4>
 set activate {enable | disable}
 set advertisement-interval <seconds_integer>
 set allowas-in <max_num_AS_integer>
 set allowas-in-enable {enable | disable}
 set as-override {enable | disable}
 set attribute-unchanged [as-path] [med] [next-hop]
 set bfd {enable | disable}
 set capability-default-originate {enable | disable}
 set capability-dynamic {enable | disable}
 set capability-graceful-restart {enable | disable}
 set capability-orf {both | none | receive | send}
 set capability-route-refresh {enable | disable}
 set connect-timer <seconds_integer>
 set description <text_str>
 set distribute-list-in <access-list-name_str>
 set distribute-list-out <access-list-name_str>
 set dont-capability-negotiate {enable | disable}
 set ebgp-enforce-multihop {enable | disable}
 set ebgp-multihop {enable | disable}
 set ebgp-multihop-ttl <seconds_integer>
 set filter-list-in <aspath-list-name_str>
 set filter-list-out <aspath-list-name_str>
 set holdtime-timer <seconds_integer>
 set interface <interface-name_str>
 set keep-alive-timer <seconds_integer>
 set maximum-prefix <prefix_integer>
 set maximum-prefix-threshold <percentage_integer>
 set maximum-prefix-warning-only {enable | disable}
 set next-hop-self {enable | disable}
 set passive {enable | disable}
 set password <string>
 set prefix-list-in <prefix-list-name_str>
 set prefix-list-out <prefix-list-name_str>
 set remote-as <id_integer>
 set remove-private-as {enable | disable}
 set retain-stale-time <seconds_integer>
 set route-map-in <routemap-name_str>
 set route-map-out <routemap-name_str>
 set route-reflector-client {enable | disable}
 set route-server-client {enable | disable}
 set send-community {both | disable | extended | standard}
 set shutdown {enable | disable}
 set soft-reconfiguration {enable | disable}
 set strict-capability-match {enable | disable}
```

```
 set unsuppress-map <route-map-name_str>
 set update-source <interface-name_str>
 set weight <weight_integer>
 end
end
end
```

### **Route reflectors (RR)**

Route reflectors in BGP concentrate route updates so other routers need only talk to the route reflectors to get all the updates. This results in smaller routing tables, fewer connections between routers, faster responses to network topology changes, and less administration bandwidth. BGP route reflectors are defined in RFC 1966.


In a BGP route reflector configuration, the AS is divided into different clusters that each include client and reflector routers. The client routers supply the reflector routers with the client's route updates. The reflectors pass this information along to other route reflectors and border routers. Only the reflectors need to be configured, not the clients — the clients will find the closest reflector and communicate with it automatically. The reflectors communicate with each other as peers. FortiGate units can be configured as either reflectors or clients.

Since route reflectors are processing more than the client routers, the reflectors should have more resources to handle the extra workload.

Smaller networks running BGP typically don't require route reflectors (RR). However, RR is a useful feature for large companies, where their AS may include 100 routers or more. For example, for a full mesh 20 router configuration within an AS, there would have to be 190 unique BGP sessions — just for routing updates within the AS. The number of sessions jumps to 435 sessions for just 30 routers, or 4950 sessions for 100 routers. From these numbers, it's plain that updating this many sessions will quickly consume the limited bandwidth and processing resources of the routers involved.

The following diagram illustrates how route reflectors can improve the situation when only six routers are involved. The AS without route reflectors requires 15 sessions between the routers. In the AS with route reflectors, the two route reflectors receive route updates from the reflector clients (unlabeled routers in the diagram) in their cluster as well as other route reflectors and pass them on to the border router. The RR configuration only requires six sessions. This example shows a reduction of 60% in the number of required sessions.

**Figure 104:** Required sessions within an AS with and without route reflectors



The BGP commands related to route reflectors includes:

```
config router bgp
 config neighbor
 set route-reflector-client {enable | disable}
 set route-server-client {enable | disable}
 end
end
```

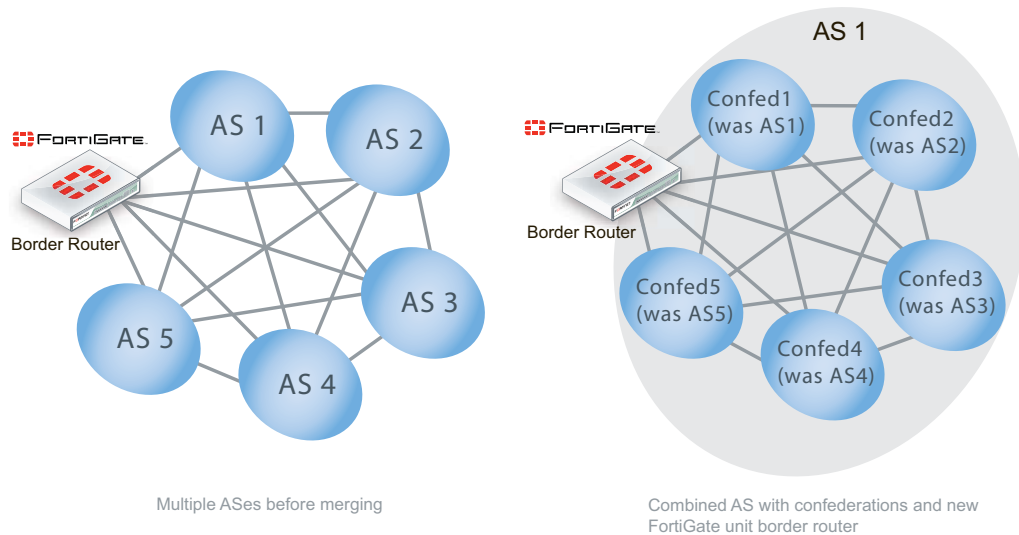
## Confederations

Confederations were introduced to reduce the number of BGP advertisements on a segment of the network, and reduce the size of the routing tables. Confederations essentially break up an AS into smaller units. Confederations are defined in RFC 3065 and RFC 1965.

Within a confederation, all routers communicate with each other in a full mesh arrangement. Communications between confederations is more like inter-AS communications in that many of the attributes are changed as they would be for BGP communications leaving the AS, or eBGP.

Confederations are useful when merging ASs. Each AS being merged can easily become a confederation, requiring few changes. Any additional permanent changes can then be implemented over time as required. The figure below shows the group of ASs before merging, and the corresponding confederations afterward as part of the single AS with the addition of a new border router. It should be noted that after merging if the border router becomes a route reflector, then each confederation only needs to communicate with one other router, instead of five others.

**Figure 105:**AS merging using confederations



Confederations and route reflectors perform similar functions — they both sub-divide large ASes for more efficient operation. They differ in that route reflector clusters can include routers that are not members of a cluster, where routers in a confederation must belong to that confederation. Also, confederations place their confederation numbers in the AS\_PATH attribute making it easier to trace.

It is important to note that while confederations essentially create sub-ASs, all the confederations within an AS appear as a single AS to external ASs.

Confederation related BGP commands include:

```
config router bgp
 set confederation-identifier <peerid_integer>
end
```

### Network Layer Reachability Information (NLRI)

Network Layer Reachability Information (NLRI) is unique to BGP-4. It is sent as part of the update messages sent between BGP routers, and contains information necessary to supernet, or aggregate route, information. The NLRI includes the length and prefix that when combined are the address of the aggregated routes referred to.

There is only one NLRI entry per BGP update message.

### BGP attributes

Each route in a BGP network has a set of attributes associated with it. These attributes define the route, and are modified as required along the route.

BGP can work well with mostly default settings, but if you are going to change settings you need to understand the roles of each attribute and how they affect those settings.

The BGP attributes include:

<b>AS_PATH</b>	A list of ASes a route has passed through. See <a href="#">“AS_PATH” on page 365</a> .
<b>MULTI_EXIT_DESC (MED)</b>	Which router to use to exit an AS with more than one external connection. See <a href="#">“MULTI_EXIT_DESC (MED)” on page 366</a> .

<b>COMMUNITY</b>	Used to apply attributes to a group of routes. See <a href="#">“COMMUNITY” on page 366.</a>
<b>NEXT_HOP</b>	Where the IP packets should be forwarded to, like a gateway in static routing. See <a href="#">“NEXT_HOP” on page 366.</a>
<b>ATOMIC_AGGREGATE</b>	Used when routes have been summarized to tell downstream routers not to de-aggregate the route. See <a href="#">“ATOMIC_AGGREGATE” on page 367.</a>
<b>ORIGIN</b>	Used to determine if the route is from the local AS or not. See <a href="#">“ORIGIN” on page 367.</a>
<b>LOCAL_PREF</b>	Used only within an AS to select the best route to a location (like MED)



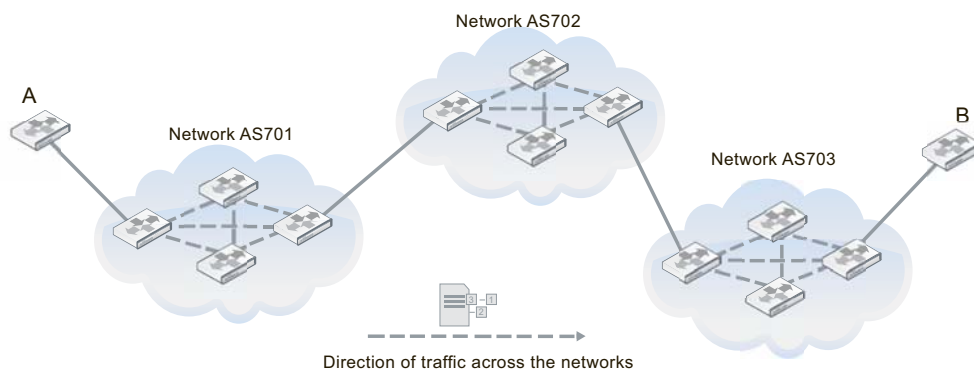
Inbound policies on FortiGate units can change the NEXT-HOP, LOCAL-PREF, MED and AS-PATH attributes of an internal BGP (iBGP) route for its local route selection purposes. However, outbound policies on the unit cannot affect these attributes.

## AS\_PATH

AS\_PATH is the BGP attribute that keeps track of each AS a route advertisement has passed through. AS\_PATH is used by confederations and by exterior BGP (EBGP) to help prevent routing loops. A router knows there is a loop if it receives an AS\_PATH with that routers AS in it. The figure below shows the route between router A and router B. The AS\_PATH from A to B would read 701,702,703 for each AS the route passes through.

As of the start of 2010, the industry upgraded from 2-byte to 4-byte AS\_PATHs. This upgrade was due to the imminent exhaustion of 2-byte AS\_PATH numbers. FortiOS supports 4-byte AS\_PATHs in its BGP implementation.

**Figure 106:**AS\_PATH of 701,702, 703 between routers A and B



The BGP commands related to AS\_PATH include:

```
config router bgp
 set bestpath-as-path-ignore {enable | disable}
end
```

## MULTI\_EXIT\_DESC (MED)

BGP AS systems can have one or more routers that connect them to other ASes. For ASes with more than one connecting router, the Multi-Exit Discriminator (MED) lists which router is best to use when leaving the AS. The MED is based on attributes such as delay. It is a recommendation only, as some networks may have different priorities.

BGP updates advertise the best path to a destination network. When the FortiGate unit receives a BGP update, the FortiGate unit examines the Multi-Exit Discriminator (MED) attribute of potential routes to determine the best path to a destination network before recording the path in the local FortiGate unit routing table.

FortiGate units have the option to treat any routes without an MED attribute as the worst possible routing choice. This can be useful because a lack of MED information is a lack of routing information which can be suspicious — possibly a hacking attempt or an attack on the network. At best it signifies an unreliable route to select.

The BGP commands related to MED include:

```
config router bgp
 set always-compare-med {enable | disable}
 set bestpath-med-confed {enable | disable}
 set bestpath-med-missing-as-worst {enable | disable}
 set deterministic-med {enable | disable}
 config neighbor
 set attribute-unchanged [as-path] [med] [next-hop]
 end
end
```

## COMMUNITY

A community is a group of routes that have the same routing policies applied to them. This saves time and resources. A community is defined by the COMMUNITY attribute of a BGP route.

The FortiGate unit can set the COMMUNITY attribute of a route to assign the route to predefined paths (see RFC 1997). The FortiGate unit can examine the COMMUNITY attribute of learned routes to perform local filtering and/or redistribution.

The BGP commands related to COMMUNITY include:

```
config router bgp
 set send-community {both | disable | extended | standard}
end
```

## NEXT\_HOP

The NEXT\_HOP attribute says what IP address the packets should be forwarded to next. Each time the route is advertised, this value is updated. The NEXT\_HOP attribute is much like a gateway in static routing.

FortiGate units allow you to change the advertising of the FortiGate unit's IP address (instead of the neighbor's IP address) in the NEXT\_HOP information that is sent to IBGP peers. This is changed with the config neighbor, set next-hop-self command.

The BGP commands related to NEXT\_HOP include:

```
config router bgp
 config neighbor
 set attribute-unchanged [as-path] [med] [next-hop]
 set next-hop-self {enable | disable}
 end
end
```

## ATOMIC\_AGGREGATE

The ATOMIC\_AGGREGATE attribute is used when routes have been summarized. It indicates which AS and which router summarize the routes. It also tells downstream routers not to de-aggregate the route. Summarized routes are routes with similar information that have been combined, or aggregated, into one route that is easier to send in updates. When it reaches its destination, the summarized routes are split back up into the individual routes.

Your FortiGate unit doesn't specifically set this attribute in the BGP router command, but it is used in the route map command.

The commands related to ATOMIC\_AGGREGATE include:

```
config router route-map
 edit <route_map_name>
 config rule
 edit <route_map_rule_id>
 set set-aggregator-as <id_integer>
 set set-aggregator-ip <address_ipv4>
 set set-atomic-aggregate {enable | disable}
 end
 end
 end
end
```

## ORIGIN

The ORIGIN attribute records where the route came from. The options can be IBGP, EBGP, or incomplete. This information is important because internal routes (IBGP) are by default higher priority than external routes (EBGP). However incomplete ORIGINS are the lowest priority of the three.

The commands related to ORIGIN include:

```
config router route-map
 edit <route_map_name>
 set comments <string>
 config rule
 edit <route_map_rule_id>
 set match-origin {egp | igp | incomplete | none}
 end
 end
 end
end
```

## How BGP works

BGP is a link-state routing protocol and keeps link-state information about the status of each network link it has connected. A BGP router receives information from its peer routers that have

been defined as neighbors. BGP routers listen for updates from these configured neighboring routers on TCP port 179.

A BGP router is a finite state machine with six various states for each connection. As two BGP routers discover each other, and establish a connection they go from the idle state, through the various states until they reach the established state. An error can cause the connection to be dropped and the state of the router to be reset to either active or idle. These errors can be caused by: TCP port 179 not being open, a random TCP port above port 1023 not being open, the peer address being incorrect, or the AS number being incorrect.

When BGP routers start a connection, they negotiate which (if any) optional features will be used such as multiprotocol extensions that can include IPv6 and VPNs.

## IBGP versus EBGP

When you read about BGP, often you see EBGP or IBGP mentioned. These are both BGP routing, but BGP used in different roles. Exterior BGP (EBGP) involves packets crossing multiple autonomous systems (ASes) where interior BGP (IBGP) involves packets that stay within a single AS. For example the AS\_PATH attribute is only useful for EBGP where routes pass through multiple ASes.

These two modes are important because some features of BGP are only used for one of EBGP or IBGP. For example confederations are used in EBGP, and route reflectors are only used in IBGP. Also routes learned from IBGP have priority over EBGP learned routes.

FortiGate units have some commands specific to EBGP. These include:

- automatically resetting the session information to external peers if the connection goes down — `set fast-external-failover {enable | disable}`
- setting an administrative distance for all routes learned from external peers (must also configure local and internal distances if this is set) — `set distance-external <distance_integer>`
- enforcing EBGP multihops and their TTL (number of hops) — `set ebgp-enforce-multihop {enable | disable}` and `set ebgp-multihop-ttl <seconds_integer>`

## BGP path determination — which route to use

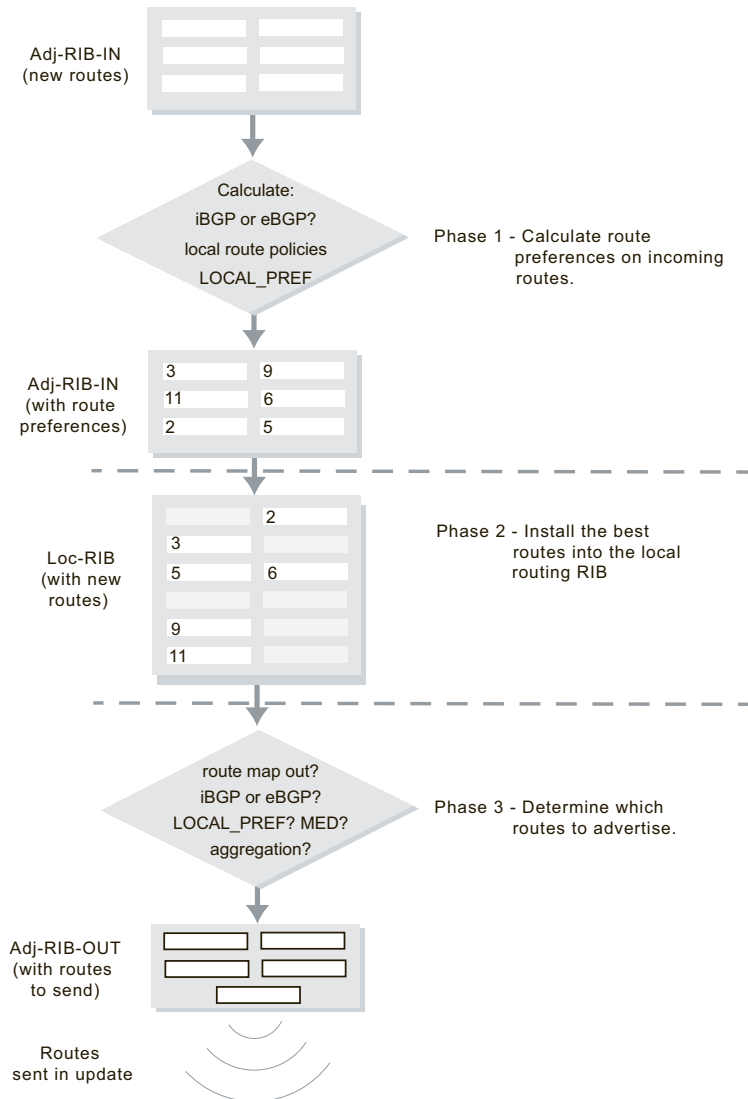
Firstly, recall that the number of available or supported routes is not set by the configuration but depends on your FortiGate's available memory. All learned routes and their attributes come into the BGP router in raw form. Before routes are installed in the routing table or are advertised to other routers, three levels of decisions must be made.

The three phases of BGP best path determination do not change. However, some manufacturers have added more information to the process, such as Cisco's WEIGHT attribute to enable an administrator to force one route's selection over another.

There is one Adj-RIB-IN and Adj-RIB-OUT for each configured neighbor. They are updated when the FortiGate unit receives BGP updates, or when the FortiGate unit sends out BGP updates.



**Figure 107:**Three phases of BGP routing decision



### Decision phase 1

At this phase, the decision is to calculate how preferred each route and its NRI are the Adjacent Routing Information Base Incoming (Adj-RIBs-In) compared to the other routes. For internal routes (IBGP), policy information or LOCAL\_PREF is used. For external peer learned routes, it is based strictly on policy. These rules set up a list of which routes are most preferred going into Phase 2.

### Decision phase 2

Phase 2 involves installing the best route to each destination into the local Routing Information Base (Loc-RIB). Effectively, the Loc-RIB is the master routing table. Each route from Phase 1 has their NEXT\_HOP checked to ensure the destination is reachable. If it is reachable, the

AS\_PATH is checked for loops. After that, routes are installed based on the following decision process:

- If there is only one route to a location, it is installed.
- If multiple routes to the same location, use the most preferred route from Level 1.
- If there is a tie, break the tie based on the following in descending order of importance: shortest AS\_PATH, smallest ORIGIN number, smallest MED, EBGP over IBGP, smallest metric or cost for reaching the NEXT\_HOP, BGP identifier, and lowest IP address.

Note that the new routes that are installed into the Loc-RIB are in addition to any existing routes in the table. Once Phase 2 is completed the Loc-RIB will consist of the best of both the new and older routes.

### Decision phase 3

Phase 3 is route distribution or dissemination. This is the process of deciding which routes the router will advertise. If there is any route aggregation or summarizing, it happens here. Also any route filtering from route maps happens here.

Once Phase 3 is complete, an update can be sent out to update the neighbor of new routes.

### Aggregate routes and addresses

BGP4 allows classless routing, which uses netmasks as well as IP addresses. This classless routing enables the configuration of aggregate routes by stating the address bits the aggregated addresses have in common. For more information, see [“Aggregated routes and addresses” on page 313](#).

The ATOMIC\_AGGREGATE attribute informs routers that the route has been aggregated, and should not be de-aggregated. An associated AGGREGATOR attribute include the information about the router that did the aggregating including its AS.

The BGP commands associated with aggregate routes and addresses are:

```
config router bgp
 config aggregate-address
 edit <aggr_addr_id>
 set as-set {enable | disable}
 set prefix <address_ipv4mask>
 set summary-only {enable | disable}
 end
 config aggregate-address6
 edit <aggr_addr_id>
 set as-set {enable | disable}
 set prefix6 <address_ipv6mask>
 set summary-only {enable | disable}
 end
 end
```

## Troubleshooting BGP

There are some features in BGP that are used to deal with problems that may arise. Typically the problems with a BGP network that has been configured, involve routes going offline frequently. This is called route flap and causes problems for the routers using that route.

## Clearing routing table entries

To see if a new route is being properly added to the routing table, you can clear all or some BGP neighbor connections (sessions) using the `execute router clear bgp` command.

For example, if you have 10 routes in the BGP routing table and you want to clear the specific route to IP address 10.10.10.1, enter the command:

```
execute router clear bgp ip 10.10.10.1
```

To remove all routes for AS number 650001, enter the command:

```
execute router clear bgp as 650001
```

## Route flap

When routers or hardware along a route go offline and back online that is called a route flap. Flapping is the term if these outages continue, especially if they occur frequently.

Route flap is a problem in BGP because each time a peer or a route goes down, all the peer routers that are connected to that out-of-service router advertise the change in their routing tables which creates a lot of administration traffic on the network. And the same traffic happens again when that router comes back online. If the problem is something like a faulty network cable that wobbles on and offline every 10 seconds, there could easily be overwhelming amounts of routing updates sent out unnecessarily.

Another possible reason for route flap occurs with multiple FortiGate units in HA mode. When an HA cluster fails over to the secondary unit, other routers on the network may see the HA cluster as being offline resulting in route flap. While this doesn't occur often, or more than once at a time, it can still result in an interruption in traffic which is unpleasant for network users. The easy solution for this problem is to increase the timers on the HA cluster, such as TTL timers, so they do not expire during the failover process. Also configuring graceful restart on the HA cluster will help with a smooth failover.

The first method of dealing with route flap should be to check your hardware. If a cable is loose or bad, it can easily be replaced and eliminate the problem. If an interface on the router is bad, either avoid using that interface or swap in a functioning router. If the power source is bad on a router, either replace the power supply or use a power conditioning backup power supply. These quick and easy fixes can save you from configuring more complex BGP options. However if the route flap is from another source, configuring BGP to deal with the outages will ensure your network users uninterrupted service.

Some methods of dealing with route flap in BGP include:

- [Holddown timer](#)
- [Dampening](#)
- [Graceful restart](#)
- [Bi-directional forwarding detection \(BFD\)](#)

### Holddown timer

The first line of defence to a flapping route is the hold down timer. This timer reduces how frequently a route going down will cause a routing update to be broadcast.

Once activated, the holddown timer won't allow the FortiGate unit to accept any changes to that route for the duration of the timer. If the route flaps five times during the timer period, only the first outage will be recognized by the FortiGate unit — for the duration of the other outages there will be no changes because the Fortigate unit is essentially treating this router as down. After the timer expires, if the route is still flapping it will happen all over again.

Even if the route isn't flapping — if it goes down, comes up, and stays back up — the timer still counts down and the route is ignored for the duration of the timer. In this situation the route will be seen as down longer than it really is, but there will be only the one set of route updates. This is not a problem in normal operation because updates are not frequent.

Also the potential for a route to be treated as down when it is really up can be viewed as a robustness feature. Typically you do not want most of your traffic being routed over an unreliable route. So if there is route flap going on, it is best to avoid that route if you can. This is enforced by the holddown timer.

### How to configure the holddown timer

There are three different route flapping situations that can occur: the route goes up and down frequently, the route goes down and back up once over a long period of time, or the route goes down and stays down for a long period of time. These can all be handled using the holddown timer.

For example, your network has two routes that you want to set the holddown timer for. One is your main route ( to 10.12.101.4) that all your Internet traffic goes through, and it can't be down for long if its down. The second is a low speed connection to a custom network that is used infrequently ( to 10.13.101.4). The holddown timer for the main route should be fairly short, lets say 60 seconds instead of the default 180 seconds. The second route timer can be left at the default or even longer since it is rarely used. In your BGP configuration this looks like:

```
config router bgp
 config neighbor
 edit 10.12.101.4
 set holddown-timer 60
 next
 edit 10.13.101.4
 set holddown-timer 180
 next
end
end
```

### Dampening

Dampening is a method used to limit the amount of network problems due to flapping routes. With dampening the flapping still occurs, but the peer routers pay less and less attention to that route as it flaps more often. One flap doesn't start dampening, but the second starts a timer where the router will not use that route — it is considered unstable. If the route flaps again before the timer expires, the timer continues to increase. There is a period of time called the reachability half-life after which a route flap will only be suppressed for half the time. This half-life comes into effect when a route has been stable for a while but not long enough to clear all the dampening completely. For the flapping route to be included in the routing table again, the suppression time must expire.

If the route flapping was temporary, you can clear the flapping or dampening from the FortiGate units cache by using one of the `execute router clear bgp` commands:

```
execute router clear bgp dampening {<ip_address> | <ip/netmask>}
```

or

```
execute router clear bgp flap-statistics {<ip> | <ip/netmask>}
```

For example, to remove route flap dampening information for the 10.10.0.0/16 subnet, enter the command:

```
execute router clear bgp dampening 10.10.0.0/16
```

The BGP commands related to route dampening are:

```
config router bgp
 set dampening {enable | disable}
 set dampening-max-suppress-time <minutes_integer>
 set dampening-reachability-half-life <minutes_integer>
 set dampening-reuse <reuse_integer>
 set dampening-route-map <routemap-name_str>
 set dampening-suppress <limit_integer>
 set dampening-unreachability-half-life <minutes_integer>
end
```

## Graceful restart

BGP4 has the capability to gracefully restart.

In some situations, route flap is caused by routers that appear to be offline but the hardware portion of the router (control plane) can continue to function normally. One example of this is when some software is restarting or being upgraded, but the hardware can still function normally.

Graceful restart is best used for these situations where routing will not be interrupted, but the router is unresponsive to routing update advertisements. Graceful restart does not have to be supported by all routers in a network, but the network will benefit when more routers support it.



FortiGate HA clusters can benefit from graceful restart. When a failover takes place, the HA cluster will advertise it is going offline, and will not appear as a route flap. It will also enable the new HA main unit to come online with an updated and usable routing table — if there is a flap the HA cluster routing table will be out of date.

---

For example, your FortiGate unit is one of four BGP routers that send updates to each other. Any of those routers may support graceful starting—when a router plans to go offline, it will send out a message to its neighbors how long it expects to be before being back online. That way its neighbor routers don't remove it from their routing tables. However if that router isn't back online when expected, the routers will mark it offline. This prevents routing flap and its associated problems.

### Scheduled time offline

Graceful restart is a means for a router to advertise it is going to have a scheduled shutdown for a very short period of time. When neighboring routers receive this notice, they will not remove that router from their routing table until after a set time elapses. During that time if the router comes back online, everything continues to function as normal. If that router remains offline longer than expected, then the neighboring routers will update their routing tables as they assume that router will be offline for a long time.

FortiGate units support both graceful restart of their own BGP routing software, and also neighboring BGP routers.

For example, if a neighbor of your FortiGate unit, with an IP address of 172.20.120.120, supports graceful restart, enter the command:

```
config router bgp
 config neighbor
 edit 172.20.120.120
 set capability-graceful-restart enable
 end
end
```

If you want to configure graceful restart on your FortiGate unit where you expect the Fortigate unit to be offline for no more than 2 minutes, and after 3 minutes the BGP network should consider the FortiGate unit offline, enter the command:

```
config router bgp
 set graceful-restart enable
 set graceful-restart-time 120
 set graceful-stalepath-time 180
end
```

The BGP commands related to BGP graceful restart are:

```
config router bgp
 set graceful-restart { disable | enable }
 set graceful-restart-time <seconds_integer>
 set graceful-stalepath-time <seconds_integer>
 set graceful-update-delay <seconds_integer>
 config neighbor
 set capability-graceful-restart { enable | disable }
 end
end

execute router restart
```

Before the restart, the router sends its peers a message to say it is restarting. The peers mark all the restarting router's routes as stale, but they continue to use the routes. The peers assume the router will restart and check its routes and take care of them if needed after the restart is complete. The peers also know what services the restarting router can maintain during its restart. After the router completes the restart, the router sends its peers a message to say it is done restarting.

## Bi-directional forwarding detection (BFD)

Bi-directional Forwarding Detection (BFD) is a protocol used to quickly locate hardware failures in the network. Routers running BFD communicate with each other, and if a timer runs out on a connection then that router is declared down. BFD then communicates this information to the routing protocol and the routing information is updated.

While BGP can detect route failures, BFD can be configured to detect these failures more quickly allowing faster responses and improved convergence. This can be balanced with the bandwidth BFD uses in its frequent route checking.

### Configurable granularity

BFD can run on the entire FortiGate unit, selected interfaces, or on BGP for all configured interfaces. The hierarchy allows each lower level to override the upper level's BFD setting. For example, if BFD was enabled for the FortiGate unit, it could be disabled only for a single

interface or for BGP. For information about FortiGate-wide BFD options, see config system settings in the [FortiGate CLI Reference](#).

BFD can only be configured through the CLI.

The BGP commands related to BFD are:

```
config system {setting | interface}
 set bfd {enable | disable | global}
 set bfd-desired-mix-tx <milliseconds>
 set bfd-detect-mult <multiplier>
 set bfd-required-mix-rx <milliseconds>
 set bfd-dont-enforce-src-port {enable | disable}

config router bgp
 config neighbor
 edit <neighbor_address_ipv4>
 set bfd {enable | disable}
 end
 end

get router info bfd neighbor
execute router clear bfd session <src_ipv4> <dst_ipv4> <interface>
```

The `config system` commands allow you to configure whether BFD is enabled in a particular unit/vdom or individual interface, and how often the interface requires sending and receiving of BFD information.

The `config router bgp` commands allow you to set the addresses of the neighbor units that are also running BFD. Both units must be configured with BFD in order to make use of it.

## Dual-homed BGP example

This is an example of a small network that uses BGP routing connections to two ISPs. This is a common configuration for companies that need redundant connections to the Internet for their business.

This configuration is for a small company connected to two ISPs. The company has one main office, the Head Office, and uses static routing for internal routing on that network.

Both ISPs use BGP routing, and connect to the Internet directly. They want the company to connect to the ISP networks using BGP. They also use graceful restart to prevent unneeded updates, and use smaller timer values to detect network failures faster.

As can be expected, the company wants to keep their BGP configuration relatively simple and easy to manage. The current configuration has only 3 routers to worry about — the 2 ISP border routers, and the FortiGate unit. This means the FortiGate unit will only have two neighbour routers to configure.

This configuration has the added benefit of being easy to expand if the Company wants to add a remote office in the future.

To keep the configuration simple, the Company is allowing only HTTP, HTTPS, FTP, and DNS traffic out of the local network. This will allow employees access to the Internet and their web-mail.

This section includes the following topics:

- [Network layout and assumptions](#)
- [Configuring the FortiGate unit](#)
- [Configuring other networking devices](#)
- [Testing this configuration](#)

## Why dual home?

Dual homing means having two separate independent connections to the Internet. Servers in this configuration have also been called bastion hosts and can include DNS servers which require multiple connections.

Benefits of dual homing can include:

- redundant Internet connection that essentially never fails
- faster connections through one ISP or the other for some destinations, such as other clients of those ISPs
- load balancing traffic to your Company network
- easier to enable more traffic through two connections than upgrading one connection to bigger bandwidth
- easier to create protection policies for different traffic through a specific ISP

Some companies require reliable internet access at all times as part of their business. Consider a doctor operating remotely who has their Internet connection fail — the consequences could easily be life or death.

Dual homing is extra expense for the second ISP connection, and more work to configure and maintain the more complex network topology.

## Potential dual homing issues

BGP comes with load balancing issues, and dual homing is the same category. BGP does not inherently deal well with load balancing, or getting default routes through BGP. Ideally one connect may be best for certain destinations, but it may not have that traffic routed to it making the load balancing less than perfect. This kind of fine tuning can be very time consuming, and usually results in a best effort situation.

When dual homing is not configured properly, your network may become a link between your ISPs and result in very high traffic between the ISPs that does not originate from your network. The problems with this situation are that your traffic may not have the bandwidth it needs, and you will be paying for a large volume of traffic that is not yours. This problem can be solved by not broadcasting or redistributing BGP routes between the ISPs.

If you learn your default routes from the ISPs in this example, you may run into an asymmetric routing problem where your traffic loops out one ISP and back to you through the other ISP. If you think this may be happening you can turn on asymmetric routing on the FortiGate unit (config system settings, set asymmetric enable) to verify that really is the problem. Turn this feature off once this is established since it disables many features on the FortiGate by disabling stateful inspection. Solutions for this problem can include using static routes for default routes instead of learning them through BGP, or configuring VDOMs on your FortiGate unit to provide a slightly different path back that is not a true loop.

## Network layout and assumptions


The network layout for the basic BGP example involves the company network being connected to both ISPs as shown below. In this configuration the FortiGate unit is the BGP border router between the Company AS, ISP1's AS, and ISP2's AS.



The components of the layout include:

- The Company AS (AS number 1) is connected to ISP1 and ISP2 through the FortiGate unit.
- The Company has one internal network — the Head Office network at 10.11.101.0/24.
- The FortiGate unit internal interface is on the the Company internal network with an IP address of 10.11.101.110.
- The FortiGate unit external1 interface is connected to ISP1's network with an IP address of 172.20.111.5, an address supplied by the ISP.
- The FortiGate unit external2 interface is connected to IPS2's network with an IP address of 172.20.222.5, an address supplied by the ISP.
- ISP1 AS has an AS number of 650001, and ISP2 has an AS number of 650002.
- Both ISPs are connected to the Internet.
- The ISP1 border router is a neighbor (peer) of the FortiGate unit. It has an address of 172.21.111.4.
- The ISP2 border router is a neighbor (peer) of the FortiGate unit. It has an address of 172.22.222.4.
- Apart from graceful restart, and shorter timers (holdtimer, and keepalive) default settings are to be used whenever possible.

**Figure 108:**Basic BGP network topology



### Assumptions

The basic BGP configuration procedure follows these assumptions:

- ISP1 is the preferred route, and ISP2 is the secondary route
- all basic configuration can be completed in both GUI and CLI
- only one AS is used for the Company

For these reasons this example configuration does not include:

- [Bi-directional forwarding detection \(BFD\)](#)
- [Route maps](#)
- [Access lists](#)
- changing redistribution defaults — make link when example is set up
- IPv6

For more information on these features, see the corresponding section.

## Configuring the FortiGate unit

In this topology, the FortiGate unit is the link between the Company Network and the ISP network. The FortiGate unit is the only BGP router on the Company Network, but there is at least one other BGP router on the ISP Network — there may be more but we don't have that information.

As mentioned in the general configuration steps, the ISP must be notified of the Company's BGP router configuration when complete as it will need to add the FortiGate BGP router as a neighbor router on its domain. This step is required for the FortiGate unit to receive BGP routing updates from the ISP network and outside networks.

If the ISP has any special BGP features enabled such as graceful restart, or route dampening that should be determined up front so those features can be enabled on the FortiGate unit.

### To configure the FortiGate unit as a BGP router

1. [Configure interfaces and default routes](#)
2. [Configure firewall services, addresses, and policies](#)
3. [Set the FortiGate BGP information](#)
4. [Add the internal network to the AS](#)
5. [Additional FortiGate BGP configuration](#)

### Configure interfaces and default routes

The FortiGate unit is connected to three networks — Company Network on the internal interface, ISP1 Network on external1 interface, and ISP2 on external2 interface.

This example uses basic interface settings. Check with your ISP to determine if additional settings are required such as setting the maximum MTU size, or if gateway detection is supported.

High end FortiGate units do not have interfaces labeled Internal, or External. Instead, for clarity's sake, we are using the alias feature to name interfaces for these roles.

Default routes to both external interfaces are configured here as well. Both are needed in case one goes offline. ISP1 is the primary connection and has a smaller administrative distance so it will be preferred over ISP2. Both distances are set low so they will be preferred over any learned routes.

### To configure the FortiGate interfaces - web-based manager

1. Go to *System > Network > Interface*.
2. Edit port 1 (internal) interface.

3. Set the following information, and select *OK*.

---

<b>Alias</b>	internal
<b>IP/Netmask</b>	10.11.101.110/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Company internal network
<b>Administrative Status</b>	Up

---

4. Edit port 2 (external1) interface.  
5. Set the following information, and select *OK*.

---

<b>Alias</b>	external1
<b>IP/Netmask</b>	172.21.111.5/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	ISP1 External BGP network
<b>Administrative Status</b>	Up

---

6. Edit port 3 (external2) interface.  
7. Set the following information, and select *OK*.

---

<b>Alias</b>	external2
<b>IP/Netmask</b>	172.22.222.5/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	ISP2 External BGP network
<b>Administrative Status</b>	Up

---

## To configure the FortiGate interfaces - CLI

```
config system interface
 edit port1
 set alias internal
 set ip 10.11.101.110 255.255.255.0
 set allowaccess http https ssh
 set description "Company internal network"
 set status up
 next
 edit port2
 set alias external1
 set ip 172.21.111.5 255.255.255.0
 set allowaccess https ssh
 set description "ISP1 External BGP network"
 set status up
 next
 edit port3
 set alias external2
 set ip 172.22.222.5 255.255.255.0
 set allowaccess https ssh
 set description "ISP2 External BGP network"
 set status up
 next
end
```

## To configure default routes for both ISPs - web-based manager

1. Go to *Router > Static > Static Routes*.
2. Delete any existing routes with a IP/Mask of address of 0.0.0.0/0.0.0.0
3. Select *Create New*, and set the following information.

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Device</b>	port2
<b>Gateway</b>	172.21.111.5
<b>Distance</b>	10

4. Select OK.
5. Select *Create New*, and set the following information.

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Device</b>	port3
<b>Gateway</b>	172.22.222.5
<b>Distance</b>	15

6. Select OK.

## To configure default routes for both ISPs - CLI

```
config router static
 edit 1
 set device "port2"
 set distance 10
 set gateway 172.21.111.5
 next
 edit 2
 set device "port3"
 set distance 15
 set gateway 172.22.222.5
 next
end
```

## Configure firewall services, addresses, and policies

To create the security policies, first you must create the firewall services group that will include all the services that will be allowed, then you must define the addresses that will be used in the security policies, and lastly you configure the security policies themselves.

To keep the configuration simple, the Company is allowing only HTTP traffic out of the local network. This will allow employees access to the Internet and their web-mail. DNS services will also be allowed through the firewall.

The security policies will allow HTTP traffic (port 80 and port 8080), HTTPS traffic (port 443), FTP traffic (port 21), and DNS traffic (port 53 and port 953) in both directions. Also BGP (port 179) may need access through the firewall.



For added security, you may want to define a smaller range of addresses for the internal network. For example if only 20 addresses are used, only allow those addresses in the range.

---

In the interest of keeping things simple, a zone will be used to group the two ISP interfaces together. This will allow using one security policy to apply to both ISPs at the same time. Remember to block intra-zone traffic as this will help prevent one ISP sending traffic to the other ISP through your FortiGate unit using your bandwidth. The zone keeps configuration simple, and in the future if there is a need for separate policies for each ISP, they can be created and the zone can be deleted.

The addresses that will be used are the addresses of the FortiGate unit internal and external ports, and the internal network.

More policies or services can be added in the future as applications are added to the network. For more information on security policies, see the firewall chapter of the [FortiGate Administration Guide](#).



When configuring security policies always enable logging to help you track and debug your traffic flow.

---

## To create a firewall services group - web-based manager

1. Go to *Firewall Objects > Service > Groups*, and select *Create New*.

2. For *Group Name*, enter “Basic\_Services”.
3. From *Available Services*, move the following six services over to the *Member* list – BGP, FTP, FTP\_GET, FTP\_PUT, DNS, HTTP, and HTTPS.
4. Select OK.

#### To create a firewall services group - CLI

```
config firewall service group
 edit "Basic_Services"
 set member "BGP" "DNS" "FTP" "FTP_GET" "FTP_PUT" "HTTP" "HTTPS"
 next
end
```

#### To create a zone for the ISP interfaces - web-based manager

1. Go to *System > Network > Interfaces*.
2. Select the caret to the right of *Create New* and then select *Zone*.
3. Enter the following information.

<b>Zone Name</b>	ISPs
<b>Block Intra-zone traffic</b>	enable
<b>interface members</b>	port2 port3

4. Select OK.

#### To create a zone for the ISP interfaces - CLI

```
config system zone
 edit "ISPs"
 set interface "port2" "port3"
 set intrazone block
 next
end
```

#### To add the firewall addresses - web-based manager

1. Go to *Firewall Objects > Address > Addresses*.
2. Select *Create New*, and set the following information.

<b>Address Name</b>	Internal_network
<b>Type</b>	Subnet / IP Range
<b>Subnet / IP Range</b>	10.11.101.0 255.255.255.0
<b>Interface</b>	port1

3. Select OK.

### To add the firewall addresses - CLI

```
config firewall address
 edit "Internal_network"
 set associated-interface "port1"
 set subnet 10.11.101.0 255.255.255.0
 next
end
```

### To add the HTTP and DNS security policies - web-based manager

1. Go to *Policy > Policy > Policy*, and select *Create New*.
2. Set the following information.

<b>Source Interface/Zone</b>	port1(internal)
<b>Source Address</b>	Internal_network
<b>Destination Interface/Zone</b>	ISPs
<b>Destination Address</b>	All
<b>Schedule</b>	always
<b>Service</b>	Basic_services
<b>Action</b>	ACCEPT
<b>Log Allowed Traffic</b>	enable
<b>Enable NAT</b>	Enable
<b>Comments</b>	ISPs basic services out policy

3. Select *OK*.
4. Select *Create New*, and set the following information.

<b>Source Interface/Zone</b>	ISPs
<b>Source Address</b>	all
<b>Destination Interface/Zone</b>	port1(internal)
<b>Destination Address</b>	Internal_network
<b>Schedule</b>	always
<b>Service</b>	Basic_services
<b>Action</b>	ACCEPT
<b>Log Allowed Traffic</b>	enable
<b>NAT</b>	Enable
<b>Comments</b>	ISPs basic services in policy

## To add the security policies - CLI

```
config firewall policy
 edit 1
 set srcintf "port1"
 set srcaddr "Internal_network"
 set dstintf "ISPs"
 set dstaddr "all"
 set schedule "always"
 set service "Basic_services"
 set action accept
 set nat enable
 set profile-status enable
 set logtraffic enable
 set comments "ISPs basic services out policy"
 next
 edit 2
 set srcintf "ISPs"
 set srcaddr "all"
 set dstintf "port1"
 set dstaddr "Internal_network"
 set schedule "always"
 set service "Basic_services"
 set action accept
 set nat enable
 set profile-status enable
 set logtraffic enable
 set comments "ISPs basic services in policy"
 next
end
```

## Set the FortiGate BGP information

When using the default information, there are only two fields to set to configure the FortiGate unit as a BGP router.

For this configuration the FortiGate unit will be in a stub area with one route out — the ISP BGP router. Until you configure the ISP router as a neighbour, even that route out is not available. So while after this part of the configuration is complete your FortiGate unit will be running BGP, it won't know about any other routers running BGP until the next part of the configuration is complete.

### To set the BGP router information - web-based manager

1. Go to *Router > Dynamic > BGP*.
2. Set the following information, and select OK.

<b>Local AS</b>	1
<b>Router ID</b>	10.11.101.110



### To set the BGP router information - CLI

```
config router BGP
 set as 1
 set router-id 10.11.101.110
end
```

### Add the internal network to the AS

The Company is one AS with the FortiGate unit configured as the BGP border router connecting that AS to the two ISPs ASes. The internal network in the Company's AS must be defined. If there were other networks in the company such as regional offices, they would be added here as well.

### To set the networks in the AS - web-based manager

1. Go to *Router > Dynamic > BGP*.
2. In *Networks*, set the following information and select *OK*.

---

<b>IP/Netmask</b>	10.11.101.0/255.255.255.0
-------------------	---------------------------

---

### To set the networks in the AS - CLI

```
config router bgp
 config network
 edit 1
 set prefix 10.11.101.0 255.255.255.0
 next
end
end
```

### Add BGP neighbor information

The configuration will not work unless you set *Remote AS* neighbors. This can be done in either the web-based manager or the CLI.

### To configure the BGP neighbors - web-based manager

1. Go to *Router > Dynamic > BGP*.
2. Add a *Neighbors IP* of 172.21.111.4 with the *Remote AS* set to 650001, then click *Add/Edit*.
3. Add another *Neighbors IP* of 172.22.222.4 with the *Remote AS* set to 650002, then click *Add/Edit*.

### To configure the BGP neighbors - CLI

```
config router BGP
 set as 1
 config neighbor
 edit "172.21.111.4"
 set remote-as 650001
 next
 edit "172.22.222.4"
 set remote-as 650002
 next
end
end
```

## Additional FortiGate BGP configuration

At this point that is all the settings that can be done in both the web-based manger and the CLI. The remaining configuration must be completed in the CLI.

These additional settings are mainly determined by your ISP requirements. They will determine your timers such as keep alive timers, if extended features like BFD and graceful restart are being used, and so on. For this example, some common simply features are being used to promote faster detections of network failures which will result in better service for the Company's internal network users.

The ISPs do not require authentication between peer routers.

These commands will enable or modify the following features on the FortiGate unit, and where possible on neighboring routers as well:

- `bestpath-med-missing-as-worst` — treats a route without an MED as the worst possible available route due to expected unreliability
- `fast-external-failover` — immediately reset the session information associated with BGP external peers if the link used to reach them goes down
- `graceful-restart*` — advertise reboots to neighbors so they do not see the router as offline, wait before declaring them offline, and how long to wait when they reboot before advertising updates. These commands applies to neighbors and are part of the BGP capabilities. This prevents unneeded routing updates.
- `holdtime-timer` — how long the router will wait for a keepalive message before declaring a router offline. A shorter time will find an offline router faster.
- `keepalive-timer` — how often the router sends out keepalive messages to neighbor routers to maintain those sessions.
- `log-neighbor-changes` — log changes to neighbor routers' status. This can be useful for troubleshooting from both internal and external networks.
- `connect-timer` — how long in seconds the FortiGate unit will try to reach this neighbor before declaring it offline.
- `weight` — used to prefer routes from one neighbor over the other. In this example ISP1 is the primary connection so it is weighted higher than ISP2

## To configure additional BGP options - CLI

```
config router bgp
 set bestpath-med-missing-as-worst enable
 set fast-external-failover enable
 set graceful-restart enable
 set graceful-restart-time 120
 set graceful-stalepath-time 180
 set graceful-update-delay 180
 set holdtime-timer 120
 set keepalive-timer 45
 set log-neighbor-changes enable
config neighbor
 edit 172.21.111.4
 set connect-timer 60
 set description "ISP1"
 set holdtime-timer 120
 set keepalive-timer 45
 set weight 250
 next
 edit 172.22.222.4
 set connect-timer 60
 set description "ISP2"
 set holdtime-timer 120
 set keepalive-timer 45
 set weight 100
 next
end
end
```

## Configuring other networking devices

There are two other networking devices that need to be configured: both ISPs' BGP routers.

The ISPs' routers must add the FortiGate unit as a neighbor so route updates can be sent in both directions. Note that ISP1 is not directly connected to ISP2 that we are aware of.

Inform both of your ISPs of your FortiGate unit's BGP information. Once they have configured their router, you can test your BGP connection to the Internet.

They will require your FortiGate unit's:

- IP address of the connected interface
- Router ID
- your Company's AS number

## Testing this configuration

With the dual-homed BGP configuration in place, you should be able to send and receive traffic, send and receive routes, and not have any routing loops. Testing the networks will confirm things are working as expected.

In general for routing you need to look at the routing table on different routers to see what routes are being installed. You also need to sniff packets to see how traffic is being routed in real time. These two sources of information will normally tell you what you need to know.

Testing of this example's network configuration should be completed in two parts:

- [Testing network connectivity](#)
- [Verifying the FortiGate unit's routing tables](#)
- [Verifying traffic routing](#)
- [Verifying the dual-homed side of the configuration](#)

## Testing network connectivity

A common first step in testing a new network topology is to test if you can reach the internet and other locations as you expect you should. If not, you may be prevented by cabling issues, software or other issues.

The easiest way to test connections is to use ping, once you ensure that all the FortiGate unit's interfaces and ISP routers have ping support enabled. Also ensure that the security policies allow ping through the firewall.

Connections to test in this example are the internal network to ISP1's router or the internet, and the same for ISP2. If you can connect on the external side of the Fortinet unit, try to ping the internal network. Those three tests should prove your basic network connections are working.



Once you have completed testing the network connectivity, turn off ping support on the external interfaces for additional security.

---

## Verifying the FortiGate unit's routing tables

The FortiGate routing table contains the routes stored for future use. If you are expecting certain routes to be there and they are not, that is a good indicator that your configuration is not what you expected.

The CLI command `get router info routing-table details` will provide you with every route's routing protocol, destination address, gateway address, interface, weighting, and if the address is directly connected or not.

If you want to limit the display to BGP routes only, use the CLI command `get router info routing-table bgp`. If there are no BGP routes in the routing table, nothing will be displayed. In the CLI command you can replace BGP with static, or other routing protocols to only display those routes.

If you want to see the contents of the routing information database (RIB), use the CLI command `get router info routing-table database`. This will display the incoming routes that may or may not make it into the routing table.

## Verifying traffic routing

Traffic may be reaching the internal network, but it may be using a different route than you think to get there.

Use a browser to try and access the Internet.

If needed, allow traceroute and other diag ports to be opened until things are working properly. Then remove access for them again.

Look for slow hops on the traceroute, or pings to a location, as they may indicate network loops that need to be fixed.

Any locations that have an unresolved traceroute or ping must be examined and fixed.

Use network packet sniffing to ensure traffic is being routed as you expect.

### Verifying the dual-homed side of the configuration

Since there are two connections to the internet in this example, theoretically you can pull the plug on one of the ISP connections, and all traffic will go through the other connection. Alternately, you may choose to remove a default route to one ISP, remove that ISP's neighbor settings, or change the weightings to prefer other other ISP. These alternate ways to test dual-homing do not change physical cabling, which may be preferred in some situations.

If this does not work as expected, things to check include:

- default static routes — if these are wrong or don't exist, the traffic can't get out.
- BGP neighbor information — If the ISP router information is incorrect, the FortiGate unit won't be able to talk to it.

## Redistributing and blocking routes in BGP

During normal BGP operation, peer routers redistribute routes from each other. However, in some specific situations it may be best to not advertise routes from one peer, such as if the peer is redundant with another peer (they share the same routes exactly), if it might be unreliable in some way, or some other reason. The FortiGate can also take routes it learns from other protocols and advertise them in BGP, for example OSPF or RIP. If your Company hosts its own web or email servers, external locations will require routes to your networks to reach those services.

In this example the Company has an internal network in an OSPF area, and is connected to a BGP AS and two BGP peers. Company goes through these two peers to reach the Internet. However, Peer 1 routes will not be advertised to Peer 2. The Company internal user and server networks are running OSPF, and will redistribute those routes to BGP so external locations can reach the web and email servers.

This section includes the following topics:

- [Network layout and assumptions](#)
- [Configuring the FortiGate unit](#)
- [Testing network configuration](#)

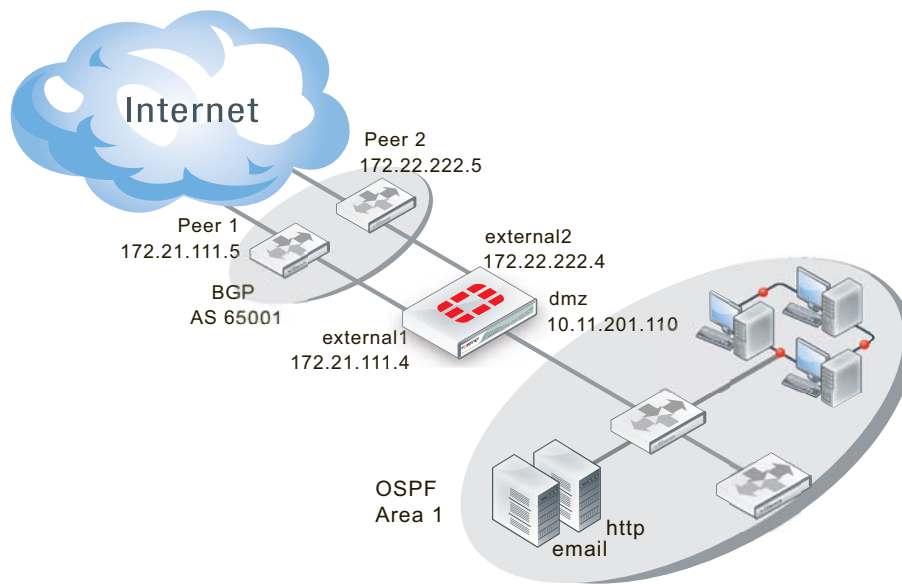
### Network layout and assumptions

The network layout for the BGP redistributing routes example involves the company network being connected to two BGP peers as shown below. In this configuration the FortiGate unit is the BGP border router between the Company AS, and the peer routers.

The components of the layout include:

- There is only one BGP AS in this example — AS 65001, shared by the FortiGate unit and both peers.
- The Company's FortiGate unit connects to the Internet through two BGP peers.
- The Company internal networks on the dmz interface of the FortiGate unit with an IP of 10.11.201.0/24.
- The FortiGate units' interfaces are connected as follows:
  - port1 (dmz) has IP 10.11.201.110 and is the internal user and server network
  - port2 (external1) has IP 172.21.111.4 and is connected to Peer 1's network
  - port3 (external2) has IP 172.22.222.4 and is connected to Peer 2's network
- Peer 1 has IP 172.21.111.5, and Peer 2 has IP 172.22.222.5.
- OSPF Area 1 is configured on the dmz interface of the FortiGate unit, and is the routing protocol used by the internal users and servers.

**Figure 109:**BGP network topology



## Assumptions

The the BGP redistributing routes configuration procedure follows these assumptions:

- the FortiGate unit has been configured following the Install Guide
- interfaces port1, port2, and port 3 exist on the FortiGate unit
- we don't know the router manufacturers of Peer 1 and Peer 2
- we don't know what other devices are on the BGP AS or OSPF Area
- all basic configuration can be completed in both GUI and CLI
- access lists and route maps will only be configured in CLI
- VDOMs are not enabled on the FortiGate unit

## Configuring the FortiGate unit

1. [Configuring the FortiGate unit — networks and firewalls](#)
2. [Configuring the FortiGate unit - BGP](#)

3. [Configuring the FortiGate unit - OSPF](#)
4. [Configuring other networking devices](#)

## Configuring the FortiGate unit – networks and firewalls

The FortiGate unit has three interfaces connected to networks — two external and one dmz.

Security policies must be in place to allow traffic to flow between these networks.

Firewall services will change depending on which routing protocol is being used on that network — either BGP or OSPF. Beyond that, all services that are allowed will be allowed in both directions due to the internal servers. The services allowed are web-server services (DNS, HTTP, HTTPS, SSH, NTP, FTP\*, SYSLOG, and MYSQL), email services (POP3, IMAP, and SMTP), and general troubleshooting services (PING, TRACEROUTE). Those last two can be removed once the network is up and working properly to increase security. Other services can be added later as needed.

### To configure the interfaces - GUI

1. Go to *System > Network > Interfaces*.
2. Edit port1 (dmz) interface.
3. Set the following information, and select *OK*.

<b>Alias</b>	dmz
<b>IP/Netmask</b>	10.11.201.110/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	OSPF internal networks
<b>Administrative Status</b>	Up

4. Edit port2 (external1) interface.
5. Set the following information, and select *OK*.

<b>Alias</b>	external1
<b>IP/Netmask</b>	172.21.111.4/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH
<b>Description</b>	BGP external1 Peer 1
<b>Administrative Status</b>	Up

6. Edit port3 (external2) interface.
7. Set the following information, and select *OK*.

<b>Alias</b>	external2
<b>IP/Netmask</b>	172.22.222.4/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH
<b>Description</b>	BGP external2 Peer2
<b>Administrative Status</b>	Up

## To configure the FortiGate interfaces (CLI)

```
config system interface
 edit port1
 set alias dmz
 set ip 10.11.201.110 255.255.255.0
 set allowaccess https ssh ping
 set description "OSPF internal networks"
 set status up
 next
 edit port2
 set alias external1
 set ip 172.21.111.4 255.255.255.0
 set allowaccess https ssh
 set description "BGP external1 Peer 1"
 set status up
 next
 edit port3
 set alias external2
 set ip 172.22.222.4 255.255.255.0
 set allowaccess https ssh
 set description "BGP external2 Peer 2"
 set status up
 next
end
```

## To configure the firewall addresses - GUI

1. Go to *Firewall Objects > Address > Addresses*.
2. Select *Create New*, and set the following information.

<b>Address Name</b>	BGP_services
<b>Type</b>	Subnet / IP Range
<b>Subnet / IP Range</b>	10.11.201.0 255.255.255.0
<b>Interface</b>	port1

3. Select *OK*.
4. Select *Create New*, and enter the following information:
5. Select *OK*.

## To configure the firewall addresses - CLI

```
config firewall address
 edit "BGP_services"
 set associated-interface "port1"
 set subnet 10.11.201.0 255.255.255.0
 next
end
```



### To configure firewall service groups - GUI

1. Go to *Firewall Objects > Service > Groups*.
2. Select *Create New*.
3. Name the group BGP\_Services.
4. Move the following services to the right list: BGP, DHCP, DNS, FTP, FTP\_GET, FTP\_PUT, HTTP, HTTPS, IMAP, MYSQL, NTP, PING, POP3, SMTP, SSH, SYSLOG, and TRACEROUTE.
5. Select *OK*.
6. Select *Create New*.
7. Name the group OSPF\_Services.
8. Move the following services to the right list: DNS, FTP, FTP\_GET, FTP\_PUT, HTTP, HTTPS, IMAP, MYSQL, NTP, OSPF, PING, POP3, SMTP, SSH, SYSLOG, and TRACEROUTE.
9. Select *OK*.

### To configure firewall service groups - CLI

```
config firewall service group
 edit "BGP_services"
 set member "BGP", "DHCP" "DNS" "FTP" "FTP_GET" "FTP_PUT" "HTTP"
 "HTTPS" "IMAP" "MYSQL" "NTP" "PING" "POP3" "SMTP" "SSH"
 "TRACEROUTE" "SYSLOG"
 next
 edit "OSPF_services"
 set member "DHCP" "DNS" "FTP" "FTP_GET" "FTP_PUT" "HTTP" "HTTPS"
 "IMAP" "MYSQL" "NTP" "PING" "POP3" "SMTP" "SSH"
 "TRACEROUTE" "SYSLOG" "OSPF"
 next
end
```

## Configuring the FortiGate unit - BGP

The only change from the standard BGP configuration for this example is configuring the blocking Peer 1's routes from being advertised to Peer 2. From the network topology you can guess that both of these peers likely share many routes in common and it makes no sense to advertise unneeded routes.

Blocking Peer 1's routes to Peer 2 is done with distribute-list-out keyword. They allow you to select which routes you will advertise to a neighbor using an access list. In this case we will block all incoming routes from Peer 1 when we send updates to Peer 2. Otherwise Peer 1 and Peer 2 are regular neighbors.

The FortiGate unit will redistribute routes learned from OSPF into BGP.

This is advanced configuration and the commands are only available in the CLI.

### To create access list to block Peer 1 - CLI

```
config access-list
 edit "block_peer1"
 config rule
 edit 1
 set prefix 172.21.111.0 255.255.255.0
 set action deny
 set exact-match enable
 end
 end
 end
end
```

### To configure BGP on the FortiGate unit - CLI

```
config router bgp
 set as 65001
 set router-id 10.11.201.110
 config redistribute ospf
 set status enable
 end
 config neighbor
 edit 172.22.222.5
 set remote-as 65001
 set distribute-list-out "block_peer1"
 next
 edit 172.21.111.5
 set remote-as 65001
 end
 end
end
```

## Configuring the FortiGate unit - OSPF

This configuration involves only one OSPF Area, so all traffic will be intra-area. If there were two or more areas with traffic going between them it would be inter-area traffic. These two types are comparable to BGP's traffic within one AS (iBGP) or between multiple ASes (eBGP). Redistributing routes from OSPF to BGP is considered external because either the start or end point is a different routing protocol.

The OSPF configuration is basic apart from redistributing BGP routes learned.

### To configure OSPF on the FortiGate unit - web-based manager

1. Go to *Router > Dynamic > OSPF*.
2. For Router ID enter `10.11.201.110` and then select *Apply*.
3. Under *Advanced Options* and *Redistribute*, select *BGP* and set BGP metric to 1.
4. For *Areas*, select *Create New*, enter the following information and then select *OK*.

<b>Area (IP)</b>	0.0.0.0
<b>Type</b>	Regular
<b>Authentication</b>	None

5. For *Networks*, select *Create New*.

6. Enter 10.11.201.0/255.255.255.0 for *IP/Netmask*, and select *OK*.
7. For *Interfaces*, select *Create New*.
8. Enter OSPF\_dmz\_network for *Name*.
9. Select port1(dmz) for *Interface*, and then select *OK*.

### To configure OSPF on the FortiGate unit - CLI

```
config router ospf
 set router-id 10.11.201.110
 config area
 edit 0.0.0.0
 set type regular
 set authentication none
 end
 config network
 edit 1
 set area 0.0.0.0
 set prefix 10.11.201.0 255.255.255.0
 end
 config interface
 edit "OSPF_dmz_network"
 set interface port1(dmz)
 set status enable
 end
 config redistribute bgp
 set status enable
 set metric 1
 end
end
```

### Configuring other networking devices

As with all BGP configurations, the peer routers will need to be updated with the FortiGate unit's BGP information including IP address, AS number, and what capabilities are being used such as IPv6, graceful restart, BFD, and so on.

### Testing network configuration

Testing this configuration involves the standard connectivity checks, but also ensuring that routes are being passed between protocols as expected.

Check the routing table on the FortiGate unit to ensure that routes from both OSPF and BGP are present.

Check the routing table on devices on the OSPF network for routes redistributed from BGP. Also check those devices for connectivity to the Internet.

Check the routing table on Peer 2 to ensure no routes from Peer 1 are present, but routes from the internal OSPF network are present.

For help with troubleshooting, see [“Troubleshooting BGP” on page 370](#).

# Open Shortest Path First (OSPF)

This section describes OSPF routing.

The following topics are included in this section:

- [OSPF Background and concepts](#)
- [Troubleshooting OSPF](#)
- [Basic OSPF example](#)
- [Advanced inter-area OSPF example](#)
- [Controlling redundant links by cost](#)

## OSPF Background and concepts

OSPF (Open Shortest Path First) is a link-state interior routing protocol, that is widely used in large enterprise organizations. It only routes packets within a single autonomous system (AS). This is different from BGP as BGP can communicate between ASes.

This section includes:

- [Background](#)
- [The parts and terminology of OSPF](#)
- [How OSPF works](#)

### Background

OSPF version 2 was defined in 1998 in RFC 2328. OSPF was designed to support classless IP addressing, and variable subnet masks. This was a shortcoming of the earlier RIP protocols.

Updates to OSPF version 2 are included in OSPF version 3 defined in 2008 in RFC 5340. OSPF3 includes support for IPv6 addressing where previously OSPF2 only supports IPv4 addressing.

The main benefit of OSPF is that it detects link failures in the network quickly and within seconds has converged network traffic successfully without any networking loops. Also OSPF has many features to control which routes are propagated and which are not, maintaining smaller routing tables. OSPF can also provide better load-balancing on external links than other interior routing protocols.

### The parts and terminology of OSPF

Parts and terminology of OSPF includes:

- [OSPFv3 and IPv6](#)
- [Router ID](#)
- [Adjacency](#)
- [Designated router \(DR\) and backup router \(BDR\)](#)
- [Area](#)
- [Authentication](#)
- [Hello and dead intervals](#)

## OSPFv3 and IPv6

OSPFv3 (OSPF version 3) includes support for IPv6. Generally, all IP addresses are in IPv6 format instead of IPv4. However, OSPFv3 area numbers use the same 32-bit numbering system as OSPFv2, as described in [RFC 2740](#). Likewise, the router ID and area ID are in the same format as OSPFv2.

As with most advanced routing features on your FortiGate unit, IPv6 settings for dynamic routing protocols must be enabled before they will be visible in the GUI. To enable IPv6 configuration in the GUI, enable it in *System > Admin > Settings*.

For IPv6, the main difference in OSPFv3 is that, rather than using a network statement to enable OSPFv3 on an interface, you define OSPF6 (OSPF for IPv6) interfaces, which are bound to interface and area. This configuration must be done in the CLI, as follows (with sample interfaces and addresses):

```
config router ospf6
 config area
 edit 0.0.0.0
 next
end
config ospf6-interface
 edit "tunnel"
 set interface "to_FGT300A-7"
next
 edit "internal_lan"
 set interface "port1"
next
 set router-id 10.174.0.113
end
```

Note that OSPFv3 neighbors use link-local IPv6 addresses, but with broadcast and point-to-point network types, neighbors are automatically discovered. You only have to manually configure neighbors when using non-broadcast network types.

## Router ID

In OSPF, each router has a unique 32-bit number called its Router ID. Often this 32-bit number is written the same as a 32-bit IPv4 address would be written in dotted decimal notation. However some brands of routers, such as Cisco routers, support a router ID entered as an integer instead of an IP address.

It is a good idea to not use IP address in use on the router for the router ID number. The router ID does not have to be a particular IP address on the router. By choosing a different number, it will be harder to get confused which number you are looking at. A good idea can be to use the as much of the area's number as possible. For example if you have 15 routers in area 0.0.0.0 they could be numbered from 0.0.0.1 to 0.0.0.15. If you have an area 1.1.1.1, then routers in that area could start at 1.1.1.10 for example.

You can manually set the router ID on your FortiGate unit.

### To manually set an OSPF router ID of 0.0.1.1 - web-based manager

1. Go to *Router > Dynamic > OSPF*.
2. For *Router ID*, enter 0.0.1.1.
3. Select *Apply*.

## To manually set an OSPF router ID of 0.0.1.1 - CLI

```
config router ospf
 set router-id 0.0.1.1
end
```

## Adjacency

In an OSPF routing network, when an OSPF router boots up it sends out OSPF Hello packets to find any neighbors, routers that have access to the same network as the router booting up. Once neighbors are discovered and Hello packets are exchanged, updates are sent, and the Link State databases of both neighbors are synchronized. At this point these neighbors are said to be adjacent.

For two OSPF routers to become neighbors, the following conditions must be met.

- The subnet mask used on both routers must be the same subnet.
- The subnet number derived using the subnet mask and each router's interface IP address must match.
- The Hello interval & The Dead interval must match.
- The routers must have the same OSPF area ID. If they are in different areas, they are not neighbors.
- If authentication is used, they must pass authentication checks.

If any of these parameters are different between the two routers, the routers do not become OSPF neighbors and cannot be adjacent. If the routers become neighbors, they are adjacent.

## Adjacency and neighbors

Neighbor routers can be in a Two-Way state, and not be adjacent. Adjacent routers normally have a neighbour state of FULL. Neighbors only exchange Hello packets, and do not exchange routing updates. Adjacent routers exchange LSAs (LSDB information) as well as Hello packets. A good example of an adjacent pair of routers is the DR and BDR.

You can check on the state of an OSPF neighbor using the CLI command `get router info ospf neighbor all`. See [“Checking the state of OSPF neighbors” on page 409](#).

## Why adjacency is important

It is important to have adjacent pairs of routers in the OSPF routing domain because routing protocol packets are only passed between adjacent routers. This means adjacency is required for two OSPF routers to exchange routes. If there is no adjacency between two routers, such as one on the 172.20.120.0 network and another on the 10.11.101.0 network, the routers do not exchange routes. This makes sense because if all OSPF routers on the OSPF domain exchanged updates it would flood the network. Also, it is better for updates to progress through adjacent routers to ensure there are no outages along the way. Otherwise, updates could skip over routers that are potentially offline, causing longer routing outages and delays while the OSPF domain learns of this outage later on.

If the OSPF network has multiple border routers and multiple connections to external networks, the designated router (DR) determines which router pairs become adjacent. The DR can accomplish this because it maintains the complete topology of the OSPF domain, including which router pairs are adjacent. The BDR also has this information in case the DR goes offline.

## Designated router (DR) and backup router (BDR)

In OSPF a router can have a number of different roles to play.

A designated router (DR) is the designated broadcasting router interface for an AS. It looks after all the initial contact and other routing administration traffic. Having only one router do all this greatly reduces the network traffic and collisions.

If something happens and the designated router goes offline, the backup designated router (BDR) takes over. An OSPF FortiGate unit interface can become either a DR or BDR. Both the DR and the BDR cover the same area, and are elected at the same time. The election process doesn't have many rules, but the exceptions can become complex.

## Benefits

The OSPF concept of the designated router is a big step above RIP. With all RIP routers doing their own updates all the time, RIP suffers from frequent and sometimes unnecessary updates that can slow down your network. With OSPF, not only do routing changes only happen when a link-state changes instead of any tiny change to the routing table, but the designated router reduces this overhead traffic even more.

However, smaller network topologies may only have a couple routers besides the designated router. This may seem excessive, but it maintains the proper OSPF form and it will still reduce the administration traffic but to a lesser extent than on a large network. Also, your network topology will be ready whenever you choose to expand your network.

## DR and BDR election

An election chooses the DR and BDR from all the available routers. The election is primarily based on the priority setting of the routers—the highest priority becomes the DR, and the second highest becomes BDR. To resolve any ties, the router with the highest router ID wins. For example 192.168.0.1 would win over 10.1.1.2.

The router priority can vary from 0 to 255, but at 0 a router will never become a DR or BDR. If a router with a higher priority comes on line after the election, it must wait until after the DR and BDR go offline before it would become the DR.

If the original DR goes offline, but then is available when the BDR goes offline later on, the original DR will be promoted back to DR without an election leaving the new BDR as it is.

With your FortiGate unit, to configure the port1 interface to be a potential OSPF designated router or backup designated router called `ospf_DR` on the network, you need to raise the priority of the router to a very high number such as 250 out of 255. This will ensure the interface has a chance to be a DR, but will not guarantee that it will be one. Give the interface a low numbered IP address—such as 10.1.1.1 instead of 192.168.1.1—to help ensure it becomes a DR, but that is not part of this example. Enter the following command:

```
config router ospf
 config ospf-interface
 edit "ospf_DR"
 set priority 250
 end
 end
```

## Area

An OSPF area is a smaller part of the larger OSPF AS. Areas are used to limit the link-state updates that are sent out. The flooding used for these updates would overwhelm a large network, so it is divided into these smaller areas for manageability.

Within an area if there are two or more routers that are viable, there will always be a designated router (DR) and a backup DR (BDR). For more on these router roles, see [“Designated router \(DR\) and backup router \(BDR\)” on page 398](#).

Defining a private OSPF area, involves:

- assigning a 32-bit number to the area that is unique on your network
- defining the characteristics of one or more OSPF areas
- creating associations between the OSPF areas that you defined and the local networks to include in the OSPF area
- if required, adjusting the settings of OSPF-enabled interfaces.



IPv6 OSPF area numbers use the same 32-bit number notation as IPv4 OSPF.

---

If you are using the web-based manager to perform these tasks, follow the procedures summarized below.

FortiGate units support the four main types of OSPF area:

- [Backbone area](#)
- [NSSA](#)
- [Stub area](#)
- [Regular area](#)

### **Backbone area**

Every OSPF network has at least one AS, and every OSPF network has a backbone area. The backbone is the main area, or possibly the only area. All other OSPF areas are connected to a backbone area. This means if two areas want to pass routing information back and forth, that routing information will go through the backbone on its way between those areas. For this reason the backbone not only has to connect to all other areas in the network, but also be uninterrupted to be able to pass traffic to all points of the network.

The backbone area is referred to as area 0 because it has an IP address of 0.0.0.0.

### **Stub area**

A stub area is an OSPF area that receives no outside routes advertised into it, and all routing in it is based on a default route. This essentially isolates it from outside areas.

Stub areas are useful for small networks that are part of a larger organization, especially if the networking equipment can't handle routing large amounts of traffic passing through, or there are other reasons to prevent outside traffic, such as security. For example most organizations don't want their accounting department to be the center of their network with everyone's traffic passing through there. It would increase the security risks, slow down their network, and it generally doesn't make sense.

A variation on the stub area is the totally stubby area. It is a stub area that does not allow summarized routes.

### **NSSA**

A not-so-stubby-area (NSSA) is a stub area that allows for external routes to be injected into it. While it still does not allow routes from external areas, it is not limited to only using the default route for internal routing.



## Regular area

A regular area is what all the other ASes are, all the non-backbone, non-stub, non-NSSA areas. A regular area generally has a connection to the backbone, does receive advertisements of outside routes, and does not have an area number of 0.0.0.0.

## Authentication

In the OSPF packet header are two authentication related fields —AuType, and Authentication.

All OSPF packet traffic is authenticated. Multiple types of authentication are supported in OSPFv2. However in OSPFv3, there is no authentication built-in but it is assumed that IPsec will be used for authentication instead.

Packets that fail authentication are discarded.

## Null authentication

Null authentication indicates there is no authentication being used. In this case the 16-byte Authentication field is not checked, and can be any value. However checksumming is still used to locate errors. On your FortiGate this is the `none` option for authentication.

## Simple Password authentication

Simple password refers to a standard plain text string of characters. The same password is used for all transactions on a network. The main use of this type of authentication is to prevent routers from accidentally joining the network. Simple password authentication is vulnerable to many forms of attack, and is not recommended as a secure form of authentication.

## Cryptographic authentication

Cryptographic authentication involves the use of a shared secret key to authenticate all router traffic on a network. The key is never sent over the network in the clear—a packet is sent and a condensed and encrypted form of the packet is appended to the end of the packet. A non-repeating sequence number is included in the OSPF packet to protect against replay attacks that could try to use already sent packets to disrupt the network. When a packet is accepted as authentic the authentication sequence number is set to the packet sequence number. If a replay attack is attempted, the packet sent will be out of sequence and ignored.

Your FortiGate unit supports all three levels of authentication through the authentication keyword associated with creating an OSPF interface .

For example to create an OSPF interface called `Accounting` on the `port1` interface that is a broadcast interface, has a hello interval of 10 seconds, has a dead interval of 40 seconds, uses text authentication (simple password) with a password of “`ospf_test`”, enter the command:

```
config router ospf
 config ospf-interface
 edit Accounting
 set interface port1
 set network-type broadcast
 set hello-interval 10
 set dead-interval 40
 set authentication text
 set authentication-key "ospf_test"
 end
 end
```

## Hello and dead intervals

The OSPF Hello protocol is used to discover and maintain communications with neighboring routers.

Hello packets are sent out at a regular interval for this purpose. The DR sends out the Hello packets. In a broadcast network, the multicast address of 224.0.0.5 is used to send out Hello packets. New routers on the network listen for and reply to these packets to join the OSPF area. If a new router never receives a Hello packet, other routers will not know it is there and will not communicate with it. However, once a new router is discovered the DR adds it to the list of routers in that area and it is integrated into the routing calculations.

Dead interval is the time other routers will wait before declaring a neighbor dead (offline). Setting a reasonable dead interval is very important. If this interval is too short, routers will be declared offline when they are just slow or momentarily inaccessible, and link-state updates will happen more than they need to, using more bandwidth. If the dead interval is too long, it will slow down network traffic overall if online routers attempt to contact offline ones instead of re-routing traffic. The CLI syntax for OSPF dead interval follows:

```
config ospf-interface
 edit ospf1
 set interface port1
 set network-type broadcast
 set dead-interval 1
end
```

## Access Lists

Access lists are filters used by FortiGate unit OSPF routing. An access list provides a list of IP addresses and the action to take for them — essentially an access list makes it easy to group addresses that will be treated the same into the same group, independent of their subnets or other matching qualities. You add a rule for each address or subnet that you want to include, specifying the action to take for it. For example if you wanted all traffic from one department to be routed a particular way, even in different buildings, you can add all the addresses to an access list and then handle that list all at once.

Each rule in an access list consists of a prefix (IP address and netmask), the action to take for this prefix (permit or deny), and whether to match the prefix exactly or to match the prefix and any more specific prefix.

The FortiGate unit attempts to match a packet against the rules in an access list starting at the top of the list. If it finds a match for the prefix, it takes the action specified for that prefix. If no match is found the default action is deny.

Access lists greatly speed up configuration and network management. When there is a problem, you can check each list instead of individual addresses. Also, it eases troubleshooting since if all addresses on one list have problems, it eliminates many possible causes right away.

If you are using the OSPF+ IPv6 protocols you will need to use access-list6, the IPv6 version of access list. The only difference is that access-list6 uses IPv6 addresses.

For example, if you want to create an access list called `test_list` that only allows an exact match of `10.10.10.10` and `11.11.11.11`, enter the command:

```
config router access-list
 edit test_list
 config rule
 edit 1
 set prefix 10.10.10.10 255.255.255.255
 set action allow
 set exact-match enable
 next
 edit 2
 set prefix 11.11.11.11 255.255.255.255
 set action allow
 set exact-match enable
 end
 end
```

Another example is if you want to deny ranges of addresses in IPv6 that start with the IPv6 equivalents of `10.10.10.10` and `11.11.11.11`, enter the command `access-list6` as follows:

```
config router access-list6
 edit test_list_ip6
 config rule
 edit 1
 set prefix6 2002:A0A:A0A:0:0:0:0:0/48
 set action deny
 next
 edit 2
 set prefix6 2002:B0B:B0B:0:0:0:0:0/48
 set action deny
 end
```

To use an `access_list`, you must call it from a routing protocol such as RIP. The following example uses the `access_list` from the earlier example called `test_list` to match routes coming in on the `port1` interface. When there is a match, it will add 3 to the hop count metric for those routes to artificially decrease their priority. Enter the following command:

```
config router ospf
 config distribute-list
 edit 5
 set access-list test_list
 set protocol connected
 end
```

If you are setting a prefix of `128.0.0.0`, use the format `128.0.0.0/1`. The default route `0.0.0.0/0` can not be exactly matched with an access-list. A prefix-list must be used for this purpose.

## How OSPF works

An OSPF installation consists of one or more areas. An OSPF area is typically divided into logical areas linked by Area Border Routers. A group of contiguous networks form an area. An Area Border Router (ABR) links one or more areas to the OSPF network backbone (area ID 0). See [“Area border router \(ABR\)” on page 315](#).

OSPF is an interior routing protocol. It includes a backbone AS, and possibly additional ASes. The DR and BDR are elected from potential routers with the highest priorities. The DR handles much of the administration to lower the network traffic required. New routers are discovered through hello packets sent from the DR using the multicast address of 224.0.0.5. If the DR goes offline at any time, the BDR has a complete table of routes that it uses when it takes over as the DR router.

OSPF does not use UDP or TCP, but is encapsulated directly in IP datagrams as protocol 89. This is in contrast to RIP, or BGP. OSPF handles its own error detection and correction functions.

The OSPF protocol, when running on IPv4, can operate securely between routers, optionally using a variety of authentication methods to allow only trusted routers to participate in routing. OSPFv3, running on IPv6, no longer supports protocol-internal authentication. Instead, it relies on IPv6 protocol security (IPsec).

Other important parts of how OSPF works includes:

- [OSPF router discovery](#)
- [How OSPF works on FortiGate units](#)
- [External routes](#)
- [Link-state Database \(LSDB\) and route updates](#)
- [OSPF packets](#)

## OSPF router discovery

OSPF-enabled routers generate Link-State Advertisements (LSA) and send them to their neighbors whenever the status of a neighbor changes or a new neighbor comes online. As long as the OSPF network is stable, LSAs between OSPF neighbors do not occur. An LSA identifies the interfaces of all OSPF-enabled routers in an area, and provides information that enables OSPF-enabled routers to select the shortest path to a destination. All LSA exchanges between OSPF-enabled routers are authenticated.

When a network of OSPF routers comes online, the follow steps occur.

1. When OSPF routers come online, they send out Hello packets to find other OSPF routers on their network segment.
2. When they discover other routers on their network segment, generally they become adjacent. Adjacent routers can exchange routing updates. See [“Adjacency” on page 398](#).
3. A DR and BDR are elected from the available routers using priority settings, and router ID. See [“Designated router \(DR\) and backup router \(BDR\)” on page 398](#), and [“DR and BDR election issues” on page 410](#).
4. Link state updates are sent between adjacent routers to map the topology of the OSPF area.
5. Once complete, the DR floods the network with the updates to ensure all OSPF routers in the area have the same OSPF route database. After the initial update, there are very few required updates if the network is stable.

## How OSPF works on FortiGate units

When a FortiGate unit interface is connected to an OSPF area, that unit can participate in OSPF communications. FortiGate units use the OSPF Hello protocol to acquire neighbors in an area. A neighbor is any router that is directly connected to the same area as the FortiGate unit, and ideally is adjacent with a state of Full. After initial contact, the FortiGate unit exchanges Hello packets with its OSPF neighbors regularly to confirm that the neighbors can be reached.

The number of routes that a FortiGate unit can learn through OSPF depends on the network topology. A single unit can support tens of thousands of routes if the OSPF network is configured properly.

## External routes

OSPF is an internal routing protocol. OSPF external routes are routes where the destination is using a routing protocol other than OSPF. OSPF handles external routes by adjusting the cost of the route to include the cost of the other routing protocol. There are two methods of calculating this cost, used for OSPF E1 and OSPF E2.

### OSPF external1 (E1)

In OSPF E1 the destination is outside of the OSPF domain. This requires a different metric to be used beyond the normal OSPF metrics. The new metric of a redistributed route is calculated by adding the external cost and the OSPF cost together.

### OSPF external2 (E2)

OSPF E2 is the default external type when routes are redistributed outside of OSPF. With OSPF E2, the metric of the redistributed route is equivalent to the external cost only, expressed as an OSPF cost. Dropping the OSPF portion can be useful in a number of situations, on border routers that have no OSPF portion for example or where the OSPF routing cost is negligible compared to the external routing cost.

### Comparing E1 and E2

The best way to understand OSPF E1 and E2 routes is to check routing tables on OSPF routers. If you look at the routes on an OSPF border router, the redistributed routes will have an associated cost that represents only the external route, as there is no OSPF cost to the route due to it already being on the edge of the OSPF domain. However, if you look at that same route on a different OSPF router inside the OSPF routing domain, it will have a higher associated cost - essentially the external cost plus the cost over the OSPF domain to that border router. The border router uses OSPF E2, where the internal OSPF router uses OSPF E1 for the same route.

### Viewing external routes

When you are trying to determine the costs for routes in your network to predict how traffic will be routed, you need to see the external OSPF routes and their associated costs. On your FortiGate unit, you find this information through your CLI.

#### To view external routes - CLI

You can view the whole routing table using `get router info routing-table all` to see all the routes including the OSPF external routes, or for a shorter list you can use the command `get router info routing-table ospf`. The letter at the left will be either E1 or E2 for external OSPF routes. The output of will look similar to the following, depending on what routes are in your routing table.

```
FGT620B# get router info routing-table all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
* - candidate default

O*E2 0.0.0.0/0 [110/10] via 10.1.1.3, tunnel_wan2, 00:02:11
O 10.0.0.1/32 [110/300] via 10.1.1.3, tunnel_wan2, 00:02:11
S 0.0.0.0/0 [10/0] via 192.168.183.254, port2
S 1.0.0.0/8 [10/0] via 192.168.183.254, port2
```

## Link-state Database (LSDB) and route updates

OSPF is based on links. The links between adjacent neighbor routers allow updates to be passed along the network. Network links allow the DR to flood the area with Link-state database (LSDB) updates. External links allow the OSPF area to connect to destinations outside the OSPF autonomous system. Information about these links is passed throughout the OSPF network as link-state updates.

The LSDB contains the information that defines the complete OSPF area, but the LSDB is not the routing table. It contains the information from all the link-state updates passed along the network. When there are no more changes required, and the network is stable then the LSDB on each router in the network will be the same. The DR will flood the LSDB to the area to ensure each router has the same LSDB.

To calculate the best route (shortest path) to a destination, the FortiGate unit applies the Shortest Path First (SPF) algorithm, based on Dijkstra's algorithm, to the accumulated link-state information. OSPF uses relative path cost metric for choosing the best route. The path cost can be any metric, but is typically the bandwidth of the path, how fast traffic will get from one point to another.

The path cost, similar to "distance" for RIP, imposes a penalty on the outgoing direction of a FortiGate unit interface. The path cost of a route is calculated by adding together all of the costs associated with the outgoing interfaces along the path to the destination. The lowest overall path cost indicates the best route, and generally the fastest route. Some brands of OSPF routers, such as Cisco, implement cost as a direct result of bandwidth between the routers. Generally this is a good cost metric because larger bandwidth means more traffic can travel without slowing down. To achieve this type of cost metric on FortiGate units, you need to set the cost for each interface manually in the CLI.



The inter-area routes may not be calculated when a Cisco type ABR has no fully adjacent neighbor in the backbone area. In this situation, the router considers summary-LSAs from all Actively summary-LSAs from all Actively Attached areas (RFC 3509).

---

The FortiGate unit dynamically updates its routing table based on the results of the SPF calculation to ensure that an OSPF packet will be routed using the shortest path to its destination. Depending on the network topology, the entries in the FortiGate unit routing table may include:

- the addresses of networks in the local OSPF area (to which packets are sent directly)
- routes to OSPF area border routers (to which packets destined for another area are sent)
- if the network contains OSPF areas and non-OSPF domains, routes to area boundary routers, which reside on the OSPF network backbone and are configured to forward packets to destinations outside the OSPF AS.

### OSPF Route updates

Once the OSPF domain is established, there should be few updates required on a stable network. When updates occur and a decision is required concerning a new route, this is the general procedure.

1. Our router gets a new route, and needs to decide if it should go in the routing table.
2. The router has an up to date LSDB of the entire area, containing information about each router, the next hop to it, and most importantly the cost to get there.
3. Our router, turns the LSDB into a shortest path first (SPF) tree using Dijkstra's algorithm. It doesn't matter if there is more than one path to a router on the network, the SPF tree only cares about the shortest path to that router.

4. Once the SPF tree has been created, and shows the shortest paths to all the OSPF routers on the network, the work is done. If the new route is the best route, it will be part of that tree. If it is not the shortest route, it will not be included in the LSDB.
5. If there has been a change from the initial LSDB to the new SPF tree, a link state update will be sent out to let the other routers know about the change so they can update their LSDBs as well. This is vital since all routers on the OSPF area must have the same LSDB.
6. If there was no change between the LSDB and the SPF tree, no action is taken.

## OSPF packets

Every OSPF packet starts with a standard 24-byte header, and another 24 bytes of information or more. The header contains all the information necessary to determine whether the packet should be accepted for further processing.

**Table 14:** OSPF packet

1-byte Version field	1-byte Type field	2-byte Packet length	3-byte Router ID
4-byte Area ID	2-byte Checksum	2-byte Auth Type	8-byte Authentication
4-byte Network Mask	2-byte Hello interval	1-byte Options field	1-byte Router Priority
4-byte Dead Router interval	4-byte DR field	4-byte BDR field	4-byte Neighbor ID

The following descriptions summarize the OSPF packet header fields.

**Version field**— The OSPF version number. This specification documents version 2 of the protocol.

**Type field**—There are 5 OSPF packet types. From one to five, respectively, they are Hello, Database Description, Link State Request, Link State Update, and Link State Acknowledgment.

**Packet length**—The length of the OSPF protocol packet in bytes. This length includes the standard OSPF 24-byte header, so all OSPF packets are at 24-bytes long.

**Router ID**—The Router ID of the packet's source.

**Area ID**—A 32-bit number identifying the area that this packet belongs to. All OSPF packets are associated with a single area. Most travel a single hop only. Packets travelling over a virtual link are labelled with the backbone Area ID of 0.0.0.0.

**Checksum**—The standard IP checksum of the entire contents of the packet, starting with the OSPF packet header but excluding the 64-bit authentication field. This checksum is calculated as the 16-bit one's complement of the one's complement sum of all the 16-bit words in the packet, excepting the authentication field. If the packet's length is not an integral number of 16-bit words, the packet is padded with a byte of zero before checksumming. The checksum is considered to be part of the packet authentication procedure; for some authentication types the checksum calculation is omitted.

**Auth Type**—Identifies the authentication procedure to be used for the packet. Authentication types include Null authentication (0), Simple password (1), Cryptographic authentication (2), and all others are reserved for future use.

**Authentication**—A 64-bit field for use by the authentication scheme. When AuType indicates no authentication is being used, the Authentication fields is not checked and can be any value.

When AuType is set to 2 (Cryptographic authentication), the 64-bit authentication field is split into the following four fields: Zero field, Key ID field, Authentication data length field, and Cryptographic sequence field.

The Key ID field indicates the key and algorithm used to create the message digest appended to the packet. The authentication data length field indicates how many bytes long the message digest is, and the cryptographic sequence number is at non-decreasing number that is set when the packet is received and authenticated to prevent replay attacks.

**Network Mask**—The subnet where this packet is valid.

**Hello interval**—The period of time between sending out Hello packets. See [“Hello and dead intervals” on page 402](#).

**Options field**— The OSPF protocol defines several optional capabilities. A router indicates the optional capabilities that it supports in its OSPF Hello packets, Database Description packets and in its LSAs. This enables routers supporting a mix of optional capabilities to coexist in a single Autonomous System.

**Router priority**—The priority between 0 and 255 that determines which routers become the DR and BDR. See [“Designated router \(DR\) and backup router \(BDR\)” on page 398](#).

**Dead router interval**—The period of time when there is no response from a router before it is declared dead. See [“Hello and dead intervals” on page 402](#).

**DR and BDR fields**—The DR and BDR fields each list the router that fills that role on this network, generally the routers with the highest priorities. See [“Designated router \(DR\) and backup router \(BDR\)” on page 398](#).

**Neighbor ID**—The ID number of a neighboring router. This ID is used to discover new routers and respond to them.

## Troubleshooting OSPF

As with other dynamic routing protocols, OSPF has some issues that may need troubleshooting from time to time. For basic troubleshooting, see the Troubleshooting chapter.

The more common issues include:

- [Clearing OSPF routes from the routing table](#)
- [Checking the state of OSPF neighbors](#)
- [Passive interface problems](#)
- [Timer problems](#)
- [Authentication issues](#)
- [DR and BDR election issues](#)

### Clearing OSPF routes from the routing table

If you think the wrong route has been added to your routing table and you want to check it out, you first have to remove that route from your table before seeing if it is added back in or not. You can clear all or some OSPF neighbor connections (sessions) using the `execute router clear ospf` command. The `exec router clear` command is much more limiting for OSPF than it is for BGP. See [“Clearing routing table entries” on page 371](#).

For example, if you have routes in the OSPF routing table and you want to clear the specific route to IP address 10.10.10.1, you will have to clear all the OSPF entries. Enter the command:

```
execute router clear ospf process
```



## Checking the state of OSPF neighbors

In OSPF each router sends out link state advertisements to find other routers on its network segment, and to create adjacencies with some of those routers. This is important because routing updates are only passed between adjacent routers. If two routers you believe to be adjacent are not, that can be the source of routing failures.

To identify this problem, you need to check the state of the OSPF neighbors of your FortiGate unit. Use the CLI command `get router info ospf neighbor all` to see all the neighbors for your FortiGate unit. You will see output in the form of:

```
FGT1 # get router info ospf neighbor
OSPF process 0:
Neighbor ID Pri State Dead Time Address Interface
10.0.0.2 1 Full/ - 00:00:39 10.1.1.2 tunnel_wan1
10.0.0.2 1 Full/ - 00:00:34 10.1.1.4 tunnel_wan2
```

The important information here is the `State` column. Any neighbors that are not adjacent to your FortiGate unit will be reported in this column as something other than `Full`. If the state is `Down`, that router is offline.

## Passive interface problems

A passive OSPF interface doesn't send out any updates. This means it can't be a DR, BDR, or an area border router among other things. It will depend on other neighbor routers to update its link-state table.

Passive interfaces can cause problems when they aren't receiving the routing updates you expect from their neighbors. This will result in the passive OSPF FortiGate unit interface having an incomplete or out-of-date link-state database, and it will not be able to properly route its traffic. It is possible that the passive interface is causing a hole in the network where no routers are passing updates to each other, however this is a rare situation.

If a passive interface is causing problems, there are simple methods to determine it is the cause. The easiest method is to make it an active interface, and if the issues disappear, then that was the cause. Another method is to examine the OSPF routing table and related information to see if it is incomplete compared to other neighbor routers. If this is the case, you can clear the routing table, reset the device and allow it to repopulate the table.

If you cannot make the interface active for some reason, you will have to change your network to fix the "hole" by adding more routers, or changing the relationship between the passive router's neighbors to provide better coverage.

## Timer problems

A timer mismatch is when two routers have different values set for the same timer. For example if one router declares a router dead after 45 seconds and another waits for 4 minutes that difference in time will result in those two routers being out of synch for that period of time—one will still see that offline router as being online.

The easiest method to check the timers is to check the configuration on each router. Another method is to sniff some packets, and read the timer values in the packets themselves from different routers. Each packet contains the hello interval, and dead interval periods, so you can compare them easily enough.

## Bi-directional Forwarding Detection (BFD)

Bi-directional Forwarding Detection (BFD) is a protocol used to quickly locate hardware failures in the network. Routers running BFD communicate with each other, and if a timer runs out on a

connection then that router is declared down. BFD then communicates this information to the routing protocol and the routing information is updated.

## Authentication issues

OSPF has a number of authentication methods you can choose from. You may encounter problems with routers not authenticating as you expect. This will likely appear simply as one or more routers that have a blind spot in their routing - they won't acknowledge a router. This can be a problem if that router connects areas to the backbone as it will appear to be offline and unusable.

To confirm this is the issue, the easiest method is to turn off authentication on the neighboring routers. With no authentication between any routers, everything should flow normally.

Another method to confirm that authentication is the problem is to sniff packets, and look at their contents. The authentication type and password are right in the packets which makes it easy to confirm they are what you expect during real time. Its possible one or more routers is not configured as you expect and may be using the wrong authentication. This method is especially useful if there are a group of routers with these problems—it may only be one router causing the problem that is seen in multiple routers.

Once you have confirmed the problem is authentication related, you can decide how to handle it. You can turn off authentication and take your time to determine how to get your preferred authentication type back online. You can try another type of authentication, such as text instead of md5, which may have more success and still provide some level of protection. The important part is that once you confirm the problem, you can decide how to fix it properly.

## DR and BDR election issues

You can force a particular router to become the DR and BDR by setting their priorities higher than any other OSPF routers in the area. This is a good idea when those routers have more resources to handle the traffic and extra work of the DR and BDR roles, since not all routers may be able to handle all that traffic.

However, if you set all the other routers to not have a chance at being elected, a priority of zero, you can run into problems if the DR and BDR go offline. The good part is that you will have some warning generally as the DR goes offline and the BDR is promoted to the DR position. But if the network segment with both the DR and BDR goes down, your network will have no way to send hello packets, send updates, or the other tasks the DR performs.

The solution to this is to always allow routers to have a chance at being promoted, even if you set their priority to one. In that case they would be the last choice, but if there are no other candidates you want that router to become the DR. Most networks would have already alerted you to the equipment problems, so this would be a temporary measure to keep the network traffic moving until you can find and fix the problem to get the real DR back online.

## Basic OSPF example

This example sets up an OSPF network at a small office. There are 3 routers, all running OSPF v2. The border router connects to a BGP network.

All three routers in this example are FortiGate units. Router1 will be the designated router (DR) and router2 will be the backup DR (BDR) due to their priorities. Router3 will not be considered for either the DR or BDR elections. Instead, Router3 is the area border router (ASBR) routing all traffic to the ISP's BGP router on its way to the Internet.

Router2 has a modem connected that provides dialup access to the Internet as well, at a reduced bandwidth. This is a PPPoE connection to a DSL modem. This provides an alternate

route to the Internet if the other route goes down. The DSL connection is slow, and is charged by the amount of traffic. For these reasons OSPF will highly favor Router3's Internet access.

The DSL connection connects to an OSPF network with the ISP, so no redistribution of routes is required. The ISP network does have to be added to that router's configuration however.

This section includes the following topics:

- [Network layout and assumptions](#)
- [Configuring the FortiGate units](#)
- [Configuring OSPF on the FortiGate units](#)
- [Configuring other networking devices](#)
- [Testing network configuration](#)

## Network layout and assumptions

There are three FortiGate units acting as OSPF v2 routers on the network—Router1, Router2, and Router3. Router1 will be the designated router (DR), and Router 2 the BDR. Router3 is the area border router (ASBR) that connects to the external ISP router running BGP. Router2 has a PPPoE DSL connection that can access the Internet.

The Head Office network is connected to Router1 and Router2 on the 10.11.101.0 subnet.

Router1 and Router3 are connected over the 10.11.103.0 subnet.


Router2 and Router3 are connected over the 10.11.102.0 subnet.

The following table lists the router, interface, address, and role it is assigned.

**Table 15:** Routers, interfaces, and IP addresses for basic OSPF example network

Router name	Interface	IP address	Interface is connected to:
<b>Router1 (DR)</b>	Internal (port1)	10.11.101.1	Head office network, and Router2
	External (port2)	10.11.102.1	Router3
<b>Router2 (BDR)</b>	Internal (port1)	10.11.101.2	Head office network, and Router1
	External (port2)	10.11.103.2	Router3
	DSL (port3)	10.12.101.2	PPPoE DSL access
<b>Router3 (ASBR)</b>	Internal1 (port1)	10.11.102.3	Router1
	Internal2 (port2)	10.11.103.3	Router2
	External (port3)	172.20.120.3	ISP's BGP network

**Figure 110:**Basic OSPF network topology



Note that other subnets can be added to the internal interfaces without changing the configuration.

### Assumptions

- The FortiGate units used in this example have interfaces named port1, port2, and port3.
- All FortiGate units in this example have factory default configuration, and are in NAT/Route mode.
- Basic firewalls are in place to allow unfiltered traffic between all connected interfaces in both directions.
- This OSPF network is not connected to any other OSPF networks.
- Both Internet connections are always available.
- The modem connection is very slow and expensive.
- Other devices may be on the network, but do not affect this basic configuration.
- Router3 is responsible for redistributing all routes into and out of the OSPF AS.

### Configuring the FortiGate units

Each FortiGate unit needs the interfaces, and basic system information such as hostname configured.

This section includes:

- [Configuring Router1](#)
- [Configuring Router2](#)
- [Configuring Router3](#)

## Configuring Router1

Router1 has two interfaces connected to the network—internal (port1) and external (port2). Its host name must be changed to Router1.

### To configure Router1 interfaces - web-based manager

1. Go to *System > Dashboard > Status*.
2. Beside the host name, select *Change*.
3. Enter a hostname of `Router1`, and select *OK*.
4. Go to *System > Network > Interfaces*, edit port1, set the following information, and select *OK*.

<b>Alias</b>	internal
<b>IP/Netmask</b>	10.11.101.1/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Head office and Router2
<b>Administrative Status</b>	Up

5. Edit port2, set the following information, and select *OK*.

<b>Alias</b>	External
<b>IP/Netmask</b>	10.11.102.1/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Router3
<b>Administrative Status</b>	Up

## Configuring Router2

Router2 configuration is the same as Router1, except Router2 also has the DSL interface to configure.

The DSL interface is configured with a username of “user1” and a password of “ospf\_example”. The default gateway will be retrieved from the ISP, and the defaults will be used for the rest of the PPPoE settings.

### To configure Router2 interfaces - web-based manager

1. Go to *System > Dashboard > Status*.
2. Beside the host name, select *Change*.
3. Enter a hostname of `Router2`, and select *OK*.
4. Go to *System > Network > Interfaces*, edit port1, set the following information, and select *OK*.

<b>Alias</b>	internal
<b>IP/Netmask</b>	10.11.101.2/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING

<b>Description</b>	Head office and Router1
<b>Administrative Status</b>	Up

5. Edit port2, set the following information, and select *OK*.

<b>Alias</b>	External
<b>IP/Netmask</b>	10.11.103.2/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Router3
<b>Administrative Status</b>	Up

6. Edit DSL (port3), set the following information, and select *OK*.

<b>Alias</b>	DSL
<b>Addressing Mode</b>	PPPoE
<b>Username</b>	user1
<b>Password</b>	ospf_example
<b>Unnumbered IP address</b>	10.12.101.2/255.255.255.0
<b>Retrieve default gateway from server</b>	Enable
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	DSL
<b>Administrative Status</b>	Up

### Configuring Router3

Router3 is similar to Router1 and Router2 configurations. The main difference is the External (port3) interface connected to the ISP BGP network which has no administration access enabled for security reasons.

#### To configure Router3 interfaces - web-based manager

1. Go to *System > Status > Dashboard*.
2. Next to hostname, select *Change*.
3. Enter a hostname of *Router3*, and select *OK*.
4. Go to *System > Network > Interfaces*, edit port1, set the following information, and select *OK*.

<b>Alias</b>	internal
<b>IP/Netmask</b>	10.11.102.3/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING

<b>Description</b>	Router1
<b>Administrative Status</b>	Up

5. Edit port2, set the following information, and select *OK*.

<b>Alias</b>	Internal2
<b>IP/Netmask</b>	10.11.103.3/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Router2
<b>Administrative Status</b>	Up

6. Edit port3, set the following information, and select *OK*.

<b>Alias</b>	External
<b>IP/Netmask</b>	172.20.120.3/255.255.255.0
<b>Administrative Access</b>	
<b>Description</b>	ISP BGP
<b>Administrative Status</b>	Up

## Configuring OSPF on the FortiGate units

With the interfaces configured, now the FortiGate units can be configured for OSPF on those interfaces. All routers are part of the backbone 0.0.0.0 area, so there is no inter-area communications needed.

For a simple configuration there will be no authentication, no graceful restart or other advanced features, and timers will be left at their defaults. Also the costs for all interfaces will be left at 10, except for the modem and ISP interfaces where cost will be used to load balance traffic. Nearly all advanced features of OSPF are only available from the CLI.

The network that is defined covers all the subnets used in this example - 10.11.101.0, 10.11.102.0, and 10.11.103.0. All routes for these subnets will be advertised. If there are other interfaces on the FortiGate units that you do not want included in the OSPF routes, ensure those interfaces use a different subnet outside of the 10.11.0.0 network. If you want all interfaces to be advertised you can use an OSPF network of 0.0.0.0 .

Each router will configure:

- router ID
- area
- network
- two or three interfaces depending on the router
- priority for DR (Router1) and BDR (Router2)
- redistribute for ASBR (Router3)

This section includes:

- [Configuring OSPF on Router1](#)
- [Configuring OSPF on Router2](#)
- [Configuring OSPF on Router3](#)

## Configuring OSPF on Router1

Router1 has a very high priority to ensure it becomes the DR for this area. Also Router1 has the lowest IP address to help ensure it will win in case there is a tie at some point. Otherwise it is a standard OSPF configuration. Setting the priority can only be done in the CLI, and it is for a specific OSPF interface.

### To configure OSPF on Router1 - web-based manager

1. Go to *Router > Dynamic > OSPF*.
2. Set *Router ID* to `10.11.101.1` and select *Apply*.
3. In *Areas*, select *Create New*, set the following information, and select *OK*.

<b>Area</b>	0.0.0.0
<b>Type</b>	Regular
<b>Authentication</b>	none

4. In *Networks*, select *Create New*, set the following information, and select *OK*.

<b>IP/Netmask</b>	10.11.0.0/255.255.0.0
<b>Area</b>	0.0.0.0

5. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

<b>Name</b>	Router1-Internal-DR
<b>Interface</b>	port1 (Internal)
<b>IP</b>	0.0.0.0
<b>Authentication</b>	none
<b>Timers (seconds)</b>	
<b>Hello Interval</b>	10
<b>Dead Interval</b>	40

6. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

<b>Name</b>	Router1-External
<b>Interface</b>	port2 (External)
<b>IP</b>	0.0.0.0
<b>Authentication</b>	none



---

**Timers (seconds)**

---

<b>Hello Interval</b>	10
<b>Dead Interval</b>	40

---

7. Using the CLI, enter the following commands to set the priority for the Router1-Internal OSPF interface to maximum, ensuring this interface becomes the DR.

```
config router ospf
 config ospf-interface
 edit Router1-Internal-DR
 set priority 255
 end
```

**To configure OSPF on Router1 - CLI**

```
config router ospf
 set router-id 10.11.101.1
 config area
 edit 0.0.0.0
 next
 end
 config network
 edit 1
 set prefix 10.11.0.0/255.255.255.0
 next
 end
 config ospf-interface
 edit "Router1-Internal"
 set interface "port1"
 set priority 255
 next
 edit "Router1-External"
 set interface "port2"
 next
 end
end
```

**Configuring OSPF on Router2**

Router2 has a high priority to ensure it becomes the BDR for this area, and configures the DSL interface slightly differently—assume this will be a slower connection resulting in the need for longer timers, and a higher cost for this route.

Otherwise it is a standard OSPF configuration.

**To configure OSPF on Router2 - web-based manager**

1. Go to *Router > Dynamic > OSPF*.
2. Set *Router ID* to `10.11.101.2` and select *Apply*.

3. In *Areas*, select *Create New*, set the following information, and select *OK*.

<b>Area</b>	0.0.0.0
<b>Type</b>	Regular
<b>Authentication</b>	none

4. In *Networks*, select *Create New*, set the following information, and select *OK*.

<b>IP/Netmask</b>	10.11.0.0/255.255.0.0
<b>Area</b>	0.0.0.0

5. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

<b>Name</b>	Router2-Internal
<b>Interface</b>	port1 (Internal)
<b>IP</b>	0.0.0.0
<b>Authentication</b>	none
<b>Timers (seconds)</b>	
<b>Hello Interval</b>	10
<b>Dead Interval</b>	40

6. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

<b>Name</b>	Router2-External
<b>Interface</b>	port2 (External)
<b>IP</b>	0.0.0.0
<b>Authentication</b>	none
<b>Timers (seconds)</b>	
<b>Hello Interval</b>	10
<b>Dead Interval</b>	40

7. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

<b>Name</b>	Router2-DSL
<b>Interface</b>	port3 (DSL)
<b>IP</b>	0.0.0.0
<b>Authentication</b>	none
<b>Cost</b>	50

---

**Timers (seconds)**

---

<b>Hello Interval</b>	20
<b>Dead Interval</b>	80

---

8. Using the CLI, enter the following commands to set the priority for the Router2-Internal OSPF interface to ensure this interface will become the BDR.

```
config router ospf
 config ospf-interface
 edit Router2-Internal
 set priority 250
 next
 end
```

### To configure OSPF on Router2 - CLI

```
config router ospf
 set router-id 10.11.101.2
 config area
 edit 0.0.0.0
 next
 end
 config network
 edit 1
 set prefix 10.11.0.0/255.255.0.0
 next
 end
 config ospf-interface
 edit "Router2-Internal"
 set interface "port1"
 set priority 255
 next
 edit "Router2-External"
 set interface "port2"
 next
 edit "Router2-DSL"
 set interface "port3"
 set cost 50
 next
 end
end
```

### Configuring OSPF on Router3

Router3 is more complex than the other two routers. The interfaces are straightforward, but this router has to import and export routes between OSPF and BGP. That requirement makes Router3 a border router or ASBR. Also Router3 needs a lower cost on its route to encourage all traffic to the Internet to route through it.

In the advanced OSPF options, Redistribute is enabled for Router3. It allows different types of routes, learned outside of OSPF, to be used in OSPF. Different metrics are assigned to these other types of routes to make them more or less preferred to regular OSPF routes.

### To configure OSPF on Router3 - web-based manager

1. Go to *Router > Dynamic > OSPF*.
2. Set *Router ID* to `10.11.101.2` and select *Apply*.
3. Expand *Advanced Options*.
4. In *Redistribute*, set the following information, and select *OK*.

Route type	Redistribute	Metric
Connected	Enable	15
Static	Enable	15
RIP	Disable	n/a
BGP	Enable	5

5. In *Areas*, select *Create New*, set the following information, and select *OK*.

Area	0.0.0.0
Type	Regular
Authentication	none

6. In *Networks*, select *Create New*, set the following information, and select *OK*.

IP/Netmask	10.11.0.0/255.255.0.0
Area	0.0.0.0

7. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	Router3-Internal
Interface	port1 (Internal)
IP	0.0.0.0
Authentication	none
Timers (seconds)	
Hello Interval	10
Dead Interval	40

8. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	Router3-Internal2
Interface	port2 (Internal2)
IP	0.0.0.0

<b>Authentication</b>	none
<b>Timers (seconds)</b>	
<b>Hello Interval</b>	10
<b>Dead Interval</b>	40

9. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

<b>Name</b>	Router3-ISP-BGP
<b>Interface</b>	port3 (ISP-BGP)
<b>IP</b>	0.0.0.0
<b>Authentication</b>	none
<b>Cost</b>	2
<b>Timers (seconds)</b>	
<b>Hello Interval</b>	20
<b>Dead Interval</b>	80

10. Using the CLI, enter the following commands to set the priority for the Router3-Internal OSPF interface to ensure this interface will become the BDR.

```
config router ospf
 config ospf-interface
 edit Router3-Internal
 set priority 250
 next
end
```

### To configure OSPF on Router3 - CLI

```
config router ospf
 set router-id 10.11.102.3
 config area
 edit 0.0.0.0
 next
 end
 config network
 edit 1
 set prefix 10.11.0.0/255.255.255.0
 next
 edit 2
 set prefix 172.20.120.0/255.255.255.0
 next
 end
 config ospf-interface
 edit "Router3-Internal"
 set interface "port1"
 set priority 255
 next
 edit "Router3-External"
 set interface "port2"
 next
 edit "Router3-ISP-BGP"
 set interface "port3"
 set cost 2
 next
 end
end
```

## Configuring other networking devices

The other networking devices required in this configuration are on the two ISP networks, the BGP network for the main Internet connection, and the DSL backup connection.

In both cases, the ISPs need to be notified of the OSPF network settings including router IP addresses, timer settings, and so on. The ISP will use this information to configure its routers that connect to this OSPF network.

## Testing network configuration

Testing the network configuration involves two parts: testing the network connectivity, and testing the OSPF routing.

To test the network connectivity use ping, traceroute, and other network tools.

To test the OSPF routing in this example, refer to the troubleshooting outlined in [“Troubleshooting OSPF” on page 408](#).

## Advanced inter-area OSPF example

This example sets up an OSPF network at a large office. There are three areas, each with two routers. Typically OSPF areas would not be this small, and if they were the areas would be combined into one bigger area. However, the stub area services the accounting department which is very sensitive about their network and do not want any of their network information broadcast through the rest of the company. The backbone area contains the bulk of the company network devices. The regular area was established for various reasons such as hosting the company servers on a separate area with extra security.

One area is a small stub area that has no independent Internet connection, and only one connection to the backbone area. That connection between the stub area and the backbone area is only through a default route. No routes outside the stub area are advertised into that area. Another area is the backbone, which is connected to the other two areas. The third area has the Internet connection, and all traffic to and from the Internet must use that area's connection. If that traffic comes from the stub area, then that traffic is treating the backbone like a transit area that only uses it to get to another area.

In the stub area, a subnet of computers is running the RIP routing protocol and those routes must be redistributed into the OSPF areas.

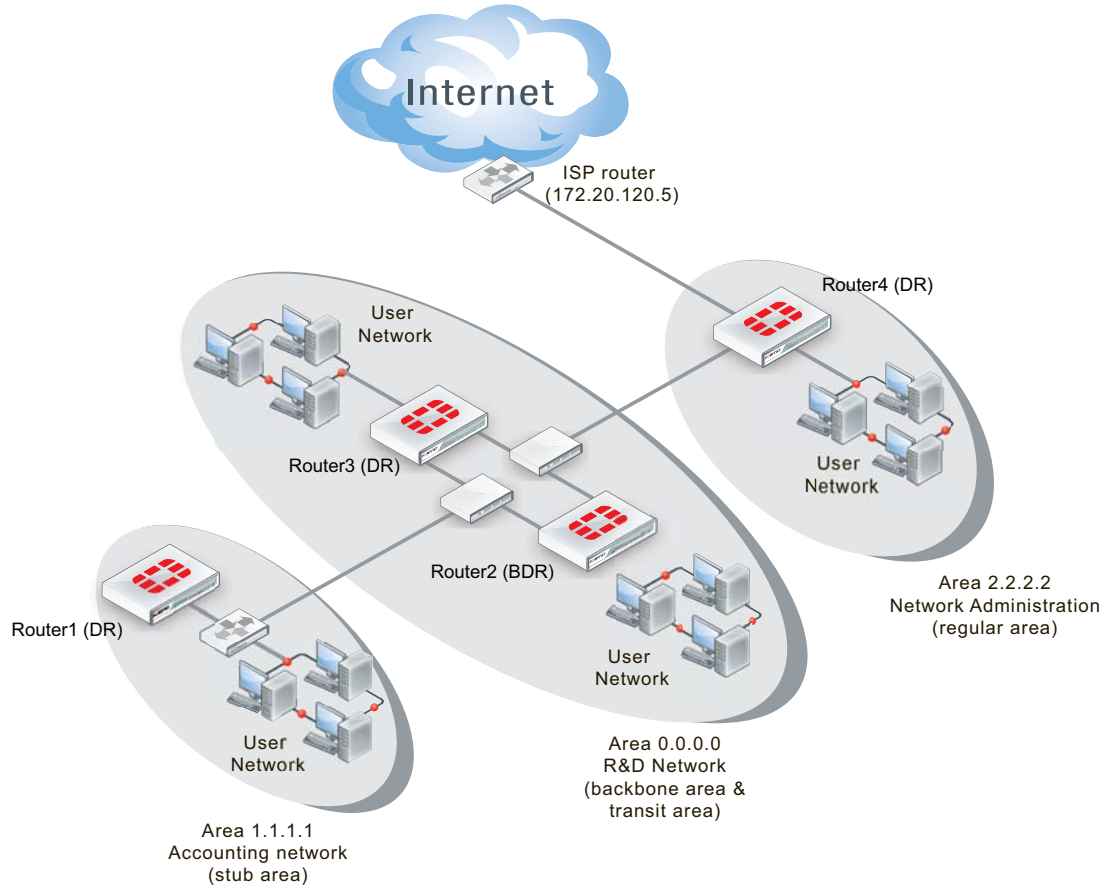
This section includes the following topics:

- [Network layout and assumptions](#)
- [Configuring the FortiGate units](#)
- [Configuring OSPF on the FortiGate units](#)
- [Configuring other networking devices](#)
- [Testing network configuration](#)

### Network layout and assumptions

There are four FortiGate units in this network topology acting as OSPF routers.

**Figure 111:**Advanced inter-area OSPF network topology



Area 1.1.1.1 is a stub area with one FortiGate unit OSPF router called Router1 (DR). Its only access outside of that area is a default route to the backbone area, which is how it accesses the Internet. Traffic must go from the stub area, through the backbone, to the third area to reach the Internet. The backbone area in this configuration is called a transit area. Also in area 1.1.1.1 there is a RIP router that will be providing routes to the OSPF area through redistribution.

Area 0.0.0.0 is the backbone area, and has two FortiGate unit routers named Router2 (BDR) and Router3 (DR).

Area 2.2.2.2 is a regular area that has an Internet connection accessed by both the other two OSPF areas. There is only one FortiGate unit router in this area called Router4 (DR). This area is more secure and requires MD5 authentication by routers.

All areas have user networks connected, but they are not important for configuring the network layout for this example.

Internal interfaces are connected to internal user networks only. External1 interfaces are connected to the 10.11.110.0 network, joining Area 1.1.1.1 and Area 0.0.0.0.

External2 interfaces are connected to the 10.11.111.0 network, joining Area 0.0.0.0 and Area 2.2.2.2. The ISP interface is called ISP.

**Table 16:** Routers, areas, interfaces, IP addresses for advanced OSPF network

Router name	Area number and type	Interface	IP address
Router1 (DR)	1.1.1.1 - stub area (Accounting)	port1 (internal)	10.11.101.1
		port2 (external1)	10.11.110.1



**Table 16:** Routers, areas, interfaces, IP addresses for advanced OSPF network

<b>Router2 (BDR)</b>	0.0.0.0 - backbone area ( R&D Network)	port1 (internal)	10.11.102.2
		port2 (external1)	10.11.110.2
		port3 (external2)	10.11.111.2
<b>Router3 (DR)</b>	0.0.0.0 - backbone area (R&D Network)	port1 (internal)	10.11.103.3
		port2 (external1)	10.11.110.3
		port3 (external2)	10.11.111.3
<b>Router4 (DR)</b>	2.2.2.2 - regular area (Network Admin)	port1 (internal)	10.11.104.4
		port2 (external2)	10.11.111.4
		port3 (ISP)	172.20.120.4

Note that other subnets can be added to the internal interfaces without changing the configuration.

### Assumptions

- The FortiGate units used in this example have interfaces named port1, port2, and port3.
- All FortiGate units in this example have factory default configuration, and are in NAT/Route mode.
- During configuration, if settings are not directly referred to they will be left at default settings.
- Basic firewalls are in place to allow unfiltered traffic between all connected interfaces in both directions.
- This OSPF network is not connected to any other OSPF areas outside of this example.
- The Internet connection is always available.
- Other devices may be on the network, but do not affect this configuration.

## Configuring the FortiGate units

This section configures the basic settings on the FortiGate units to be OSPF routers. These configurations include multiple interface settings, and hostname.

There are four FortiGate units in this example. The two units in the backbone area can be configured exactly the same except for IP addresses, so only router3 (the DR) configuration will be given with notes indicating router2 (the BDR) IP addresses.

Configuring the FortiGate units includes:

- [Configuring Router1](#)
- [Configuring Router2](#)
- [Configuring Router3](#)
- [Configuring Router4](#)

### Configuring Router1

Router1 is part of the Accounting network stub area (1.1.1.1).

### To configure Router1 interfaces - web-based manager

1. Go to *System > Dashboard > Status*.
2. Next to hostname, select *Change*.
3. Enter a hostname of `Router1`, and select *OK*.
4. Go to *System > Network > Interfaces*, edit port1, set the following information, and select *OK*.

<b>Alias</b>	internal
<b>IP/Netmask</b>	10.11.101.1/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Accounting network
<b>Administrative Status</b>	Up

5. Edit port2, set the following information, and select *OK*.

<b>Alias</b>	External1
<b>IP/Netmask</b>	10.11.110.1/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Backbone network and internet
<b>Administrative Status</b>	Up

### Configuring Router2

Router2 is part of the R&D network backbone area (0.0.0.0). Router2 and Router3 are in this area. They provide a redundant connection between area 1.1.1.1 and area 2.2.2.2.

Router2 has three interfaces configured; one to the internal network, and two to Router3 for redundancy.

### To configure Router2 interfaces - web-based manager

1. Go to *System > Dashboard > Status*.
2. Next to hostname, select *Change*.
3. Enter a hostname of `Router2`, and select *OK*.
4. Go to *System > Network > Interfaces*, edit port1 (internal), set the following information, and select *OK*.

<b>Alias</b>	internal
<b>IP/Netmask</b>	10.11.102.2/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Internal RnD network
<b>Administrative Status</b>	Up

5. Edit port2 (external1), set the following information, and select *OK*.

---

<b>Alias</b>	external1
<b>IP/Netmask</b>	10.11.110.2/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Router3 first connection
<b>Administrative Status</b>	Up

---

6. Edit port3 (external2), set the following information, and select *OK*.

---

<b>Alias</b>	external2
<b>IP/Netmask</b>	10.11.111.2/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Router3 second connection
<b>Administrative Status</b>	Up

---

### Configuring Router3

Router3 is part of the R&D network backbone area (0.0.0.0). Router2 and Router3 are in this area. They provide a redundant connection between area 1.1.1.1 and area 2.2.2.2.

#### To configure Router3 interfaces - web-based manager

1. Go to *System > Dashboard > Status*.
2. Next to hostname, select *Change*.
3. Enter a hostname of `Router3`, and select *OK*.
4. Go to *System > Network > Interfaces*, edit port1 (internal), set the following information, and select *OK*.

---

<b>Alias</b>	internal
<b>IP/Netmask</b>	10.11.103.3/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Internal RnD network
<b>Administrative Status</b>	Up

---

5. Edit port2 (external1), set the following information, and select *OK*.

---

<b>Alias</b>	external1
<b>IP/Netmask</b>	10.11.110.3/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING

---

<b>Description</b>	Router2 first connection
<b>Administrative Status</b>	Up

6. Edit port3 (external2), set the following information, and select *OK*.

<b>Alias</b>	external2
<b>IP/Netmask</b>	10.11.111.3/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Router2 second connection
<b>Administrative Status</b>	Up

## Configuring Router4

Router4 is part of the Network Administration regular area (2.2.2.2). This area provides internet access for both area 1.1.1.1 and the backbone area.

This section configures interfaces and hostname.

### To configure Router4 interfaces - web-based manager

1. Go to *System > Dashboard > Status*.
2. Next to hostname, select *Change*.
3. Enter a hostname of Router4, and select *OK*.
1. Go to *System > Network > Interfaces*.
2. Edit port1 (internal).
3. Set the following information, and select *OK*.

<b>Alias</b>	internal
<b>IP/Netmask</b>	10.11.101.4/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Accounting network
<b>Administrative Status</b>	Up

4. Edit port2 (external2).
5. Set the following information, and select *OK*.

<b>Alias</b>	external2
<b>IP/Netmask</b>	10.11.110.4/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Backbone and Accounting network
<b>Administrative Status</b>	Up

6. Edit port3 (ISP).

7. Set the following information, and select *OK*.

<b>Alias</b>	ISP
<b>IP/Netmask</b>	172.20.120.4/255.255.255.0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	ISP and internet
<b>Administrative Status</b>	Up

## Configuring OSPF on the FortiGate units

Three of the routers are designated routers (DR) and one is a backup DR (BDR). This is achieved through the lowest router ID numbers, or OSPF priority settings.

Also each area needs to be configured as each respective type of area - stub, backbone, or regular. This affects how routes are advertised into the area.

### To configure OSPF on Router1 - web-based manager

1. Go to *Router > Dynamic > OSPF*.
2. Enter 10.11.101.1 for the *Router ID* and select *Apply*.
3. In *Areas*, select *Create New*, set the following information, and select *OK*.

<b>Area</b>	1.1.1.1
<b>Type</b>	Stub
<b>Authentication</b>	None

4. In *Networks*, select *Create New*, set the following information, and select *OK*.

<b>IP/Netmask</b>	10.11.101.0/255.255.255.0
<b>Area</b>	1.1.1.1

5. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

<b>Name</b>	Accounting
<b>Interface</b>	port1 (internal)
<b>IP</b>	10.11.101.1
<b>Authentication</b>	None

6. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

<b>Name</b>	Backbone1
<b>Interface</b>	port2 (external1)
<b>IP</b>	10.11.110.1
<b>Authentication</b>	None

### To configure OSPF on Router2 - web-based manager

1. Go to *Router > Dynamic > OSPF*.
2. Enter 10.11.102.2 for the *Router ID* and select *Apply*.
3. In *Areas*, select *Create New*, set the following information, and select *OK*.

---

<b>Area</b>	0.0.0.0
-------------	---------

---

<b>Type</b>	Regular
-------------	---------

---

<b>Authentication</b>	None
-----------------------	------

---

4. In *Networks*, select *Create New*, set the following information, and select *OK*.

---

<b>IP/Netmask</b>	10.11.102.2/255.255.255.0
-------------------	---------------------------

---

<b>Area</b>	0.0.0.0
-------------	---------

---

5. In *Networks*, select *Create New*, set the following information, and select *OK*.

---

<b>IP/Netmask</b>	10.11.110.2/255.255.255.0
-------------------	---------------------------

---

<b>Area</b>	0.0.0.0
-------------	---------

---

6. In *Networks*, select *Create New*, set the following information, and select *OK*.

---

<b>IP/Netmask</b>	10.11.111.2/255.255.255.0
-------------------	---------------------------

---

<b>Area</b>	0.0.0.0
-------------	---------

---

7. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

---

<b>Name</b>	RnD network
-------------	-------------

---

<b>Interface</b>	port1 (internal)
------------------	------------------

---

<b>IP</b>	10.11.102.2
-----------	-------------

---

<b>Authentication</b>	None
-----------------------	------

---

8. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

---

<b>Name</b>	Backbone1
-------------	-----------

---

<b>Interface</b>	port2 (external1)
------------------	-------------------

---

<b>IP</b>	10.11.110.2
-----------	-------------

---

<b>Authentication</b>	None
-----------------------	------

---

9. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

---

<b>Name</b>	Backbone2
-------------	-----------

---

<b>Interface</b>	port3 (external2)
------------------	-------------------

---

<b>IP</b>	10.11.111.2
<b>Authentication</b>	None

**To configure OSPF on Router3 - web-based manager**

1. Go to *Router > Dynamic > OSPF*.
2. Enter 10.11.103.3 for the *Router ID* and then select *Apply*.
3. In *Areas*, select *Create New*, set the following information, and then select *OK*.

<b>Area</b>	0.0.0.0
<b>Type</b>	Regular
<b>Authentication</b>	None

4. In *Networks*, select *Create New*, set the following information, and select *OK*.

<b>IP/Netmask</b>	10.11.102.3/255.255.255.0
<b>Area</b>	0.0.0.0

5. In *Networks*, select *Create New*, set the following information, and select *OK*.

<b>IP/Netmask</b>	10.11.110.3/255.255.255.0
<b>Area</b>	0.0.0.0

6. In *Networks*, select *Create New*, set the following information, and select *OK*.

<b>IP/Netmask</b>	10.11.111.3/255.255.255.0
<b>Area</b>	0.0.0.0

7. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

<b>Name</b>	RnD network
<b>Interface</b>	port1 (internal)
<b>IP</b>	10.11.103.3
<b>Authentication</b>	None

8. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

<b>Name</b>	Backbone1
<b>Interface</b>	port2 (external1)
<b>IP</b>	10.11.110.3
<b>Authentication</b>	None

9. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

<b>Name</b>	Backbone2
<b>Interface</b>	port3 (external2)
<b>IP</b>	10.11.111.3
<b>Authentication</b>	None

**To configure OSPF on Router4 - web-based manager**

1. Go to *Router > Dynamic > OSPF*.
2. Enter 10.11.104.4 for the *Router ID* and then select *Apply*.
3. In *Areas*, select *Create New*.
4. Set the following information, and select *OK*.

<b>Area</b>	2.2.2.2
<b>Type</b>	Regular
<b>Authentication</b>	None

5. In *Networks*, select *Create New*, set the following information, and select *OK*.

<b>IP/Netmask</b>	10.11.104.0/255.255.255.0
<b>Area</b>	0.0.0.0

6. In *Networks*, select *Create New*, set the following information, and select *OK*.

<b>IP/Netmask</b>	10.11.111.0/255.255.255.0
<b>Area</b>	0.0.0.0

7. In *Networks*, select *Create New*, set the following information, and select *OK*.

<b>IP/Netmask</b>	172.20.120.0/255.255.255.0
<b>Area</b>	0.0.0.0

8. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

<b>Name</b>	Network Admin network
<b>Interface</b>	port1 (internal)
<b>IP</b>	10.11.104.4
<b>Authentication</b>	None



9. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

<b>Name</b>	Backbone2
<b>Interface</b>	port2 (external2)
<b>IP</b>	10.11.111.4
<b>Authentication</b>	None

10. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

<b>Name</b>	ISP
<b>Interface</b>	port3 (ISP)
<b>IP</b>	172.20.120.4
<b>Authentication</b>	None

## Configuring other networking devices

All network devices on this network are running OSPF routing. The user networks (Accounting, R&D, and Network Administration) are part of one of the three areas.

The ISP needs to be notified of your network configuration for area 2.2.2.2. Your ISP will not advertise your areas externally as they are intended as internal areas. External areas have assigned unique numbers. The area numbers used in this example are similar to the 10.0.0.0 and 192.168.0.0 subnets used in internal networking.

## Testing network configuration

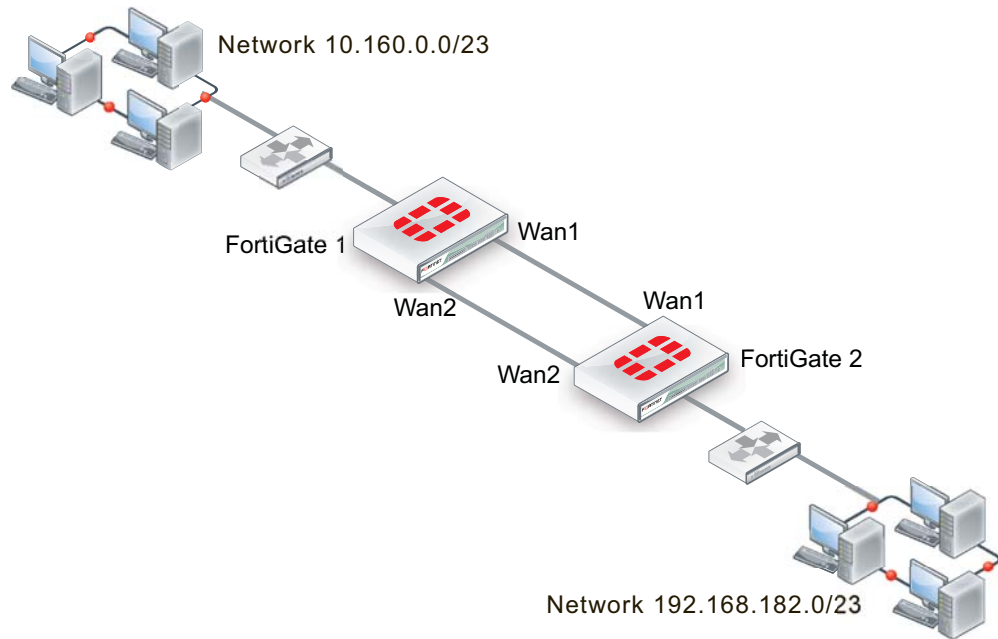
There are two main areas to test in this network configuration; network connectivity, and OSPF routing.

To test the network connectivity, see if computers on the Accounting or R&D networks can access the internet. If you need to troubleshoot network connectivity, see the [Troubleshooting](#) chapter.

To test the OSPF routing, check the routing tables on the FortiGate units to ensure the expected OSPF routes are present. If you need help troubleshooting OSPF routing, see [“Troubleshooting OSPF” on page 408](#).

## Controlling redundant links by cost

In this scenario, two FortiGate units have redundant links: one link between their WAN1 interfaces and another between their WAN2 interfaces.



FortiGate 1 should learn the route to network 192.168.182.0 and FortiGate 2 should learn the route to network 10.160.0.0. Under normal conditions, they should learn these routes through the WAN1 link. The WAN2 link should be used only as a backup.

With the default settings, each FortiGate unit learns these routes from both WAN1 and WAN2.

#### **FortiGate 1:**

```
FGT1 # get router info ospf neighbor
OSPF process 0:
Neighbor ID Pri State Dead Time Address Interface
10.2.2.2 1 Full/Backup 00:00:33 10.182.0.187 wan1
10.2.2.2 1 Full/Backup 00:00:31 10.183.0.187 wan2

FGT1 # get router info routing-table ospf
O*E2 0.0.0.0/0 [110/10] via 10.183.0.187, wan2, 00:00:01
[110/10] via 10.182.0.187, wan1, 00:00:01
O 192.168.182.0/23 [110/20] via 10.183.0.187, wan2, 00:02:04
[110/20] via 10.182.0.187, wan1, 00:02:04
```

#### **FortiGate 2:**

```
FGT2 # get router info ospf neighbor
OSPF process 0:
Neighbor ID Pri State Dead Time Address Interface
10.1.1.1 1 Full/DR 00:00:38 10.182.0.57 wan1
10.1.1.1 1 Full/DR 00:00:38 10.183.0.57 wan2

FGT2 # get router info routing-table ospf
O 10.160.0.0/23 [110/20] via 10.183.0.57, wan2, 00:00:39
[110/20] via 10.182.0.57, wan1, 00:00:39
```

## Adjusting the route costs

On both FortiGate units, the cost of the route through WAN2 is adjusted higher so that this route will only be used if the route through WAN1 is unavailable. The default cost is 10. The WAN2 route will be changed to a cost of 200.

### **On both FortiGate units:**

```
config router ospf
 config ospf-interface
 edit "WAN2_higher_cost"
 set cost 200
 set interface "wan2"
 end
```

Now both FortiGate units use only the WAN1 route:

### **FortiGate 1:**

```
FGT1 # get router info routing-table ospf
O*E2 0.0.0.0/0 [110/10] via 10.182.0.187, wan1, 00:00:40
O 192.168.182.0/23 [110/20] via 10.182.0.187, wan1, 00:00:40
```

### **FortiGate 2:**

```
FGT2 # get router info routing-table ospf
O 10.160.0.0/23 [110/20] via 10.182.0.57, wan1, 00:09:37
```

### **LSDB check on FortiGate 1:**

```
FGT1 # get router info ospf database router lsa
Router Link States (Area 0.0.0.0)
LS age: 81
Options: 0x2 (*|-|-|-|-|E|-)
Flags: 0x0
LS Type: router-LSA
Link State ID: 10.1.1.1
Advertising Router: 10.1.1.1
LS Seq Number: 8000000b
Checksum: 0xe637
Length: 60
Number of Links: 3
```

```
Link connected to: Stub Network
(Link ID) Network/subnet number: 10.160.0.0
(Link Data) Network Mask: 255.255.254.0
Number of TOS metrics: 0
TOS 0 Metric: 10
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 10.183.0.187
(Link Data) Router Interface address: 10.183.0.57
Number of TOS metrics: 0
TOS 0 Metric: 200
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 10.182.0.57
(Link Data) Router Interface address: 10.182.0.57
Number of TOS metrics: 0
TOS 0 Metric: 10
```

```
LS age: 83
Options: 0x2 (*|-|-|-|-|E|-)
Flags: 0x2 : ASBR
LS Type: router-LSA
Link State ID: 10.2.2.2
Advertising Router: 10.2.2.2
LS Seq Number: 8000000e
Checksum: 0xfc9b
Length: 60
Number of Links: 3
```

```
Link connected to: Stub Network
(Link ID) Network/subnet number: 192.168.182.0
(Link Data) Network Mask: 255.255.254.0
Number of TOS metrics: 0
TOS 0 Metric: 10
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 10.183.0.187
(Link Data) Router Interface address: 10.183.0.187
Number of TOS metrics: 0
TOS 0 Metric: 200
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 10.182.0.57
(Link Data) Router Interface address: 10.182.0.187
Number of TOS metrics: 0
TOS 0 Metric: 10
```

## Verifying route redundancy

Bring down WAN1 and then check the routes on the two FortiGate units.

### **FortiGate 1:**

```
FGT1 # get router info routing-table ospf
FGT1 # get router info routing-table ospf
O*E2 0.0.0.0/0 [110/10] via 10.183.0.187, wan2, 00:00:06
O 192.168.182.0/23 [110/210] via 10.183.0.187, wan2, 00:00:06
```

### **FortiGate 2:**

```
FGT2 # get router info routing-table ospf
O 10.160.0.0/23 [110/210] via 10.183.0.57, wan2, 00:00:14
```

The WAN2 interface is now in use on both units.

# Intermediate System to Intermediate System Protocol (IS-IS)

This section describes the Intermediate System to Intermediate System Protocol (IS-IS).

The following topics are included in this section:

- [IS-IS background and concepts](#)
- [Troubleshooting IS-IS](#)
- [Simple IS-IS example](#)

## IS-IS background and concepts

Intermediate System to Intermediate System Protocol (IS-IS) allows routing of ISO's OSI protocol stack Connectionless Network Service (CLNS). IS-IS is an Interior Gateway Protocol (IGP) not intended to be used between Autonomous Systems (ASes).

This section contains:

- [Background](#)
- [How IS-IS works](#)
- [Parts and terminology of IS-IS](#)

### Background

IS-IS was developed by Digital Equipment Corporation and later standardized by ISO in 1992 as ISO 19589 (see [RFC 1142](#)—note this RFC is different from the ISO version). At roughly the same time, the Internet Engineering Task Force developed OSPF (see "[Open Shortest Path First \(OSPF\)](#)" on [page 396](#)). After the initial version, IP support was added to IS-IS and this version was called Integrated IS-IS (see [RFC 1195](#)). Its widespread use started when an early version of IS-IS was included with BSD v4.3 Linux as the routed daemon. The routing algorithm used by IS-IS, the Bellman–Ford algorithm, first saw widespread use as the initial routing algorithm of the ARPANET.

IS-IS is a link state protocol well-suited to smaller networks that is in widespread use and has near universal support on routing hardware. It is quick to configure, and works well if there are no redundant paths. However, IS-IS updates are sent out node-by-node, so it can be slow to find a path around network outages. IS-IS also lacks good authentication, can not choose routes based on different quality of service methods, and can create network loops if you are not careful. IS-IS uses Dijkstra's algorithm to find the best path, like OSPF.

While OSPF is more widely known, IS-IS is a viable alternative to OSPF in enterprise networks and ISP infrastructures, largely due to its native support for IPv6 and its non-disruptive methods for splitting, merging, migrating, and renumbering network areas.

The FortiGate implementation supports both IS-IS (see [RFCs 1142](#) and [1162](#)) and Integrated IS-IS (see [RFCs 1195](#) and [5308](#)).

## How IS-IS works

As one of the original modern dynamic routing protocols, IS-IS is straightforward. Its routing algorithm is not complex, there are some options to allow fine tuning, and it is straightforward to configure IS-IS on FortiGate units.

From [RFC 1142](#):

*The routing algorithm used by the Decision Process is a shortest path first (SPF) algorithm. Instances of the algorithm are run independently and concurrently by all intermediate systems in a routing domain. IntraDomain routing of a PDU occurs on a hop-by-hop basis: that is, the algorithm determines only the next hop, not the complete path, that a data PDU will take to reach its destination.*

### IS-IS versus static routing

IS-IS was one of the earliest dynamic routing protocols to work with IP addresses. As such, it is not as complex as more recent protocols. However, IS-IS is a big step forward from simple static routing.

While IS-IS may be slow in response to network outages, static routing has zero response. The same is true for convergence—static routing has zero convergence. Both IS-IS and static routing have the limited hop count, so it is neither a strength nor a weakness.

### TLV

IS-IS uses *type-length-variable (TLV)* parameters to carry information in Link-State PDUs (LSPs). Each IS-IS LSP consists of a variable-length header to which TLVs are appended in order to extend IS-IS for IP routing. The TLV field consists of one octet of type (T), one octet of length (L), and “L” octets of Value (V). They are included in all of the IS-IS [Packet types](#). For a complete breakdown of the LSP, refer to [“LSP structure” on page 439](#).

In IS-IS, TLVs are used to determine route-leaking and authentication, and are also used for IPv4 and IPv6 awareness and reachability.

- To determine which TLVs are responsible for route-leaking, refer to [“Default routing” on page 442](#).
- To determine which TLVs are responsible for authentication, refer to [“Authentication” on page 443](#).
- To determine which TLVs are responsible for IPv4 and IPv6 awareness and reachability, refer to [“Integrated IS-IS” on page 444](#).

For a complete list of reserved TLV codepoints, refer to [RFC 3359](#).

### LSP structure

It is difficult to fully understand a routing protocol without knowing what information is carried in its packets. Knowing how routers exchange each type of information will help you better understand the IS-IS protocol and will allow you to configure your network more appropriately.

This section provides information on the contents of the IS-IS LSP. LSPs describe the network topology and can include IP routes and checksums.

### NSAP and NET

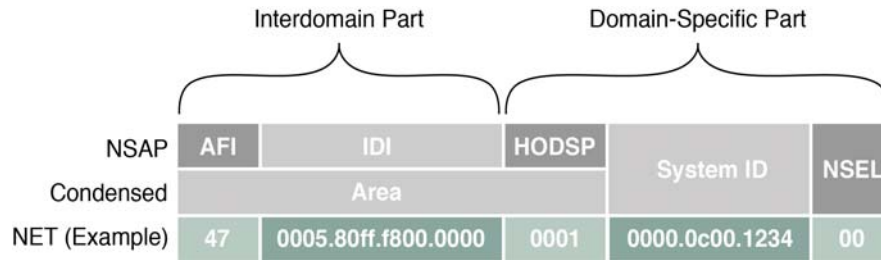
IS-IS routing protocol utilizes ISO network addressing to identify network interfaces. The addresses are known as Network Service Access Points (NSAPs). In general, IS-IS routers consist of only one NSAP, whereas IP addressing requires one IP address per interface.

In IS-IS, the NSAP address is translated into a Network Entity Title (NET), which is the same as the NSAP but can differentiate end systems by way of a byte called the *n-selector* (NSEL). In

order for adjacencies to form in IS-IS, the NSEL must necessarily be set to zero, to indicate “this system”. The total NET can be anywhere between 8 and 20 bytes long due to the support for variable length area addressing.

The following diagram identifies the individual parts of the NSAP, with explanations below.

**Table 17:** NSAP and NET example



**AFI** — The *Authority and Format Identifier (AFI)* specifies the format of the addressing family used. IS-IS is designed to carry routing information for several different protocols. Each entry has an address family identifier that identifies the globally unique Interdomain Part (IDP). For example, 49 is the AFI for private addresses, whereas 47 is the AFI for international organizations.

**IDI** — The *Initial Domain Identifier (IDI)* identifies the routing domain within an interconnected network. The length of the IDI is typically determined by the AFI. If you are using an AFI of 49, you do not need to specify an IDI, since the network is private.

**HODSP** — The *High Order Domain-Specific Part (HODSP)* identifies the unique address within a specific routing domain. Together, the AFI, IDI, and HODSP define the area address. All of the nodes within an area must have the same area address.

**System ID** — The *System ID* represents the 6-8 byte router identifier. The ID could be Media Access Control (MAC) format, as in the example above, or a static length IP address expressed in binary-coded decimal (BCD) format.

**NSEL** — The *n-selector (NSEL)*, as previously described, identifies the network layer transport service and must always be set to zero for IS-IS NETs.

## Parts and terminology of IS-IS

Before you can understand how IS-IS functions, you need to understand some of the main concepts and parts of IS-IS.

This section includes:

- [DIS election and pseudonode LSP](#)
- [Packet types](#)
- [Default routing](#)
- [Timer options](#)
- [Authentication](#)
- [Integrated IS-IS](#)

### DIS election and pseudonode LSP

In IS-IS routing protocol, a single router is chosen to be the designated intermediate system (DIS). The election of the DIS is determined automatically and dynamically on the LAN depending on highest interface priority and the subnetwork point of attachment (SNPA). The FortiGate is typically the DIS, and each router in its LAN is an intermediate system (IS).



Unlike OSPF, which elects a designated router (DR) and backup designated router (BDR), the DIS has no backup and determines the election of a new DIS whenever a router is added to the LAN or whenever the current DIS drops. A backup DIS is irrelevant since all of the routers on an IS-IS system are synchronized, and the short Hello interval used by the DIS quickly detects failures and the subsequent replacement of the DIS.

Synchronization of all the nodes in an IS-IS area could prove troublesome when updating the network infrastructure, and would demand ever-increasing resources each time a new router is added (at an exponential scale). For this purpose the DIS creates a pseudonode, which is essentially a virtual, logical node representing the LAN. The pseudonode requests adjacency status from all the routers in a multi-access network by sending IS-IS Hello (IIH) PDUs to Level 1 and Level 2 routers (where Level 1 routers share the same address as the DIS and Level 2 routers do not). Using a pseudonode to alter the representation of the LAN in the link-state database (LSD) greatly reduces the amount of adjacencies that area routers have to report. In essence, a pseudonode *collapses* a LAN topology, which allows a more linear scale to link-state advertising.

In order to maintain the database synchronization, the DIS periodically sends complete sequence number packets (CSNPs) to all participating routers.

## Packet types

Four general packet types (PDUs) are communicated through IS-IS, appearing at both Level 1 and Level 2. They are described below.

**Intermediate System-to-Intermediate System Hello (IIH) PDU** — As mentioned previously, the IIH PDU, or Hello packet, detects neighboring routers and indicates to the pseudonode the area's adjacency mesh. The Hello packet, flooded to the multicast address, contains the system ID of the sending router, the holding time, the circuit type of the interface on which the PDU was sent, the PDU length, the DIS identifier, and the interface priority (used in DIS election). The Hello packet also informs its area routers that it is the DIS.

Hello packets are padded to the maximum IS-IS PDU size of 1492 bytes (the full MTU size) to assist in the detection of transmission errors with large frames or with MTU mismatches between adjacencies.

The DIS typically floods Hello packets to the entire LAN every three seconds.

**Link-state PDU (LSP)** — The LSP contains information about each router in an area and its connected interfaces. LSPs are refreshed periodically and acknowledged on the network by way of sequence number PDUs. If new LSP information is found, based on the most recent complete sequence number PDU (CSNP), then out-of-date entries in the link-state database (LSDB) are removed and the LSDB is updated.

For a more detailed breakdown of the LSP, see [“LSP structure” on page 439](#).

**Complete sequence number PDU (CSNP)** — CSNPs contain a list of all LSPs in the current LSDB. The CSNP informs other area routers of missing or outdated links in the adjacency mesh. The receiving routers then use this information to update their own database to ensure that all area routers converge.

In contrast to Hello packets, CSNPs are sent every ten seconds and only between neighbor. In other words, they are never flooded.

**Partial sequence number PDU (PSNP)** — PSNPs are used to request and acknowledge LSP information from an adjacency. When a router compares a CSNP with its local database and determines a discrepancy, the router requests an updated LSP using a PSNP. Once received, the router stores the LSP in its local database and responds to the DIS with acknowledgement.

## Default routing

The default route is used if either there are no other routes in the routing table or if none of the other routes apply to a destination. Including the gateway in the default route gives all traffic a next-hop address to use when leaving the local network. The gateway address is normally another router on the edge of the local network.

FortiGate units come with a default static route with an IPv4 address of 0.0.0.0, an administration distance of 10, and a gateway IPv4 address. Beginner administrators can use the default route settings until a more advanced configuration is warranted.

By default, all routes are displayed in the Routing Monitor list. To display the routes in the routing table, go to *Route > Monitor > Routing Monitor*.

## Route leaking

Route leaking is the term used to describe the bi-directional flow of information between internal and external routing interfaces. By default, IS-IS leaks routing information from a Level 1 area into a Level 2 area. In order to leak Level 2 routing information into a Level 1 area, you must configure an export policy. Whether or not a route is leaked is determined by the ATT bit, using TLV 128 (for internal IP reachability) and TLV 130 (for external IP address information). For more information on TLVs, see [“Troubleshooting IS-IS” on page 445](#).

To configure IS-IS route leaking, use the following CLI commands.

1. On a Level 1-2 router:

```
config router isis
 set redistribute-l2 enable
end
```

2. On a Level 1 router:

```
config router isis
 get router info routing-table isis
 get router info isis route
end
```

## Default information originate option

Enabling default-information-originate generates and advertises a default route into the FortiGate unit's IS-IS-enabled networks. The generated route may be based on routes learned through a dynamic routing protocol, routes in the routing table, or both. IS-IS does not create the default route unless you use the `always` option.

Select *Disable* if you experience any issues or if you wish to advertise your own static routes into IS-IS updates.

The CLI commands associated with default information originate include:

```
config router isis
 set default-originate
end
```

## Timer options

IS-IS uses various timers to regulate its performance including a garbage timer, update timer, and timeout timer. The FortiGate unit default timer settings (30, 180, and 120 seconds respectively) are effective in most configurations—if you change these settings, ensure that the new settings are compatible with local routers and access servers.

You can configure the three IS-IS timers in the CLI, using the following commands:

```
config router isis
 set garbage-timer
 set update-timer
 set timeout-timer
end
```

You will find more information on each timer below.

### **Garbage timer**

The garbage timer is the amount of time (in seconds) that the DIS will advertise a route as being unreachable before deleting the route from the routing table. If this timer is shorter, it will keep more up-to-date routes in the routing table and remove old ones faster. This results in a smaller routing table which is useful if you have a very large network, or if your network changes frequently.

### **Update timer**

The update timer determines the interval between routing updates. Generally, this value is set to 30 seconds. There is some randomness added to help prevent network traffic congestion, which could result from all routers simultaneously attempting to update their neighbors. The update timer should be at least three times smaller than the timeout timer, otherwise you will experience an error.

If you are experiencing significant traffic on your network, you can increase this interval to send fewer updates per minute. However, ensure you increase the interval for all the routers on your network or you will experience timeouts that will degrade your network speed.

### **Timeout timer**

The timeout timer is the maximum amount of time (in seconds) that a route is considered reachable while no updates are received for the route. This is the maximum time the DIS will keep a reachable route in the routing table while no updates for that route are received. If the DIS receives an update for the route before the timeout period expires, the timer is restarted. The timeout period should be at least three times longer than the update period, otherwise you will experience an error.

If you are experiencing problems with routers not responding in time to updates, increase this timer. However, remember that longer timeout intervals result in longer overall update periods — it may be considerable time before the DIS is done waiting for all the timers to expire on unresponsive routes.

## **Authentication**

In routing protocols, it is typically desirable to establish authentication rules that prevent malicious and otherwise unwanted information from being injected into the routing table. IS-IS routing protocol utilizes TLV 10 to establish authentication. For more information on TLVs, see [“TLV” on page 439](#).

Initially, IS-IS used plain Clear Text to navigate the authentication rules, but this was found to be insecure since the Clear Text packets were unencrypted and could be exposed to packet sniffers. As per [RFC 3567](#), HMAC-MD5 and Enhanced Clear Text authentication features were introduced to IS-IS, both of which encrypt authentication data, making them considerably more secure than using plain Clear Text authentication.

### **HMAC-MD5 authentication**

Hashed Message Authentication Codes - Message Digest 5 (HMAC-MD5) is a mechanism for applying a cryptographic hash function to the message authentication process. It is applied at

both Level 1 and Level 2 routing. In IS-IS, an HMAC-MD5 can be applied to each type of LSP, on different interfaces, and with different passwords.

Authentication data is hashed using an AH (Authentication Header) key. From [RFC 2085](#):

*The “AH Key” is used as a shared secret between two communicating parties. The Key is not a “cryptographic key” as used in a traditional sense. Instead, the AH key (shared secret) is hashed with the transmitted data and thus, assures that an intervening party cannot duplicate the authentication data. [...] Implementation should, and as frequently as possible, change the AH key. Keys need to be chosen at random, or generated using a cryptographically strong pseudo-random generator seeded with a random seed.”*

Clear Text authentication uses the configuration commands `area-password` and `domain-password` for authentication, but when migrating from Clear Text authentication to HMAC-MD5, these command settings are automatically overwritten.

By the year 2005, the MD5 hash function had been identified as vulnerable to collision search attacks and various weaknesses. While such vulnerabilities do not compromise the use of MD5 within HMAC, administrators need to be aware of potential developments in cryptanalysis and cryptographic hash functions in the likely event that the underlying hash function needs to be replaced.

### Enhanced Clear Text authentication

Enhanced Clear Text authentication is an extension to Clear Text authentication that allows the encryption of passwords as they are displayed in the configuration. It includes a series of authentication mode commands and an authentication key chain, and allows for more simple password modification and password management. Enhanced Clear Text authentication also provides for smoother migration to and from changing authentication types. Intermediate systems continue to use the original authentication method until all the area routers are updated to use the new method.

### Authentication key chain

A key chain is a list of one or more authentication keys including the send and receive lifetimes for each key. Keys are used for authenticating routing packets only during the specified lifetimes. A router migrates from one key to the next according to the scheduled send and receive lifetimes. If an active key is unavailable, then the PDU is automatically discarded.

From [RFC 5310](#):

*It should be noted that the cryptographic strength of the HMAC depends upon the cryptographic strength of the underlying hash function and on the size and quality of the key.*

### Integrated IS-IS

Integrated IS-IS is an extended version of IS-IS that includes support for both IPv4 and IPv6. IPv4 and IPv6 interface addresses are determined by TLVs 132 and 232, respectively. The parameter responsible for IPv6 reachability is TLV 236. For more information on TLVs, see [“Troubleshooting IS-IS” on page 445](#).

The FortiGate unit command `config router isis` is almost the same except that IPv6 addresses are used. Also, if you are going to use prefix or access lists with Integrated IS-IS, you must use the `config router access-list6` or `config prefix-list6` versions of those commands.

If you want to route IPv4 traffic over an IPv6 network, you can use the command `config system ipv6-tunnel` to configure the FortiGate unit to do this. The IPv6 interface is configured under `config system interface`. All subnets between the source and destination addresses must support IPv6. This command is not supported in Transparent mode.

For example, you want to set up a tunnel on the port1 interface starting at 2002:C0A8:3201:: on your local network and tunnel it to address 2002:A0A:A01:: where it will need access to an IPv4 network again. Use the following command:

```
config system ipv6-tunnel
 edit test_tunnel
 set destination 2002:A0A:A01::
 set interface port1
 set source 2002:C0A8:3201::
 end
end
```

The CLI commands associated with Integrated IS-IS include:

```
config router isis
config router access-list6
config router prefix-list6
config system ipv6-tunnel
get router info6 *
```

## Troubleshooting IS-IS

If you want to troubleshoot Integrated IS-IS, it is the same as with IS-IS but you must specify the different protocol, and use IPv6 addresses. This applies to commands such as `get router info6` when you want to see the routing table, or other related information.

This section includes:

- [Routing loops](#)
- [Split horizon and Poison reverse updates](#)

### Routing loops

Normally in routing, a path between two addresses is chosen and traffic is routed along that path from one address to the other. When there is a routing loop, that normal path doubles back on itself creating a loop. When there are loops, the network has problems.

A routing loop happens when a normally functioning network has an outage, and one or more routers are offline. When packets encounter this, an alternate route is attempted to maneuver around the outage. During this phase it is possible for a route to be attempted that involves going back a hop, and trying a different hop forward. If that hop forward is blocked by the outage as well, a hop back and possibly the original hop forward may be selected. You can see if this continues, how it can consume not only network bandwidth but also many resources on those routers affected. The worst part is this situation will continue until the network administrator changes the router settings, or the downed routers come back online.

### Routing loop effect on the network

In addition to this “traffic jam” of routed packets, every time the routing table for a router changes that router sends an update out to all of the IS-IS routers connected to it. In a network loop, its possible for a router to change its routes very quickly as it tries and fails along these new routes. This can quickly result in a flood of updates being sent out, which can effectively grind the network to a halt until the problem is fixed.

## How can you spot a routing loop

Any time network traffic slows down, you will be asking yourself if it is a network loop or not. Often slowdowns are normal, they are not a full stoppage, and normal traffic resumes in a short period of time.

If the slow down is a full halt of traffic or a major slowdown does not return to normal quickly, you need to do serious troubleshooting quickly.

Some methods to troubleshoot your outage include:

- [Checking your logs](#)
- [Using SNMP network monitoring](#)
- [Using dead gateway detection and e-mail alerts](#)
- [Looking at the packet flow](#)

If you aren't running SNMP, dead gateway detection, or you have non-Fortinet routers in your network, you can use networking tools such as ping and traceroute to define the outage on your network and begin to fix it.

### Checking your logs

If your routers log events to a central location, it can be easy to check the logs for your network for any outages.

On your FortiGate unit, go to *Log & Report > Log & Archive Access*. You will want to look at both event logs and traffic logs. Events to look for will generally fall under CPU and memory usage, interfaces going offline (due to dead gateway detection), and other similar system events.

Once you have found and fixed your network problem, you can go back to the logs and create a report to better see how things developed during the problem. This type of forensics analysis can better help you prepare for next time.

### Using SNMP network monitoring

If your network had no problems one minute and slows to a halt the next, chances are something changed to cause that problem. Most of the time an offline router is the cause, and once you find that router and bring it back online, things will return to normal.

If you can enable a hardware monitoring system such as SNMP or sFlow on your routers, you can be notified of the outage and where it is exactly as soon as it happens.

Ideally you can configure SNMP on all your FortiGate routers and be alerted to all outages as they occur.

### To use SNMP to detect potential routing loops

1. Go to *System > Config > SNMP*.
2. Enable *SNMP Agent*.
3. Optionally enter the *Description*, *Location*, and *Contact* information for this device for easier location of the problem report.
4. In either *SNMP v1/v2c* section or *SNMP v3* section, as appropriate, select *Create New*.
5. Enter the *Community Name* that you want to use.
6. In *Hosts*, select *Add* to add an IP address where you will be monitoring the FortiGate unit. You can add up to 8 different addresses.
7. Ensure that ports 161 and 162 (SNMP queries and traps) are allowed through your security policies.
8. In *SNMP Event*, select the events you want to be notified of. For routing loops this should include *CPU Overusage*, *Memory Low*, and possibly *Log disk space low*. If there are

problems, the log will be filling up quickly, and the FortiGate unit's resources will be overused.

9. Select *OK*.
10. Configure SNMP host (manager) software on your administration computer. This will monitor the SNMP information sent out by the FortiGate unit. Typically you can configure this software to alert you to outages or CPU spikes that may indicate a routing loop.

### Using dead gateway detection and e-mail alerts

Another tool available to you on FortiGate units is the dead gateway detection. This feature allows the FortiGate unit to ping a gateway at regular intervals to ensure it is online and working. When the gateway is not accessible, that interface is marked as down.

#### To detect possible routing loops with dead gateway detection and e-mail alerts

1. To configure dead gateway detection, go to *Router > Static > Settings* and select *Create New*.
2. Set the *Ping Interval* (how often to send a ping), and *Failover Threshold* (how many lost pings is considered a failure). A smaller interval and smaller number of lost pings will result in faster detection, but will create more traffic on your network.
3. To configure interface status change notification, go to *Log & Report > Log Config > Alert E-mail*.
4. After you enter your email details, select the events you want to be alerted about — in our case *Configuration changes*. You may also want to log CPU and Memory usage as a network outage will cause your CPU activity to spike.



If you have VDOMs configured, you will have to enter the basic SMTP server information in the Global section, and the rest of the configuration within the VDOM that includes this interface.

---

After this configuration, when this interface on the FortiGate unit cannot connect to the next router, the FortiGate unit will bring down the interface and alert you with an email to the outage.

### Looking at the packet flow

If you want to see what is happening on your network, look at the packets travelling on the network. In this situation, you are looking for routes that have metrics higher than 15 as that indicates they are unreachable. Ideally if you debug the flow of the packets, and record the routes that are unreachable, you can create an accurate picture of the network outage.

### Action to take on discovering a routing loop

Once you have mapped the problem on your network, and determined it is in fact a routing loop there are a number of steps to take in correcting it.

1. Get any offline routers back online. This may be a simple reboot, or you may have to replace hardware. Often this first step will restore your network to its normal operation, once the routing tables finish being updated.
2. Change your routing configuration on the edges of the outage. Even if step 1 brought your network back online, you should consider making changes to improve your network before the next outage occurs. These changes can include configuring features like hold downs and triggers for updates, split horizon, and poison reverse updates.

## Split horizon and Poison reverse updates

Split horizon is best explained with an example. You have three routers linked serially, let's call them A, B, and C. A is only linked to B, C is only linked to B, and B is linked to both A and C. To get to C, A must go through B. If the link to C goes down, it is possible that B will try to use A's route to get to C. This route is A-B-C, so it will not work. However, if B tries to use it this begins an endless loop.

This situation is called a split horizon because from B's point of view the horizon stretches out in each direction, but in reality it only is on one side.

Poison reverse is the method used to prevent routes from running into split horizon problems. Poison reverse "poisons" routes away from the destination that use the current router in their route to the destination. This "poisoned" route is marked as unreachable for routers that cannot use it. In IS-IS this means that route is marked with a distance of 16.

## Simple IS-IS example

This is an example of a typical medium-sized network configuration using IS-IS routing.

Imagine a company with four FortiGate devices connected to one another. A FortiGate at one end of the network connects to two routers, each with its own local subnet; one of these routers uses OSPF and the other uses RIP.

Your task is to configure the four FortiGates to route traffic and process network updates using IS-IS, such that the farthest FortiGate (see 'FGT4' in [Network layout and assumptions](#)) receives route updates for the two routers at the opposite end of the network. Furthermore, FGT4 has been given a loopback subnet that must be identified by the router running RIP.

Since the internal networks use OSPF and RIP, those protocols will need to be redistributed through the IS-IS network. To keep the example simple, there will be no authentication of router traffic.

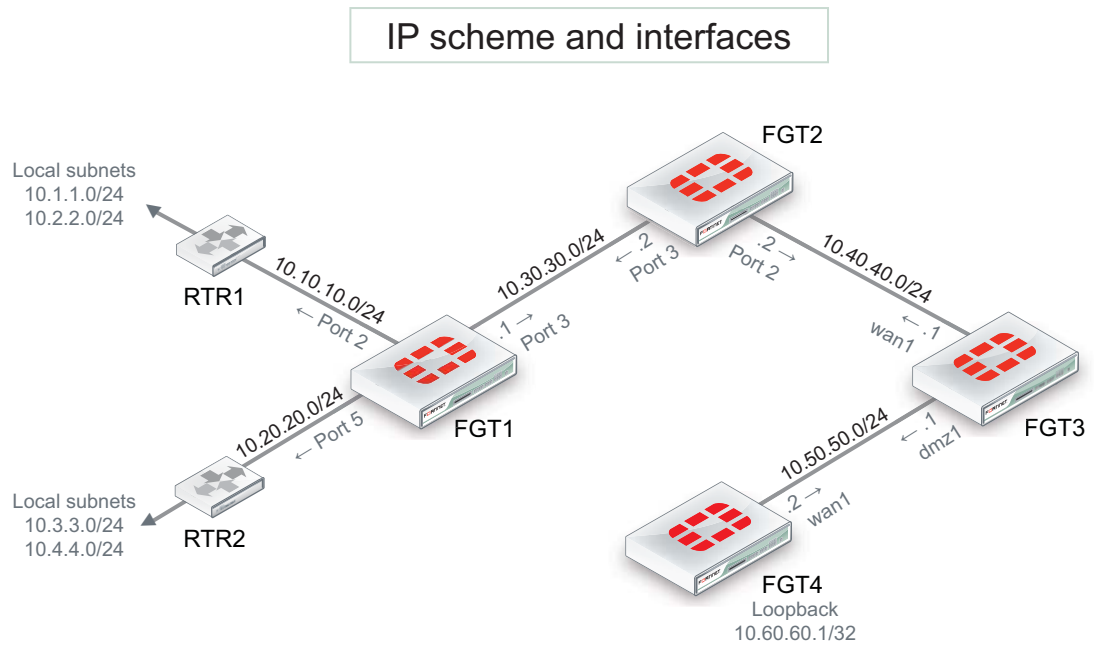
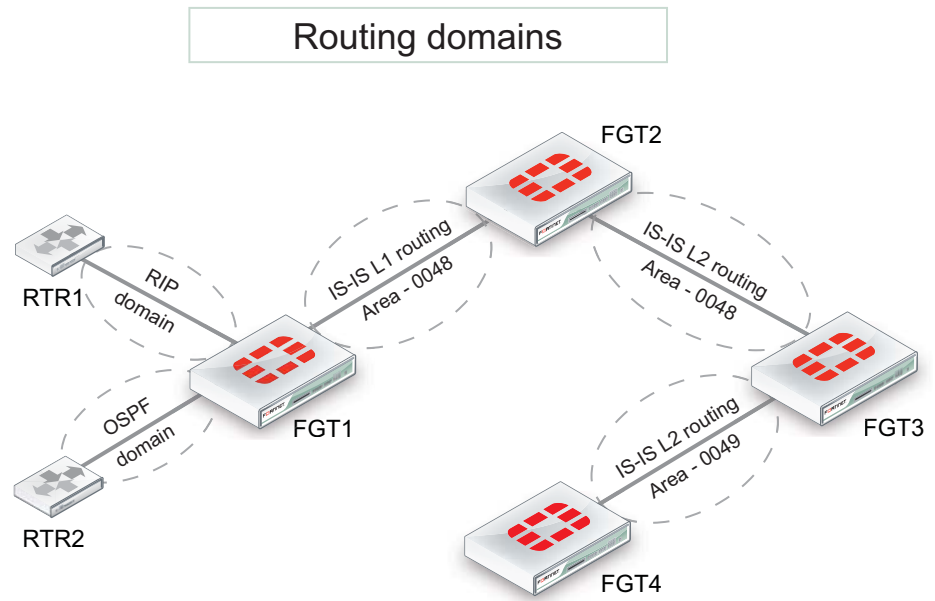
With IS-IS properly configured in this example, if a router fails or temporarily goes offline, the route change will propagate throughout the system.

This section includes the following topics:

- [Network layout and assumptions](#)
- [Expectations](#)
- [CLI configuration](#)
- [Verification](#)
- [Troubleshooting](#)



## Network layout and assumptions



- It is assumed that each FortiGate is operating in NAT mode.
- All interfaces have been previously assigned and no static routes are required.
- The AFI (Authority and Format Identifier) used is 49 : Locally administered (private).
- The Area identifiers are 0048 and 0049.

## Expectations

- FGT4 must get the IS-IS route updates for RTR1 and RTR2 local subnets (10.1.1.0, 10.2.2.0, 10.3.3.0, 10.4.4.0).
- RTR1 must receive (via RIP2) the loopback subnet of FGT4 (10.60.60.1/32).

## CLI configuration

The following CLI configuration occurs on each FortiGate (as identified), including only the relevant parts.

### FGT1

```
config router isis
 config isis-interface
 edit "port3"
 set circuit-type level-1
 set network-type broadcast
 set status enable
 next
 end
config isis-net
 edit 1
 set net 49.0048.1921.6818.2136.00
 next
end
config redistribute "connected"
end
config redistribute "rip"
 set status enable
 set level level-1
end
config redistribute "ospf"
 set status enable
 set level level-1
end
end
config router rip
 config interface
 edit "port2"
 set receive-version 2
 set send-version 2
 next
 end
config network
 edit 1
 set prefix 10.10.10.0 255.255.255.0
 next
end
config redistribute "isis"
 set status enable
end
end
```

## FGT2

```
config router isis
 config isis-interface
 edit "port3"
 set circuit-type level-1
 set network-type broadcast
 set status enable
 next
 edit "port2"
 set network-type broadcast
 set status enable
 next
 end
 config isis-net
 edit 1
 set net 49.0048.1221.6818.2110.00
 next
 end
 set redistribute-l1 enable
 set redistribute-l2 enable
end
```

## FGT3

```
config router isis
 set is-type level-2-only
 config isis-interface
 edit "wan1"
 set network-type broadcast
 set status enable
 next
 edit "dmz1"
 set network-type broadcast
 set status enable
 next
 end
 config isis-net
 edit 1
 set net 49.0048.1921.6818.2108.00
 next
 edit 2
 set net 49.0049.1921.6818.2108.00
 next
 end
end
```

## FGT4

```
config router isis
 set is-type level-2-only
 config isis-interface
 edit "wan1"
 set network-type broadcast
 set status enable
 next
 end
config isis-net
 edit 1
 set net 49.0049.1721.0160.1004.00
 next
end
config redistribute "connected"
 set status enable
end
end
```

## Verification

Once the network has been configured, you need to test that it works as expected. Use the following CLI commands on the devices indicated.

### Verifying if RTR1 receives loopback subnet of FGT4

```
(RTR1) # get router info routing-table all
```

*Result:*

```
C 10.1.1.0/24 is directly connected, vlan1
C 10.2.2.0/24 is directly connected, vlan2
C 10.10.10.0/24 is directly connected, dmz1
R 10.40.40.0/24 [120/2] via 10.10.10.1, dmz1, 00:04:07
R 10.50.50.0/24 [120/2] via 10.10.10.1, dmz1, 00:04:07
R 10.60.60.1/32 [120/2] via 10.10.10.1, dmz1, 00:04:07
```

(\*) If required, filtering out 10.50.50.0 and 10.40.40.0 from the routing table could be done with a route-map.

### Verification on FGT2, which is the border between L1 and L2 routing levels; looking at IS-IS information

```
FGT2 # get router info isis interface
```

*Result:*

```
port2 is up, line protocol is up
 Routing Protocol: IS-IS ((null))
 Network Type: Broadcast
 Circuit Type: level-1-2
 Local circuit ID: 0x01
 Extended Local circuit ID: 0x00000003
 Local SNPA: 0009.0f85.ad8c
```

```

IP interface address:
 10.40.40.2/24
IPv6 interface address:
Level-1 Metric: 10/10, Priority: 64, Circuit ID: 1221.6818.2110.01
Number of active level-1 adjacencies: 0
Level-2 Metric: 10/10, Priority: 64, Circuit ID: 1221.6818.2110.01
Number of active level-2 adjacencies: 1
Next IS-IS LAN Level-1 Hello in 6 seconds
Next IS-IS LAN Level-2 Hello in 1 seconds
port3 is up, line protocol is up
 Routing Protocol: IS-IS ((null))
 Network Type: Broadcast
 Circuit Type: level-1
 Local circuit ID: 0x02
 Extended Local circuit ID: 0x00000004
 Local SNPA: 0009.0f85.ad8d
 IP interface address:
 10.30.30.2/24
 IPv6 interface address:
Level-1 Metric: 10/10, Priority: 64, Circuit ID: 1221.6818.2110.02
Number of active level-1 adjacencies: 1
Next IS-IS LAN Level-1 Hello in 2 seconds

```

```
FGT2 # get router info isis neighbour
```

*Result:*

System Id	Interface	SNPA	State	Holdtime	Type	Protocol
1921.6818.2108	port2	0009.0f04.0794	Up	22	L2	IS-IS
1921.6818.2136	port3	0009.0f85.acf7	Up	29	L1	IS-IS

### Verification on FGT3, which is border between 2 areas; looking at IS-IS information

IS-IS router CLI commands available:

```
FGT3 # get router info isis ?
```

*Result:*

```

interface show isis interfaces
neighbor show CLNS neighbor adjacencies
is-neighbor show IS neighbor adjacencies
database show IS-IS link state database
route show IS-IS IP routing table
topology show IS-IS paths

```

Example of interface status and neighbors:

```
FGT3 # get router info isis interface
```

*Result:*

```

wan1 is up, line protocol is up
 Routing Protocol: IS-IS ((null))
 Network Type: Broadcast
 Circuit Type: level-1-2
 Local circuit ID: 0x01
 Extended Local circuit ID: 0x00000003

```

Local SNPA: 0009.0f04.0794  
IP interface address:  
10.40.40.1/24  
IPv6 interface address:  
Level-2 Metric: 10/10, Priority: 64, Circuit ID: 1221.6818.2110.01  
Number of active level-2 adjacencies: 1  
Next IS-IS LAN Level-2 Hello in 3 seconds

dmz1 is up, line protocol is up  
Routing Protocol: IS-IS ((null))  
Network Type: Broadcast  
Circuit Type: level-1-2  
Local circuit ID: 0x02  
Extended Local circuit ID: 0x00000005  
Local SNPA: 0009.0f04.0792  
IP interface address:  
10.50.50.1/24  
IPv6 interface address:  
Level-2 Metric: 10/10, Priority: 64, Circuit ID: 1721.0160.1004.01  
Number of active level-2 adjacencies: 1  
Next IS-IS LAN Level-2 Hello in 7 seconds

FGT3 # **get router info isis neighbour**

*Result:*

System Id	Interface	SNPA	State	Holdtime	Type	Protocol
1221.6818.2110	wan1	0009.0f85.ad8c	Up	8	L2	IS-IS
1721.0160.1004	dmz1	0009.0f52.7704	Up	8	L2	IS-IS

### **Verification on FGT4 that the remote subnets from RTR1 and RTR2 are in the routing table and learned with IS-IS**

FGT4 # **get router info routing-table all**

*Result:*

Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP  
O - OSPF, IA - OSPF inter area  
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2  
E1 - OSPF external type 1, E2 - OSPF external type 2  
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area  
\* - candidate default

**i L2 10.1.1.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46**  
**i L2 10.2.2.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46**  
**i L2 10.3.3.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46**  
**i L2 10.4.4.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46**  
i L2 10.10.10.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46  
i L2 10.11.11.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46  
i L2 10.20.20.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46  
i L2 10.30.30.0/24 [115/30] via 10.50.50.1, wan1, 00:13:55  
i L2 10.40.40.0/24 [115/20] via 10.50.50.1, wan1, 00:15:30  
C 10.50.50.0/24 is directly connected, wan1  
C 10.60.60.1/32 is directly connected, loopback

## Troubleshooting

The following diagnose commands are available for further IS-IS troubleshooting and will display all IS-IS activity (sent and received packets):

```
FGT # diagnose ip router isis level info
FGT # diagnose ip router isis all enable
FGT # diagnose debug enable
```

...to stop the debug type output :

```
FGT # diagnose ip router isis level none
```

Output and interpretation depends on the issue faced. You can provide this information to TAC if you open a support ticket.

### Debugging IPv6 on IS-ISng

The debug command is very useful to see what is happening on the network at the packet level. The following CLI commands specify both IPv6 and IS-IS, so only IS-ISng packets will be reported. The output from these commands will show you the IS-ISng traffic on your FortiGate unit including RECV, SEND, and UPDATE actions.

```
FGT # diagnose ipv6 router isis level info
FGT # diagnose ipv6 router isis all enable
FGT # diagnose debug enable
```

These three commands will:

- turn on debugging in general
- set the debug level to information, a verbose reporting level
- turn on all IS-IS router settings

Part of the information displayed from the debugging is the metric (hop count). If the metric is 16, then that destination is unreachable since the maximum hop count is 15.

In general, you should see an update announcement, followed by the routing table being sent out, and a received reply in response.

# Chapter 3 Authentication for FortiOS 5.0

This FortiOS Handbook chapter contains the following sections:

[Introduction to authentication](#) describes some basic elements and concepts of authentication.

[Authentication servers](#) describes external authentication servers, where a FortiGate unit fits into the topology, and how to configure a FortiGate unit to work with that type of authentication server.

[Users and user groups](#) describes the different types of user accounts and user groups. Authenticated access to resources is based on user identities and user group membership. Two-factor authentication methods, including FortiToken, provide additional security.

[Managing Guest Access](#) explains how to manage temporary accounts for visitors to your premises.

[Configuring authenticated access](#) provides detailed procedures for setting up authenticated access in security policies and authenticated access to VPNs.

[Certificate-based authentication](#) describes authentication by means of X.509 certificates.

[SSO using a FortiAuthenticator unit](#) describes how to use a FortiAuthenticator unit as an SSO agent that can integrate with external network authentication systems such as RADIUS and LDAP to gather user logon information and send it to the FortiGate unit. Users can also log on through a FortiAuthenticator-based web portal or the FortiClient SSO Mobility Agent.

[Single Sign-On to Windows AD](#) describes how to set up Single Sign-On in a Windows AD network by configuring the FortiGate unit to poll domain controllers for information user logons and user privileges.

[Agent-based FSSO](#) describes how to set up Single Sign-On in Windows AD, Citrix, or Novell networks by installing Fortinet Single Sign On (FSSO) agents on domain controllers. The FortiGate unit receives information about user logons and allows access to network resources based on user group memberships.

[SSO using RADIUS accounting records](#) describes how to set up Single Sign-On in a network that uses RADIUS authentication. In this configuration, the RADIUS server send RADIUS accounting records to the FortiGate unit when users log on or off the network. The record includes a user group name that can be used in FortiGate security policies to determine which resources each user can access.

[Monitoring authenticated users](#) describes FortiOS authenticated user monitor screens.

[Examples and Troubleshooting](#) provides configuration examples and troubleshooting suggestions.



# Introduction to authentication

Identifying users and other computers—authentication—is a key part of network security. This section describes some basic elements and concepts of authentication.

The following topics are included in this section:

- [What is authentication?](#)
- [Methods of authentication](#)
- [Types of authentication](#)
- [User's view of authentication](#)
- [FortiGate administrator's view of authentication](#)

## What is authentication?

Businesses need to authenticate people who have access to company resources. In the physical world this may be a swipe card to enter the building, or a code to enter a locked door. If a person has this swipe card or code, they have been authenticated as someone allowed in that building or room.

Authentication is the act of confirming the identity of a person or other entity. In the context of a private computer network, the identities of users or host computers must be established to ensure that only authorized parties can access the network. The FortiGate unit enables controlled network access and applies authentication to users of security policies and VPN clients.

## Methods of authentication

FortiGate unit authentication is divided into three basic types: password authentication for people, certificate authentication for hosts or endpoints, and two-factor authentication for additional security beyond just passwords. An exception to this is that FortiGate units in an HA cluster and FortiManager units use password authentication.

Password authentication verifies individual user identities, but access to network resources is based on membership in user groups. For example, a security policy can be configured to permit access only to the members of one or more user groups. Any user who attempts to access the network through that policy is then authenticated through a request for their username and password.

Methods of authentication include:

- [Local password authentication](#)
- [Server-based password authentication](#)
- [Certificate-based authentication](#)
- [Two-factor authentication](#)

## Local password authentication

The simplest authentication is based on user accounts stored locally on the FortiGate unit. For each account, a username and password is stored. The account also has a disable option so that you can suspend the account without deleting it.

Local user accounts work well for a single-FortiGate installation. If your network has multiple FortiGate units that will use the same accounts, the use of an external authentication server can simplify account configuration and maintenance.

You create local user accounts in the web-based manager under *User & Device > User > User Definition*. This page is also used to create accounts where an external authentication server stores and verifies the password.

## Server-based password authentication

Using external LDAP, RADIUS, or TACACS+ authentication servers is desirable when multiple FortiGate units need to authenticate the same users, or where the FortiGate unit is added to a network that already contains an authentication server.

When you use an external authentication server to authenticate users, the FortiGate unit sends the user's entered credentials to the external server. The password is encrypted. The server's response indicates whether the supplied credentials are valid or not.

You must configure the FortiGate unit to access the external authentication servers that you want to use. The configuration includes the parameters that authenticate the FortiGate unit to the authentication server.

You can use external authentication servers in two ways:

- Create user accounts on the FortiGate unit, but instead of storing each user's password, specify the server used to authenticate that user. As with accounts that store the password locally, you add these users to appropriate user groups.
- Add the authentication server to user groups. Any user who has an account on the server can be authenticated and have the access privileges of the FortiGate user group. Optionally, when an LDAP server is a FortiGate user group member, you can limit access to users who belong to specific groups defined on the LDAP server.

## Certificate-based authentication

An RSA X.509 server certificate is a small file issued by a Certificate Authority (CA) that is installed on a computer or FortiGate unit to authenticate itself to other devices on the network. When one party on a network presents the certificate as authentication, the other party can validate that the certificate was issued by the CA. The identification is therefore as trustworthy as the Certificate Authority (CA) that issued the certificate.

To protect against compromised or misused certificates, CAs can revoke any certificate by adding it to a Certificate Revocation List (CRL). Certificate status can also be checked online using Online Certificate Status Protocol (OCSP).

RSA X.509 certificates are based on public-key cryptography, in which there are two keys: the private key and the public key. Data encrypted with the private key can be decrypted only with the public key and vice versa. As the names suggest, the private key is never revealed to anyone and the public key can be freely distributed. Encryption with the recipient's public key creates a message that only the intended recipient can read. Encryption with the sender's private key creates a message whose authenticity is proven because it can be decrypted only with the sender's public key.

Server certificates contain a signature string encrypted with the CA's private key. The CA's public key is contained in a CA root certificate. If the signature string can be decrypted with the CA's public key, the certificate is genuine.

## Certificate authorities

A certificate authority can be:

- an organization, such as VeriSign Inc., that provides certificate services
- a software application, such as Microsoft Certificate Services or OpenSSH

For a company web portal or customer-facing SSL VPN, a third-party certificate service has some advantages. The CA certificates are already included in popular web browsers and customers trust the third-party. On the other hand, third-party services have a cost.

For administrators and for employee VPN users, the local CA based on a software application provides the required security at low cost. You can generate and distribute certificates as needed. If an employee leaves the organization, you can simply revoke their certificate.

## Certificates for users

FortiGate unit administrators and SSL VPN users can install certificates in their web browsers to authenticate themselves. If the FortiGate unit uses a CA-issued certificate to authenticate itself to the clients, the browser will also need the appropriate CA certificate.

FortiGate IPsec VPN users can install server and CA certificates according to the instructions for their IPsec VPN client software. The FortiClient Endpoint Security application, for example, can import and store the certificates required by VPN connections.

FortiGate units are also compatible with some Public Key Infrastructure systems. For an example of this type of system, see [“RSA ACE \(SecurID\) servers” on page 483](#).

## Two-factor authentication

A user can be required to provide both something they know (their username and password combination) and something they have (certificate or a random token code). Certificates are installed on the user's computer.

Two-factor authentication is available for PKI users. For more information, see [“Certificate” on page 495](#).

Another type of two-factor authentication is to use a randomly generated token (multi-digit number) along with the username and password combination. One method is a FortiToken — a one time passcode (OTP) generator that generates a unique code every 60 seconds. Others use email or SMS text messaging to deliver the random token code to the user or administrator.

When one of these methods is configured, the user enters this code at login after the username and password have been verified. The FortiGate unit verifies the token code after as well as the password and username. For more information, see [“Two-factor authentication” on page 494](#)

## Types of authentication

FortiOS supports two different types of authentication based on your situation and needs.

Security policy authentication, or identity-based policies, is easily applied to all users logging on to a network, or network service. For example if a group of users on your network such as the accounting department who have access to sensitive data need to access the Internet, it is a good idea to make sure the user is a valid user and not someone trying to send company secrets to the Internet. Security policy authentication can be applied to as many or as few users

as needed, and it supports a number of authentication protocols to easily fit with your existing network.

VPN authentication can be for both the remote VPN device as well as the VPN users. VPNs are used to communicate with locations outside the company network as if they were part of the company network. This level of trust, once a VPN is established, is easily established with authentication to verify the remote user is in fact a valid user. In this situation without authentication, anyone malicious or otherwise could connect to the company network with potentially full access.

## Firewall authentication (identity-based policies)

Security policies enable traffic to flow between networks. If you want to limit which users have access to particular resources, you create identity-based policies (IBP) that allow access only to members of specific user groups. Authentication, a request for username and password, is triggered when a user attempts to access a resource for which data must pass through an identity-based policy.

The user's authentication expires if the connection is idle for too long, 5 minutes by default but that can be customized.

Identity-based policies are the mechanism for FSSO, NTLM, certificate based, and RADIUS SSO authentication.

### FSSO

Fortinet Single Sign on (FSSO) provides seamless authentication support for Microsoft Windows Active Directory (AD) and Novell eDirectory users in a FortiGate environment.

On a Microsoft Windows or Novell network, users authenticate with the Active Directory or Novell eDirectory at logon. FSSO provides authentication information to the FortiGate unit so that users automatically get access to permitted resources. See [“Introduction to FSSO agents” on page 564](#).

### NTLM

The NT LAN Manager (NTLM) protocol is used when the MS Windows Active Directory (AD) domain controller can not be contacted. NTLM is a browser-based method of authentication.

The FSSO software is installed on each AD server and the FortiGate unit is configured to communicate with each FSSO client. When a user successfully logs into their Windows PC (and is authenticated by the AD Server), the FSSO client communicates the user's name, IP address, and group login information to the FortiGate unit. The FortiGate unit sets up a temporary access policy for the user, so when they attempt access through the firewall they do not need to re-authenticate. This model works well in environments where the FSSO client can be installed on all AD servers.

In system configurations where it is not possible to install FSSO clients on all AD servers, the FortiGate unit must be able to query the AD servers to find out if a user has been properly authenticated. This is achieved using the NTLM messaging features of Active Directory and Internet Explorer.

Even when NTLM authentication is used, the user is not asked again for their username and password. Internet Explorer stores the user's credentials and the FortiGate unit uses NTLM messaging to validate them in the Windows AD environment.

Note that if the authentication reaches the timeout period, the NTLM message exchange restarts. For more information on NTLM, see [“NTLM authentication” on page 524](#) and [“FSSO NTLM authentication support” on page 569](#).

## Certificates

Certificates can be used as part of an identity-based policy. All users being authenticated against the policy are required to have the proper certificate. See [“Certificate-based authentication” on page 458](#)

## RADIUS SSO

RADIUS Single Sign-On (RSSO) is a remote authentication method that does not require any local users to be configured, and relies on RADIUS Start records to provide the FortiGate unit with authentication information. That information identifies the user and user group, which is then matched using a security policy. See [“SSO using RADIUS accounting records” on page 602](#).

## FortiGuard Web Filter override authentication

Optionally, users can be allowed the privilege of overriding FortiGuard Web Filtering to view blocked web sites. Depending on the override settings, the override can apply to the user who requested it, the entire user group to which the user belongs, or all users who share the same web filter profile. As with other FortiGate features, access to FortiGuard overrides is controlled through user groups. Firewall and Directory Services user groups are eligible for the override privilege. For more information about web filtering and overrides, see the UTM chapter of this FortiOS Handbook.

## VPN authentication

Authentication involves authenticating the user. In IPsec VPNs authenticating the user is optional, but authentication of the peer device is required.

This section includes:

- [Authenticating IPsec VPN peers \(devices\)](#)
- [Authenticating IPsec VPN users](#)
- [Authenticating SSL VPN users](#)
- [Authenticating PPTP and L2TP VPN users](#)

### Authenticating IPsec VPN peers (devices)

A VPN tunnel has one end on a local trusted network, and the other end is at a remote location. The remote peer (device) must be authenticated to be able to trust the VPN tunnel. Without that authentication, it is possible for a malicious hacker to masquerade as a valid VPN tunnel device and gain access to the trusted local network.

The three ways to authenticate VPN peers are with a preshared key, RSA X.509 certificate, an a specific peer ID value.

The simplest way for IPsec VPN peers to authenticate each other is through the use of a preshared key, also called a shared secret. The preshared key is a text string used to encrypt the data exchanges that establish the VPN tunnel. The preshared key must be six or more characters. The VPN tunnel cannot be established if the two peers do not use the same key. The disadvantage of preshared key authentication is that it can be difficult to securely distribute and update the preshared keys. See [“Authenticating the FortiGate unit with a pre-shared key” on page 1640](#).

RSA X.509 certificates are a better way for VPN peers to authenticate each other. Each peer offers a certificate signed by a Certificate Authority (CA) which the other peer can validate with the appropriate CA root certificate. For more information about certificates, see [“Certificate-based authentication” on page 532](#).

You can supplement either preshared key or certificate authentication by requiring the other peer to provide a specific peer ID value. The peer ID is a text string configured on the peer device. On a FortiGate peer or FortiClient Endpoint Security peer, the peer ID provided to the remote peer is called the Local ID.

### Authenticating IPsec VPN users

An IPsec VPN can be configured to accept connections from multiple dynamically addressed peers. You would do this to enable employees to connect to the corporate network while traveling or from home. On a FortiGate unit, you create this configuration by setting the *Remote Gateway* to *Dialup User*.

It is possible to have an IPsec VPN in which remote peer devices authenticate using a common preshared key or a certificate, but there is no attempt to identify the user at the remote peer. To add user authentication, you can do one of the following:

- require a unique preshared key for each peer
- require a unique peer ID for each peer
- require a unique peer certificate for each peer
- require additional user authentication (XAuth)

The peer ID is a text string configured on the peer device. On a FortiGate peer or FortiClient Endpoint Security peer, the peer ID provided to the remote peer is called the Local ID.

### Authenticating SSL VPN users

SSL VPN users can be

- user accounts with passwords stored on the FortiGate unit
- user accounts authenticated by an external RADIUS, LDAP or TACACS+ server
- PKI users authenticated by certificate

You need to create a user group for your SSL VPN. Simply create a firewall user group, enable SSL VPN access for the group, and select the web portal the users will access.

SSL VPN access requires an SSL VPN security policy that permits access to members of your user group.

### Authenticating PPTP and L2TP VPN users

PPTP and L2TP are older VPN tunneling protocols that do not provide authentication themselves. FortiGate units restrict PPTP and L2TP access to users who belong to one specified user group. Users authenticate themselves to the FortiGate unit by username/password. You can configure PPTP and L2TP VPNs only in the CLI. Before you configure the VPN, create a firewall user group and add to it the users who are permitted to use the VPN. Users are authenticated when they attempt to connect to the VPN. For more information about configuring PPTP or L2TP VPNs, see the [FortiGate CLI Reference](#).

## Single Sign On authentication for users

“Single Sign-On” means that users logged on to a computer network are authenticated for access to network resources through the FortiGate unit without having to enter their username and password again. FortiGate units directly provide Single Sign On capability for:

- Microsoft Windows networks using either Active Directory or NTLM authentication
- Novell networks, using eDirectory

In combination with a FortiAuthenticator unit, the FortiGate unit can provide Single Sign-On capability that integrates multiple external network authentication systems such as Windows Active Directory, Novell e-Directory, RADIUS and LDAP. The FortiAuthenticator unit gathers user logon information from all of these sources and sends it to the FortiGate unit.

Through the SSO feature, the FortiGate unit knows the username, IP address, and external user groups to which the user belongs. When the user tries to access network resources, the FortiGate unit selects the appropriate security policy for the destination. If the user belongs to one of the permitted user groups, the connection is allowed.

For detailed information about SSO, see

- [“SSO using a FortiAuthenticator unit” on page 550](#)
- [“Agent-based FSSO” on page 563](#)

## User’s view of authentication

From the user’s point of view, they see a request for authentication when they try to access a protected resource, such as an FTP repository of intellectual property or simply access a website on the Internet. The way the request is presented to the user depends on the method of access to that resource.

VPN authentication usually controls remote access to a private network.

### Web-based user authentication

Security policies usually control browsing access to an external network that provides connection to the Internet. In this case, the FortiGate unit requests authentication through the web browser.

The user types a username and password and then selects *Continue* or *Login*. If the credentials are incorrect, the authentication screen is redisplayed with blank fields so that the user can try again. When the user enters valid credentials, access is granted to the required resource. In some cases, if a user tries to authenticate several times without success, a message appears, such as: “Too many bad login attempts. Please try again in a few minutes.” This indicates the user is locked out for a period of time. This prevents automated brute force password hacking attempts. The administrator can customize these settings if required.



After a defined period of user inactivity (the authentication timeout, defined by the FortiGate administrator), the user’s access expires. The default is 5 minutes. To access the resource, the user will have to authenticate again.

---

### VPN client-based authentication

A VPN provides remote clients with access to a private network for a variety of services that include web browsing, email, and file sharing. A client program such as FortiClient negotiates the connection to the VPN and manages the user authentication challenge from the FortiGate unit.

FortiClient can store the username and password for a VPN as part of the configuration for the VPN connection and pass them to the FortiGate unit as needed. Or, FortiClient can request the username and password from the user when the FortiGate unit requests them.

SSL VPN is a form of VPN that can be used with a standard Web browser. There are two modes of SSL VPN operation (supported in NAT/Route mode only):

- web-only mode, for remote clients equipped with a web-browser only
- tunnel mode, for remote computers that run a variety of client and server applications.



After a defined period of user inactivity on the VPN connection (the idle timeout, defined by the FortiGate administrator), the user's access expires. The default is 30 minutes. To access the resource, the user will have to authenticate again.

---

## FortiGate administrator's view of authentication

Authentication is based on user groups. The FortiGate administrator configures authentication for security policies and VPN tunnels by specifying the user groups whose members can use the resource. Some planning is required to determine how many different user groups need to be created. Individual user accounts can belong to multiple groups, making allocation of user privileges very flexible.

A member of a user group can be:

- a user whose username and password are stored on the FortiGate unit
- a user whose name is stored on the FortiGate unit and whose password is stored on a remote or external authentication server
- a remote or external authentication server with a database that contains the username and password of each person who is permitted access

The general process of setting up authentication is as follows:

1. If remote or external authentication is needed, configure the required servers.
2. Configure local and peer (PKI) user identities. For each local user, you can choose whether the FortiGate unit or a remote authentication server verifies the password. Peer members can be included in user groups for use in security policies.
3. Create user groups.
4. Add local/peer user members to each user group as appropriate. You can also add an authentication server to a user group. In this case, all users in the server's database can authenticate. You can only configure peer user groups through the CLI.
5. Configure security policies and VPN tunnels that require authenticated access.

For authentication troubleshooting, see the specific chapter for the topic or for general issues see [“Troubleshooting” on page 629](#).

## General authentication settings

Go to *User & Device > Authentication > Settings* to configure authentication timeout, protocol support, and authentication certificates.

When user authentication is enabled within a security policy, the authentication challenge is normally issued for any of the four protocols (depending on the connection protocol):

- HTTP (can also be set to redirect to HTTPS)
- HTTPS
- FTP
- Telnet.



The selections made in the *Protocol Support* list of the Authentication Settings screen control which protocols support the authentication challenge. Users must connect with a supported protocol first so they can subsequently connect with other protocols. If HTTPS is selected as a method of protocol support, it allows the user to authenticate with a customized Local certificate.

When you enable user authentication within a security policy, the security policy user will be challenged to authenticate. For user ID and password authentication, users must provide their user names and passwords. For certificate authentication (HTTPS or HTTP redirected to HTTPS only), you can install customized certificates on the unit and the users can also have customized certificates installed on their browsers. Otherwise, users will see a warning message and have to accept a default Fortinet certificate.

<b>Authentication Timeout</b>	Enter a length of time in minutes, from 1 to 480. Authentication timeout controls how long an authenticated firewall connection can be idle before the user must authenticate again. The default value is 30
<b>Protocol Support</b>	Select the protocols to challenge during firewall user authentication.
<b>Certificate</b>	If using HTTPS protocol support, select the local certificate to use for authentication. Available only if HTTPS protocol support is selected.
<b>Apply</b>	Select to apply the selections for user authentication settings.



When you use certificate authentication, if you do not specify any certificate when you create the security policy, the global settings will be used. If you specify a certificate, the per-policy setting will overwrite the global setting.

---

# Authentication servers

FortiGate units support the use of external authentication servers. An authentication server can provide password checking for selected FortiGate users or it can be added as a member of a FortiGate user group.

If you are going to use authentication servers, you must configure the servers before you configure FortiGate users or user groups that require them.



MAC OS and iOS devices, including iPhones and iPads, can perform user authentication with FortiOS units using RADIUS servers, but not with LDAP or TACACS+ servers.

---

This section includes the following topics:

- [FortiAuthenticator servers](#)
- [RADIUS servers](#)
- [LDAP servers](#)
- [TACACS+ servers](#)
- [SSO servers](#)
- [RSA ACE \(SecurID\) servers](#)

## FortiAuthenticator servers

FortiAuthenticator is an Authentication, Authorization, and Accounting (AAA) server, that includes a RADIUS server, an LDAP server, and can replace the FSSO Collector Agent on a Windows AD network. Multiple FortiGate units can use a single FortiAuthenticator for FSSO, remote authentication, and FortiToken management.

For more information, see the [FortiAuthenticator Administration Guide](#).

## RADIUS servers

Remote Authentication and Dial-in User Service (RADIUS) is a broadly supported client-server protocol that provides centralized authentication, authorization, and accounting functions. RADIUS clients are built into gateways that allow access to networks such as Virtual Private Network servers, Network Access Servers (NAS), as well as network switches and firewalls that use authentication. FortiGate units fall into the last category.

RADIUS servers use UDP packets to communicate with the RADIUS clients on the network to authenticate users before allowing them access to the network, to authorize access to resources by appropriate users, and to account or bill for those resources that are used. RADIUS servers are currently defined by RFC 2865 (RADIUS) and RFC 2866 (Accounting), and listen on either UDP ports 1812 (authentication) and 1813 (accounting) or ports 1645 (authentication) and 1646 (accounting) requests. RADIUS servers exist for all major operating systems.

You must configure the RADIUS server to accept the FortiGate unit as a client. FortiGate units use the authentication and accounting functions of the RADIUS server.



FortiOS does not accept all characters from auto generated keys from MS Windows 2008. These keys are very long and as a result RADIUS authentication will not work. Maximum key length for MS Windows 2008 is 128 bytes. In older versions of FSAE, it was 40 bytes.

### Microsoft RADIUS servers

Microsoft Windows Server 2000, 2003, and 2008 have RADIUS support built-in. Microsoft specific RADIUS features are defined in RFC 2548. The Microsoft RADIUS implementation can use Active Directory for user credentials.

For details on Microsoft RADIUS server configurations, refer to Microsoft documentation.

### RADIUS user database

The RADIUS user database is commonly an SQL or LDAP database, but can also be any combination of:

- usernames and passwords defined in a configuration file
- user account names and passwords configured on the computer where the RADIUS server is installed.

If users are members of multiple RADIUS groups, then the user group authentication timeout value does not apply. See [“Membership in multiple groups” on page 506](#).

### RADIUS authentication with a FortiGate unit

To use RADIUS authentication with a FortiGate unit

- configure one or more RADIUS servers on the FortiGate unit
- assign users to a RADIUS server

When a configured user attempts to access the network, the FortiGate unit will forward the authentication request to the RADIUS server which will match the username and password remotely. Once authenticated the RADIUS server passes the authorization granted message to the FortiGate unit which grants the user permission to access the network.

The RADIUS server uses a “shared secret” key along with MD5 hashing to encrypt information passed between RADIUS servers and clients, including the FortiGate unit. Typically only user credentials are encrypted. Additional security can be configured through IPsec tunnels.

## RADIUS attribute value pairs

RADIUS packets include a set of attribute value pairs (AVP) to identify information about the user, their location and other information. The FortiGate unit sends the following RADIUS attributes.

**Table 18:** FortiOS supported RADIUS attributes

RADIUS Attribute	Name	Description	AVP type
1	Acct-Session-ID	Unique number assigned to each start and stop record to make it easy to match them, and to eliminate duplicate records.	44
2	username	Name of the user being authenticated	1
3	NAS-Identifier	Identifier or IP address of the Network Access Server (NAS) that is requesting authentication. In this case, the NAS is the FortiGate unit.	32
4	Framed-IP-Address	Address to be configured for the user.	8
5	Fortinet-VSA	See <a href="#">“Vendor-specific attributes” on page 469</a>	26
6	Acct-Input-Octets	Number of octets received from the port over the course of this service being provided.  Used to charge the user for the amount of traffic they used.	42
7	Acct-Output-Octets	Number of octets sent to the port while delivering this service.  Used to charge the user for the amount of traffic they used.	43

[Table 19](#) describes the supported authentication events and the RADIUS attributes that are sent in the RADIUS accounting message.

**Table 19:** RADIUS attributes sent in RADIUS accounting message

Authentication Method	RADIUS Attributes						
	1	2	3	4	5	6	7
Web	a	a	a		a		
XAuth of IPsec (without DHCP)	a	a	a		a		
XAuth of IPsec (with DHCP)	a	a	a	a	a		
PPTP/L2TP (in PPP)	a	a	a	a	a	a	a
SSL-VPN	a	a	a		a		

## Vendor-specific attributes

Vendor specific attributes (VSA) are the method RADIUS servers and client companies use to extend the basic functionality of RADIUS. Some major vendors, such as Microsoft, have published their VSAs, however many do not.

In order to support vendor-specific attributes (VSA), the RADIUS server requires a dictionary to define which VSAs to support. This dictionary is typically supplied by the client or server vendor.

The Fortinet RADIUS vendor ID is 12356.

The FortiGate unit RADIUS VSA dictionary is supplied by Fortinet and is available through the Fortinet Knowledge Base (<http://kb.forticare.com>) or through Technical Support. Fortinet's dictionary for FortiOS 4.0 and up is configured this way:

```
##
Fortinet's VSA's
#
VENDOR fortinet 12356
BEGIN-VENDOR fortinet
ATTRIBUTE Fortinet-Group-Name 1 string
ATTRIBUTE Fortinet-Client-IP-Address 2 ipaddr
ATTRIBUTE Fortinet-Vdom-Name 3 string
ATTRIBUTE Fortinet-Client-IPv6-Address 4 octets
ATTRIBUTE Fortinet-Interface-Name 5 string
ATTRIBUTE Fortinet-Access-Profile 6 string
#
Integer Translations
#
END-VENDOR Fortinet
```

Note that using the Fortinet-Vdom-Name, users can be tied to a specific VDOM on the FortiGate unit. See the documentation provided with your RADIUS server for configuration details.

## Role Based Access Control

In Role Based Access Control (RBAC), network administrators and users have varying levels of access to network resources based on their role, and that role's requirement for access specific resources. For example, a junior accountant does not require access to the sales presentations, or network user account information.

There are three main parts to RBAC: role assignment, role authorization, and transaction authorization. Role assignment is accomplished when someone in an organization is assigned a specific role by a manager or HR. Role authorization is accomplished when a network administrator creates that user's RADIUS account and assigns them to the required groups for that role. Transaction authorization occurs when that user logs on and authenticates before performing a task.

RBAC is enforced when FortiOS network users are remotely authenticated via a RADIUS server. For users to authenticate, an identity-based security policy must be matched. That policy only matches a specific group of users. If VDOMs are enabled, the matched group will be limited to a specific VDOM. Using this method network administrators can separate users into groups that match resources, protocols, or VDOMs. It is even possible to limit users to specific FortiGate units if the RADIUS servers serve multiple FortiOS units.

For more information on identity-based policies, see [“Authentication in security policies” on page 517](#).

## Configuring the FortiGate unit to use a RADIUS server

The information you need to configure the FortiGate unit to use a RADIUS server includes

- the RADIUS server's domain name or IP address
- the RADIUS server's shared secret key.

You can optionally specify the NAS IP or Called Station ID. When configuring the FortiGate to use a RADIUS server, the FortiGate is a Network Access Server (NAS). If the FortiGate interface has multiple IP addresses, or you want the RADIUS requests to come from a different address you can specify it here. Called Station ID applies to carrier networks. However, if the NAS IP is not included in the RADIUS configuration, the IP of the FortiGate unit interface that communicates with the RADIUS server is used instead.

A maximum of 10 remote RADIUS servers can be configured on the FortiGate unit. One or more servers must be configured on FortiGate before remote users can be configured. To configure remote users, see [“Creating users” on page 491](#).

On the FortiGate unit, the default port for RADIUS traffic is 1812. Some RADIUS servers use port 1645. If this is the case with your server, you can either:

- Re-configure the RADIUS server to use port 1812. See your RADIUS server documentation for more information on this procedure.

or

- Change the FortiGate unit default RADIUS port to 1645 using the CLI:

```
config system global
 set radius-port 1645
end
```

One wildcard admin account can be added to the FortiGate unit when using RADIUS authentication. This uses the wildcard character to allow multiple admin accounts on RADIUS to use a single account on the FortiGate unit. See [“Example — wildcard admin accounts - CLI” on page 476](#).

### To configure the FortiGate unit for RADIUS authentication - web-based manager

1. Go to *User & Device > Authentication > RADIUS Servers* and select *Create New*.
2. Enter the following information and select OK.

<b>Name</b>	A name to identify the RADIUS server on the FortiGate unit.
<b>Primary Server Name/IP</b>	Enter the domain name (such as fgt.exmample.com) or the IP address of the RADIUS server.
<b>Primary Server Secret</b>	Enter the server secret key, such as radiusSecret. This can be a maximum of 16 characters long.  This must match the secret on the RADIUS primary server.
<b>Secondary Server Name/IP</b>	Optionally enter the domain name (such as fgt.exmample.com) or the IP address of the secondary RADIUS server.
<b>Secondary Server Secret</b>	Optionally, enter the secondary server secret key, such as radiusSecret2. This can be a maximum of 16 characters long.  This must match the secret on the RADIUS secondary server.
<b>Authentication Scheme</b>	If you know the RADIUS server uses a specific authentication protocol, select it from the list. Otherwise select <i>Use Default Authentication Scheme</i> . The Default option will usually work.

<b>NAS IP/ Called Station ID</b>	Enter the IP address to be used as an attribute in RADIUS access requests.  <b>NAS-IP-Address</b> is RADIUS setting or IP address of FortiGate interface used to talk to RADIUS server, if not configured.  <b>Called Station ID</b> is same value as NAS-IP Address but in text format.
<b>Include in every User Group</b>	When enabled this RADIUS server will automatically be included in all user groups. This is useful if all users will be authenticating with the remote RADIUS server.



For MAC OS and iOS devices to authenticate, you must use MS-CHAP-v2 authentication. In the CLI, the command is `set auth-type ms_chap_v2`.

3. Select OK.

### To configure the FortiGate unit for RADIUS authentication - CLI example

```
config user radius
 edit ourRADIUS
 set auth-type auto
 set server 10.11.102.100
 set secret radiusSecret
 end
```

For more information about RADIUS server options, refer to the [FortiGate CLI Reference](#).

## Troubleshooting RADIUS

To test the connection to the RADIUS server use the following command:

```
diagnose test authserver radius-direct <server_name or IP> <port number> <secret>
```

For the port number, enter -1 to use the default port. Otherwise enter the port number to check.

Additional RADIUS related troubleshooting is located at [“Troubleshooting FSSO” on page 596](#)

## LDAP servers

Lightweight Directory Access Protocol (LDAP) is an Internet protocol used to maintain authentication data that may include departments, people, groups of people, passwords, email addresses, and printers. LDAP consists of a data-representation scheme, a set of defined operations, and a request/response network.

The scale of LDAP servers range from big public servers such as BigFoot and Infospace, to large organizational servers at universities and corporations, to small LDAP servers for workgroups that may be using OpenLDAP. This document focuses on the institutional and workgroup applications of LDAP.

This section includes:

- [Components and topology](#)
- [LDAP directory organization](#)
- [Configuring the FortiGate unit to use an LDAP server](#)
- [Example — wildcard admin accounts - CLI](#)
- [Example of LDAP to allow Dial-in through member-attribute - CLI](#)
- [Troubleshooting LDAP](#)

## Components and topology

LDAP organization starts with directories. A directory is a set of objects with similar attributes organized in a logical and hierarchical way. Generally, an LDAP directory tree reflects geographic and organizational boundaries, with the Domain name system (DNS) names to structure the top level of the hierarchy. The common name identifier for most LDAP servers is `cn`, however some servers use other common name identifiers such as `uid`.

When LDAP is configured and a user is required to authenticate the general steps are:

1. The FortiGate unit contacts the LDAP server for authentication.
2. To authenticate with the FortiGate unit, the user enters a username and password.
3. The FortiGate unit sends this username and password to the LDAP server.
4. If the LDAP server can authenticate the user, the user is successfully authenticated with the FortiGate unit.
5. If the LDAP server cannot authenticate the user, the connection is refused by the FortiGate unit.

## Binding

Binding is the step where the LDAP server authenticates the user. If the user is successfully authenticated, binding allows the user access to the LDAP server based on that user's permissions.

The FortiGate unit can be configured to use one of three types of binding:

- anonymous - bind using anonymous user search
- regular - bind using username/password and then search
- simple - bind using a simple password authentication without a search

You can use simple authentication if the user records all fall under one domain name (`dn`). If the users are under more than one `dn`, use the anonymous or regular type, which can search the entire LDAP database for the required username.

If your LDAP server requires authentication to perform searches, use the regular type and provide values for username and password.

## Supported versions

The FortiGate unit supports LDAP protocol functionality defined in RFC 2251: Lightweight Directory Access Protocol v3, for looking up and validating user names and passwords. FortiGate LDAP supports all LDAP servers compliant with LDAP v3, including FortiAuthenticator. In addition, FortiGate LDAP supports LDAP over SSL/TLS, which can be configured only in the CLI.

FortiGate LDAP does not support proprietary functionality, such as notification of password expiration, which is available from some LDAP servers. FortiGate LDAP does not supply information to the user about why authentication failed.




## LDAP directory organization

To configure your FortiGate unit to work with an LDAP server, you need to understand the organization of the information on the server.

The top of the hierarchy is the organization itself. Usually this is defined as Domain Component (DC), a DNS domain. If the name contains a dot, such as `example.com`, it is written as two parts separated by a comma: `dc=example,dc=com`.

In this example, Common Name (CN) identifiers reside at the Organization Unit (OU) level, just below DC. The Distinguished Name (DN) is `ou=People,dc=example,dc=com`.

**Figure 112:**LDAP object hierarchy



In addition to the DN, the FortiGate unit needs an identifier for the individual person. Although the FortiGate unit GUI calls this the Common Name (CN), the identifier you use is not necessarily CN. On some servers, CN is the full name of a person. It might be more convenient to use the same identifier used on the local computer network. In this example, User ID (UID) is used.

### Locating your identifier in the hierarchy

You need to determine the levels of the hierarchy from the top to the level that contain the identifier you want to use. This defines the DN that the FortiGate unit uses to search the LDAP database. Frequently used distinguished name elements include:

- uid (user identification)
- pw (password)
- cn (common name)
- ou (organizational unit)
- o (organization)
- c (country)

One way to test this is with a text-based LDAP client program. For example, OpenLDAP includes a client, `ldapsearch`, that you can use for this purpose.

Enter the following at the command line:

```
ldapsearch -x '(objectclass=*)'
```

The output is lengthy, but the information you need is in the first few lines:

```
version: 2
#
filter: (objectclass=*)
requesting: ALL

dn: dc=example,dc=com
dc: example
objectClass: top
objectClass: domain

dn: ou=People,dc=example,dc=com
ou: People
objectClass: top
objectClass: organizationalUnit
...
dn: uid=tbrown,ou=People,dc=example,dc=com
uid: tbrown
cn: Tom Brown
```

In the output above, you can see `tbrown` (uid) and `Tom Brown`(cn). Also note the dn is `ou=People, dc=example, dc=com`.

## Configuring the FortiGate unit to use an LDAP server

After you determine the common name and distinguished name identifiers and the domain name or IP address of the LDAP server, you can configure the server on the FortiGate unit. The maximum number of remote LDAP servers that can be configured is 10.

One or more servers must be configured on FortiGate before remote users can be configured. To configure remote users, see [“Creating users” on page 491](#).

### To configure the FortiGate unit for LDAP authentication - web-based manager

1. Go to *User & Device > Authentication > LDAP Servers* and select *Create New*.
2. Enter a *Name* for the LDAP server.
3. In *Server Name/IP* enter the server’s FQDN or IP address.
4. If necessary, change the *Server Port* number. The default is port 389.
5. Enter the *Common Name Identifier* (20 characters maximum).
6. `cn` is the default, and is used by most LDAP servers.
7. In the *Distinguished Name* field, enter the base distinguished name for the server using the correct X.500 or LDAP format.

The FortiGate unit passes this distinguished name unchanged to the server. The maximum number of characters is 512.

If you don’t know the distinguished name, leave the field blank and select the Query icon to the right of the field. See the [“Using the Query icon” on page 475](#).

8. In *Bind Type*, select *Regular*.
9. In *User DN*, enter the LDAP administrator’s distinguished name.
10. In *Password*, enter the LDAP administrator’s password.

## 11. Select OK.



To verify your Distinguished Name field is correct, you can select the *Test* button. If your DN field entry is valid, you will see the part of the LDAP database it defines. If your DN field entry is not valid, it will display an error message and return no information.

For detailed information about configuration options for LDAP servers, see the Online Help on your FortiGate unit or the FortiGate CLI Reference.

### To configure the FortiGate unit for LDAP authentication - CLI example

```
config user ldap
 edit ourLDAPsrv
 set server 10.11.101.160
 set cnid cn
 set dn cn=users,dc=office,dc=example,dc=com
 set type regular
 set username
 cn=administrator,cn=users,dc=office,dc=example,dc=com
 set password w5AiGVMLkgyPQ
 set password-expiry-warning enable
 set password-renewal enable
 end
```

### password-expiry-warning and password-renewal

In SSLVPN, when an LDAP user is connecting to the LDAP server it is possible for them to receive any pending password expiry or renewal warnings. When the password renewal or expiry warning exists, SSLVPN users will see a prompt allowing them to change their password.

`password-expiry-warning` allows FortiOS to detect from the LDAP server when a password is expiring or has expired using server controls or error codes.

`password-renewal` allows FortiOS to perform the online LDAP password renewal operations the LDAP server expects.

On an OpenLDAP server, when a user attempts to logon with an expired password they are allowed to logon on but only to change their password.

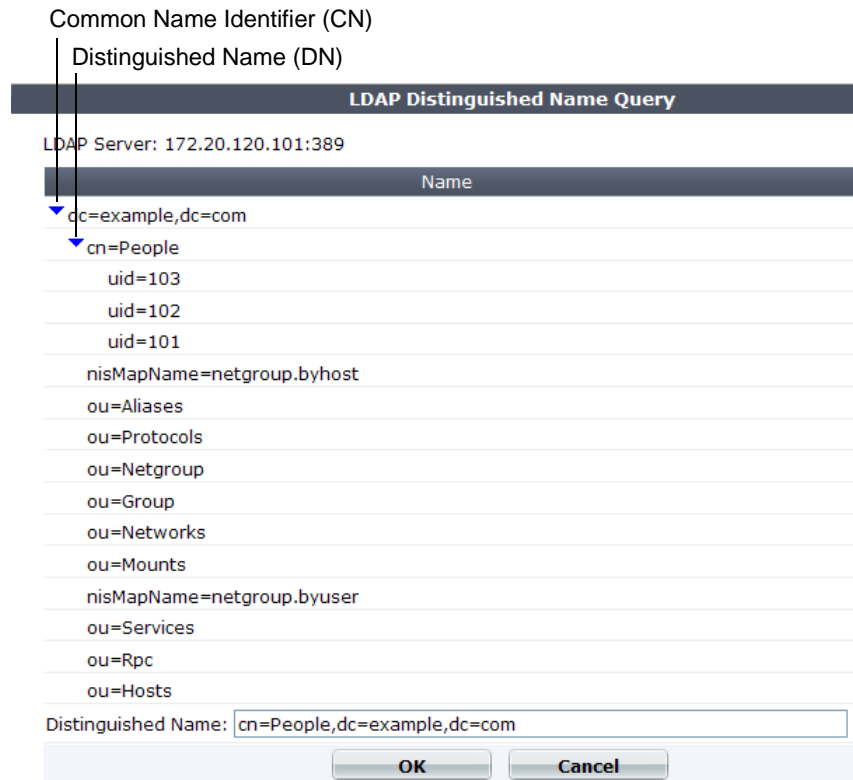
When changing passwords on a Windows AD system, the connection must be SSL-protected.

### Using the Query icon

The LDAP Distinguished Name Query list displays the LDAP directory tree for the LDAP server connected to the FortiGate unit. This helps you to determine the appropriate entry for the DN field. To see the distinguished name associated with the Common Name identifier, select the Expand icon next to the CN identifier. Select the DN from the list. The DN you select is displayed in the Distinguished Name field. Select OK and the Distinguished Name you selected will be saved in the Distinguished Name field of the LDAP Server configuration.

To see the users within the LDAP Server user group for the selected Distinguished Name, expand the Distinguished Name in the LDAP Distinguished Name Query tree.

**Figure 113:**LDAP server Distinguished Name Query tree



### Example — wildcard admin accounts - CLI

A wildcard admin account is an administrator account with the wildcard option enabled. This option allows multiple different remote administration accounts to match one local administration account, avoiding the need to set up individual admin accounts on the FortiGate unit. Instead multiple LDAP admin accounts will all be able to use one FortiGate admin account.

The initial benefit of wildcard admin accounts is fast configuration of the FortiGate unit's administration account to work with your LDAP network. The many to one ratio saves on effort, and potential errors.

The ongoing benefit is that as long as the users on the LDAP system belong to that group, and the test admin user settings don't change on the FortiGate unit, no other work is required. This point is important as it can help avoid system updates or changes that would otherwise require changes to the LDAP administrator account configuration. Even if a user is added to or removed from the LDAP group, no changes are required on the FortiGate unit.

Two potential issues with wildcard admin accounts are that multiple users may be logged on to the same account at the same time. This becomes an issue if they are changing the same information at the same time. The other potential issue is that security is reduced because multiple people have login access for the same account. If each user was assigned their own account, a hijacking of one account would not affect the other users.

Note that wildcard admin configuration also applies to RADIUS. When configuring for RADIUS, configure the RADIUS server, and RADIUS user group instead of LDAP. When using web-based management, wildcard admin is the only type of remote administrator account that does not require you to enter a password on account creation. That password is normally used when the remote authentication server is unavailable during authentication.

In this example, default values are used where possible. If a specific value is not mentioned, it is set to its default value.

## Configuring the LDAP server

The important parts of this configuration are the username and group lines. The username is the domain administrator account. The group binding allows only the group with the name `GRP` to access.



The dn used here is as an example only. On your network use your own domain name.

### To configure LDAP server - CLI

```
config user ldap
 edit "ldap_server"
 set server "192.168.201.3"
 set cnid "sAMAccountName"
 set dn "DC=example,DC=com,DC=au"
 set type regular
 set username "CN=Administrator,CN=Users,DC=example,DC=COM"
 set password *
 set group "CN=GRP,OU=training,DC=example,DC=COM"
 set filter ""
 next
end
```

### To configure the user group and add the LDAP server - CLI

```
config user group
 edit "ldap_grp"
 set member "ldap"
 config match
 edit 1
 set server-name "ldap_server"
 set group-name "TRUE"
 next
 end
next
end
```

## Configuring the admin account

The wildcard part of this example is only available in the CLI for admin configuration. When enabled, this allows all LDAP group members to login to the FortiGate unit without the need to create a separate admin account for each user. In effect the members of that group will each be able to login as "test".

### To configure the admin account - CLI

```
config system admin
 edit "test"
 set remote-auth enable
 set accprofile "super_admin"
 set wildcard enable
 set remote-group "ldap_grp"
 next
end
```

For troubleshooting, test that the admin account is operational, and see [“Troubleshooting LDAP” on page 479](#).

### Example of LDAP to allow Dial-in through member-attribute - CLI

In this example, users defined in MicroSoft Windows Active Directory (AD) are allowed to setup a VPN connection simply based on an attribute that is set to TRUE, instead of based on being part of a specific group.

In AD, the “Allow Dial-In” property is activated in the user properties, and this sets the `msNPAllowDialin` attribute to “TRUE”.

This same procedure can be used for other member attributes, as your system requires.

### Configuring LDAP member-attribute settings

To accomplish this with a FortiGate unit, the member attribute must be set. Setting member attributes can only be accomplished through the CLI using the `member-attr` keyword - the option is not available through the web-based manager.

Before configuring the FortiGate unit, the AD server must be configured and have the `msNPAllowDialin` attribute set to “TRUE” for the users in question. If not, those users will not be able to properly authenticate.

The dn used here is as an example only. On your network use your own domain name.

### To configure user LDAP member-attribute settings - CLI

```
config user ldap
 edit "ldap_server"
 set server "192.168.201.3"
 set cnid "sAMAccountName"
 set dn "DC=fortinet,DC=com,DC=au"
 set type regular
 set username "fortigate@example.com"
 set password *****
 set member-attr "msNPAllowDialin"
 next
end
```

### Configuring LDAP group settings

A user group that will use LDAP must be configured. This example adds the member `ldap` to the group which is the LDAP server name that was configured earlier.

## To configure LDAP group settings - CLI

```
config user group
 edit "ldap_grp"
 set member "ldap"
 config match
 edit 1
 set server-name "ldap"
 set group-name "TRUE"
 next
 end
end
```

Once these settings are in place, users can authenticate.

## Troubleshooting LDAP

The examples in this section use the values from the previous example.

### LDAP user test

A quick way to see if the LDAP configuration is correct is to run a diagnose CLI command with LDAP user information. The following command tests with a user called `netAdmin` and a password of `fortinet`. If the configuration is correct the test will be successful.

```
FGT# diag test authserver ldap ldap_server netAdmin fortinet
'ldap_server' is not a valid ldap server name — an LDAP server by that name
has not been configured on the FortiGate unit, check your spelling.

authenticate 'netAdmin' against 'ldap_server' failed! — the user netAdmin
does not exist on ldap_server, check your spelling of both the user and sever and ensure the
user has been configured on the FortiGate unit.
```

### LDAP authentication debugging

For a more in-depth test, you can use a `diag debug` command. The sample output from a shows more information about the authentication process that may prove useful if there are any problems.

Ensure the “Allow Dial-in” attribute is still set to “TRUE” and run the following CLI command. `fnbamd` is the Fortinet non-blocking authentication daemon.

```
FGT# diag debug enable
FGT# diag debug reset
FGT# diag debug application fnbamd -1
FGT# diag debug enable
```

The output will look similar to:

```
get_member_of_groups-Get the memberOf groups.
get_member_of_groups- attr='msNPAllowDialin', found 1 values
get_member_of_groups-val[0]='TRUE'
fnbamd_ldap_get_result-Auth accepted
fnbamd_ldap_get_result-Going to DONE state res=0
fnbamd_auth_poll_ldap-Result for ldap svr 192.168.201.3 is SUCCESS
fnbamd_auth_poll_ldap-Passed group matching
```

If the “Allow Dial-in” attribute is not set but it is expected, the last line of the above output will instead be:

```
fnbamd_auth_poll_ldap-Failed group matching
```

## TACACS+ servers

When users connect to their corporate network remotely, they do so through a remote access server. As remote access technology has evolved, the need for security when accessing networks has become increasingly important. This need can be filled using a Terminal Access Controller Access-Control System (TACACS+) server.

TACACS+ is a remote authentication protocol that provides access control for routers, network access servers, and other networked computing devices via one or more centralized servers. TACACS+ allows a client to accept a username and password and send a query to a TACACS+ authentication server. The server host determines whether to accept or deny the request and sends a response back that allows or denies the user access to the network.

TACACS+ offers fully encrypted packet bodies, and supports both IP and AppleTalk protocols. TACACS+ uses TCP port 49, which is seen as more reliable than RADIUS’s UDP protocol.

There are several different authentication protocols that TACACS+ can use during the authentication process:

**Table 20:** Authentication protocols

Protocol	Definition
<b>ASCII</b>	Machine-independent technique that uses representations of English characters. Requires user to type a username and password that are sent in clear text (unencrypted) and matched with an entry in the user database stored in ASCII format.
<b>PAP</b>	Password Authentication Protocol (PAP) Used to authenticate PPP connections. Transmits passwords and other user information in clear text.
<b>CHAP</b>	Challenge-Handshake Authentication Protocol (CHAP) Provides the same functionality as PAP, but is more secure as it does not send the password and other user information over the network to the security server.
<b>MS-CHAP</b>	MicroSoft Challenge-Handshake Authentication Protocol v1 (MSCHAP) Microsoft-specific version of CHAP.
<b>default</b>	The default protocol configuration, Auto, uses PAP, MS-CHAP, and CHAP, in that order.

## Configuring a TACACS+ server on the FortiGate unit

A maximum of 10 remote TACACS+ servers can be configured for authentication.

One or more servers must be configured on FortiGate before remote users can be configured. To configure remote users, see [“Creating users” on page 491](#).

The TACACS+ page in the web-based manager is not available until a TACACS+ server has been configured in the CLI. For more information see the CLI Reference.

### To configure the FortiGate unit for TACACS+ authentication - web-based manager

1. Go to *User & Device > Authentication > TACACS+ Servers* and select *Create New*.



2. Enter the following information, and select *OK*.

<b>Name</b>	Enter the name of the TACACS+ server.
<b>Server Name/IP</b>	Enter the server domain name or IP address of the TACACS+ server.
<b>Server Key</b>	Enter the key to access the TACACS+ server.
<b>Authentication Type</b>	Select the authentication type to use for the TACACS+ server. <i>Auto</i> tries PAP, MSCHAP, and CHAP (in that order).

### To configure the FortiGate unit for TACACS+ authentication - CLI

```
config user tacacs+
 edit tacacs1
 set authen-type auto
 set key abcdef
 set port 49
 set server 192.168.0.101
 end
```

## SSO servers

Novell and Microsoft Windows networks provide user authentication based on directory services: eDirectory for Novell, Active Directory for Windows. Users can log on at any computer in the domain and have access to resources as defined in their user account. The Fortinet Single Sign On (FSSO) agent enables FortiGate units to authenticate these network users for security policy or VPN access without asking them again for their username and password.

When a user logs in to the Windows or Novell domain, the FSSO agent sends the FortiGate unit the user's IP address and the names of the user groups to which the user belongs. The FortiGate unit uses this information to maintain a copy of the domain controller user group database. Because the domain controller authenticates users, the FortiGate unit does not perform authentication. It recognizes group members by their IP address.

In the FortiOS FSSO configuration, you specify the server where the FSSO Collector agent is installed. The Collector agent retrieves the names of the Novell or Active Directory user groups from the domain controllers on the domains, and then the FortiGate unit gets them from the Collector agent. You cannot use these groups directly. You must define FSSO type user groups on your FortiGate unit and then add the Novell or Active Directory user groups to them. The FSSO user groups that you created are used in security policies and VPN configurations to provide access to different services and resources.

FortiAuthenticator servers can replace the Collector agent when FSSO is using polling mode. The benefits of this is that FortiAuthenticator is a stand-alone server that has the necessary FSSO software pre-installed. For more information, see the [FortiAuthenticator Administration Guide](#).

## Single Sign-on Agent configuration settings

The following are SSO configuration settings in *User & Device > Authentication > Single Sign-On*.

---

### SSO Server List

Lists all the collector agents' lists that you have configured. On this page, you can create, edit or delete FSSO agents.

**Note:** You can create a redundant configuration on your unit if you install a collector agent on two or more domain controllers. If the current (or first) collector agent fails, the Fortinet unit switches to the next one in its list of up to five collector agents.

---

<b>Create New</b>	Creates a new agent. When you select <i>Create New</i> , you are automatically redirected to the New page.
<b>Edit</b>	Modifies the settings for the selected SSO server.  To remove multiple entries from the list, for each servers you want removed, select the check box and then select <i>Delete</i> .  To remove all agents from the list, on the FSSO Agent page, select the check box at the top of the check box column and then select <i>Delete</i> .
<b>Delete</b>	Removes an agent from the list on the page.

---

### Settings when *Type is Poll Active Directory Server*

---

<b>Server</b>	The IP address of the domain controller (DC).
<b>User</b>	The user ID used to access the domain controller.
<b>Password</b>	Enter the password for the account used to access the DC.
<b>LDAP Server</b>	Select the check box and select an LDAP server to access the Directory Service.
<b>Enable Polling</b>	Enable to allow the FortiGate unit to poll this DC.
<b>Users/Groups</b>	A list of user and user group names retrieved from the DC.

---

### Settings when *Type is Fortinet Single Sign On Agent*

---

<b>Name</b>	Enter a name for the SSO server.
<b>Primary Agent IP/Name</b> <b>Secondary Agent IP/Name</b>	Enter the IP address or name of the Directory Service server where this SSO agent is installed. The maximum number of characters is 63.
<b>Password</b>	Enter the password for the collector agent. This is required only if you configured your Fortinet Single Sign On Agent collector agent to require authenticated access.
<b>More FSSO agents</b>	Select to add up to three additional SSO agents.
<b>Users/Groups</b>	A list of user and user group names retrieved from the server.

---

## RSA ACE (SecurID) servers

SecurID is a two-factor system that uses one-time password (OTP) authentication. It is produced by the company RSA. This system includes portable tokens carried by users, an RSA ACE/Server, and an Agent Host. In our configuration, the FortiGate unit is the Agent Host.

### Components

When using SecurID, users carry a small device or “token” that generates and displays a random password. According to RSA, each SecurID authenticator token has a unique 64-bit symmetric key that is combined with a powerful algorithm to generate a new code every 60 seconds. The token is time-synchronized with the SecurID RSA ACE/Server.

The RSA ACE/Server is the management component of the SecurID system. It stores and validates the information about the SecurID tokens allowed on your network. Alternately the server could be an RSA SecurID 130 Appliance.

The Agent Host is the server on your network, in this case it is the FortiGate unit, that intercepts user logon attempts. The Agent Host gathers the user ID and password entered from their SecurID token, and sends that information to the RSA ACE/Server to be validated. If valid, a reply comes back indicating it is a valid logon and the FortiGate unit allows the user access to the network resources specified in the associated security policy.

### Configuring the SecurID system

To use SecurID with a FortiGate unit, you need:

- to configure the RSA server and the RADIUS server to work with each other (see RSA server documentation)
- [To configure the RSA SecurID 130 Appliance](#)
- or
- [To configure the FortiGate unit as an Agent Host on the RSA ACE/Server](#)
- [To configure the FortiGate unit to use the RADIUS server](#)
- [To create a SecurID user group and user](#)
- [To configure a security policy with SecurID authentication](#)

The following instructions are based on RSA ACE/Server version 5.1, or RSA SecurID 130 Appliance, and assume that you have successfully completed all the external RSA and RADIUS server configuration steps listed above.

For this example, the RSA server is on the internal network, with an IP address of 192.128.100.100. The FortiGate unit internal interface address is 192.168.100.3, RADIUS shared secret is fortinet123, RADIUS server is at IP address 192.168.100.102.

#### To configure the RSA SecurID 130 Appliance

1. Go to the IMS Console for SecurID and logon.
2. Go to *RADIUS > RADIUS Clients*, and select *Add New*.
3. Enter the following information to configure your FortiGate as a SecurID Client, and select *Save*.

---

#### RADIUS Client Basics

---

<b>Client Name</b>	FortiGate
--------------------	-----------

---

<b>Associated RSA Agent</b>	FortiGate
<b>RADIUS Client Settings</b>	
<b>IP Address</b>	192.168.100.3 The IP address of the FortiGate unit internal interface.
<b>Make / Model</b>	Select Standard Radius
<b>Shared Secret</b>	fortinet123 The RADIUS shared secret.
<b>Accounting</b>	Leave unselected
<b>Client Status</b>	Leave unselected

### To configure the FortiGate unit as an Agent Host on the RSA ACE/Server

1. On the RSA ACE/Server computer, go to *Start > Programs > RSA ACE/Server*, and then *Database Administration - Host Mode*.
2. On the *Agent Host* menu, select *Add Agent Host*.
3. Enter and save the following information.

<b>Name</b>	FortiGate
<b>Network Address</b>	192.168.100.3 The IP address of the FortiGate unit.
<b>Secondary Nodes</b>	Optionally enter other IP addresses that resolve to the FortiGate unit.

If needed, refer to the RSA ACE/Server documentation for more information.

### To configure the FortiGate unit to use the RADIUS server

1. Go to *User & Device > Authentication > RADIUS Servers* and select *Create New*.
2. Enter the following information, and select OK.

<b>Name</b>	RSA
<b>Type</b>	Query
<b>Primary Server Address</b>	192.168.100.102 Optionally select Test to ensure the IP address is correct and the FortiGate can contact the RADIUS server.
<b>Primary Server Secret</b>	fortinet123
<b>Authentication Scheme</b>	Select <i>Use Default Authentication Scheme</i> .

### To create a SecurID user group and user

1. Go to *User & Device > User > User Groups*, and select *Create New*.
2. Enter the following information, and select OK.

<b>Name</b>	RSA_group
<b>Remote Authentication servers</b>	Select the RSA server.

3. Go to *User & Device > User > User Definition*, and select *Create New*.
4. Enter the following information, and select OK.

<b>User Name</b>	wloman
<b>Match User on RADIUS server</b>	RSA
<b>Add this user to groups</b>	Select RSA_group

To test this configuration, on your FortiGate unit use the CLI command:

```
diag test auth rad RSA auto wloman 111111111
```

The series of 1s is the one time password that your RSA SecurID token generates and you enter.

### Using the SecurID user group for authentication

You can use the SecurID user group in several FortiOS features that authenticate by user group including

- [Security policy](#)
- [IPsec VPN XAuth](#)
- [PPTP VPN](#)
- [SSL VPN](#)

The following sections assume the SecurID user group is called `securIDgrp` and has already been configured. Unless otherwise states, default values are used.

### Security policy

To use SecurID in a security policy, you must include the SecurID user group in an identity-based security policy. This procedure will create a security policy that allows HTTP, FTP, and POP3 traffic from the `internal` interface to `wan1`. If these interfaces are not available on your FortiGate unit, substitute other similar interfaces.

### To configure a security policy with SecurID authentication

1. Go to *Policy > Policy > Policy*.
2. Select *Create New*.
3. In *Policy Subtype*, select *User Identity*.

4. Enter

<b>Incoming Interface</b>	internal
<b>Source Address</b>	all
<b>Outgoing Interface</b>	wan1
<b>Enable NAT</b>	Selected.

5. In *Configure Authentication Rules*, select *Create New*.

6. Enter

<b>Destination Address</b>	all
<b>Group(s)</b>	securIDgrp
<b>Schedule</b>	always
<b>Services</b>	HTTP, FTP, POP3
<b>Action</b>	ACCEPT

7. To generate usage reports on traffic authenticated with this policy, enable *Log Allowed Traffic*.
8. To either limit traffic or guarantee minimum bandwidth for traffic that uses the SecurID security policy, enable *Traffic Shaping* and *Shared Traffic Shaper* and then select one of the default shapers from the list such as *guarantee-100kbps*.
9. Select OK.

You are returned to the security policy creation page, with the information you just entered in the *Configure Authentication Rules* table.

10. Optionally, you can modify any challenge pages or logon pages users will see. Select *Customize Authentication Messages* and select the Edit icon that appears.
11. Select OK.

The SecurID security policy is configured.

For more detail on configuring security policies, see the [FortiOS Handbook FortiGate Fundamentals chapter](#).

## IPsec VPN XAuth

Extended Authentication (XAuth) increases security by requiring additional user authentication information in a separate exchange at the end of the VPN Phase 1 negotiation. If the SecurID user group is used, this extended information will require users to enter their SecurID code. For more on XAuth, see [“Configuring XAuth authentication” on page 529](#).

This Phase 1 configuration will be named `securIDxAuth` and it will connect with IP address `10.11.101.155` on the `wan1` interface.

### To configure IPsec VPN XAuth with SecurID authentication - web-based manager

1. Go to *VPN > IPsec > Auto Key (IKE)*.

2. Select *Create Phase 1* and enter

<b>Name</b>	securIDxAAuth
<b>Remote Gateway</b>	Dialup User
<b>Local Interface</b>	wan1
<b>Mode</b>	Main (ID protection)
<b>Authentication Method</b>	Preshared Key
<b>Pre-shared Key</b>	fortinet
<b>Peer Options</b>	Accept any peer ID.

3. Select *Advanced...* and enter

<b>XAUTH</b>	Enable as Server
<b>Server Type</b>	AUTO
<b>User Group</b>	securIDgrp

4. Select *OK*.

## PPTP VPN

PPTP VPN is configured in the CLI. In the PPTP configuration (`config vpn pptp`), set `usrgrp` to the SecurID user group.

## SSL VPN

In the SecurID user group, select the appropriate web portal for these users. In the security policy for the SSL VPN, include the SecurID user group in the list of selected user groups.

# Users and user groups

FortiGate authentication controls system access by user group. By assigning individual users to the appropriate user groups you can control each user's access to network resources. The members of user groups are user accounts, of which there are several types. Local users and peer users are defined on the FortiGate unit. User accounts can also be defined on remote authentication servers.

This section describes how to configure local users and peer users and then how to configure user groups. For information about configuration of authentication servers see [“Authentication servers” on page 466](#).

This section contains the following topics:

- [Users](#)
- [User groups](#)

## Users

A user is a user account consisting of username, password, and in some cases other information, configured on the FortiGate unit or on an external authentication server. Users can access resources that require authentication only if they are members of an allowed user group. There are several different types of user accounts with slightly different methods of authentication:

User type	Authentication
Local user, password stored on the FortiGate unit	The username and password must match a user account stored on the FortiGate unit. Authentication by FortiGate security policy.
Local user, password stored on a remote server	The username must match a user account stored on the FortiGate unit and the username and password must match a user account stored on the remote authentication server.
Authentication server user	A FortiGate user group can include user accounts or groups that exist on a remote authentication server.
FSSO user	With Fortinet Single Sign On (FSSO), users on a Microsoft Windows or Novell network can use their network authentication to access resources through the FortiGate unit. Access is controlled through FSSO user groups which contain Windows or Novell user groups as their members.
Peer user with certificate authentication	A peer user is a digital certificate holder that authenticates using a client certificate. No password is required, unless two-factor authentication is enabled.



User type	Authentication
IM Users	IM users are not authenticated. The FortiGate unit can allow or block each IM user name from accessing the IM protocols. A global policy for each IM protocol governs access to these protocols by unknown users.
Guest Users	Guest user accounts are temporary. The account expires after a selected period of time.

This section includes:

- [Local users](#)
- [PKI or peer users](#)
- [Two-factor authentication](#)
- [FortiToken](#)
- [IM users](#)
- [Monitoring users](#)

## Local users

Local users are defined on the FortiGate unit in *User & Device > User > User Definition*.

---

### User page

Lists each individual local user's list that you created. On this page, you can edit, delete or create a new local users list.

**Note:** If you want to have users always authenticate whenever their time expires, use the `hard-timeout` value in the `auth-type` command. This is available only in the CLI.

---

**Create New** Creates a new local user account. When you select *Create New*, you are automatically redirected to New User page.

---

**Edit** Modifies a user's account settings. When you select *Edit*, you are automatically redirected to the Edit User page.

---

**Delete** Removes a user from the list. Removing the user name removes the authentication configured for the user.

The *Delete* icon is not available if the user belongs to a user group.

To remove multiple local user accounts from within the list, on the User page, in each of the rows of user accounts you want removed, select the check box and then select *Delete*.

To remove all local user accounts from the list, on the User page, select the check box in the check box column and then select *Delete*.

---

**User Name** The local user name. If the user is authenticated externally, the username on the FortiGate unit must be identical to the username on the authentication server.

---

**Type** The authentication type to use for this user. The authentication types are Local (user and password stored on Fortinet unit), LDAP, RADIUS, and TACACS+ (user and password matches a user account stored on the authentication server).

---

<b>Two-factor Authentication</b>	Indicates whether two-factor authentication is configured for the user. Gray “X” — not enabled Green check mark — enabled
<b>Ref.</b>	Displays the number of times this object is referenced by other objects. Select the number to open the Object Usage window and view the list of referring objects. The list is grouped into expandable categories, such as Firewall Policy. Numbers of objects are shown in parentheses.  To view more information about the referring object, use the icons: <ul style="list-style-type: none"> <li>• <b>View the list page for these objects</b> – available for object categories. Goes to the page where the object is listed. For example, if the category is User Groups, opens User Groups list.</li> <li>• <b>Edit this object</b> – opens the object for editing. modifies</li> <li>• <b>View the details for this object</b> – displays current settings for the object.</li> </ul>
<b>New User or Edit User page</b>	
Provides settings for a new or existing local user.	
<b>User Name</b>	A name that identifies the user.
<b>Disable</b>	Select to prevent this user from authenticating.
<b>Password</b>	Select to authenticate this user using a password stored on the FortiGate unit. Enter the password. Best practice is to create a password at least six characters long.
<b>Match users on LDAP servers</b>	Select to authenticate this user using a password stored on an LDAP server. Select the LDAP server from the list.  You can select only an LDAP server that has been added to the Fortinet LDAP configuration. For more information, see <a href="#">“Configuring the FortiGate unit to use an LDAP server” on page 474.</a>
<b>Match users on RADIUS server</b>	Select to authenticate this user using a password stored on a RADIUS server. Select the RADIUS server from the list.  You can select only a RADIUS server that has been added to the Fortinet RADIUS configuration. For more information, see <a href="#">“Configuring the FortiGate unit to use a RADIUS server” on page 470.</a>
<b>Match users on TACACS+ server</b>	Select to authenticate this user using a password stored on a TACACS+ server. Select the TACACS+ server from the list.  You can select only a TACACS+ server that has been added to the Fortinet TACACS+ configuration. For more information, see <a href="#">“TACACS+ servers” on page 480.</a>

---

<b>Contact Info</b>	Provide the email address or SMS cell number at which the user will receive token password codes. For custom SMS service, you must first enter the SMS service provider in <i>System &gt; Config &gt; Messaging Servers</i> before you can select it from the drop-down list. See <a href="#">“FortiToken” on page 497</a> .
<b>Enable Two-factor Authentication</b>	Select to enable two-factor authentication. Then select the Token (FortiToken or FortiToken Mobile) for this user account. See <a href="#">“Associating FortiTokens with accounts” on page 500</a> .

---

## Creating users

Before configuring any authentication, except RADIUS SSO, you must first create local users. For more about RADIUS SSO, see [“SSO using RADIUS accounting records” on page 602](#).

When creating a new user, there are only two differences between a local and a remote user:

- local users require a password to be configured
- remote users do not require a password, but do require a remote authentication server to be configured

### To create a local user - web-based manager

1. Go to *User & Device > User > User Definition* and select *Create New*.
2. Enter the username in the *username* field.
  - Select *Password* and type a password. Best practices dictate that the password be at least six characters long.



To authenticate this user using an external authentication server, select the *Match user* option for the appropriate type of server and select the server name. Password is not required. You must configure the remote server access first. See [“Authentication servers” on page 466](#).

3. Optionally select *Enable Two-factor Authentication* to use that option with this user.  
When enabled, additional options will be displayed. Select one of the following options and configure it as stated.
  - Select *FortiToken*, and choose the FortiToken serial number to associate with this user.
  - Select *Email to* and enter the user’s email address to email them the token code.
  - Select *SMS* and enter the Mobile Provider from the list, and enter the user’s mobile phone number that will receive the token code in a text message.
4. Select *OK*.



The Mobile Provider for SMS must be entered in the CLI using the `config system sms-server` command before it will be available to select in the web-based manager.

---

## To create a local user - CLI examples

### Locally authenticated user

```
config user local
 edit user1
 set type password
 set passwd ljt_pj2gpepfdw
 end
```

### User authenticated on an LDAP server

```
config user local
 edit user2
 set type ldap
 set ldap_server ourLDAPsrv
 end
```

### User authenticated on a RADIUS server

```
config user local
 edit user3
 set type radius
 set radius_server ourRADIUSsrv
 end
```

### User authenticated on a TACACS+ server

```
config user local
 edit user4
 set type tacacs+
 set tacacs+_server ourTACACS+srv
 end
```

### User authenticated with a FortiToken

```
config user local
 edit user5
 set type password
 set passwd ljt_pj2gpepfdw
 set two_factor fortitoken
 set fortitoken 182937197
 end
```

### User authenticated using email

```
config user local
 edit user6
 set type password
 set passwd ljt_pj4h7epfdw
 set two_factor email
 set email-to user6@sample.com
 end
```

User authenticated using SMS text message

```
config system sms-server
 edit "Sample Mobile Inc"
 set mail-server mail.sample.com
 end

config user local
 edit user7
 set type password
 set passwd 3ww_pjt68dw
 set two_factor sms
 set sms-server custom
 set sms-custom-server "Sample Mobile Inc"
 set sms-phone 2025551234
 end
```

## Removing users

Best practices dictate that when a user account is no longer in use, it be deleted. Removing local and remote users from FortiOS involve the same steps.

If the user account is referenced by any configuration objects those references must be removed before the user can be deleted. See [“Removing references to users” on page 493](#).

### To remove a user from the FortiOS configuration - web-based manager

1. Go to *User & Device > User > User Definition*.
2. Select the check box of the user that you want to remove.
3. Select *Delete*.
4. Select *OK*.

### To remove a user from the FortiOS configuration - CLI example

```
config user local
 delete user4444
end
```

## Removing references to users

You cannot remove a user that belongs to a user group. Remove the user from the user group first, and then delete the user.

### To remove references to a user - web-based manager

1. Go to *User & Device > User > User Definition*.
2. If the number in the far right column for the selected user contains any number other than zero, select it.
3. A more detailed list of object references to this user is displayed. Use its information to find and remove these references to allow you to delete this user.

## PKI or peer users

A PKI, or peer user, is a digital certificate holder. A PKI user account on the FortiGate unit contains the information required to determine which CA certificate to use to validate the user's

certificate. Peer users can be included in firewall user groups or peer certificate groups used in IPsec VPNs. For more on certificates, see [“Certificates overview” on page 533](#).

To define a peer user you need:

- a peer username
- the text from the subject field of the user’s certificate, or the name of the CA certificate used to validate the user’s certificate

## Creating a peer user

The configuration page for PKI users in the web-based manager. Follow the CLI-based instructions.

### To create a peer user for PKI authentication - CLI example

```
config user peer
 edit peer1
 set subject peer1@mail.example.com
 set ca CA_Cert_1
 end
```

There are other configuration settings that can be added or modified for PKI authentication. For example, you can configure the use of an LDAP server to check access rights for client certificates. For information about the detailed PKI configuration settings, see the [FortiGate CLI Reference](#).

## Two-factor authentication

The standard logon requires a username and password. This is one factor authentication—your password is one piece of information you need to know to gain access to the system.

Two factor authentication adds the requirement for another piece of information for your logon. Generally the two factors are something you know (password) and something you have (certificate, token, etc.). This makes it harder for a hacker to steal your logon information. For example if you have a FortiToken device, the hacker would need to both use it and know your password to gain entry to your account.

Two-factor authentication is available on both user and admin accounts. But before you enable two-factor authentication on an administrator account, you need to ensure you have a second administrator account configured to guarantee administrator access to the FortiGate unit if you are unable to authenticate on the main admin account for some reason.



Two-factor authentication does not work with explicit proxies.

---

The methods of two-factor authentication include:

- [Certificate](#)
- [Email](#)
- [SMS](#)
- [FortiToken](#)

## Certificate

You can increase security by requiring both certificate and password authentication for PKI users. Certificates are installed on the user's computer. Requiring a password also protects against unauthorized use of that computer.

Optionally peer users can enter the code from their FortiToken instead of the certificate.

### To create a peer user with two-factor authentication - CLI example

```
config user peer
 edit peer1
 set subject E=peer1@mail.example.com
 set ca CA_Cert_1
 set two-factor enable
 set passwd fdktguefheygfe
 end
```

For more information on certificates, see [“Certificates overview” on page 533](#).

## Email

Two-factor email authentication sends a randomly generated six digit numeric code to the specified email address. Enter that code when prompted at logon. This token code is valid for 60 seconds. If you enter this code after that time, it will not be accepted.

A benefit is that you do not require mobile service to authenticate. However, a potential issue is if your email server does not deliver the email before the 60 second life of the token expires.

The code will be generated and emailed at the time of logon, so you must have email access at that time to be able to receive the code.

### To configure an email provider - web-based manager

1. Go to *System > Config > Messaging Servers*.
2. Enter the *SMTP Server* and *Default Reply To* address.
3. If applicable, enable *Authentication* and enter the *SMTP User* and *Password* to use.
4. Select *Apply*.

### To configure an email provider - CLI

```
config system email-server
 edit <provider_name>
 set server <server_domain-name>
 next
end
```

### To enable email two-factor authentication - web-based manager

1. To modify an administrator account, go to *System > Admin > Administrators*. To modify a user account go to *User & Device > User > User Definition*.
2. Edit the user account.
3. Enable and enter the user's *Email Address*.
4. Select *Enable Two-factor Authentication*.
5. Select the *Token* that the user has.
6. Select *OK*.

### To enable email two-factor authentication - CLI

```
config user local
 edit <user_name>
 set email-to <user_email>
 set two-factor email
 end
```

## SMS

SMS two-factor authentication sends the token code in an SMS text message to the mobile device indicated when this user attempts to logon. This token code is valid for 60 seconds. If you enter this code after that time, it will not be accepted. Enter this code when prompted at logon to be authenticated.

SMS two-factor authentication has the benefit that you do not require email service before logging on. A potential issue is if the mobile service provider does not send the SMS text message before the 60 second life of the token expires.

If you do not use the FortiGuard Messaging Service, you need to configure an SMS service.

### To configure an SMS service for your FortiGate unit - web-based manager

1. Go to *System > Config > Messaging Servers*.
2. In *SMS Service*, select *Create New*.
3. Enter a *Name* for the SMS service and the service *Address* (domain name), then select *OK*.
4. Select *Apply*.

### To configure an SMS service - CLI

```
config system sms-server
 edit <provider_name>
 set mail-server <server_domain-name>
 next
end
```

### To configure SMS two-factor authentication - web-based manager

1. To modify an:
  - administrator account, go to *System > Admin > Administrators*, or
  - user account go to *User & Device > User > User Definition*.
2. Edit the user account.
3. Select SMS and either:
  - Select *FortiGuard Messaging Service*or
  - Select *Custom* and then choose the *SMS Provider* to use.
4. Enter the phone number of the mobile device that will receive the SMS text messages.
5. Select *Enable Two-factor Authentication*.
6. Select the *Token* that the user has.
7. Select *OK*.



## To enable SMS two-factor authentication - CLI

```
config user local
 edit <user_name>
 set sms-phone <user_phone>
 set sms-server fortiguard
 set two-factor sms
 end
```

If you have problems receiving the token codes via SMS messaging, contact your mobile provider to ensure you are using the correct phone number format to receive text messages and that your current mobile plan allows text messages.

## FortiToken

FortiToken is a disconnected one-time password (OTP) generator. It is a small physical device with a button that when pressed displays a six digit authentication code. This code is entered with a user's username and password as two-factor authentication. The code displayed changes every 60 seconds, and when not in use the LCD screen is blanked to extend the battery life.

There is also a mobile phone application, FortiToken Mobile, that performs much the same function.

FortiTokens have a small hole in one end. This is intended for a lanyard to be inserted so the device can be worn around the neck, or easily stored with other electronic devices. Do not put the FortiToken on a key ring as the metal ring and other metal objects can damage it. The FortiToken is an electronic device like a cell phone and must be treated with similar care.

Any time information about the FortiToken is transmitted, it is encrypted. When the FortiGate unit receives the code that matches the serial number for a particular FortiToken, it is delivered and stored encrypted. This is in keeping with the Fortinet's commitment to keeping your network highly secured.

FortiTokens can be added to user accounts that are local, IPsec VPN, SSL VPN, and even Administrators. See [“Associating FortiTokens with accounts” on page 500](#).

A FortiToken can be associated with only one account on one FortiGate unit.

If a user loses their FortiToken, it can be locked out using the FortiGate so it will not be used to falsely access the network. Later if found, that FortiToken can be unlocked on the FortiGate to allow access once again. See [“FortiToken maintenance” on page 501](#).

There are three tasks to complete before FortiTokens can be used to authenticate accounts:

1. [Adding FortiTokens to the FortiGate](#)
2. [Activating a FortiToken on the FortiGate](#)
3. [Associating FortiTokens with accounts](#)

## The FortiToken authentication process

The steps during FortiToken two-factor authentication are as follows.

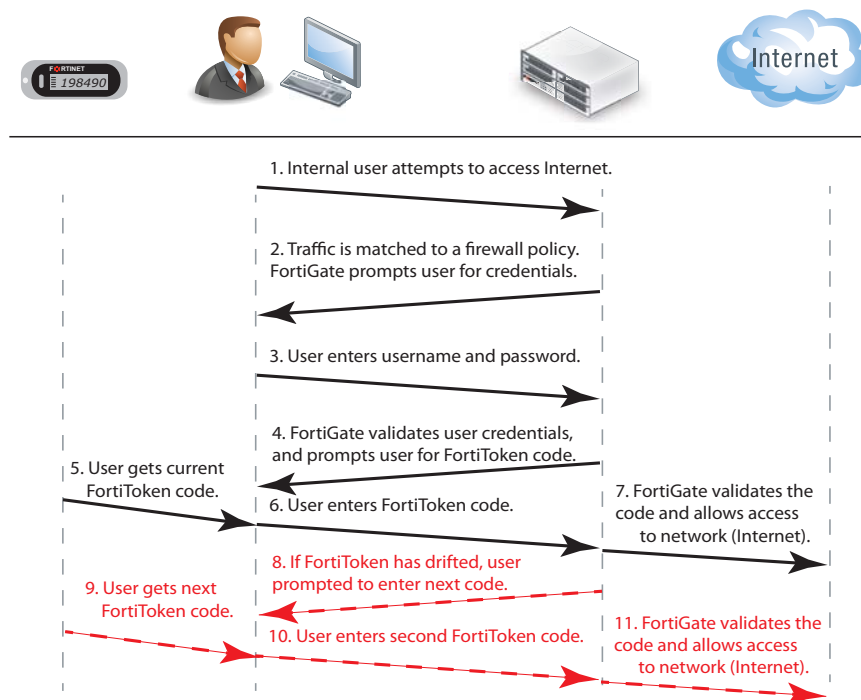
1. User attempts to access a network resource.
2. FortiGate unit matches the traffic to an authentication security policy, and FortiGate unit prompts the user for username and password.
3. User enters their username and password.
4. FortiGate unit verifies their information, and if valid prompts the user for the FortiToken code.
5. User gets the current code from their FortiToken device.

6. User enters current code at the prompt.
7. FortiGate unit verifies the FortiToken code, and if valid allows access to the network resources such as the Internet.

The following steps are only if the time on the FortiToken has drifted from the time on the FortiGate unit and needs to be synchronized.

8. If time on FortiToken has drifted, FortiGate unit will prompt user to enter a second code to confirm.
9. User gets the next code from their FortiToken device
10. User enters the second code at the prompt.
11. FortiGate unit uses both codes to update its clock to match the FortiToken and then proceeds as in step 7.

**Figure 114:**FortiToken authentication process



When configured the FortiGate unit accepts the username and password, authenticates them either locally or remotely, and prompts the user for the FortiToken code. The FortiGate then authenticates the FortiToken code. When FortiToken authentication is enabled, the prompt field for entering the FortiToken code is automatically added to the authentication screens.

Even when an Administrator is logging in through a serial or Telnet connection and their account is linked to a FortiToken, that Administrator will be prompted for the token's code at each login.



If you have attempted to add invalid FortiToken serial numbers, there will be no error message. The serial numbers will simply not be added to the list.

## Adding FortiTokens to the FortiGate

Before one or more FortiTokens can be used to authenticate logons, they must be added to the FortiGate. The import feature is used to enter many FortiToken serial numbers at one time. The serial number file must be a text file with one FortiToken serial number per line.

One FortiToken can be added to multiple FortiGate units. This is useful for maintaining two-factor authentication for employees over multiple office locations, such as for employees who travel frequently between offices.

### To manually add a FortiToken to the FortiGate - web-based manager

1. Go to *User & Device > Two-factor Authentication > FortiTokens*.
2. Select *Create New*.
3. In *Type*, select *Hard Token* or *Mobile Token*.
4. Enter one or more FortiToken serial numbers (hard token) or activation codes (mobile token).
5. Select *OK*.

### To import multiple FortiTokens to the FortiGate - web-based manager

1. Go to *User & Device > Two-factor Authentication > FortiTokens*.
2. Select *Create New*.
3. In *Type*, select *Hard Token*.
4. Select *Import*.
5. Select *Serial Number File* or *Seed File*, depending on which file you have.
6. Browse to the local file location on your local computer.
7. Select *OK*.  
The file is imported.
8. Select *OK*.

### To add two FortiTokens to the FortiGate - CLI

```
config user fortitoken
 edit <serial_number>
 next
 edit <serial_number2>
 next
end
```

## Activating a FortiToken on the FortiGate

Once one or more FortiTokens have been added to the FortiGate unit, they must be activated before being available to be associated with accounts. The process of activation involves the FortiGate querying FortiGuard servers about the validity of each FortiToken. The serial number and information is encrypted before it is sent for added security.



A FortiGate unit requires a connection to FortiGuard servers to activate a FortiToken.

### To activate a FortiToken on the FortiGate unit - web-based manager

1. Go to *User & Device > Two-factor Authentication > FortiTokens*.

2. Select one or more FortiTokens with a status of New.
3. Select *Activate*.
4. Refresh web browser. The status of selected FortiTokens will change to Activated.

The selected FortiTokens are now available for use with user and admin accounts.

#### To activate a FortiToken on the FortiGate unit - CLI

```
config user fortitoken
 edit <token_serial_num>
 set status activate
 next
end
```

### Associating FortiTokens with accounts

The final step before using the FortiTokens to authenticate logons is associating a FortiToken with an account. The accounts can be local user or administrator accounts.

#### To add a FortiToken to a local user account - web-based manager

1. Ensure that your FortiToken serial number has been added to the FortiGate successfully, and its status is Activated.
2. Go to *User & Device > User > User Definition*, and select *Create New*.
3. Enter the username and password for this user account.
4. Select *Enable Two-factor Authentication*.
5. Select FortiToken, and select the serial number from the list that matches that user's FortiToken.
6. Select OK.

#### To add a FortiToken to a local user account - CLI

```
config user local
 edit <username>
 set type password
 set passwd "myPassword"
 set two-factor fortitoken
 set fortitoken <serial_number>
 set status enable
 next
end
```

#### To add a FortiToken to an administrator account - web-based manager

1. Ensure that your FortiToken serial number has been added to the FortiGate successfully, and its status is Activated.
2. Go to *System > Admin > Administrators*, and select an admin account.  
This account is assumed to be configured except for two-factor authentication.
3. Select *Enable Two-factor Authentication*.
4. Select FortiToken, and select the serial number from the list that matches that user's FortiToken.
5. Select OK.

### To add a FortiToken to a local user account - CLI

```
config user local
 edit <username>
 set type password
 set passwd "myPassword"
 set two-factor fortitoken
 set fortitoken <serial_number>
 set status enable
 next
end
```

The `fortitoken` keyword will not be visible until `fortitoken` is selected for the `two-factor` keyword.



Before a new FortiToken can be used, it may need to be synchronized due to clock drift.

---

### FortiToken maintenance

Once FortiTokens are entered into the FortiGate unit, there are only two tasks to maintain them – changing the status,

To change the status of a FortiToken between Activated and Locked - CLI

```
config user fortitoken
 edit <token_serial_num>
 set status lock
 next
end
```

Any user attempting to login using this FortiToken will not be able to authenticate.

### To list the drift on all FortiTokens configured on this FortiGate unit - CLI

```
diag fortitoken drift
FORTITOKEN DRIFT
```

This command lists the serial number and drift for each FortiToken configured on this FortiGate unit. This command is useful to check if it is necessary to synchronize the FortiGate and any particular FortiTokens.

## IM users

Instant Messenger (IM) protocols are gaining in popularity as an essential way to communicate between two or more individuals in real time. Some companies even rely on IM protocols for critical business applications such as Customer/Technical Support.

The most common IM protocols in use today include AOL Instant Messenger (AIM), Yahoo Instant Messenger, MSN messenger, and ICQ. You configure each IM user as either allowed or denied use of IM applications. For each protocol, a global allow or deny policy configured in `config imp2p policy` governs unknown IM users. IM User settings are not available in the web-based manager.

## Monitoring users

To monitor user activity in the web-based manager, go to *User & Device > Monitor > Firewall*. The list of users who are logged on is displayed with some information about them such as their user group, security policy ID, how long they have been logged on, their IP address, traffic volume, and their authentication method as one of FSSO, NTLM, or firewall (FW-auth).

From this screen you can de-authenticate all users who are logged on. The de-authenticate button is at the top left of this screen.

To see information about banned users go to *User & Device > Monitor > Banned User*. Displayed information about users who have been banned includes what application the triggered the ban (Application Protocol), the reason for the ban (Cause or rule), Created, and when the ban expires.

### Filtering the list of users

When there are many users logged on, it can be difficult to locate a specific user or multiple users to analyze. Applying filters to the list allows you to organize the user list to meet your needs, or only display some the users that meet your current requirements.

Select *Column Settings* at the bottom of the screen to adjust columns that are displayed for users, including what order they are displayed in. This can be very helpful in locating information you are looking for.

The *username* column includes a green arrow to the right of the title. Select this arrow to sort the list of users by ordering them in ascending (down arrow) or descending order. This is the only column that allows this.

Each column heading has a grey filter icon. Click on the filter icon to configure a filter for the data displayed in that column. Each column has similar options including a field to enter the filtering information, a check box to select the negative of the text in the field, and the options to add more fields, apply the filter, clear all filters, or cancel without saving. To enter multiple terms in the field, separate each of them with a comma. To filter entries that contain a specific prefix, use an \* (asterisk).

For example, to create a filter to display only users with an IP address of 10.11.101.x who authenticated using one of security policies five through eight, and who belong to the user group *Accounting*.

1. Go to *User & Device > Monitor > Firewall*.
2. Select the filter icon beside *IP address*.
3. Enter 10.11.101.0. and select *Apply*.
4. Select the filter icon beside *Policy ID*.
5. Enter 5-8 and select *Apply*.
6. Select *Add new filter*.
7. Select the filter icon beside *User Group*.
8. Enter *Accounting* and select *Apply*.

## User groups

A user group is a list of user identities. An identity can be:

- a local user account (username/password) stored on the FortiGate unit
- a local user account with the password stored on a RADIUS, LDAP, or TACACS+ server
- a PKI user account with digital client authentication certificate stored on the FortiGate unit
- a RADIUS, LDAP, or TACACS+ server, optionally specifying particular user groups on that server
- a user group defined on an FSSO server.

Identity-based policies and some types of VPN configurations allow access to specified user groups only. This restricted access enforces Role Based Access Control (RBAC) to your organization's network and its resources. Users must be in a group and that group must be part of the security policy.



You cannot change the type of a group unless the group is empty.

---

In most cases, the FortiGate unit authenticates users by requesting their username and password. The FortiGate unit checks local user accounts first. If a match is not found, the FortiGate unit checks the RADIUS, LDAP, or TACACS+ servers that belong to the user group. Authentication succeeds when a matching username and password are found. If the user belongs to multiple groups on a server, those groups will be matched as well.



FortiOS does not allow username overlaps between RADIUS, LDAP, or TACACS+ servers.

---

There are two types of FortiGate user groups: Firewall user groups, and FSSO user groups.

### Firewall user groups

Firewall user groups are used locally as part of authentication and can contain any type of user identity except an FSSO group. When a user attempts to access resources controlled by an Identity-Based Policy (IBP), the FortiGate unit requires authentication from that user. If the user authenticates successfully and is a member of one of the permitted groups, the session is allowed to proceed.

This section includes:

- [SSL VPN access](#)
- [IPsec VPN access](#)
- [Configuring a firewall user group](#)
- [User group timeouts](#)
- [Viewing, editing and deleting user groups](#)

## SSL VPN access

In any firewall user group, you can enable SSL VPN access and select the web-portal that the users can access. When the user connects to the FortiGate unit via HTTPS on the SSL VPN port (default 10443), the FortiGate unit requests a username and password.

SSL VPN access also requires an SSL VPN security policy (*Action* is *SSL VPN*) with an identity-based rule enabling access for the user group. For more information, see the [FortiOS Handbook SSL VPN chapter](#).

## IPsec VPN access

A firewall user group can provide access for dialup users of an IPsec VPN. In this case, the IPsec VPN phase 1 configuration uses the *Accept peer ID in dialup group* peer option. The user's VPN client is configured with the username as peer ID and the password as pre-shared key. The user can connect successfully to the IPsec VPN only if the username is a member of the allowed user group and the password matches the one stored on the FortiGate unit.



A user group cannot be used as a dialup group if any member of the group is authenticated using an external authentication server.

---

For more information, see the [FortiOS Handbook IPsec VPN chapter](#).

## Configuring a firewall user group

A user group can contain:

- local users, whether authenticated by the FortiGate unit or an authentication server
- PKI users
- authentication servers, optionally specifying particular user groups on the server

### To create a Firewall user group - web-based manager

1. Go to *User & Device > User > User Groups* and select *Create New*.
2. Enter a name for the user group.
3. In *Type*, select *Firewall*.
4. From the *Available Users* list, select users and then select the right arrow button to move the names to the *Members* list.

If you select an authentication server as a group member, by default all user accounts on the authentication server are members of this FortiGate user group. Follow steps 5 through 8 if you want to include only specific user groups from the authentication server. Otherwise, select *OK*.

5. Select *Add*.
6. To add a remote authentication server, select *Add* and select the authentication server from the drop down *Remote Server* list.

The option to add remote servers is available only if at least one remote server has been configured.

7. In the *Group Name* field, either select *Any* to match all possible groups, or select *Specify* and enter the group name in the appropriate format for the type of server.

For example, an LDAP server requires LDAP format, such as:  
cn=users,dn=office,dn=example,dn=com

8. Repeat steps 5 through 7 to add all the authentication server user groups that are required.



9. Select *OK*.

### To create a firewall user group - CLI example

In this example, the members of `accounting_group` are `User1` and all of the members of `rad_accounting_group` on `myRADIUS` external RADIUS server.

```
config user group
 edit accounting_group
 set group-type firewall
 set member User1 myRADIUS
 config match
 edit 0
 set server-name myRADIUS
 set group-name rad_accounting_group
 end
 end
end
```



Matching user group names from an external authentication server might not work if the list of group memberships for the user is longer than 8000 bytes. Group names beyond this limit are ignored.

---

`server_name` is the name of the RADIUS, LDAP, or TACACS+ server, but it must be a member of this group first and must also be a configured remote server on the FortiGate unit.

`group_name` is the name of the group on the RADIUS, LDAP, or TACACS+ server such as “engineering” or “cn=users,dc=test,dc=com”.

Before using group matching with TACACS+, you must first enable authentication. For example if you have a configured TACACS+ server called `myTACS`, use the following CLI commands.

```
config user tacacs+
 edit myTACS
 set authorization enable
 next
end
```

For more information about user group CLI commands, see the [Fortinet CLI Guide](#).

### Multiple group enforcement support

Previously, when a user belonged to multiple user groups, this user could only access the group services that were within one group. With multiple group enforcement, a user can access the services within the groups that the user is part of.

For example, `userA` belongs to `user_group1`, `user_group2`, `user_group3`, and `user_group4`; previously `userA` could only access services within one of those four groups, typically the group that matches the first security policy. This can be annoying if HTTP access is in `user_group1`, FTP access is in `user_group2`, and email access is in `user_group3`. Now `userA` can access services within `user_group1`, `user_group2`, `user_group3`, and `user_group4`.

This feature is available only in the CLI and is enabled by default. It applies to RADIUS, LDAP, and TACACS+ servers. The new command for this feature is `auth-multi-group` found in `config user settings` and checks all groups a user belongs to for authentication.

## User group timeouts

User groups can have timeout values per group in addition to FortiGate-wide timeouts. There are essentially three different types of timeouts that are configurable for user authentication on the FortiGate unit — idle timeout, hard timeout, and session timeout. These are in addition to any external timeouts such as those associated with RADIUS servers.

If VDOMs are enabled, the global level user setting `authtimeout` is the default all VDOMs inherit. If VDOMs are not enabled, user settings `authtimeout` is the default. The default timeout value is used when the `authtimeout` keyword for a user group is set to zero.

Each type of timeout will be demonstrated using the existing user group `example_group`. Timeout units are minutes. A value of zero indicates the global timeout is used.

### Membership in multiple groups

When a user belongs to multiple groups in RADIUS groups, the group auth-timeout values are ignored. Instead the global timeout value is used. The default value is 5 minutes, but it can be set from 1 to 480 minutes.

```
config user setting
 set auth-timeout-type idle-timeout
 set auth-timeout 300
end
```

### Idle timeout

The default type of timeout is idle timeout. When a user initiates a session, it starts a timer. As long as data is transferred in this session, the timer continually resets. If data flow stops, the timer is allowed to advance until it reaches its limit. At that time the user has been idle for too long, and the user is forced to re-authenticate before traffic is allowed to continue in that session.

### To configure user group authentication idle timeout - CLI

```
config user settings
 set auth-timeout-type idle-timeout
end
config user group
 edit example_group
 set auth-timeout 480
 next
end
```

### Hard timeout

Where the idle timeout is reset with traffic, the hard timeout is absolute. From the time the first session a user establishes starts, the hard timeout counter starts. When the timeout is reached, all the sessions for that user must be re-authenticated. This timeout is not affected by any event.

### To configure user group authentication hard timeout - CLI

```
config user settings
 set auth-timeout-type hard-timeout
end
config user group
 edit example_group
 set auth-timeout 480
 next
end
```

### Session timeout

The session timeout works much like the hard timeout in that its an absolute timer that can not be affected by events. However, when the timeout is reached existing sessions may continue but new sessions are not allowed until re-authentication takes place. The timeout can be set from 1 to 480 minutes. Setting the timeout value to zero removes the timeout value allowing the user to remain logged on without limit.

### To configure a user group authentication new session hard timeout - CLI

```
config user setting
 set auth-timeout-type new-session
end

config user group
 edit example_group
 set authtimeout 30
 next
end
```

## SSO user groups

SSO user groups are part of FSSO authentication and contain only Windows or Novell network users. No other user types are permitted as members. Information about the Windows or Novell user groups and the logon activities of their members is provided by the Fortinet Single Sign On (FSSO) which is installed on the network domain controllers.

You can specify FSSO user groups in identity-based security policies in the same way as you specify firewall user groups. FSSO user groups cannot have SSL VPN or dialup IPsec VPN access.

For information about configuring FSSO user groups, see [“Creating Fortinet Single Sign-On \(FSSO\) user groups” on page 591](#). For complete information about installing and configuring FSSO, see [“Agent-based FSSO” on page 563](#).

## Configuring Peer user groups

Peer user groups can only be configured using the CLI. Peers are digital certificate holders defined using the `config user peer` command. The peer groups you define here are used in dialup IPsec VPN configurations that accept RSA certificate authentication from members of a peer certificate group. For more information, see [“Authenticating IPsec VPN users with security certificates” on page 545](#).

### To create a peer group - CLI example

```
config user peergrp
 edit vpn_peergrp1
 set member pki_user1 pki_user2 pki_user3
 end
```

## Viewing, editing and deleting user groups

To view the list of FortiGate user groups, go to *User & Device > User > User Groups*.

### Editing a user group

When editing a user group in the CLI you must set the type of group this will be — either a firewall group, or a Fortinet Single Sign-On Service group. Once the type of group is set, and members are added you cannot change the group type without removing the members.

In the web-based manager, if you change the type of the group any members will be removed automatically.

### To edit a user group - web-based manager

1. Go to *User & Device > User > User Groups*.
2. Select the check box for the user group that you want to edit.
3. Select the *Edit* button.
4. Modify the user group as needed.
5. Select *OK*.

### To edit a user group - CLI example

This example adds user3 to Group1. Note that you must re-specify the full list of users:

```
config user group
 edit Group1
 set group-type firewall
 set member user2 user4 user3
 end
```

### Deleting a user group

Before you delete a user group, you must ensure there are no objects referring to, it such as security policies. If there are, you must remove those references before you are able to delete the user group.

### To remove a user group - web-based manager

1. Go to *User & Device > User > User Groups*.
2. Select the check box for the user group that you want to remove.
3. Select the *Delete* button.
4. Select *OK*.

### To remove a user group - CLI example

```
config user group
 delete Group2
end
```

# Managing Guest Access

## Introduction

Visitors to your premises might need user accounts on your network for the duration of their stay. If you are hosting a large event such as a conference, you might need to create many such temporary accounts. The FortiOS Guest Management feature is designed for this purpose.

A guest user account User ID can be the user's email address, a randomly generated string, or an ID that the administrator assigns. Similarly, the password can be administrator-assigned or randomly generated.

You can create many guest accounts at once using randomly-generated User IDs and passwords. This reduces administrator workload for large events.

## User's view of guest access

1. The user receives an email, SMS message, or printout from a FortiOS administrator listing a User ID and password.
2. The user logs onto the network with the provided credentials.
3. After the expiry time, the credentials are no longer valid.

## Administrator's view of guest access

1. Create one or more guest user groups.  
All members of the group have the same characteristics: type of User ID, type of password, information fields used, type and time of expiry.
2. Create guest accounts using Guest Management.
3. Use captive portal authentication and select the appropriate guest group.

## Configuring guest user access

### Creating guest management administrators

The guest management administrator can be a regular FortiGate administrator. Optionally, you can create administrator accounts that can perform only guest management. This type of administrator is also limited to specific guest user groups.

#### To create a guest management administrator

1. Go to *System > Admin > Administrators* and create a regular administrator account.  
For detailed information see the System Administration chapter.
2. Select *Restrict to Provision Guest Accounts*.
3. In *Guest Groups*, add the guest groups that this administrator manages.

## Creating guest user groups

The guest group configuration determines the fields that are provided when you create a guest user account.

### To create a guest user group

1. Go to *User & Device > User > User Groups* and select *Create New*.
2. Enter

<b>Name</b>	Enter a name for the group.
<b>Type</b>	Guest
<b>Enable Batch Account Creation</b>	<p>Create multiple accounts automatically. When this is enabled:</p> <ul style="list-style-type: none"><li>• <i>User ID</i> and <i>Password</i> are set to <i>Auto-Generate</i>.</li><li>• The user accounts have only <i>User ID</i>, <i>Password</i>, and <i>Expiration</i> fields. Only the <i>Expiration</i> field is editable. If the expiry time is a duration, such as “8 hours”, this is the time after first login.</li><li>• You can print the account information. Users do not receive email or SMS notification.</li></ul> <p>See <a href="#">“To create multiple guest user accounts automatically” on page 511</a>.</p>
<b>User ID</b>	<p>Select one of:</p> <ul style="list-style-type: none"><li>• Email — User’s email address</li><li>• Specify — Administrator assigns user ID</li><li>• Auto-Generate — FortiGate unit creates a random user ID</li></ul>
<b>Password</b>	<p>Select one of:</p> <ul style="list-style-type: none"><li>• Specify — Administrator assigns user ID</li><li>• Auto-Generate — FortiGate unit creates a random password</li><li>• Disable — no password</li></ul>
<b>Expire Type</b>	<p>Choose one of:</p> <p>Immediately — expiry time is counted from creation of account</p> <p>After first login — expiry time is counted from user’s first login</p>
<b>Default Expire Time</b>	Set the expire time. The administrator can change this for individual users.
<b>Enable Name</b>	If enabled, user must provide a name.
<b>Enable Sponsor</b>	If enabled, user form has Sponsor field. Select <i>Required</i> or <i>Optional</i> .
<b>Enable Company</b>	If enabled, user form has Company field. Select <i>Required</i> or <i>Optional</i> .
<b>Enable Email</b>	If enabled, user is notified by email.
<b>Enable SMS</b>	If enabled, user is notified by SMS. Select whether FortiGuard Messaging Service or a another SMS provider is used. You can add SMS providers in <i>System &gt; Config &gt; Messaging Servers</i> .

## Creating guest user accounts

Guest user accounts are not the same as local user accounts created in *User & Device > User > User Definition*. Guest accounts are not permanent; they expire after a defined time period. You create guest accounts in *User & Device > User > Guest Management*.

### To create a guest user account

1. Go to *User & Device > User > Guest Management*.
2. In *Guest Groups*, select the guest group to manage.
3. Select *Create New* and fill in the fields in the *New User* form.  
Fields marked *Optional* can be left blank. The guest group configuration determines the fields that are available.
4. Select *OK*.

### To create multiple guest user accounts automatically

1. Go to *User & Device > User > Guest Management*.
2. In *Guest Groups*, select the guest group to manage.  
The guest group must have the *Enable Batch Guest Account Creation* option enabled.
3. Select *Create New > Multiple Users*.  
Use the down-pointing caret to the right of *Create New*.
4. Enter *Number of Accounts*.
5. Optionally, change the *Expiration*.
6. Select *OK*.

## Guest Management Account List

Go to *User & Device > User > Guest Management* to create, view, edit or delete guest user accounts.

<b>Create New</b>	Creates a new guest user account.
<b>Edit</b>	Edit the selected guest user account.
<b>Delete</b>	Delete the selected guest user account.
<b>Purge</b>	Remove all accounts from the list.
<b>Send</b>	Send the user account information to a printer or to the guest. Depending on the group settings and user information, the information can be sent to the user by email or SMS.
<b>Refresh</b>	Update the list.
<b>Guest Groups</b>	Select the guest group to list. New accounts are added to this group.
<b>User ID</b>	The user ID. Depending on the guest group settings, this can be the user's email address, an ID that the administrator specified, or a randomly-generated ID.
<b>Expires</b>	Indicates a duration such as "3 hours". A duration on its own is relative to the present time. Or, the duration is listed as "after first login."

## Guest access in a retail environment

Some retail businesses such as coffee shops provide free WiFi Internet access for their customers. For this type of application, the FortiOS guest management feature is not required; the WiFi access point is open and customers do not need logon credentials. However, the business might want to contact its customers later with promotional offers to encourage further patronage. Using FortiOS device-based security policies, it is possible to collect customer email addresses for this purpose.

The first time a customer's device attempts to use the WiFi connection, FortiOS requests an email address, which it validates. The customer's subsequent connections go directly to the Internet without interruption.

### Implementing email harvesting

The customer's first contact with your network will be with a captive portal which presents a web page requesting an email address. When FortiOS has validated the email address, the customer's device MAC address is added to the Collected Emails device group.

You need configure a device policy that allows traffic to flow from the WiFi SSID to the Internet interface. Within that policy, you need an authentication rule to allow members of the Collected Emails device group to access the Internet. This rule must be listed first. Unknown devices are not members of the Collected Emails device group, so they don't match the rule.

You also need to select *Prompt E-mail collection Portal for all devices*.

#### To create the device policy

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Enter the following information:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Device Identity
<b>Incoming Interface</b>	wifi
<b>Source Address</b>	all
<b>Outgoing Interface</b>	wan1
<b>Enable NAT</b>	Enable.

You are now ready to create the authentication rule.

#### To create the authentication rule - web-based manager

1. In *Configure Authentication Rules*, select *Create New* and enter:

<b>Destination Address</b>	all
<b>Device</b>	Collected Emails
<b>Compliant with Endpoint Profile</b>	not selected
<b>Schedule</b>	always
<b>Service</b>	ALL



---

<b>Action</b>	ACCEPT
---------------	--------

---

2. Select *OK*.
3. If asked, confirm that you accept FortiOS will enable device identification on the source interface.
4. In *Device Policy Options*, select *Prompt E-mail Collection Portal for all devices*.
5. Optionally, customize authentication messages.
6. Select *OK* to complete configuration of the security policy.

#### To create the authentication rule - CLI

```
config firewall policy
 edit 3
 set srcintf "wifi"
 set dstintf "wan1"
 set srcaddr "all"
 set action accept
 set email-collection-portal enable
 set identity-based enable
 set identity-from device
 set nat enable
 config identity-based-policy
 edit 1
 set schedule "always"
 set dstaddr "abc"
 set service "ALL"
 set devices "collected-emails"
 end
 end
end
```

### Checking for harvested emails

In the web-based manager, go to *User & device > Device > Device Definitions*. In the CLI you can use the `diagnose user device list` command. For example,

```
FGT-100D # diagnose user device list
hosts
vd 0 d8:d1:cb:ab:61:0f gen 35 req 30 redir 1 last 43634s
7-11_2-int
ip 10.0.2.101 ip6 fe80::dad1:cbff:feab:610f
type 2 'iPhone' src http c 1 gen 29
os 'iPhone' version 'iOS 6.0.1' src http id 358 c 1
email 'yo@yourdomain.com'
vd 0 74:e1:b6:dd:69:f9 gen 36 req 20 redir 0 last 39369s
7-11_2-int
ip 10.0.2.100 ip6 fe80::76e1:b6ff:fedd:69f9
type 1 'iPad' src http c 1 gen 5
os 'iPad' version 'iOS 6.0' src http id 293 c 1
host 'Joes's-iPad' src dhcp
email 'you@fortinet.com'
```

# Configuring authenticated access

When you have configured authentication servers, users, and user groups, you are ready to configure security policies and certain types of VPNs to require user authentication.

This section describes:

- [Authentication timeout](#)
- [Password policy](#)
- [Authentication protocols](#)
- [Authentication in security policies](#)
- [Limited access for unauthenticated users](#)
- [VPN authentication](#)

## Authentication timeout

An important feature of the security provided by authentication is that it is temporary—a user must re-authenticate after logging out. Also if a user is logged on and authenticated for an extended period of time, it is a good policy to have them re-authenticate at set periods. This ensures a user's session is cannot be spoofed and used maliciously for extended periods of time — re-authentication will cut any spoof attempts short. Shorter timeout values are more secure.

### Security authentication timeout

You set the security user authentication timeout to control how long an authenticated connection can be idle before the user must authenticate again. The maximum timeout is 480 minutes (8 hours).

#### To set the security authentication timeout - web-based manager

1. Go to *User & Device > Authentication > Settings*.
2. Enter the *Authentication Timeout* value in minutes.  
The default authentication timeout is 5 minutes.
3. Select *Apply*.

### SSL VPN authentication timeout

You set the SSL VPN user authentication timeout (*Idle Timeout*) to control how long an authenticated connection can be idle before the user must authenticate again. The maximum timeout is 28 800 seconds. The default timeout is 300 seconds.

#### To set the SSL VPN authentication timeout - web-based manager

1. Go to *VPN > SSL > Config*.
2. Enter the *Idle Timeout* value (seconds).
3. Select *Apply*.

## Password policy

Password authentication is effective only if the password is sufficiently strong and is changed periodically. By default, the FortiGate unit requires only that passwords be at least eight characters in length. You can set a password policy to enforce higher standards for both length and complexity of passwords. Password policies can apply to administrator passwords or IPsec VPN preshared keys.

To set a password policy in the web-based manager, go to *System > Admin > Settings*. In the CLI, use the `config system password-policy` command.

The default minimum password length on the FortiGate unit is eight characters, but up to 32 characters is permitted. Fortinet suggests a minimum length of 14 characters.

Users usually create passwords composed of alphabetic characters and perhaps some numbers. Password policy can require the inclusion of uppercase letters, lowercase letters, numerals or punctuation characters.

### Configuring password minimum requirement policy

Best practices dictate that passwords include:

- one or more uppercase characters
- one or more lower care characters
- one or more of the numerals
- one or more non alphanumeric characters, such as punctuation marks.

The minimum number of each of these types of characters can be set in both the web-based manager and the CLI.

The following procedures show how to force administrator passwords to contain at least two uppercase, four lower care, two digits, and one non-alphanumeric characters. Leave the minimum length at the default of eight characters.

#### To change administrator password minimum requirements - web-based manager

1. Go to *System > Admin > Settings*.
2. Select *Enable Password Policy*.
3. Select *Must Contain*.
4. Enter the following information:

<b>uppercase Letters</b>	2
<b>lower case Letters</b>	4
<b>Numerical Digits</b>	2
<b>Non-alphanumeric Letters</b>	1

5. Under *Apply Password Policy to*, select *Admin Password*.
6. Select *Apply*.

## To change administrator password minimum requirements - CLI

```
config system password-policy
 set status enable
 set apply-to admin-password
 set min-upper-case-letter 2
 set min-lower-case-letter 4
 set min-number 2
 set min-non-alphanumeric 1
 set change-4-characters enable
end
```

The `change-4-characters` option forces new passwords to change a minimum of four characters in the old password. Changing fewer characters results in the new password being rejected. This option is only available in the CLI.

## Password best practices

In addition to length and complexity, there are security factors that cannot be enforced in a policy. Guidelines issued to users will encourage proper password habits.

Best practices dictate that password expiration also be enabled. This forces passwords to be changed on a regular basis. You can set the interval in days. The more sensitive the information this account has access to, the shorter the password expiration interval should be. For example 180 days for guest accounts, 90 days for users, and 60 days for administrators.

Avoid:

- real words found in any language dictionary
- numeric sequences, such as “12345”
- sequences of adjacent keyboard characters, such as “qwerty”
- adding numbers on the end of a word, such as “hello39”
- adding characters to the end of the old password, such as “hello39” to “hello3900”
- repeated characters
- personal information, such as your name, birthday, or telephone number.

## Maximum logon attempts and blackout period

When you logon and fail to enter the correct password you could be a valid user, or a hacker attempting to gain access. For this reason, best practices dictate to limit the number of failed attempts to logon before a blackout period where you cannot logon.

To set a maximum of five failed authentication attempts before the blackout, using the following CLI command:

```
config user setting
 set auth-invalid-max 5
end
```

To set the length of the blackout period to five minutes, or 300 seconds, once the maximum number of failed logon attempts has been reached, use the following CLI command:

```
config user setting
 set auth-blackout-time 300
end
```

## Authentication protocols

When user authentication is enabled on a security policy, the authentication challenge is normally issued for any of the four protocols, HTTP, HTTPS, FTP, and Telnet, which are dependent on the connection protocol. By making selections in the Protocol Support list, the user controls which protocols support the authentication challenge. The user must connect with a supported protocol first, so that they can subsequently connect with other protocols.

For example, if you have selected HTTP, FTP, or Telnet, a username and password-based authentication occurs. The FortiGate unit then prompts network users to input their security username and password. If you have selected HTTPS, certificate-based authentication (HTTPS, or HTTP redirected to HTTPS only) occurs.



FTP and Telnet authentication replacement messages cannot be customized. For HTTP and HTTPS replacement messages see [“Authentication replacement messages” on page 518](#).

---

For certificate-based authentication, you must install customized certificates on the FortiGate unit and on the browsers of network users. If you do not install certificates on the network user’s web browser, the network users may see an SSL certificate warning message and have to manually accept the default FortiGate certificate. The network user’s web browser may deem the default certificate as invalid.

When you use certificate authentication, if you do not specify any certificate when you create the security policy, the global settings are used. If you specify a certificate, the per-policy setting will overwrite the global setting. For more information about the use of certification authentication see [“Certificate-based authentication” on page 532](#).

### To set the authentication protocols

1. Go to *User & Device > Authentication > Settings*.
2. In *Protocol Support*, select the required authentication protocols.
3. If using HTTPS protocol support, in *Certificate*, select a Local certificate from the drop-down list.
4. Select *Apply*.

## Authentication in security policies

Security policies control traffic between FortiGate interfaces, both physical interfaces and VLAN subinterfaces. Without authentication, a security policy enables access from one network to another for all users on the source network. Authentication enables you to allow access only for users who are members of selected user groups. To include authentication in a security policy, you must create an identity-based policy.



You can configure user authentication for security policies only when *Action* is set to *Accept*. If the policy is set to *Deny*, *IPsec*, or *SSL VPN* the options will be different.

---

The style of the authentication method varies by the authentication protocol. If you have selected HTTP, FTP or Telnet, a username and password-based authentication occurs. The FortiGate unit prompts network users to input their security username and password. If you

have selected HTTPS, certificate-based authentication (HTTPS or HTTP redirected to HTTPS only) occurs. You must install customized certificates on the FortiGate unit and on the browsers of network users, which the FortiGate unit matches.

This section includes:

- [Enabling authentication protocols](#)
- [Authentication replacement messages](#)
- [Access to the Internet](#)
- [Configuring authentication security policies](#)
- [Identity-based policy](#)
- [NTLM authentication](#)
- [Certificate authentication](#)
- [Restricting number of concurrent user logons](#)

## Enabling authentication protocols

Users can authenticate using FTP, HTTP, HTTPS, and Telnet. However, these protocols must be enabled first.

Another authentication option is to redirect any attempts to authenticate using HTTP to a more secure channel that uses HTTPS. This forces users to a more secure connection before entering their user credentials.

### To enable support for authentication protocols - web-based manager

1. Go to *User & Device > Authentication > Settings*.
2. Select one or more of HTTP, HTTPS, FTP, Telnet, or Redirect HTTP Challenge to a Secure Channel (HTTPS). Only selected protocols will be available for use in authentication.
3. Select the *Certificate* to use, for example `Fortinet_Factory`.
4. Select *Apply*.

### To enable support for authentication protocols - CLI

```
config user setting
 set auth-type ftp http https telnet
 set auth-cert Fortinet_Factory
end
```

## Authentication replacement messages

A replacement message is the body of a webpage containing a message about a blocked website message, a file too large message, a disclaimer, or even a login page for authenticating. The user is presented with this message instead of the blocked content.

Authentication replacement messages are the prompts a user sees during the security authentication process such as login page, disclaimer page, and login success or failure pages. These are different from most replacement messages because they are interactive requiring a user to enter information, instead of simply informing the user of some event as other replacement messages do.

Replacement messages have a system-wide default configuration, a per-VDOM configuration, and disclaimers can be customized for multiple security policies within a VDOM.

These replacement messages are used for authentication using HTTP and HTTPS. Authentication replacement messages are HTML messages. You cannot customize the security authentication messages for FTP and Telnet.

The authentication login page and the authentication disclaimer include replacement tags and controls not found on other replacement messages.

More information about replacement messages can be found in the `config system replacemsg` section of the [FortiOS CLI Reference](#).

**Table 21:** List of authentication replacement messages

Replacement message name (CLI name)	Description
<b>Login challenge page (auth-challenge-page)</b>	<p>This HTML page is displayed if security users are required to answer a question to complete authentication. The page displays the question and includes a field in which to type the answer. This feature is supported by RADIUS and uses the generic RADIUS challenge-access auth response. Usually, challenge-access responses contain a Reply-Message attribute that contains a message for the user (for example, "Please enter new PIN"). This message is displayed on the login challenge page. The user enters a response that is sent back to the RADIUS server to be verified.</p> <p>The Login challenge page is most often used with RSA RADIUS server for RSA SecurID authentication. The login challenge appears when the server needs the user to enter a new PIN. You can customize the replacement message to ask the user for a SecurID PIN.</p> <p>This page uses the <code>%%QUESTION%%</code> tag.</p>
<b>Disclaimer page (auth-disclaimer-page-1) (auth-disclaimer-page-2) (auth-disclaimer-page-3)</b>	<p>Prompts user to accept the displayed disclaimer when leaving protected network.</p> <p>The web-based manager refers to this as User Authentication Disclaimer, and it is enabled with a security policy that also includes at least one identity-based policy. When a security user attempts to browse a network through the FortiGate unit using HTTP or HTTPS this disclaimer page is displayed.</p> <p>The extra pages seamlessly extend the size of the page from 8 192 characters to 16 384 and 24 576 characters respectively. When configuring the disclaimer page in the web-based manager this is shown by its size being 24 576 characters.</p> <p>See <a href="#">"Disclaimer" on page 521</a>.</p>
<b>Email token page (auth-email-token-page)</b>	<p>The page prompting a user to enter their email token. See <a href="#">"Email" on page 495</a>.</p>
<b>FortiToken page (auth-fortitoken-page)</b>	<p>The page prompting a user to enter their FortiToken code. See <a href="#">"FortiToken" on page 497</a>.</p>

**Table 21:** List of authentication replacement messages

Replacement message name (CLI name)	Description
<b>Keepalive page (auth-keepalive-page)</b>	<p>The HTML page displayed with security authentication keepalive is enabled using the following CLI command:</p> <pre>config system global   set auth-keepalive enable end</pre> <p>Authentication keepalive keeps authenticated firewall sessions from ending when the authentication timeout ends. In the web-based manager, go to <i>User &amp; Device &gt; Authentication &gt; Settings</i> to set the <i>Authentication Timeout</i>.</p> <p>This page includes %%TIMEOUT%%.</p>
<b>Login failed page (auth-login-failed-page)</b>	<p>The Disclaimer page replacement message does not re-redirect the user to a redirect URL or the security policy does not include a redirect URL. When a user selects the button on the disclaimer page to decline access through the FortiGate unit, the Declined disclaimer page is displayed.</p>
<b>Login page (auth-login-page)</b>	<p>The authentication HTML page displayed when users who are required to authenticate connect through the FortiGate unit using HTTP or HTTPS.</p> <p>Prompts the user for their username and password to login.</p> <p>This page includes %%USERNAMEID%% and %%PASSWORDID%% tags.</p>
<b>Declined disclaimer page (auth-reject-page)</b>	<p>The page displayed if a user declines the disclaimer page. See <a href="#">“Disclaimer” on page 521</a>.</p>
<b>SMS Token page (auth-sms-token-page)</b>	<p>The page prompting a user to enter their SMS token. See <a href="#">“SMS” on page 496</a>.</p>
<b>Success message (auth-success-msg)</b>	<p>The page displayed when a user successfully authenticates. Prompts user to attempt their connection again (as the first was interrupted for authentication).</p>

## Access to the Internet

A policy for accessing the Internet is similar to a policy for accessing a specific network, but the destination address is set to *all*. The destination interface is the one that connects to the Internet Service Provider (ISP). For general purpose Internet access, the Service is set to ANY.

Access to HTTP, HTTPS, FTP and Telnet sites may require access to a domain name service. DNS requests do not trigger authentication. You must configure a policy to permit unauthenticated access to the appropriate DNS server, and this policy must **precede** the policy for Internet access. Failure to do this will result in the lack of a DNS connection and a corresponding lack of access to the Internet.



## Configuring authentication security policies

To include authentication in a security policy, you must create an identity-based policy. An identity-based policy can authenticate by certificate, FSSO, and NTLM. The two exceptions to this are RADIUS SSO and FSSO Agents. See [“SSO using RADIUS accounting records” on page 602](#), and [“Introduction to FSSO agents” on page 564](#).

Before creating an identity-based security policy, you need to configure one or more users and firewall user groups. For more information, see [“Users and user groups” on page 488](#).

Creating the security policy is the same as a regular security policy except you must select the action specific to your authentication method:

**Table 22:** Authentication methods allowed for each policy Action

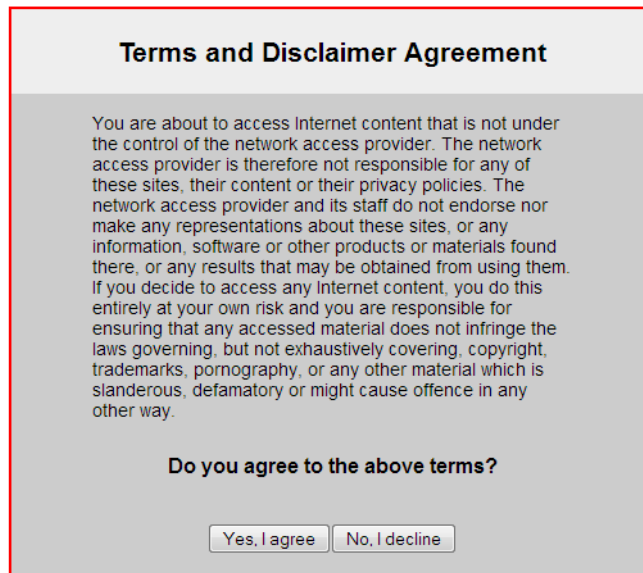
Action	Authentication method	Where authentication is used
<b>ACCEPT</b>	FSSO Agent or identity-based policy – FSSO	See <a href="#">“Agent-based FSSO” on page 563</a> .
	identity-based policy – NTLM	See <a href="#">“NTLM authentication” on page 524</a> .
	identity-based policy – Certificates	See <a href="#">“Configuring certificate-based authentication” on page 543</a> .
	RADIUS SSO	See <a href="#">“SSO using RADIUS accounting records” on page 602</a> .
<b>IPSEC</b>	IPsec Phase 1 and 2	See <a href="#">“Configuring authentication of remote IPsec VPN users” on page 528</a> .
<b>SSL-VPN</b>	SSL certificates	See <a href="#">“Configuring authentication of SSL VPN users” on page 528</a> .
<b>DENY</b>	none	none

### Disclaimer

When configuring a User Identity authentication security policy, there is an option to enable a disclaimer. The disclaimer is a replacement message that when enabled, web traffic matching this policy will be presented with the disclaimer that the user must choose to agree or decline.

The default disclaimer contains a warning that any content the user is about to access is the responsibility of the user and not the company or owner of the network. It is presented in [Figure 115](#). You can customize the text and the appearance as required.

**Figure 115:**Default disclaimer message



**Terms and Disclaimer Agreement**

You are about to access Internet content that is not under the control of the network access provider. The network access provider is therefore not responsible for any of these sites, their content or their privacy policies. The network access provider and its staff do not endorse nor make any representations about these sites, or any information, software or other products or materials found there, or any results that may be obtained from using them. If you decide to access any Internet content, you do this entirely at your own risk and you are responsible for ensuring that any accessed material does not infringe the laws governing, but not exhaustively covering, copyright, trademarks, pornography, or any other material which is slanderous, defamatory or might cause offence in any other way.

**Do you agree to the above terms?**

## Customizing authentication replacement messages

Customizing disclaimers or other authentication replacement messages involves changing the text of the disclaimer message, and possibly the overall appearance of the message.

Disclaimers are useful in many situations. Often companies find it useful to brand the disclaimers with their specific company policy text, logo, and design. One example of this is at an Internet cafe where customers pay for usage and must accept terms of usage before accessing the internet. The cafe benefits from a customized disclaimer that alerts the customer to their online policies. The same is true for other authentication replacement messages such as the login page.

Changing the disclaimer at *System > Config > Replacement messages* is not the same as selecting to customize a disclaimer used in a policy. The *System > Config* location is the default message that all disclaimers inherit. The security policy location is a customized disclaimer that inherits the default format for the disclaimer message, but then can be customized for this policy.

### To customize the disclaimer for a security policy - web-based manager

1. Go to *Policy > Policy > Policy*. Either select an existing User Identity policy or create a new one.
2. Enable *Disclaimer*, and select *Customize Authentication Messages*.
3. Select the Edit icon. You can select and edit any of the pages. Change your text or layout as needed.

## Enabling security logging

There are two types of logging that relate to authentication — event logging, and security logging.

When enabled, event logging records system events such as configuration changes, and authentication. To configure event logging, go to *Log&Report > Log Config > Log Settings* and enable *Event Logging*. Select the events you want to log, such as *User activity event*.

When enabled, security logging will log UTM and security policy traffic.

You must enable logging within a security policy, as well as the options that are applied to a security policy, such as UTM features. Event logs are enabled within the Event Log page,

For more information on logging, see the [FortiOS Log and Reporting chapter](#).

For more information on specific types of log messages, see the [FortiOS Log Message Reference](#).



You need to set the logging severity level to *Notification* when configuring a logging location to record traffic log messages.

---

### To enable logging within an existing security policy - web-based manager

1. Go to *Policy > Policy*.
2. Expand to reveal the policy list of a policy.
3. Select the security policy you want to enable logging on and then select *Edit*.
4. To log all general firewall traffic, select the check box beside *Log Allowed Traffic*.
5. On the security policy's page, select the check box beside *UTM*.
6. In *UTM Security Profiles*, select enable the UTM profiles that you want applied to the policy, then select the profile or sensor from the drop-down list as well.
7. Select *OK*.

## Identity-based policy

An identity-based policy (IBP) performs user authentication in addition to the normal security policy duties. If the user does not authenticate, access to network resources is refused. This enforces Role Based Access Control (RBAC) to your organization's network and resources.

Identity-based policies also support Single Sign-On operation. The user groups selected in the policy are of the Fortinet Single Sign-On (FSSO) type.

User authentication can occur through any of the following supported protocols, including: HTTP, HTTPS, FTP, and Telnet. The authentication style depends on which of these protocols is included in the selected security services group and which of those enabled protocols the network user applies to trigger the authentication challenge.

For username and password-based authentication (HTTP, FTP, and Telnet) the FortiGate unit prompts network users to enter their username, password, and token code if two-factor authentication is selected for that user account. See [“Two-factor authentication” on page 494](#). For certificate-based authentication, including HTTPS or HTTP redirected to HTTPS only, see [“Certificate authentication” on page 525](#).



FortiManager does not support pushing identity based policies down to FortiGate units.

Set these commands in the CLI to see the other identity-based commands that were hidden before. In the following procedure, this is policy number 7.

```
config firewall policy
 edit 7
 set action ACCEPT
 set identity-based enable
 next
end
```

With identity-based policies, once the FortiGate unit matches the source and destination addresses, it processes the identity sub-policies for the user groups and services. This means unique security policies must be placed **before** an identity-based policy to be effective.

When the identity-based policy has been configured, the option to customize authentication messages is available. This allows you to change the text, style, layout, and graphics of the replacement messages associated with this firewall policy. When enabled, customizing these messages follows the same method as changing the disclaimer. See [“Disclaimer” on page 521](#).

Types of authentication also available in identity-based policies are

- [NTLM authentication](#)
- [Certificate authentication](#)

### Identity-based sub-policies

Once IBP is enabled in a policy, a table appears. Selecting *Add* allows you to configure authentication rules which are added to this table as sub-policies.

Just as with regular security policies, with these identity-based sub-policies traffic is matched from the top of the list of sub-policies down until the criteria is met. If there is no matching policy packets are dropped, even if they have been authenticated. Each sub-policy has its own UTM profile fields, traffic shaping, logging, and so on that take effect when the User Group, Service and Schedule are matched.

The order of these sub-policies is just as important as with regular security policies. For example if a user is a member of two groups, and each group has a separate sub-policy entry, the top one in the list will be matched first.

## NTLM authentication

The NT LAN Manager (NTLM) protocol is used when the MS Windows Active Directory (AD) domain controller can not be contacted. NTLM uses web browsers to send and receive authentication information. See [“NTLM” on page 460](#) and [“FSSO NTLM authentication support” on page 569](#).

NTLM authentication is enabled when you configure FSSO and enable NTLM in the identity-based policy (IBP). There must be at least one FSSO Collector agent configured on the FortiGate. Any users and user groups associated with the security policy will use NTLM to authenticate without further configuration. However some extra configuration in the CLI may be required for certain cases including guest access, and defining NTLM enabled browsers.



If there are multiple domains, a trust relation must exist between them. This is automatic if they are in a forest. With the trust relation, only one FSSO DC agent needs to be installed. Without the trust relation, FSSO DC agents must be installed on each domain controller.

---

## NTLM guest access - CLI

Guest profile access may be granted to users failing NTLM authentication, such as visitors who have no user credentials on the network. To allow guest users in NTLM, use the following CLI command:

```
config firewall policy
 edit 8
 set action accept
 set identity-based enable
 set ntlm enable
 set ntlm-guest enable
 next
end
```

## NTLM enabled browsers - CLI

User agent strings for NTLM enabled browsers allow the inspection of initial HTTP-User-Agent values, so that non-supported browsers are able to go straight to guest access without needlessly prompting the user for credentials that will fail. `ntlm-guest` must be enabled to use this option.

```
config firewall policy
 edit 9
 set action accept
 set identity-based enable
 set ntlm enable
 set ntlm-guest enable
 set ntlm-enabled-browsers <user_agent_string>
 next
end
```

<user\_agent\_string> is the name of the browser that is NTLM enabled. Examples of these values include “MSIE”, “Mozilla” (which includes FireFox), and “Opera”.

Value strings can be up to 63 characters in length, and may not contain cross site scripting (XSS) vulnerability characters such as brackets. The FortiGate unit prevents use of these characters to prevent exploit of cross site scripting (XSS) vulnerabilities.

## Certificate authentication

Certificates can be used as part of an identity-based policy. A customized certificate must be installed on the FortiGate unit and in the web browser, which the FortiGate unit will attempt to match.

All users being authenticated against the policy are required to have the proper certificate, which must be imported into the FortiGate unit. See [“Certificate-based authentication” on page 532](#).

To require the user to accept a disclaimer to connect to the destination, select *Enable Disclaimer*. If the user is to be redirected after accepting the disclaimer, enter the URL in the *Redirect URL to* field. You can edit the User Authentication Disclaimer replacement message text in *System > Config > Replacement Messages*.

## Certificate redirect authentication

Under *User & Device > Authentication > Settings*, select *Redirect HTTP Challenge to a Secure Channel (HTTPS)*. This forces users to use secure connections to send their authentication information.

The following steps happen during a redirect:

1. User tries to access the Internet and the HTTP traffic hits the FortiGate security policy with authentication and HTTPS redirect enabled.
2. The FortiGate redirects the user with the HTTPS port and IP address of the interface connected to the user, such as internal.
3. User authenticates over the HTTPS connection as with normal authentication.
4. On successful authentication, the FortiGate provides access to the Internet as originally requested.

## Restricting number of concurrent user logons

Some users on your network may often have multiple account sessions open at one time either to the same network resource or accessing to the admin interface on the FortiGate unit.

While there are valid reasons for having multiple concurrent sessions open, hackers also do this to speed up their malicious work. Often a hacker is making multiple attempts to gain access to the internal network or the admin interface of the FortiGate unit, usually from different IP addresses to appear to the FortiGate unit as legitimate users. For this reason, the more concurrent sessions a hacker has open at once, the faster they will achieve their goal.

To help prevent this, you can disallow concurrent administrative access using the same administrator user name, but from a different IP address. This allows valid users to continue their legitimate work while limiting hackers' activity.

### To disable concurrent administrator sessions - CLI

```
config system global
 set admin-concurrent disable
end
```

## Limited access for unauthenticated users

When configuring User Identity policies, if you select the option *Skip this policy for unauthenticated user* the policy will only apply to users who have already authenticated with the FortiGate unit. This feature is intended for networks with two kinds of users:

- Single sign-on users who have authenticated when their devices connected to their network
- Other users who do not authenticate with the network so are “unauthenticated”

Sessions from authenticated users can match this policy and sessions from unauthenticated users will skip this policy and potentially be matched with policies further down the policy list. Typically, you would arrange a policy with *Skip this policy for unauthenticated user* at the top of a policy list.

You can also use the following CLI command to enable skipping policies for unauthenticated users:

```
config firewall policy
 edit <id>
 set identity-based enable
 set fall-through-unauthenticated enable
 next
```

## Use case - allowing limited access for unauthenticated users

Consider an office with open use PCs in common areas. Staff and customers do not have to log in to these PCs and can use them for limited access to the Internet. From their desks, employees of this office log into PCs which are logged into the office network. The FortiGate unit on the office network uses single sign-on to get user credentials from the network authentication server.

The open use PCs have limited access to the Internet. Employee PCs can access internal resources and have unlimited access to the Internet.

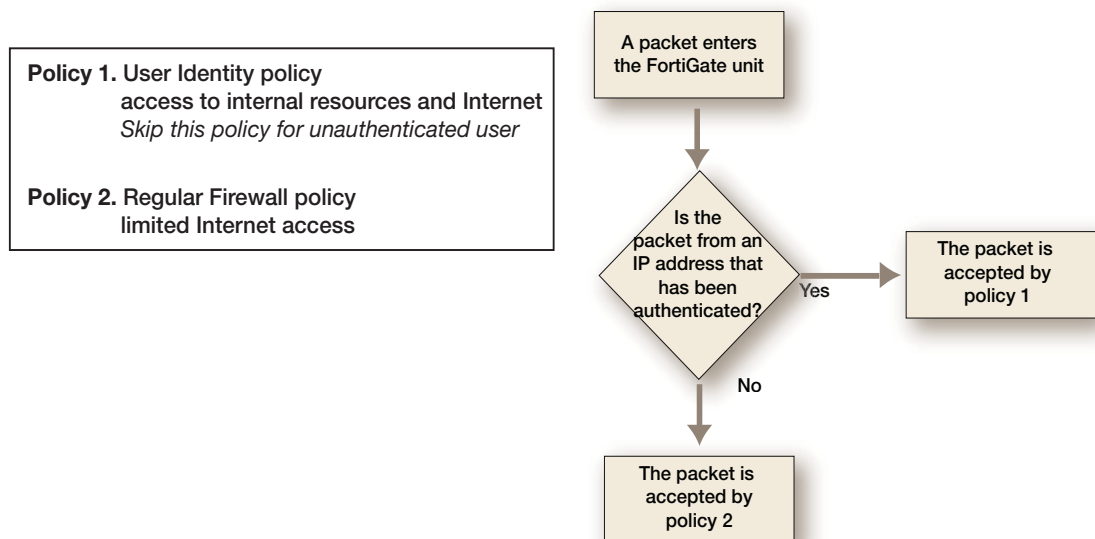
To support these different levels of access you can add a user identity policy to the top of the policy list that allows authenticated users to access internal resources and to have unlimited access to the Internet. In this policy, select *Skip this policy for unauthenticated user*.

Add a normal firewall policy below this policy that allows limited access to the Internet.

Sessions from authenticated PCs will be accepted by the User Identity policy. Sessions from unauthenticated PCs will skip the User Identity policy and be accepted by the normal firewall policy.

Figure 116 shows how the FortiGate unit handles packets received from authenticated and unauthenticated users.

**Figure 116:**Packet flow for authenticated and unauthenticated users



## Use case - multiple levels of authentication

As a variation of the above use case, Policy 2 could be a User Identity policy and *Skip this policy for unauthenticated user* would not be selected. Sessions from unauthenticated users that are accepted by Policy2 would now require users to authenticate before traffic can connect

through the FortiGate unit. The result is different levels of authentication: Single sign on for some users and firewall authentication for others.

## VPN authentication

All VPN configurations require users to authenticate. Authentication based on user groups applies to:

- SSL VPNs
- PPTP and L2TP VPNs
- an IPsec VPN that authenticates users using dialup groups
- a dialup IPsec VPN that uses XAUTH authentication (Phase 1)

You must create user accounts and user groups before performing the procedures in this section. If you create a user group for dialup IPsec clients or peers that have unique peer IDs, their user accounts must be stored locally on the FortiGate unit. You cannot authenticate these types of users using a RADIUS or LDAP server.

### Configuring authentication of SSL VPN users

The general procedure for authenticating SSL VPN users is:

1. Configure user accounts.
2. Create one or more user groups for SSL VPN users.  
See “Configuring user accounts and user groups for SSL VPN” in the *FortiOS Handbook SSL VPN chapter*.
3. Enable SSL VPN.
4. Optionally, set inactivity and authentication timeouts.
5. Configure a security policy with SSL VPN action. Add an identity-based rule to allow access for the user groups you created for SSL VPN users.  
See “Configuring security policies” in the *FortiOS Handbook SSL VPN chapter*.

### Configuring authentication timeout

By default, the SSL VPN authentication expires after 8 hours (28 800 seconds). You can change it only in the CLI, and the time entered must be in seconds. For example, to change this timeout to one hour, you would enter:

```
config vpn ssl settings
 set auth-timeout 3600
end
```

If you set the authentication timeout (`auth-timeout`) to 0 when you configure the timeout settings, the remote client does not have to re-authenticate unless they log out of the system. To fully take advantage of this setting, the value for `idle-timeout` has to be set to 0 also, so that the client does not time out if the maximum idle time is reached. If the `idle-timeout` is not set to the infinite value, the system will log out if it reaches the limit set, regardless of the `auth-timeout` setting.

### Configuring authentication of remote IPsec VPN users

An IPsec VPN on a FortiGate unit can authenticate remote users through a dialup group. The user account name is the peer ID and the password is the pre-shared key.



Authentication through user groups is supported for groups containing only local users. To authenticate users using a RADIUS or LDAP server, you must configure XAUTH settings. See [“Configuring XAuth authentication” on page 529](#).

### To configure user group authentication for dialup IPsec - web-based manager

1. Configure the dialup users who are permitted to use this VPN. Create a user group with Type:Firewall and add them to it.

For more information, see [“Users and user groups” on page 488](#).

2. Go to *VPN > IPsec > Auto Key (IKE)*, select *Create Phase 1* and enter the following information.

<b>Name</b>	Name for group of dialup users using the VPN for authentication.
<b>Remote Gateway</b>	List of the types of remote gateways for VPN. Select <i>Dialup User</i> .
<b>Authentication Method</b>	List of authentication methods available for users. Select <i>Preshared Key</i> and enter the preshared key.
<b>Peer Options</b>	Select <i>Accept peer ID in dialup group</i> . Select the user group that is to be allowed access to the VPN. The listed user groups contain only users with passwords on the FortiGate unit.

3. Select *Advanced* to reveal additional parameters and configure other VPN gateway parameters as needed.
4. Select *OK*.

### To configure user group authentication for dialup IPsec - CLI example

The `peertype` and `usrgrp` options configure user group-based authentication.

```
config vpn ipsec phase1
 edit office_vpn
 set interface port1
 set type dynamic
 set psksecret yORRAzltNGhzgtV32jend
 set proposal 3des-sha1 aes128-sha1
 set peertype dialup
 set usrgrp Group1
 end
```

## Configuring XAuth authentication

Extended Authentication (XAuth) increases security by requiring additional user authentication information in a separate exchange at the end of the VPN Phase 1 negotiation. The FortiGate unit asks the user for a username and password. It then forwards the user’s credentials (the password is encrypted) to an external RADIUS or LDAP server for verification.

XAuth can be used in addition to or in place of IPsec phase 1 peer options to provide access security through an LDAP or RADIUS authentication server. You must configure a dialup user group whose members are all externally authenticated.

### To configure authentication for a dialup IPsec VPN - web-based manager

1. Configure the users who are permitted to use this VPN. Create a user group and add the users to the group.

For more information, see [“Users and user groups” on page 488](#).

2. Go to *VPN > IPsec > Auto Key (IKE)*.
3. Select *Create Phase 1* and configure the basic VPN phase1 settings.

*Remote Gateway* must be *Dialup User*.

4. Select *Advanced* to reveal additional parameters and enter the following information.

<b>XAuth</b>	Select <i>Enable as Server</i> .
<b>Server Type</b>	Select <i>PAP, CHAP, or AUTO</i> . Use CHAP whenever possible. Use PAP with all implementations of LDAP and with other authentication servers that do not support CHAP, including some implementations of Microsoft RADIUS. Use AUTO with the Fortinet Remote VPN Client and where the authentication server supports CHAP but the XAuth client does not.
<b>User Group</b>	Select the user group that is to have access to the VPN. The list of user groups does not include any group that has members whose password is stored on the FortiGate unit.

5. Select *OK*.

For more information about XAUTH configuration, see the IPsec VPN chapter of this FortiOS Handbook.

### To configure authentication for a dialup IPsec VPN - CLI example

The `xauthtype` and `authusrgrp` fields configure XAuth authentication.

```
config vpn ipsec phase1
 edit office_vpn
 set interface port1
 set type dynamic
 set psksecret yORRAzltNGhzgtV32jend
 set proposal 3des-sha1 aes128-sha1
 set peertype dialup
 set xauthtype pap
 set authusrgrp Group1
 end
```

Some parameters specific to setting up the VPN itself are not shown here. For detailed information about configuring IPsec VPNs, see the [FortiOS Handbook IPsec VPN chapter](#).

## Configuring authentication of PPTP VPN users and user groups

Configuration of a PPTP VPN is possible only through the CLI. You can configure user groups and security policies using either CLI or web-based manager.

### To configure authentication for a PPTP VPN

1. Configure the users who are permitted to use this VPN. Create a security user group and add them to it.

For more information, see [“Users and user groups” on page 488](#).

2. Configure the PPTP VPN in the CLI as in this example.

```
config vpn pptp
 set status enable
 set sip 192.168.0.100
 set eip 192.168.0.110
 set usrgrp PPTP_Group
end
```

The `sip` and `eip` fields define a range of virtual IP addresses assigned to PPTP clients.

3. Configure a security policy. The source interface is the one through which the clients will connect. The source address is the PPTP virtual IP address range. The destination interface and address depend on the network to which the clients will connect. The policy action is ACCEPT.

## Configuring authentication of L2TP VPN users/user groups

Configuration of a L2TP VPN is possible only through the CLI. You can configure user groups and security policies using either CLI or web-based manager.

### To configure authentication for a PPTP VPN

1. Configure the users who are permitted to use this VPN. Create a user group and add them to it.

For more information, see [“Users and user groups” on page 488](#).

2. Configure the L2TP VPN in the CLI as in this example.

```
config vpn l2tp
 set status enable
 set sip 192.168.0.100
 set eip 192.168.0.110
 set usrgrp L2TP_Group
end
```

The `sip` and `eip` fields define a range of virtual IP addresses assigned to L2TP clients.

3. Configure a security policy. The source interface is the one through which the clients will connect. The source address is the L2TP virtual IP address range. The destination interface and address depend on the network to which the clients will connect. The policy action is ACCEPT.

# Certificate-based authentication

This section provides an overview of how the FortiGate unit verifies the identities of administrators, SSL VPN users, or IPsec VPN peers using X.509 security certificates.

The following topics are included in this section:

- [What is a security certificate?](#)
- [Certificates overview](#)
- [Managing X.509 certificates](#)
- [Configuring certificate-based authentication](#)
- [Example — Generate a CSR on the FortiGate unit](#)
- [Example — Generate and Import CA certificate with private key pair on OpenSSL](#)
- [Example — Generate an SSL certificate in OpenSSL](#)

## What is a security certificate?

A security certificate is a small text file that is part of a third-party generated public key infrastructure (PKI) to help guarantee the identity of both the user logging on and the web site they where they are logging in.

A certificate includes identifying information such as the company and location information for the web site, as well as the third-party company name, the expiry date of the certificate, and the encrypted public key.

FortiGate units use X.509 certificates to authenticate single sign-on (SSO) for users. The X.509 standard has been in use since before 2000, but has gained popularity with the Internet's increased popularity. X.509 v3 is defined in RFC 5280 and specifies standard formats for public key certificates, certificate revocation lists, and a certification path validation algorithm. The unused earlier X.509 version 1 was defined in RFC 1422.

The main difference between X.509 and PGP certificates is that where in PGP anyone can sign a certificate, for X.509 only a trusted authority can sign certificates. This limits the source of certificates to well known and trustworthy sources. Where PGP is well suited for one-on-one communications, the X.509 infrastructure is intended to be used in many different situations including one-to-many communications. Some common filename extensions for X.509 certificates are listed in [Table 23](#).

**Table 23:** Common certificate filename extensions

Filetype	Format name	Description
<b>.pem</b>	Privacy Enhanced Mail (PEM)	Base64 encoded DER certificate, that uses “-----BEGIN CERTIFICATE-----” and “-----END CERTIFICATE-----”
<b>.cer .crt .der</b>	Security CERTificate	Usually binary DER form, but Base64-encoded certificates are common too.

**Table 23:** Common certificate filename extensions

<b>.p7b</b> <b>.p7c</b>	PKCS#7 SignedData	Structure without data, just certificates or CRLs. PKCS#7 is a standard for signing or encrypting (officially called “enveloping”) data.
<b>.p12</b>	PKCS#12	May contain certificate(s) (public) and private keys (password protected)
<b>.pfx</b>	personal information exchange (PFX)	Older format. Came before PKCS#12. Usually today data is in PKCS#12 format.

## Certificates overview

Certificates play a major role in authentication of clients connecting to network services via HTTPS, both for administrators and SSL VPN users. Certificate authentication is optional for IPsec VPN peers.

- [Certificates and protocols](#)
- [IPsec VPNs and certificates](#)
- [Certificate types on the FortiGate unit](#)

## Certificates and protocols

There are a number of protocols that are commonly used with certificates including SSL and HTTPS, and other certificate-related protocols.

### SSL and HTTPS

The secure HTTP (HTTPS) protocol uses SSL. Certificates are an integral part of SSL. When a web browser connects to the FortiGate unit via HTTPS, a certificate is used to verify the FortiGate unit’s identity to the client. Optionally, the FortiGate unit can require the client to authenticate itself in return.

By default, the FortiGate unit uses a self-signed security certificate to authenticate itself to HTTPS clients. When the certificate is offered, the client browser displays two security messages.

- The first message prompts users to accept and optionally install the FortiGate unit’s self-signed security certificate. If the user does not accept the certificate, the FortiGate unit refuses the connection. When the user accepts the certificate, the FortiGate login page is displayed, and the credentials entered by the user are encrypted before they are sent to the FortiGate unit. If the user chooses to install the certificate, the prompt is not displayed again.
- Just before the FortiGate login page is displayed, a second message informs users that the FortiGate certificate distinguished name differs from the original request. This message is displayed because the FortiGate unit redirects the connection (away from the distinguished name recorded in the self-signed certificate) and can be ignored.

Optionally, you can install an X.509 server certificate issued by a certificate authority (CA) on the FortiGate unit. You can then configure the FortiGate unit to identify itself using the server certificate instead of the self-signed certificate. For more information, see the [FortiOS Handbook SSL VPN chapter](#). or “[Authenticating SSL VPN users with security certificates](#)” on [page 544](#).

After successful certificate authentication, communication between the client browser and the FortiGate unit is encrypted using SSL over the HTTPS link.

## Certificate-related protocols

There are multiple protocols that are required for handling certificates. These include the Online Certificate Status Protocol (OCSP), Secure Certificate Enrollment Protocol (SCEP), and Server-based Certificate Validation Protocol (SCVP).

### Online Certificate Status Protocol

Online Certificate Status Protocol (OCSP) allows the verification of X.509 certificate expiration dates. This is important to prevent hackers from changing the expiry date on an old certificate to a future date.

Normally certificate revocation lists (CRLs) are used, but OCSP is an alternate method available. However a CRL is a public list, and some companies may want to avoid the public exposure of their certificate structure even if it is only invalid certificates.

The OSCP check on the certificate's revocation status is typically carried out over HTTP with a request-response format. The authority responding can reply with a status of good, revoked, or unknown for the certificate in question.

### Secure Certificate Enrollment Protocol

Secure Certificate Enrollment Protocol (SCEP) is an automated method of signing up for certificates. Typically this involves generating a request you send directly to the SCEP service, instead of generating a file request that may or may not be signed locally.

### Server-based Certificate Validation Protocol

Server-based Certificate Validation Protocol (SCVP) is used to trace a certificate back to a valid root level certificate. This ensures that each step along the path is valid and trustworthy.

## IPsec VPNs and certificates

Certificate authentication is a more secure alternative to preshared key (shared secret) authentication for IPsec VPN peers. Unlike administrators or SSL VPN users, IPsec peers use HTTP to connect to the VPN gateway configured on the FortiGate unit. The VPN gateway configuration can require certificate authentication before it permits an IPsec tunnel to be established. See [“Authenticating IPsec VPN users with security certificates” on page 545](#).

## Certificate types on the FortiGate unit

There are different types of certificates available that vary depending on their intended use. FortiOS supports local, remote, CA, and CRL certificates.

### Local certificates

Local certificates are issued for a specific server, or web site. Generally they are very specific, and often for an internal enterprise network. For example a personal web site for John Smith at [www.example.com](http://www.example.com) (such as <http://www.example.com/home/jsmith>) would have its own local certificate.

These can optionally be just the certificate file, or also include a private key file and PEM passphrase for added security.

For information about generating a certificate request, see [“Generating a certificate signing request” on page 536](#). For information about installing a local certificate, see [“Obtaining and installing a signed server certificate from an external CA” on page 538](#).

## Remote certificates

Remote certificates are public certificates without a private key. For dynamic certificate revocation, you need to use an Online Certificate Status Protocol (OCSP) server. The OCSP is configured in the CLI only. Installed Remote (OCSP) certificates are displayed in the Remote Certificates list. You can select *Import* to install a certificate from the management PC.

## CA root certificates

CA root certificates are similar to local certificates, however they apply to a broader range of addresses or to whole company; they are one step higher up in the organizational chain. Using the local certificate example, a CA root certificate would be issued for all of `www.example.com` instead of just the smaller single web page.

## Certificate revocation list

Certificate revocation list (CRL) is a list of certificates that have been revoked and are no longer usable. This list includes certificates that have expired, been stolen, or otherwise compromised. If your certificate is on this list, it will not be accepted. CRLs are maintained by the CA that issues the certificates and includes the date and time when the next CRL will be issued as well as a sequence number to help ensure you have the most current version of the CRL.

## Certificate signing

The trust in a certificate comes from the authority that signs it. For example if VeriSign signs your CA root certificate, it is trusted by everyone. While these certificates are universally accepted, it is cumbersome and expensive to have all certificates on a corporate network signed with this level of trust.

With self-signed certificates nobody, except the other end of your communication, knows who you are and therefore they do not trust you as an authority. However this level is useful for encryption between two points — neither point may care about who signed the certificate, just that it allows both points to communicate. This is very useful for internal networks and communications.

A general rule is that CA signed certificates are accepted and sometimes required, but it is easier to self-sign certificates when you are able.

For more on the methods of certificate signing see [“Generating a certificate signing request” on page 536](#).

## Managing X.509 certificates

Managing security certificates is required due to the number of steps involved in both having a certificate request signed, and then distributing the correct files for use.

You use the FortiGate unit or CA software such as OpenSSL to generate a certificate request. That request is a text file that you send to the CA for verification, or alternately you use CA software to self-validate. Once validated, the certificate file is generated and must be imported to the FortiGate unit before it can be used. These steps are explained in more detail later in this section.

This section provides procedures for generating certificate requests, installing signed server certificates, and importing CA root certificates and CRLs to the FortiGate unit.

For information about how to install root certificates, CRLs, and personal or group certificates on a remote client browser, refer to your browser’s documentation.

This section includes:

- [Generating a certificate signing request](#)
- [Generating certificates with CA software](#)
- [Obtaining and installing a signed server certificate from an external CA](#)
- [Installing a CA root certificate and CRL to authenticate remote clients](#)
- [Troubleshooting certificates](#)

## Generating a certificate signing request

Whether you create certificates locally with a software application or obtain them from an external certificate service, you will need to generate a certificate signing request (CSR).

When you generate a CSR, a private and public key pair is created for the FortiGate unit. The generated request includes the public key of the FortiGate unit and information such as the FortiGate unit's public static IP address, domain name, or email address. The FortiGate unit's private key remains confidential on the FortiGate unit.

After you submit the request to a CA, the CA will verify the information and register the contact information on a digital certificate that contains a serial number, an expiration date, and the public key of the CA. The CA will then sign the certificate, and you install the certificate on the FortiGate unit.

The Certificate Request Standard is a public key cryptography standard (PKCS) published by RSA, specifically PKCS10 which defines the format for CSRs. This is defined in RFC 2986.

### To generate a certificate request in FortiOS - web-based manager

1. Go to *System > Certificates > Local Certificates*.
2. Select *Generate*.
3. In the *Certificate Name* field, enter a unique meaningful name for the certificate request. Typically, this would be the hostname or serial number of the FortiGate unit or the domain of the FortiGate unit such as example.com.



Do not include spaces in the certificate name. This will ensure compatibility of a signed certificate as a PKCS12 file to be exported later on if required.

- 
4. Enter values in the *Subject Information* area to identify the FortiGate unit:
    - If the FortiGate unit has a static IP address, select *Host IP* and enter the public IP address of the FortiGate unit. If the FortiGate unit does not have a public IP address, use an email address (or fully qualified domain name (FQDN) if available) instead.
    - If the FortiGate unit has a dynamic IP address and subscribes to a dynamic DNS service, use a FQDN if available to identify the FortiGate unit. If you select *Domain Name*, enter the FQDN of the FortiGate unit. Do not include the protocol specification (`http://`) or any port number or path names.



If a domain name is not available and the FortiGate unit subscribes to a dynamic DNS service, an “unable to verify certificate” type message may be displayed in the user's browser whenever the public IP address of the FortiGate unit changes.

---



- If you select *E-Mail*, enter the email address of the owner of the FortiGate unit.

5. Enter values in the *Optional Information* area to further identify the FortiGate unit.

<b>Organization Unit</b>	Name of your department. You can enter a series of OUs up to a maximum of 5. To add or remove an OU, use the plus (+) or minus (-) icon.
<b>Organization</b>	Legal name of your company or organization.
<b>Locality (City)</b>	Name of the city or town where the FortiGate unit is installed.
<b>State/Province</b>	Name of the state or province where the FortiGate unit is installed.
<b>Country</b>	Select the country where the FortiGate unit is installed.
<b>e-mail</b>	Contact email address.
<b>Subject Alternative Name</b>	<p>Optionally, enter one or more alternative names for which the certificate is also valid. Separate names with a comma. A name can be:</p> <ul style="list-style-type: none"> <li>• e-mail address</li> <li>• IP address</li> <li>• URI</li> <li>• DNS name (alternatives to the Common Name)</li> <li>• directory name (alternatives to the Distinguished Name)</li> </ul> <p>You must precede the name with the name type. Examples:</p> <p>IP:1.1.1.1</p> <p>email:test@fortinet.com</p> <p>email:my@other.address</p> <p>URI:http://my.url.here/</p>

6. From the *Key Size* list, select *1024 Bit*, *1536 Bit* or *2048 Bit*. Larger keys are slower to generate but more secure.
  7. In *Enrollment Method*, you have two methods to choose from. Select *File Based* to generate the certificate request, or *Online SCEP* to obtain a signed SCEP-based certificate automatically over the network. For the SCEP method, enter the URL of the SCEP server from which to retrieve the CA certificate, and the CA server challenge password.
  8. Select *OK*.
  9. The request is generated and displayed in the *Local Certificates* list with a status of *PENDING*.
  10. Select the *Download* button to download the request to the management computer.
  11. In the *File Download* dialog box, select *Save* and save the Certificate Signing Request on the local file system of the management computer.
  12. Name the file and save it on the local file system of the management computer.
- The certificate request is ready for the certificate to be generated.

## Generating certificates with CA software

CA software allows you to generate unmanaged certificates and CA certificates for managing other certificates locally without using an external CA service. Examples of CA software include `ssl-ca` from OpenSSL (available for Linux, Windows, and Mac) or `gensslcert` from SuSE. MS Windows Server 2000 and 2003 come with a CA as part of their certificate services, and in MS Windows 2008 CA software can be installed as part of the Active Directory installation. See [“Example — Generate and Import CA certificate with private key pair on OpenSSL” on page 546](#).

The general steps for generating certificates with CA software are

1. Install the CA software as a stand-alone root CA.
2. Provide identifying information for your self-administered CA.

While following these steps, the methods vary slightly when generating server certificates, CA certificates, and PKI certificates.

### Server certificate

1. Generate a Certificate Signing Request (CSR) on the FortiGate unit.
2. Copy the CSR base-64 encoded text (PKCS10 or PKCS7) into the CA software and generate the certificate.  
PKCS10 is the format used to send the certificate request to the signing authority. PKCS7 is the format the signing authority can use for the newly signed certificate.
3. Export the certificate as a X.509 DER encoded binary file with `.CER` extension
4. Upload the certificate file to the FortiGate unit Local Certificates page (type is Certificate).

### CA certificate

1. Retrieve the CA Certificate from the CA software as a DER encoded file.
2. Upload the CA certificate file to the FortiGate unit CA Certificates page at *System > Certificates > CA Certificates*.

### PKI certificate

1. Generate a Certificate Signing Request (CSR) on the FortiGate unit.
2. Copy the CSR base-64 encoded text (PKCS#10 or PKCS#7) into the CA software and generate the certificate.  
PKCS10 is the format used to send the certificate request to the signing authority. PKCS7 is the format the signing authority can use for the newly signed certificate.
3. Export the certificate as a X.509 DER encoded binary file with `.CER` extension.
4. Install the certificate in the user's web browser or IPsec VPN client as needed.

## Obtaining and installing a signed server certificate from an external CA

To obtain a signed server certificate for a FortiGate unit, you must send a request to a CA that provides digital certificates that adhere to the X.509 standard. The FortiGate unit provides a way for you to generate the request.

### To submit the certificate signing request (file-based enrollment)

1. Using the web browser on the management computer, browse to the CA web site.
2. Follow the CA instructions for a base-64 encoded PKCS#10 certificate request and upload your certificate request.

3. Follow the CA instructions to download their root certificate and CRL.

When you receive the signed server certificate from the CA, install the certificate on the FortiGate unit.

#### **To install or import the signed server certificate - web-based manager**

1. On the FortiGate unit, go to *System > Certificates > Local Certificates*.
2. Select *Import*.
3. From *Type*, select *Local Certificate*.
4. Select *Browse*, browse to the location on the management computer where the certificate was saved, select the certificate, and then select *Open*.
5. Select *OK*, and then select *Return*.

## **Installing a CA root certificate and CRL to authenticate remote clients**

When you apply for a signed personal or group certificate to install on remote clients, you can obtain the corresponding root certificate and CRL from the issuing CA. When you receive the signed personal or group certificate, install the signed certificate on the remote client(s) according to the browser documentation. Install the corresponding root certificate (and CRL) from the issuing CA on the FortiGate unit according to the procedures given below.

#### **To install a CA root certificate**

1. After you download the root certificate of the CA, save the certificate on the management computer. Or, you can use online SCEP to retrieve the certificate.
2. On the FortiGate unit, go to *System > Certificates > CA Certificates*.
3. Select *Import*.
4. Do one of the following:
  - To import using SCEP, select *SCEP*. Enter the URL of the SCEP server from which to retrieve the CA certificate. Optionally, enter identifying information of the CA, such as the filename.
  - To import from a file, select *Local PC*, then select *Browse* and find the location on the management computer where the certificate has been saved. Select the certificate, and then select *Open*.
5. Select *OK*, and then select *Return*.

The system assigns a unique name to each CA certificate. The names are numbered consecutively (CA\_Cert\_1, CA\_Cert\_2, CA\_Cert\_3, and so on).

#### **To import a certificate revocation list**

A Certificate Revocation List (CRL) is a list of the CA certificate subscribers paired with certificate status information. The list contains the revoked certificates and the reason(s) for revocation. It also records the certificate issue dates and the CAs that issued them.

When configured to support SSL VPNs, the FortiGate unit uses the CRL to ensure that the certificates belonging to the CA and remote peers or clients are valid. The CRL has an “effective date” and a “next update” date. The interval is typically 7 days (for Microsoft CA). FortiOS will update the CRL automatically. Also, there is a CLI command to specify an “update-interval” in seconds. Recommendation should be 24 hours (86400 seconds) but depends on company security policy.

1. After you download the CRL from the CA web site, save the CRL on the management computer.
2. Go to *System > Certificates > CRL*.
3. Select *Import*.

4. Do one of the following:
  - To import using an HTTP server, select *HTTP* and enter the URL of the HTTP server.
  - To import using an LDAP server see this KB [article](#).
  - To import using an SCEP server, select *SCEP* and select the Local Certificate from the list. Enter the URL of the SCEP server from which the CRL can be retrieved.
  - To import from a file, select *Local PC*, then select *Browse* and find the location on the management computer where the CRL has been saved. Select the CRL and then select *Open*.
5. Select *OK*, and then select *Return*.

## Troubleshooting certificates

There are times when there are problems with certificates — a certificate is seen as expired when its not, or it can't be found. Often the problem is with a third party web site, and not FortiOS. However, some problems can be traced back to FortiOS such as DNS or routing issues.

### Certificate is reported as expired when it is not

Certificates often are issued for a set period of time such as a day or a month, depending on their intended use. This ensures everyone is using up-to-date certificates. It is also more difficult for hackers to steal and use old certificates.

Reasons a certificate may be reported as expired include:

- It really has expired based on the “best before” date in the certificate
- The FortiGate unit clock is not properly set. If the FortiGate clock is fast, it will see a certificate as expired before the expiry date is really here.
- The requesting server clock is not properly set. A valid example is if your certificate is 2 hours from expiring, a server more than two time zones away would see the certificate as expired. Otherwise, if the server's clock is set wrongly it will also have the same effect.
- The certificate was revoked by the issuer before the expiry date. This may happen if the issuer believes a certificate was either stolen or misused. Its possible it is due to reasons on the issuer's side, such as a system change or such. In either case it is best to contact the certificate issuer to determine what is happening and why.

### A secure connection cannot be completed (Certificate cannot be found)

Everyone who uses a browser has encountered a message such as *This connection is untrusted*. Normally when you try to connect securely to a web site, that web site will present its valid certificate to prove their identity is valid. When the web site's certificate cannot be verified as valid, the message appears stating *This connection is untrusted* or something similar. If you usually connect to this web site without problems, this error could mean that someone is trying to impersonate or hijack the web site, and best practices dictates you not continue.

Reasons a web site's certificate cannot be validated include:

- The web site uses an unrecognized self-signed certificate. These are not secure because anyone can sign them. If you accept self-signed certificates you do so at your own risk. Best practices dictate that you must confirm the ID of the web site using some other method before you accept the certificate.
- The certificate is valid for a different domain. A certificate is valid for a specific location, domain, or sub-section of a domain such as one certificate for `support.example.com`

that is not valid for `marketing.example.com`. If you encounter this problem, contact the webmaster for the web site to inform them of the problem.

- There is a DNS or routing problem. If the web site's certificate cannot be verified, it will not be accepted. Generally to be verified, your system checks with the third party certificate signing authority to verify the certificate is valid. If you cannot reach that third party due to some DNS or routing error, the certificate will not be verified.
- Firewall is blocking required ports. Ensure that any firewalls between the requesting computer and the web site allow the secure traffic through the firewall. Otherwise a hole must be opened to allow it through. This includes ports such as 443 (HTTPS) and 22 (SSH).

## Online updates to certificates and CRLs

If you obtained your local or CA certificate using SCEP, you can configure online renewal of the certificate before it expires. Similarly, you can receive online updates to CRLs.

### Local certificates

In the `config vpn certificate local` command, you can specify automatic certificate renewal. The relevant fields are:

<code>scep-url &lt;URL_str&gt;</code>	The URL of the SCEP server. This can be HTTP or HTTPS. The following options appear after you add the <code>&lt;URL_str&gt;</code> .
<code>scep-password &lt;password_str&gt;</code>	The password for the SCEP server.
<code>auto-regenerate-days &lt;days_int&gt;</code>	How many days before expiry the FortiGate unit requests an updated local certificate. The default is 0, no auto-update.
<code>auto-regenerate-days -warning &lt;days_int&gt;</code>	How many days before local certificate expiry the FortiGate generates a warning message. The default is 0, no warning.

In this example, an updated certificate is requested three days before it expires.

```
config vpn certificate local
 edit mycert
 set scep-url http://scep.example.com/scep
 set scep-server-password my_pass_123
 set auto-regenerate-days 3
 set auto-regenerate-days-warning 2
 end
```

### CA certificates

In the `config vpn certificate ca` command, you can specify automatic certificate renewal. The relevant fields are:

<code>scep-url &lt;URL_str&gt;</code>	The URL of the SCEP server. This can be HTTP or HTTPS.
<code>auto-update-days &lt;days_int&gt;</code>	How many days before expiry the FortiGate unit requests an updated CA certificate. The default is 0, no auto-update.
<code>auto-update-days-warning &lt;days_int&gt;</code>	How many days before CA certificate expiry the FortiGate generates a warning message. The default is 0, no warning.

In this example, an updated certificate is requested three days before it expires.

```
config vpn certificate ca
 edit mycert
 set scep-url http://scep.example.com/scep
 set auto-update-days 3
 set auto-update-days-warning 2
 end
```

## Certificate Revocation Lists

If you obtained your CRL using SCEP, you can configure online updates to the CRL using the `config vpn certificate crl` command. The relevant fields are:

Variable	Description
<code>http-url &lt;http_url&gt;</code>	URL of the server used for automatic CRL certificate updates. This can be HTTP or HTTPS.
<code>scep-cert &lt;scep_certificate&gt;</code>	Local certificate used for SCEP communication for CRL auto-update.
<code>scep-url &lt;scep_url&gt;</code>	URL of the SCEP CA server used for automatic CRL certificate updates. This can be HTTP or HTTPS.
<code>update-interval &lt;seconds&gt;</code>	How frequently, in seconds, the FortiGate unit checks for an updated CRL. Enter 0 to update the CRL only when it expires. Not available for http URLs.
<code>update-vdom &lt;update_vdom&gt;</code>	VDOM used to communicate with remote SCEP server for CRL auto-update.

In this example, an updated CRL is requested only when it expires.

```
config vpn certificate crl
 edit cert_crl
 set http-url http://scep.example.com/scep
 set scep-cert my-scep-cert
 set scep-url http://scep.ca.example.com/scep
 set update-interval 0
 set update-vdom root
 end
```

## Backing up and restoring local certificates

The FortiGate unit provides a way to export and import a server certificate and the FortiGate unit's personal key through the CLI. If required (to restore the FortiGate unit configuration), you can import the exported file through the *System > Certificates > Local Certificates* page of the web-based manager.



As an alternative, you can back up and restore the entire FortiGate configuration through the *System Information* widget on the Dashboard of the web-based manager. Look for *[Backup]* and *[Restore]* in the *System Configuration* row. The backup file is created in a FortiGate-proprietary format.

### To export a server certificate and private key - CLI

This procedure exports a server (local) certificate and private key together as a password protected PKCS12 file. The export file is created through a customer-supplied TFTP server. Ensure that your TFTP server is running and accessible to the FortiGate unit before you enter the command.

1. Connect to the FortiGate unit through the CLI.

2. Type the following command:

```
execute vpn certificate local export tftp <cert_name> <exp_filename>
 <tftp_ip>
```

where:

- <cert\_name> is the name of the server certificate; typing ? displays a list of installed server certificates.
  - <exp\_filename> is a name for the output file.
  - <tftp\_ip> is the IP address assigned to the TFTP server host interface.
- 3 Move the output file from the TFTP server location to the management computer for future reference.

### To import a server certificate and private key - web-based manager

1. Go to *System > Certificates > Local Certificates* and select *Import*.

2. In *Type*, select *PKCS12 Certificate*.

3. Select *Browse*. Browse to the location on the management computer where the exported file has been saved, select the file, and then select *Open*.

4. In the *Password* field, type the password needed to upload the exported file.

5. Select *OK*, and then select *Return*.

### To import separate server certificate and private key files - web-based manager

Use the following procedure to import a server certificate and the associated private key file when the server certificate request and private key were not generated by the FortiGate unit. The two files to import must be available on the management computer.

1. Go to *System > Certificates > Local Certificates* and select *Import*.

2. In *Type*, select *Certificate*.

3. Select the *Browse* button beside the *Certificate file* field. Browse to the location on the management computer where the certificate file has been saved, select the file, and then select *Open*.

4. Select the *Browse* button beside the *Key file* field. Browse to the location on the management computer where the key file has been saved, select the file, and then select *Open*.

5. If required, in the *Password* field, type the associated password, and then select *OK*.

6. Select *Return*.

## Configuring certificate-based authentication

You can configure certificate-based authentication for FortiGate administrators, SSL VPN users, and IPsec VPN users.

In Microsoft Windows 7, you can use the certificate manager to keep track of all the different certificates on your local computer. To access certificate manager, in Windows 7 press the Windows key, enter "certmgr.msc" at the search prompt, and select the displayed match.

Remember that in addition to these system certificates, many applications require you to register certificates with them directly.

To see FortiClient certificates, open the FortiClient Console, and select VPN. The VPN menu has options for My Certificates (local or client) and CA Certificates (root or intermediary certificate authorities). Use Import on those screens to import certificate files from other sources.

## Authenticating administrators with security certificates

You can install a certificate on the management computer to support strong authentication for administrators. When a personal certificate is installed on the management computer, the FortiGate unit processes the certificate after the administrator supplies a username and password.

To enable strong administrative authentication:

- Obtain a signed personal certificate for the administrator from a CA and load the signed personal certificate into the web browser on the management computer according to the browser documentation.
- Install the root certificate and the CRL from the issuing CA on the FortiGate unit (see [“Installing a CA root certificate and CRL to authenticate remote clients” on page 539](#)).
- Create a PKI user account for the administrator.
- Add the PKI user account to a firewall user group dedicated to PKI-authenticated administrators.
- In the administrator account configuration, select *PKI* as the account *Type* and select the *User Group* to which the administrator belongs.

## Authenticating SSL VPN users with security certificates

While the default self-signed certificates can be used for HTTPS connections, it is preferable to use the X.509 server certificate to avoid the redirection as it can be misinterpreted as possible session hijacking. However, the server certificate method is more complex than self-signed security certificates. Also the warning message is typically displayed for the initial connection, and future connections will not generate these messages.

X.509 certificates can be used to authenticate IPsec VPN peers or clients, or SSL VPN clients. When configured to authenticate a VPN peer or client, the FortiGate unit prompts the VPN peer or client to authenticate itself using the X.509 certificate. The certificate supplied by the VPN peer or client must be verifiable using the root CA certificate installed on the FortiGate unit in order for a VPN tunnel to be established.

### To enable certificate authentication for an SSL VPN user group

1. Install a signed server certificate on the FortiGate unit and install the corresponding root certificate (and CRL) from the issuing CA on the remote peer or client.
2. Obtain a signed group certificate from a CA and load the signed group certificate into the web browser used by each user. Follow the browser documentation to load the certificates.
3. Install the root certificate and the CRL from the issuing CA on the FortiGate unit (see [“Installing a CA root certificate and CRL to authenticate remote clients” on page 539](#)).
4. Create a PKI user for each SSL VPN user. For each user, specify the text string that appears in the Subject field of the user’s certificate and then select the corresponding CA certificate.
5. Use the `config user peergrp` CLI command to create a peer user group. Add to this group all of the SSL VPN users who are authenticated by certificate.
6. Go to *Policy > Policy > Policy*.
7. Edit the SSL-VPN security policy.



8. Select *SSL Client Certificate Restrictive*.
9. Select *OK*.

## Authenticating IPsec VPN users with security certificates

To require VPN peers to authenticate by means of a certificate, the FortiGate unit must offer a certificate to authenticate itself to the peer.

### To enable the FortiGate unit to authenticate itself with a certificate:

1. Install a signed server certificate on the FortiGate unit.  
See [“To install or import the signed server certificate - web-based manager” on page 539](#).
2. Install the corresponding CA root certificate on the remote peer or client. If the remote peer is a FortiGate unit, see [“To install a CA root certificate” on page 539](#).
3. Install the certificate revocation list (CRL) from the issuing CA on the remote peer or client. If the remote peer is a FortiGate unit, see [“To import a certificate revocation list” on page 539](#).
4. In the VPN phase 1 configuration, set *Authentication Method* to *RSA Signature* and from the *Certificate Name* list select the certificate that you installed in Step 1.

To authenticate a VPN peer using a certificate, you must install a signed server certificate on the peer. Then, on the FortiGate unit, the configuration depends on whether there is only one VPN peer or if this is a dialup VPN that can have multiple peers.

### To configure certificate authentication of a single peer

1. Install the CA root certificate and CRL.
2. Create a PKI user to represent the peer. Specify the text string that appears in the Subject field of the user’s certificate and then select the corresponding CA certificate.
3. In the VPN phase 1 *Peer Options*, select *Accept this peer certificate only* and select the PKI user that you created.

### To configure certificate authentication of multiple peers (dialup VPN)

1. Install the corresponding CA root certificate and CRL.
2. Create a PKI user for each remote VPN peer. For each user, specify the text string that appears in the Subject field of the user’s certificate and then select the corresponding CA certificate.
3. Use the `config user peergrp` CLI command to create a peer user group. Add to this group all of the PKI users who will use the IPsec VPN.

In the VPN phase 1 *Peer Options*, select *Accept this peer certificate group only* and select the peer group that you created.

## Example – Generate a CSR on the FortiGate unit

This example follows all the steps required to create and install a local certificate on the FortiGate unit, without using CA software.

The FortiGate unit is called myFortiGate60, and is located at 10.11.101.101 (a private IP address) and <http://myfortigate.example.com>. Mr. John Smith ([john.smith@myfortigate.example.com](mailto:john.smith@myfortigate.example.com)) is the IT administrator for this FortiGate unit, and the unit belongs to the Sales department located in Greenwich, London, England.

### To generate a certificate request on the FortiGate unit - web-based manager

1. Go to *System > Certificates > Local Certificates*.
2. Select *Generate*.

3. In the *Certificate Name* field, enter `myFortiGate60`.



Do not include spaces in the certificate name. This will ensure compatibility of a signed certificate as a PKCS12 file to be exported later on if required.

Since the IP address is private, we will use the FQDN instead.

4. Select *Domain Name*, and enter `http://myfortigate.example.com`.
5. Enter values in the *Optional Information* area to further identify the FortiGate unit.

<b>Organization Unit</b>	Sales
<b>Organization</b>	Example.com
<b>Locality (City)</b>	Greenwich
<b>State/Province</b>	London
<b>Country</b>	England
<b>e-mail</b>	john.smith@myfortigate.example.com

6. From the *Key Size* list, select *2048 Bit* or the most secure option available to you.
7. In *Enrollment Method*, select *File Based* to generate the certificate request
8. Select *OK*.  
The request is generated and displayed in the *Local Certificates* list with a status of *PENDING*.
9. Select the *Download* button to download the request to the management computer.
10. In the *File Download* dialog box, select *Save* and save the Certificate Signing Request on the local file system of the management computer.
11. Name the file and save it on the local file system of the management computer.

## Example – Generate and Import CA certificate with private key pair on OpenSSL

This example explains how to generate a certificate using OpenSSL on MS Windows. OpenSSL is available for Linux and Mac OS as well, however their terminology will vary slightly from what is presented here.

### Assumptions

Before starting this procedure, ensure that you have downloaded and installed OpenSSL on Windows. One source is <http://www.slproweb.com/products/Win32OpenSSL.html>.

### Generating and importing the CA certificate and private key

The two following procedures will generate a CA certificate file and private key file, and then import it to the FortiGate unit as a local certificate.

### To generate the private key and certificate

1. At the Windows command prompt, go to the OpenSSL bin directory. If you installed to the default location this will be the following command:

```
cd c:\OpenSSL-Win32\bin
```

2. Enter the following command to generate the private key. You will be prompted to enter your PEM pass phrase. Choose something easy to remember such as fortinet123.

```
openssl genrsa -des3 -out fgtcapriv.key 2048
```

This command generates an RSA DES3 2038-bit encryption key.

3. The following command will generate the certificate using the key from the previous step.

```
openssl req -new -x509 -days 3650 -extensions v3_ca -key
fgtcapriv.key -out fgtca.crt
```

This step generates an X509 CA certificate good for 10 years that uses the key generated in the previous step. The certificate filename is `fgtca.crt`.

You will be prompted to enter information such as PEM Pass Phrase from the previous step, Country Name, State, Organization Name, Organizational Unit (such as department name), Common Name (the FQDN), and Email Address.

### To import the certificate to the FortiGate unit - web-based manager

1. Go to *System > Certificates > Local Certificates*.
2. Select *Import*.
3. Select Certificate for Type.  
Fields for Certificate file, Key file, and Password are displayed.
4. For Certificate file, enter `c:\OpenSSL-Win32\bin\fgtca.crt`.
5. For Key file, enter `c:\OpenSSL-Win32\bin\fgtcapriv.key`.
6. For Password, enter the PEM Pass Phrase you entered earlier, such as fortinet123.
7. Select OK.

The Certificate will be added to the list of Local Certificates and be ready for use. It will appear in the list as the filename you uploaded — `fgtca`. You can add comments to this certificate to make it clear where its from and how it is intended to be used. If you download the certificate from FortiOS, it is a `.CER` file.

It can now be used in [“Authenticating IPsec VPN users with security certificates” on page 545](#), and [“Authenticating SSL VPN users with security certificates” on page 544](#).

Optionally, you can install the certificate as a CA Certificate. CA certificates are used in HTTPS proxy/inspection. To do this, under *CA Certificates* select *Import*. Select *Local PC* and enter the certificate file `c:\OpenSSL-Win32\bin\fgtca.crt`. Then select *OK*. This certificate will be displayed in the CA Certificate list under the name `CA_Cert_1`.

## Example — Generate an SSL certificate in OpenSSL

This example explains how to generate a CA signed SSL certificate using OpenSSL on MS Windows. OpenSSL is available for Linux and Mac OS as well, however their terminology will vary slightly from what is presented here.

This example includes:

- [Assumptions](#)
- [Generating a CA signed SSL certificate](#)
- [Generating a self-signed SSL certificate](#)
- [Import the SSL certificate into FortiOS](#)

## Assumptions

- Before starting this procedure, ensure that you have downloaded and installed OpenSSL on MS Windows. One download source is <http://www.slproweb.com/products/Win32OpenSSL.html>.

## Generating a CA signed SSL certificate

This procedure assumes:

- you have already completed [“Example — Generate and Import CA certificate with private key pair on OpenSSL” on page 546](#) successfully.

### To generate the CA signed SSL certificate

1. At the Windows command prompt, go to the OpenSSL bin directory. If you installed to the default location this will be the following command:

```
cd c:\OpenSSL-Win32\bin
```

2. Enter the following command to generate the private key. You will be prompted to enter your PEM pass phrase. Choose something easy to remember such as fortinet.

```
openssl genrsa -des3 -out fgtssl.key 2048
```

This command generates an RSA DES3 2038-bit encryption key.

3. Create a certificate signing request for the SSL certificate. This step requires you to enter the information listed in step 3 of the previous example — [“To generate the private key and certificate” on page 547](#). You can leave the Challenge Password blank.

```
openssl req -new -key fgtssl.key -out fgtssl.csr
```

4. Using the CSR from the previous step, you can now create the SSL certificate using the CA certificate that was created in [“Example — Generate and Import CA certificate with private key pair on OpenSSL” on page 546](#).

```
openssl x509 -req -days 365 -in fgtssl.csr -CA fgtca.crt -CAkey fgtcapriv.key -set_serial 01 -out fgtssl.crt
```

This will generate an X.509 certificate good for 365 days signed by the CA certificate fgtca.crt.

## Generating a self-signed SSL certificate

This procedure does not require any existing certificates.

1. At the Windows command prompt, go to the OpenSSL bin directory. If you installed to the default location this will be the following command:

```
cd c:\OpenSSL-Win32\bin
```

2. Enter the following command to generate the private key. You will be prompted to enter your PEM pass phrase. Choose something easy to remember such as fortinet.

```
openssl genrsa -des3 -out fgtssl.key 2048
openssl req -new -key fgtssl.key -out fgtssl.csr
openssl x509 -req -days 365 -in fgtssl.csr -signkey fgtssl.key
-out fgtssl.crt
```

These commands:

- generate an RSA 3DES 2048-bit private key,
- generate an SSL certificate signing request, and
- sign the CSR to generate an SSL .CRT certificate file.

## Import the SSL certificate into FortiOS

### To import the certificate to FortiOS- web-based manager

1. Go to *System > Certificates > Local Certificates*.
2. Select *Import*.
3. Select Certificate for Type.  
Fields for Certificate file, Key file, and Password are displayed.
4. For Certificate file, enter `c:\OpenSSL-Win32\bin\fgtssl.crt`.
5. For Key file, enter `c:\OpenSSL-Win32\bin\fgtssl.key`.
6. For Password, enter the PEM Pass Phrase you entered, such as `fortinet`.
7. Select OK.

The SSL certificate you just uploaded can be found under *System > Certificates > Local Certificates* under the name of the file you uploaded — `fgtssl`.

### To confirm the certificate is uploaded properly - CLI

```
config vpn certificate local
 edit fgtssl
 get
 end
```

The `get` command will display all the certificate's information. If it is not there or the information is not correct, you will need to remove the corrupted certificate (if it is there) and upload it again from your PC.

### To use the new SSL certificate - CLI

```
config vpn ssl settings
 set servercert fgtssl
end
```

This assigns the `fgtssl` certificate as the SSL server certificate. For more information see the [FortiOS Handbook SSL VPN chapter](#).

# SSO using a FortiAuthenticator unit

If you use a FortiAuthenticator unit in your network as a single sign-on agent,

- Users can authenticate through a web portal on the FortiAuthenticator unit.
- Users with FortiClient Endpoint Security installed can be automatically authenticated by the FortiAuthenticator unit through the FortiClient SSO Mobility Agent.

The FortiAuthenticator unit can integrate with external network authentication systems such as RADIUS and LDAP to gather user logon information and send it to the FortiGate unit.

## User's view of FortiAuthenticator SSO authentication

There are two different ways users can authenticate through a FortiAuthenticator unit.

### Users without FortiClient Endpoint Security - SSO widget

To log onto the network, the user accesses the organization's web page with a web browser. Embedded on that page is a simple logon widget, like this:



User not logged in. Click *Login* to go to the FortiAuthenticator login page.



User logged in. Name displayed. *Logout* button available.

The SSO widget sets a cookie on the user's browser. When the user browses to a page containing the login widget, the FortiAuthenticator unit recognizes the user and updates its database if the user's IP address has changed. The user will not need to re-authenticate until the login timeout expires, which can be up to 30 days.

### Users with FortiClient Endpoint Security - FortiClient SSO Mobility Agent

The user simply accesses resources and all authentication is performed transparently with no request for credentials. IP address changes, such as those due to WiFi roaming, are automatically sent to the FortiAuthenticator unit. When the user logs off or otherwise disconnects from the network, the FortiAuthenticator unit is aware of this and deauthenticates the user.

The FortiClient SSO Mobility Agent, a feature of FortiClient Endpoint Security v5.0, must be configured to communicate with the appropriate FortiAuthenticator unit. After that, the agent automatically provides user name and IP address information to the FortiAuthenticator unit for transparent authentication.

## Administrator's view of FortiAuthenticator SSO authentication

You can configure either or both of these authentication types on your network.

### SSO widget

You need to configure the Single Sign-On portal on the FortiAuthenticator unit. Go to *SSO & Dynamic Policies > SSO > Login Portal* to do this. Copy the *Embeddable login widget* code for use on your organization's home page. Identity-based security policies on the FortiGate unit determine which users or groups of users can access which network resources.

### FortiClient SSO Mobility Agent

Your users must be running FortiClient Endpoint Security v5.0 to make use of this type of authentication.

On the FortiAuthenticator unit, you need to enable *FortiClient Service* when you define the unit's secret key. Go to *SSO & Dynamic Policies > SSO > Options*. You need to provide your users the FortiAuthenticator IP address and secret key so that they can configure the FortiClient SSO Mobility Agent on their computers. See "[Configuring the FortiGate unit](#)" on page 552.

## Configuring the FortiAuthenticator unit

The FortiAuthenticator unit can poll FortiGate units, Windows Active Directory, RADIUS servers, LDAP servers, and FortiClients for information about user logon activity.

### To configure FortiAuthenticator polling

1. Go to *SSO & Dynamic Policies > SSO > Options*.
2. In the *FortiGate* section, leave the Listening Port at 8000, unless your network requires you to change this. The FortiGate unit must allow traffic on this port to pass through the firewall. Optionally, you can set the Login Expiry time. This is the length of time users can remain logged in before the system logs them off automatically. The default is 480 minutes (8 hours).
3. Select *Enable Authentication* and enter the *Secret key*. Be sure to use the same secret key when configuring the FSSO Agent on FortiGate units.
4. In the *Fortinet Single Sign-On (FSSO)* section, enter

---

**Enable Windows Active Directory domain controllers** Select for integration with Windows Active Directory.

---

**Enable Radius accounting service** Select if you want to use a Remote LDAP server.

---

**Use remote LDAP server for SSO groups lookup** Optionally, you can provide SSO only to certain groups. If so, enable this option and select the remote LDAP server.

---

**Enable FortiClient service Enable Authentication** Select both options to enable single sign-on by clients running FortiClient Endpoint Security. Enter the *Secret key*. Be sure to use the same secret key in the FortiClient Single Sign-On Mobility Agent settings.

---

5. Select *OK*.

For more information, see the [FortiAuthenticator Administration Guide](#).

## Configuring the FortiGate unit

### Adding a FortiAuthenticator unit as an SSO agent

On the FortiGate unit, you need to add the FortiAuthenticator unit as a Single Sign-On agent that provides user logon information.

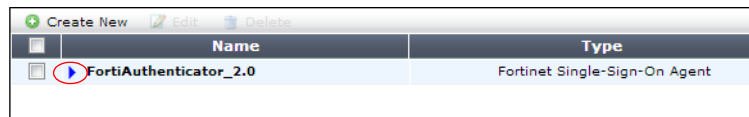
#### To add a FortiAuthenticator unit as SSO agent

1. Go to *User & Device > Authentication > Single Sign-On* and select *Create New*.
2. In *Type*, select *Fortinet Single-Sign-On Agent*.
3. Enter a *Name* for the FortiAuthenticator unit.
4. In *Primary Agent IP/Name*, enter the IP address of the FortiAuthenticator unit.
5. In *Password*, enter the secret key that you defined for the FortiAuthenticator unit.

On the FortiAuthenticator unit, you go to *SSO & Dynamic Policies > SSO > Options* to define the secret key. Select *Enable Authentication*.

6. Select *OK*.

In a few minutes, the FortiGate unit receives a list of user groups from the FortiAuthenticator unit. The entry in the Single Sign-On server list shows a blue caret.



	Name	Type
<input type="checkbox"/>	FortiAuthenticator_2.0	Fortinet Single-Sign-On Agent

When you open the server, you can see the list of groups. You can use the groups in identity-based security policies.

### Configuring an FSSO user group

You cannot use FortiAuthenticator SSO user groups directly in a security policy. Create an FSSO user group and add FortiAuthenticator SSO user groups to it. FortiGate FSSO user groups are available for selection in identity-based security policies.

#### To create an FSSO user group

1. Go to *User & Device > User > User Groups* and select *Create New*.
2. Enter a *Name* for the group.
3. In *Type*, select *Fortinet Single Sign-On (FSSO)*.
4. Add *Available Members* to the *Members* list.  
The Available Members are SSO groups provided by SSO agents.
5. Select *OK*.

### Configuring security policies

You can create identity-based policies based on FSSO groups as you do for local user groups. For more information about security policies see the Firewall chapter.



## Configuring the FortiClient SSO Mobility Agent

The user's device must have FortiClient Endpoint Security v5.0 installed. Only two pieces of information are required to set up the SSO Mobility Agent feature: the FortiAuthenticator unit IP address and the preshared secret.

The user needs to know the FortiAuthenticator IP address and preshared secret to set up the SSO Mobility Agent. Or, you could preconfigure FortiClient

### To configure FortiClient SSO Mobility Agent

1. In FortiClient Endpoint Security, go to *File > Settings*.

You must run the FortiClient application as an administrator to access these settings.

2. Select *Enable single sign-on mobility agent*. Enter the FortiAuthenticator unit IP address, including the listening port number specified on the FortiAuthenticator unit.

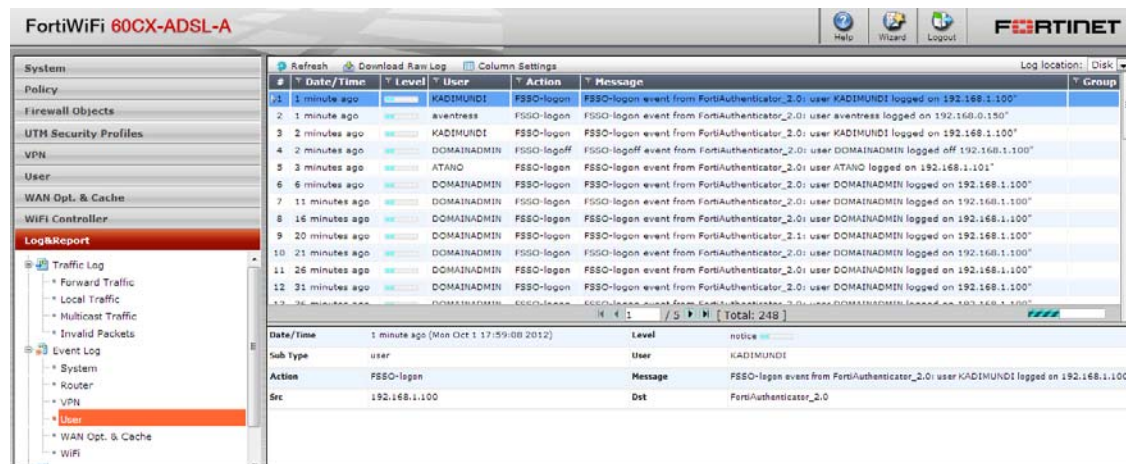
Example: 192.168.0.99:8001. You can omit the port number if it is 8005.

3. Enter the preshared key.

Select OK.

## Viewing SSO authentication events on the FortiGate unit

User authentication events are logged in the FortiGate event log. Go to *Log & Report > Event Log > User*.



The screenshot displays the FortiGate WebUI interface for a FortiWiFi 60CX-ADSL-A unit. The left sidebar shows the navigation menu with 'Log & Report' selected, and 'Event Log' > 'User' highlighted. The main panel shows a table of authentication events. Below the table, a detailed view of a selected event is shown.

#	Date/Time	Level	User	Action	Message	Group
1	1 minute ago	notice	KADIMUNDI	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user KADIMUNDI logged on 192.168.1.100*	
2	1 minute ago	notice	aventress	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user aventress logged on 192.168.0.150*	
3	2 minutes ago	notice	KADIMUNDI	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user KADIMUNDI logged on 192.168.1.100*	
4	2 minutes ago	notice	DOMAINADMIN	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user DOMAINADMIN logged on 192.168.1.100*	
5	3 minutes ago	notice	ATANO	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user ATANO logged on 192.168.1.101*	
6	6 minutes ago	notice	DOMAINADMIN	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user DOMAINADMIN logged on 192.168.1.100*	
7	11 minutes ago	notice	DOMAINADMIN	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user DOMAINADMIN logged on 192.168.1.100*	
8	16 minutes ago	notice	DOMAINADMIN	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user DOMAINADMIN logged on 192.168.1.100*	
9	20 minutes ago	notice	DOMAINADMIN	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.1: user DOMAINADMIN logged on 192.168.1.100*	
10	21 minutes ago	notice	DOMAINADMIN	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user DOMAINADMIN logged on 192.168.1.100*	
11	26 minutes ago	notice	DOMAINADMIN	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user DOMAINADMIN logged on 192.168.1.100*	
12	31 minutes ago	notice	DOMAINADMIN	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user DOMAINADMIN logged on 192.168.1.100*	
13	36 minutes ago	notice	DOMAINADMIN	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user DOMAINADMIN logged on 192.168.1.100*	

Date/Time	Level	Message
1 minute ago (Mon Oct 1 17:59:00 2012)	notice	FSSO-logout event from FortiAuthenticator_2.0: user KADIMUNDI logged on 192.168.1.100*

Sub Type	User
user	KADIMUNDI

Action	Message
FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user KADIMUNDI logged on 192.168.1.100*

Src	Dst
192.168.1.100	FortiAuthenticator_2.0

# Single Sign-On to Windows AD

The FortiGate unit can authenticate users transparently and allow them network access based on their privileges in Windows AD. This means that users who have logged on to the network are not asked again for their credentials to access network resources through the FortiGate unit, hence the term “Single Sign-On”.

The following topics are included:

- [Introduction to Single Sign-On with Windows AD](#)
- [Configuring Single Sign On to Windows AD](#)
- [FortiOS FSSO log messages](#)
- [Testing FSSO](#)
- [Troubleshooting FSSO](#)

## Introduction to Single Sign-On with Windows AD

Introduced in FortiOS 5.0, Single Sign-On (SSO) support provided by FortiGate polling of domain controllers is simpler than the earlier method that relies on agent software installed on Windows AD network servers. No Fortinet software needs to be installed on the Windows network. The FortiGate unit needs access only to the Windows AD global catalog and event log.

When a Windows AD user logs on at a workstation in a monitored domain, the FortiGate unit

- detects the logon event in the domain controller’s event log and records the workstation name, domain, and user,
- resolves the workstation name to an IP address,
- uses the domain controller’s LDAP server to determine which groups the user belongs to,
- creates one or more log entries on the FortiGate unit for this logon event as appropriate.

When the user tries to access network resources, the FortiGate unit selects the appropriate security policy for the destination. The selection consist of matching the FSSO group or groups the user belongs to with the security policy or policies that match that group. If the user belongs to one of the permitted user groups associated with that policy, the connection is allowed. Otherwise the connection is denied.

## Configuring Single Sign On to Windows AD

On the FortiGate unit, security policies control access to network resources based on user groups. With Fortinet Single Sign On, this is also true but each FortiGate user group is associated with one or more Windows AD user groups. This is how Windows AD user groups get authenticated in the FortiGate security policy.

Fortinet Single Sign On sends information about Windows user logons to FortiGate units. If there are many users on your Windows AD domains, the large amount of information might affect the performance of the FortiGate units.

To configure your FortiGate unit to operate with either a Windows AD or a Novell eDirectory FSSO install, you

- Configure LDAP access to the Windows AD global catalog. See [“Configuring LDAP server access” on page 555](#).
- Add Active Directory user groups to FortiGate FSSO user groups. See [“Creating Fortinet Single Sign-On \(FSSO\) user groups” on page 557](#).
- Configure the LDAP Server as a Single Sign-On server. See [“Configuring the LDAP Server as a Single Sign-On server” on page 557](#)
- Create security policies for FSSO-authenticated groups. See [“Creating security policies” on page 557](#).
- Optionally, specify a guest protection profile to allow guest access. See [“Enabling guest access through FSSO security policies” on page 559](#).

## Configuring LDAP server access

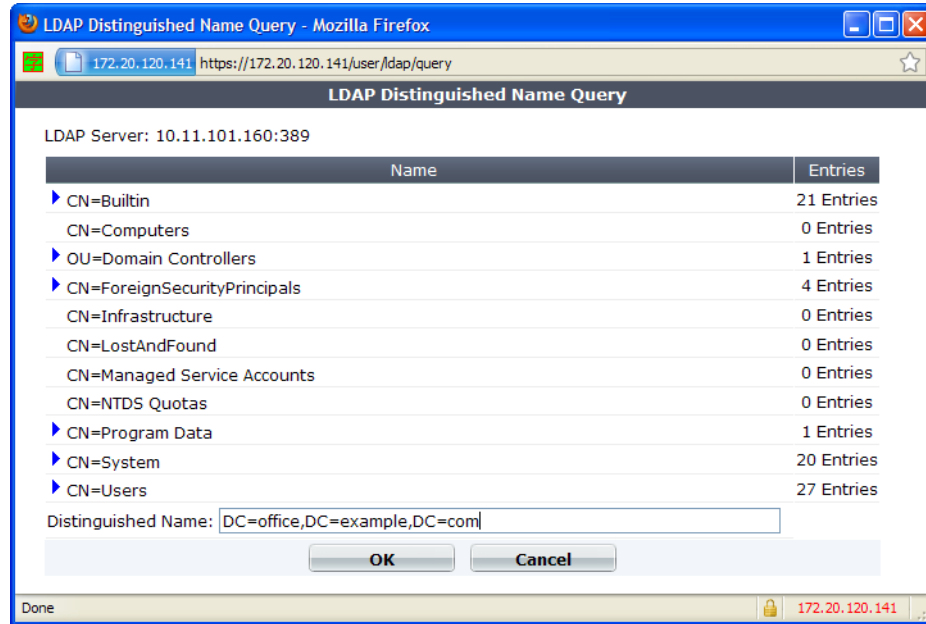
The FortiGate unit needs access to the domain controller’s LDAP server to retrieve user group information.

The LDAP configuration on the FortiGate unit not only provides access to the LDAP server, it sets up the retrieval of Windows AD user groups for you to select in FSSO. The LDAP Server configuration (in *User & Device > Authentication > LDAP Servers*) includes a function to preview the LDAP server’s response to your distinguished name query. If you already know the appropriate Distinguished Name (DN) and User DN settings, you may be able to skip some of the following steps.

### To add an LDAP server - web-based manager

1. Go to *User & Device > Authentication > LDAP Servers* and select *Create New*.
2. Enter the *Server Name/IP* and *Server Port* (default 389).
3. Select the *Query distinguished name* button to the right of the *Distinguished Name* field.  
A new window opens.
4. If more than one name is listed, you might need to explore each name following the steps below to determine which one is relevant to your needs.
5. Copy the name string to the *Distinguished Name* field and select *OK*.  
This closes the window and copies the name string to the *Distinguished Name* field of the LDAP Server configuration.
6. Set *Bind Type* to *Regular*.
7. In the *User DN* field, enter the administrative account name that you created for FSSO.  
For example, if the account is FSSO\_Admin, enter “cn=FSSO\_Admin,cn=users”.
8. Make sure that the *User DN* entry ends with a comma and append the string from the *Distinguished Name* field to the end of it.  
Example: cn=FSSO\_Admin,cn=users,dc=office,dc=example,dc=com
9. Enter the administrative account password in the *Password* field.
10. Select the *Query distinguished name* button again.  
The LDAP Distinguished Name Query window opens:

**Figure 117: Authenticated DN query**



You can expand any of the DNs that contain entries. When you select an expandable DN, the *Distinguished Name* field is updated. Look for the DN that contains the users or groups whose logon you want to monitor.

11. Select the DN that you want to monitor and then select *OK*.

This closes the window and updates the *Distinguished Name* field of the LDAP Server configuration with the selected Domain Name Identifier (DNI).

12. Check the following fields and select *OK*:

<b>Name</b>	Enter a name to identify the LDAP server.
<b>Common Name Identifier</b>	The default common name identifier is <code>cn</code> . This is correct for most LDAP servers. However some servers use other identifiers such as <code>uid</code> .
<b>Secure Connection</b>	Optional.

#### To configure LDAP for FSSO - CLI example

```

config user ldap
 edit "ADserver"
 set server "10.11.101.160"
 set cnid "cn"
 set dn "cn=users,dc=office,dc=example,dc=com"
 set type regular
 set username
 "cn=administrator,cn=users,dc=office,dc=example,dc=com"
 set password set_a_secure_password
 next
end

```

## Creating Fortinet Single Sign-On (FSSO) user groups

You cannot use Windows or Novell groups directly in FortiGate security policies. You must create FortiGate user groups of the FSSO type and add Windows or Novell groups to them.

### To create a user group for FSSO authentication - web-based manager

1. Go to *User & Device > User > User Groups*. and select *Create New*.  
The New User Group dialog box opens.
2. In the *Name* box, enter a name for the group, FSSO\_Internet\_users for example.
3. In *Type*, select *Fortinet Single Sign-On (FSSO)*.
4. From the *Available Members* list, select the required FSSO groups.  
Using the CTRL or SHIFT keys, you can select multiple groups.
5. Select the green right arrow button to move the selected groups to the *Members* list.
6. Select *OK*.

### To create the FSSO\_Internet-users user group - CLI

```
config user group
 edit FSSO_Internet_users
 set group-type fss-service
 set member CN=Engineering,cn=users,dc=office,dc=example,dc=com
 CN=Sales,cn=users,dc=office,dc=example,dc=com
 end
```

## Configuring the LDAP Server as a Single Sign-On server

The LDAP server must be added to the FortiGate Single Sign-On configuration.

### To add the LDAP server as a Single Sign-On server

1. Go to *User & Device > Authentication > Single Sign-On* and select *Create New*.
2. Enter

<b>Type</b>	Poll Active Directory Server
<b>Server Name/IP</b>	Server Name or IP address of the Domain Controller
<b>User</b>	A Domain user name
<b>Password</b>	The user's password
<b>LDAP Server</b>	Select the LDAP server you added earlier.
<b>Enable Polling</b>	Select

3. Select *OK*.

## Creating security policies

Policies that require FSSO authentication are very similar to other security policies. Using identity-based policies, you can configure access that depends on the FSSO user group. This allows each FSSO user group to have its own level of access to its own group of services

In this situation, Example.com is a company that has its employees and authentication servers on an internal network. The FortiGate unit intercepts all traffic leaving the internal network and

requires FSSO authentication to access network resources on the Internet. The following procedure configures the security policy for FSSO authentication. FSSO is installed and configured including the RADIUS server, FSSO Collector agent, and user groups on the FortiGate

For the following procedure, the internal interface is `port1` and the external interface connected to the Internet is `port2`. There is an address group for the internal network called `company_network`. The FSSO user group is called `fsso_group`, and the FSSO RADIUS server is `fsso_rad_server`.

### To configure an FSSO authentication security policy - web-based manager

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Enter the following information.

<b>Policy Type</b>	Firewall
<b>Policy subtype</b>	User Identity
<b>Incoming Interface</b>	port1
<b>Source Address</b>	company_network
<b>Outgoing Interface</b>	port2
<b>Enable NAT</b>	Select

3. In *Configure Authentication Rules*, select *Create New*.
4. Enter

<b>Destination Address</b>	all
<b>Groups</b>	Select from the FSSO user groups that you created earlier.  FSSO_Guest_users is a default user group enabled when FSSO is configured. It allows guest users on the network who do not have an FSSO account to still authenticate and have access to network resources. See <a href="#">“Enabling guest access through FSSO security policies” on page 559</a> .
<b>Schedule</b>	always
<b>Service</b>	HTTP, HTTPS, FTP, and Telnet
<b>Action</b>	ACCEPT
<b>Log Allowed Traffic</b>	Select. Logging FSSO logon events helps troubleshoot any FSSO related issues.
<b>UTM Security Profiles</b>	Enable AntiVirus, IPS, Web Filter, and Email Filter default profiles.

5. Select *OK*.  
A new line of information will appear in the identity-based policy table, listing the user groups, services, schedule, UTM, and logging selected for the rule.
6. Select *OK*.
7. Ensure the FSSO authentication policy is higher in the policy list than more general policies for the same interfaces.

## To create a security policy for FSSO authentication - CLI

```
config firewall policy
 edit 0
 set srcintf internal
 set dstintf wan1
 set srcaddr company_network
 set dstaddr all
 set action accept
 set identity-based enable
 set nat enable
 config identity-based-policy
 edit 1
 set schedule any
 set groups company_network FSSO_guest_users
 set service HTTP HTTPS FTP TELNET
 end
 end
 end
```

Here is an example of how this FSSO authentication policy is used. Example.com employee on the internal company network logs on to the internal network using their RADIUS username and password. When that user attempts to access the Internet, which requires FSSO authentication, the FortiGate authentication security policy intercepts the session, checks with the FSSO Collector agent to verify the user's identity and credentials, and then if everything is verified the user is allowed access to the Internet.

## Enabling guest access through FSSO security policies

You can enable guest users to access FSSO security policies. Guests are users who are unknown to Windows AD and servers that do not logon to a Windows AD domain.

To enable guest access in your FSSO security policy, add an identity-based policy assigned to the built-in user group `FSSO_Guest_Users`. Specify the services, schedule and UTM profiles that apply to guest users — typically guests have access to a reduced set of services. See [“Creating security policies” on page 557](#).

## FortiOS FSSO log messages

There are two types of FortiOS log messages — firewall and event. FSSO related log messages are generated from authentication events. These include user logon and log off events, and NTLM authentication events. These log messages are central to network accounting policies, and can also be useful in troubleshooting issues. For more information on firewall logging, see [“Enabling security logging” on page 522](#). For more information on logging, see the [FortiOS Handbook Log and Reporting chapter](#).

## Enabling authentication event logging

For the FortiGate unit to log events, that specific type of event must be enabled under logging.

When VDOMs are enabled certain options may not be available, such as CPU and memory usage events. You can enable event logs only when you are logged on to a VDOM; you cannot enable event logs globally.

To ensure you log all the events need, set the minimum log level to Notification or Information. Firewall logging requires Notification as a minimum. The closer to Debug level, the more information will be logged. While this extra information is useful, you must

### To enable event logging

1. Go to *Log&Report > Log Config > Log Settings*.
2. In *Event Logging*, select

<b>System activity event</b>	All system-related events, such as ping server failure and gateway status.
<b>User activity event</b>	All administration events, such as user logins, resets, and configuration updates.

Optionally you can enable any or all of the other logging event options.

3. Select *Apply*.

**Figure 118:**Authentication log messages

The screenshot shows a log viewer interface with a table of log entries and a detailed view of the first entry. The log table has columns: #, Date, Time, Level, Sub Type, ID, Action, Message, and User Interface. The detailed view shows fields: Date (2011-05-02), Time (11:51:47), Level (notice), Sub Type (auth), ID (43011), Virtual Domain (root), Src (172.20.120.230), Dst (N/A), User (test), Group (testee), Policy ID (4), User Interface (test(172.20.120.230)), Action (authentication), Status (timed\_out), Reason (Authentication timed out), and Message (User from 172.20.120.230 was timed out).

#	Date	Time	Level	Sub Type	ID	Action	Message	User Interface
1	2011-05-02	11:51:47	notice	auth	43011	authentication	User from 172.20.120.230 was timed out	test(172.20.120.230)
2	2011-05-02	11:43:31	notice	auth	43008	authentication	User test succeeded in authentication	HTTPS(172.20.120.230)
3	2011-05-02	11:43:23	notice	auth	43009	authentication	User techdoc failed in authentication	HTTPS(172.20.120.230)

Date	2011-05-02	Time	11:51:47
Level	notice	Sub Type	auth
ID	43011	Virtual Domain	root
Src	172.20.120.230	Dst	N/A
User	test	Group	testee
Policy ID	4	User Interface	test(172.20.120.230)
Action	authentication	Status	timed_out
Reason	Authentication timed out	Message	User from 172.20.120.230 was timed out

**Table 24:**List of FSSO related log messages

Message ID	Severity	Description
43008	Notification	Authentication was successful
43009	Notification	Authentication session failed
43010	Warning	Authentication locked out
43011	Notification	Authentication timed out
43012	Notification	FSSO authentication was successful
43013	Notification	FSSO authentication failed
43014	Notification	FSSO user logged on
43015	Notification	FSSO user logged off
43016	Notification	NTLM authentication was successful
43017	Notification	NTLM authentication failed



For more information on logging, see the [FortiOS Handbook Log and Reporting chapter](#).

## Testing FSSO

Once FSSO is configured, you can easily test to ensure your configuration is working as expected. For additional FSSO testing, see [“Troubleshooting FSSO” on page 561](#).

1. Logon to one of the stations on the FSSO domain, and access an Internet resource.
2. Connect to the CLI of the FortiGate unit, and if possible log the output.
3. Enter the following command:
4. Check the output. If FSSO is functioning properly you will see something similar to the following:

```
diagnose debug authd fssso list
----FSSO logons----
IP: 192.168.1.230 User: ADMINISTRATOR Groups: VLAD-AD/DOMAIN USERS
IP: 192.168.1.240 User: ADMINISTRATOR Groups: VLAD-AD/DOMAIN USERS
Total number of users logged on: 2
----end of FSSO logons----
```

The exact information will vary based on your installation.

5. Check the FortiGate event log, for FSSO-auth action or other FSSO related events with FSSO information in the message field. For a list of FSSO log message IDs, see [Table 24 on page 560](#).
6. To check server connectivity, run the following commands from the CLI:

```
FGT# diagnose debug enable
FGT# diagnose debug authd fssso server-status
FGT# Server Name Connection Status

SBS-2003 connected
```

## Troubleshooting FSSO

When installing, configuring, and working with FSSO some problems are quite common. A selection of these problems follows including explanations and solutions.

Some common Windows AD problems include:

- [General troubleshooting tips for FSSO](#)
- [Users on a particular computer \(IP address\) can not access the network](#)
- [Guest users do not have access to network](#)

### General troubleshooting tips for FSSO

The following tips are useful in many FSSO troubleshooting situations.

- Ensure all firewalls are allowing the FSSO required ports through.  
FSSO has a number of required ports that must be allowed through all firewalls or connections will fail. These include: ports 139, 389 (LDAP), 445, 636 (LDAP).
- Ensure there is at least 64kbps bandwidth between the FortiGate unit and domain controllers. If there is insufficient bandwidth, some FSSO information might not reach the FortiGate unit. The best solution is to configure traffic shaping between the FortiGate unit and the domain controllers to ensure that the minimum bandwidth is always available.

## Users on a particular computer (IP address) can not access the network

Windows AD Domain Controller agent gets the username and workstation where the logon attempt is coming from. If there are two computers with the same IP address and the same user trying to logon, it is possible for the authentication system to become confused and believe that the user on computer\_1 is actually trying to access computer\_2.

Windows AD does not track when a user logs out. It is possible that a user logs out on one computer, and immediately logs onto a second computer while the system still believes the user is logged on the original computer. While this is allowed, information that is intended for the session on one computer may mistakenly end up going to the other computer instead. The result would look similar to a hijacked session.

### Solutions

- Ensure each computer has separate IP addresses.
- Encourage users to logout on one machine before logging onto another machine.
- If multiple users have the same username, change the usernames to be unique.
- Shorten timeout timer to flush inactive sessions after a shorter time.

## Guest users do not have access to network

A group of guest users was created, but they don't have access.

### Solution

The group of the guest users was not included in a policy, so they do not fall under the guest account. To give them access, associate their group with a security policy.

Additionally, there is a default group called `FSSO_Guest_Users`. Ensure that group is part of an identity-based security policy to allow traffic.

# Agent-based FSSO

FortiOS can provide single sign-on capabilities to Windows AD, Citrix, or Novell eDirectory users with the help of agent software installed on these networks. The agent software sends information about user logons to the FortiGate unit. With user information such as IP address and user group memberships from the network, FortiGate security policies can allow authenticated network access to users who belong to the appropriate user groups without requesting their credentials again.

For Windows AD networks, FortiGate units can provide SSO capability without agent software by directly polling the Windows AD domain controllers. For information about this type of SSO, see [“Single Sign-On to Windows AD” on page 554](#).

The following topics are included:

- [Introduction to agent-based FSSO](#)
- [FSSO NTLM authentication support](#)
- [Agent installation](#)
- [Configuring the FSSO Collector agent for Windows AD](#)
- [Configuring the FSSO TS agent for Citrix](#)
- [Configuring the FSSO eDirectory agent for Novell eDirectory](#)
- [Configuring FSSO on FortiGate units](#)
- [FortiOS FSSO log messages](#)
- [Testing FSSO](#)
- [Troubleshooting FSSO](#)

## Introduction to agent-based FSSO

Fortinet Single Sign-On (FSSO), through agents installed on the network, monitors user logons and passes that information to the FortiGate unit. When a user logs on at a workstation in a monitored domain, FSSO

- detects the logon event and records the workstation name, domain, and user,
- resolves the workstation name to an IP address,
- determines which user groups the user belongs to,
- sends the user logon information, including IP address and groups list, to the FortiGate unit
- creates one or more log entries on the FortiGate unit for this logon event as appropriate.

When the user tries to access network resources, the FortiGate unit selects the appropriate security policy for the destination. If the user belongs to one of the permitted user groups associated with that policy, the connection is allowed. Otherwise the connection is denied.



FSSO can also provide NTLM authentication service for requests coming from FortiGate. SSO is very convenient for users, but may not be supported across all platforms. NTLM is not as convenient, but it enjoys wider support. See [“FSSO NTLM authentication support” on page 569](#).

---

## Introduction to FSSO agents

There are several different FSSO agents that can be used in an FSSO implementation.

- [Domain Controller \(DC\) agent](#)
- [eDirectory agent](#)
- [Citrix/Terminal Server \(TS\) agent](#)
- [Collector \(CA\) agent](#)

Consult the latest FortiOS and FSSO Release Notes for operating system compatibility information.

### Domain Controller (DC) agent

The Domain Controller (DC) agent must be installed on every domain controller if you will use DC Agent mode, but is not required if you use Polling mode. See [“FSSO for Windows AD” on page 565](#).

### eDirectory agent

The eDirectory agent is installed on a Novell network to monitor user logons and send the required information to the FortiGate unit. It functions much like the Collector agent on a Windows AD domain controller. The agent can obtain information from the Novell eDirectory using either the Novell API or LDAP.

### Citrix/Terminal Server (TS) agent

The Citrix/Terminal Server (TS) agent is installed on a Citrix terminal server to monitor user logons in real time. It functions much like the DC Agent on a Windows AD domain controller.

### Collector (CA) agent

This agent is installed as a service on a server in the Windows AD network to monitor user logons and send the required information to the FortiGate unit. The Collector agent can collect information from

- Domain Controller agent (Windows AD)
- TS agent (Citrix Terminal Server)

In a Windows AD network, the Collector agent can optionally obtain logon information by polling the AD domain controllers. In this case, DC agents are not needed.

The Collector can obtain user group information from the DC agent or Optionally, a FortiGate unit can obtain group information directly from AD using Lightweight Directory Access Protocol (LDAP).

On a Windows AD network, the FSSO software can also serve NT LAN Manager (NTLM) requests coming from client browsers (forwarded by the FortiGate unit) with only one or more Collector agents installed. See [“FSSO NTLM authentication support” on page 569](#).

The CA is responsible for DNS lookups, group verification, workstation checks, and as mentioned FortiGate updates of logon records. The FSSO Collector Agent sends Domain Local Security Group and Global Security Group information to FortiGate units. The CA communicates with the FortiGate over TCP port 8000 and it listens on UDP port 8002 for updates from the DC agents.

The FortiGate unit can have up to five CAs configured for redundancy. If the first on the list is unreachable, the next is attempted, and so on down the list until one is contacted. See [“Configuring FSSO on FortiGate units” on page 589](#).

All DC agents must point to the correct Collector agent port number and IP address on domains with multiple DCs.

See “Configuring Collector agent settings” on page 577.



A FortiAuthenticator unit can act much like a Collector agent, collecting Windows AD user logon information and sending it to the FortiGate unit. It is particularly useful in large installations with several FortiGate units. For more information, see the [FortiAuthenticator Administration Guide](#).

## FSSO for Windows AD

FSSO for Windows AD requires at least one Collector agent. Domain Controller agents may also be required depending on the Collector agent working mode. There are two working modes to monitor user logon activity: DC Agent mode or Polling mode.

**Table 25:** Collector agent DC Agent mode versus Polling mode

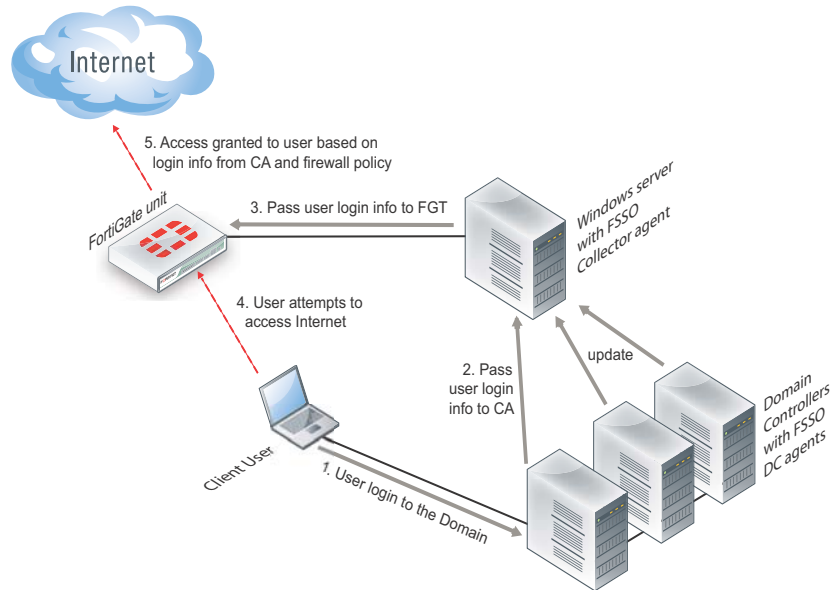
	DC Agent mode	Polling Mode
<b>Installation</b>	Complex — Multiple installations: one agent per DC plus Collector agent, requires a reboot	Easy — only Collector agent installation, no reboot required
<b>Resources</b>	Shares resources with DC system	Has own resources
<b>Network load</b>	Each DC agent requires minimum 64kpbs bandwidth, adding to network load	Increase polling period during busy period to reduce network load
<b>Level of Confidence</b>	Captures all logons	Potential to miss a login if polling period is too great

### DC Agent mode

DC Agent mode is the standard mode for FSSO. In DC Agent mode (see [Figure 119](#)), a Fortinet authentication agent is installed on each domain controller. These DC agents monitor user logon events and pass the information to the Collector agent, which stores the information and sends it to the FortiGate unit.

The DC agent installed on the domain controllers is not a service like the Collector agent — it is a DLL file called `dcagent.dll` and is installed in the `Windows\system32` directory. It must be installed on all domain controllers of the domains that are being monitored.

**Figure 119:**FSSO in DC agent mode



DC Agent mode provides reliable user logon information, however you must install a DC agent on every domain controller. A reboot is needed after the agent is installed. Each installation requires some maintenance as well. For these reasons it may not be possible to use the DC Agent mode.

Each domain controller connection needs a minimum guaranteed 64kbps bandwidth to ensure proper FSSO functionality. You can optionally configure traffic shapers on the FortiGate unit to ensure this minimum bandwidth is guaranteed for the domain controller connections.

### Polling mode


In Polling mode there are two options — NetAPI polling, and Event log polling. Both share the advantages of being transparent and agentless.

NetAPI polling is used to retrieve server logon sessions. This includes the logon event information for the Controller agent. NetAPI runs faster than Event log polling but it may miss some user logon events under heavy system load. It requires a query round trip time of less than 10 seconds.

Event log polling may run a bit slower, but will not miss events, even when the installation site has many users that require authentication. It does not have the 10 second limit or NetAPI polling. Event log polling requires fast network links. Event log polling is required if there are Mac OS users logging into Windows AD.

In Polling mode (see [Figure 120](#)), the Collector agent polls port 445 of each domain controller for user logon information every few seconds and forwards it to the FortiGate unit. There are no DC Agents installed, so the Collector agent polls the domain controllers directly.

**Figure 120:**FSSO in Polling mode



A major benefit of Polling mode is that no FSSO DC Agents are required. If it is not possible to install FSSO DC Agents on your domain controllers, this is the alternate configuration available to you. Polling mode results in a less complex install, and reduces ongoing maintenance. The minimum permissions required in Polling mode are to read the event log or call NetAPI. To install FSSO with minimum permissions, see [“Installing FSSO without using an administrator account” on page 574](#).

### Collector agent AD Access mode - Standard versus Advanced

The Collector agent has two ways to access Active Directory user information. The main difference between Standard and Advanced mode is the naming convention used when referring to username information.

Standard mode uses regular Windows convention: Domain\Username. Advanced mode uses LDAP convention: CN=User, OU=Name, DC=Domain.

If there is no special requirement to use LDAP— best practices suggest you set up FSSO in Standard mode. This mode is easier to set up, and is usually easier to maintain and troubleshoot.

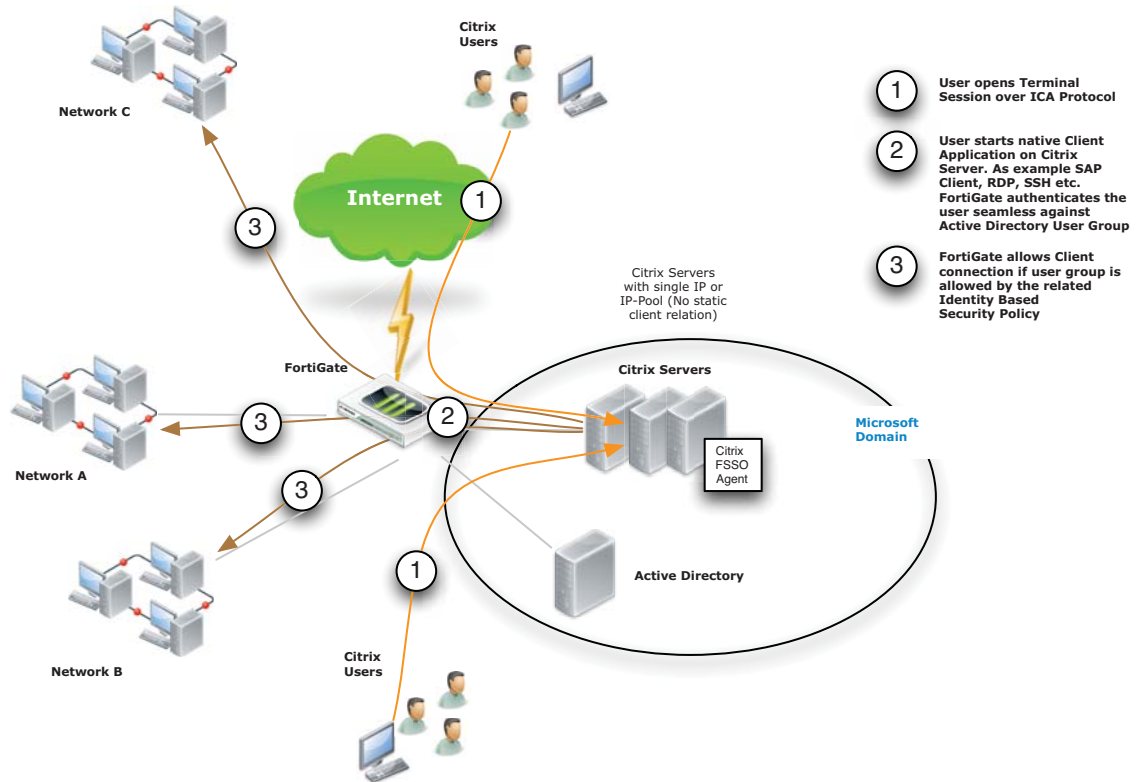
Standard and advanced modes have the same level of functionality with the following exceptions:

1. Users have to create Group filters on the Collector agent. This differs from Advanced mode where Group filters are configured from the FortiGate unit. Fortinet strongly encourages users to create filters from CA.
2. Advanced mode supports nested or inherited groups. This means that users may be a member of multiple monitored groups. Standard mode does not support nested groups so a user must be a direct member of the group being monitored.

### FSSO for Citrix

Citrix users can enjoy a similar Single Sign-On experience as Windows AD users. The FSSO TS agent installed on each Citrix server provides user logon information to the FSSO Collector agent on the network. The FortiGate unit uses this information to authenticate the user in security policies.

**Figure 121:**Citrix SSO topology



Citrix users do not have unique IP addresses. When a Citrix user logs on, the TS agent assigns that user a range of ports. By default each user has a range of 200 ports.

## FSSO for Novell eDirectory

FSSO in a Novell eDirectory environment works similar to the FSSO Polling mode in the Windows AD environment. The eDirectory agent polls the eDirectory servers for user logon information and forwards the information to the FortiGate unit. There is no need for the Collector agent.

When a user logs on at a workstation, FSSO:

- detects the logon event by polling the eDirectory server and records the IP address and user ID,
- looks up in the eDirectory which groups this user belongs to,
- sends the IP address and user groups information to the FortiGate unit.

When the user tries to access network resources, the FortiGate unit selects the appropriate security policy for the destination. If the user belongs to one of the permitted user groups, the connection is allowed.

FSSO is supported on the Novell E-Directory 8.8 operating system.

For a Novell network, there is only one FSSO component to install — the eDirectory agent. In some cases, you also need to install the Novell Client.



## FSSO security issues

When the different components of FSSO are communicating there are some inherent security features.

FSSO installation requires an account with network admin privileges. The security inherent in these types of accounts helps ensure access to FSSO configurations is not tampered with.

User passwords are never sent between FSSO components. The information that is sent is information to identify a user including the username, group or groups, and IP address.

NTLM uses base-64 encoded packets, and uses a unique randomly generated challenge nonce to avoid sending user information and password between the client and the server. For more information on NTLM, see [“FSSO NTLM authentication support” on page 569](#).

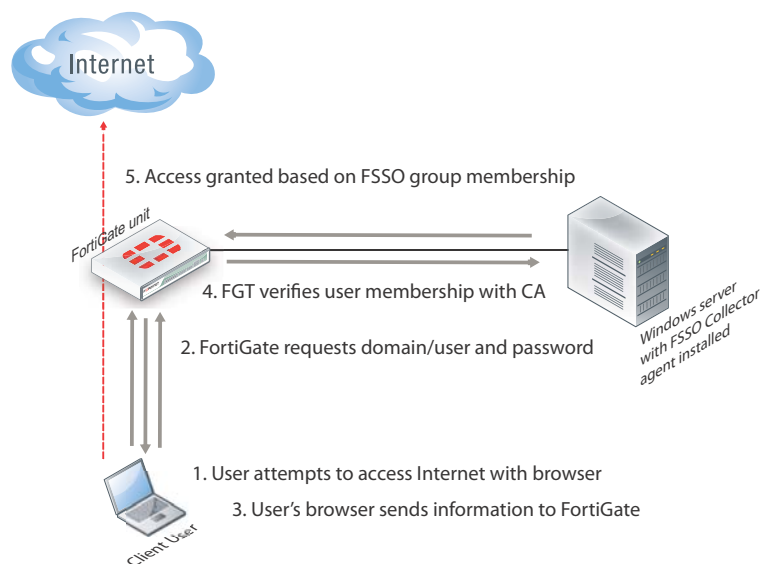
## FSSO NTLM authentication support

In a Windows AD network, FSSO can also provide NTLM authentication service to the FortiGate unit. When the user makes a request that requires authentication, the FortiGate unit initiates NTLM negotiation with the client browser. The FortiGate unit does not process the NTLM packets itself. Instead, it forwards all the NTLM packets to the FSSO service to process.

NTLM has the benefit of not requiring an FSSO agent, but it is not transparent to users, and the user’s web browser must support NTLM.

The NTLM protocol protects the user’s password by not sending it over the network. Instead, the server sends the client a random number that the client must encrypt with the hash value of the user’s password. The server compares the result of the client’s encryption with the result of its own encryption. The two will match only if both parties used the same password.

**Figure 122:**NTLM authentication



If the NTLM authentication with the Windows AD network is successful, and the user belongs to one of the groups permitted in the applicable security policy, the FortiGate unit allows the connection.

Fortinet has tested NTLM authentication with Internet Explorer and Firefox browsers.


## NTLM in a multiple domain environment

In a multiple domain environment for NTLM, the important factor is that there is a trust relation between the domains. In a forest, this relation is automatically created. So you can install FSSO agent on one of the domain controllers without worry.

But in case of multiple domains that are not in a forest, you need to create a trust relation between the domains. If you do not want to have a trust relation between your multiple domains, you need to use FSAE 4.0 MR1 and the DC agent needs to be installed once on each domain. Then you can use security policies to configure server access.

In [Figure 123](#), three domains are shown connected to the FSSO Collector agent server. The Client logs on to their local Domain Controller, which then sends the user logon event information to the Collector Agent. When the Client attempts to access the Internet, the FortiGate unit contacts the Collector Agent for the logon information, sees the Client is authenticated, and allows access to the Internet. There are multiple domains each with a domain controller agent (DCagent) that sends logon information to the Collector agent. If the multiple domains have a trust relationship, only one DCagent is required instead of one per domain.

**Figure 123:**FSSO NTLM with multiple domains not in a forest



### Understanding the NTLM authentication process

1. The user attempts to connect to an external (internet) HTTP resource. The client application (browser) on the user's computer issues an unauthenticated request through the FortiGate unit.
2. The FortiGate is aware that this client has not authenticated previously, so responds with a 401 Unauthenticated status code, and tells the client which authentication method to reply with in the header: Proxy-Authenticated: NTLM. Then the initial session is dismantled.
3. The client application connects again to the FortiGate, and issues a GET-request, with a Proxy-Authorization: NTLM <negotiate string> header. <negotiate-string> is a base64-encoded NTLM Type 1 negotiation packet.

4. The FortiGate unit replies with a 401 "proxy auth required" status code, and a `Proxy-Authenticate: NTLM <challenge string>` (a base 64-encoded NTLM Type 2 challenge packet). In this packet is the challenge nonce, a random number chosen for this negotiation that is used once and prevents replay attacks.



The TCP connection must be kept alive, as all subsequent authentication-related information is tied to the TCP connection. If it is dropped, the authentication process must start again from the beginning.

- 
5. The client sends a new GET-request with a header: `Proxy-Authenticate: NTLM <authenticate string>`, where `<authenticate string>` is a NTLM Type 3 Authentication packet that contains:
    6. username and domain
    7. the challenge nonce encoded with the client password (it may contain the challenge nonce twice using different algorithms).

If the negotiation is successful and the user belongs to one of the groups permitted in the security policy, the connection is allowed, Otherwise, the FortiGate unit denies the authentication by issuing a 401 return code and prompts for a username and password. Unless the TCP connection is broken, no further credentials are sent from the client to the proxy.



If the authentication policy reaches the authentication timeout period, a new NTLM handshake occurs.

---

## Agent installation

After reading the appropriate sections of ["Introduction to agent-based FSSO" on page 563](#) to determine which FSSO agents you need, you can proceed to perform the necessary installations.

Ensure you have administrative rights on the servers where you are installing FSSO agents. It is best practice to install FSSO agents using the built-in local administrator account. Optionally, you can install FSSO without an admin account. See ["Installing FSSO without using an administrator account" on page 574](#).



In Windows 2008 by default, you do not have administrative user rights if you are logged on as a user other than as the built-in administrator, even if you were added to the local Administrators group on the computer.

---

The FSSO installer first installs the Collector agent. You can then continue with installation of the DC agent, or you can install it later by going to *Start > Programs > Fortinet >*

*Fortinet Single Sign On Agent > Install DC Agent.* The installer will install a DC agent on the domain controllers of all of the trusted domains in your network.



Each domain controller connection needs a minimum guaranteed 64kpbs bandwidth to ensure proper FSSO functionality. Traffic shapers configured on the FortiGate can help guarantee these minimum bandwidths.

## Collector agent installation

To install FSSO, you must obtain the FSSO\_Setup file from the [Fortinet Support web site](#). This is available as either an executable (.exe) or a Microsoft Installer (.msi) file. Then you follow these two installation procedures on the server that will run the Collector agent. This can be any server or domain controller that is part of your network. These procedures also install the DC Agent on all of the domain controllers in your network.

### To install the Collector agent

1. Create an account with administrator privileges and a password that does not expire. See Microsoft Advanced Server documentation for help with this task.

To use a non-admin read only account, see [“Installing FSSO without using an administrator account” on page 574](#).

2. Log on to the account that you created in Step 1.
3. Double-click the FSSOSetup.exe file.
4. The Fortinet SSO Collector Agent Setup Wizard starts.
5. Select *Next*.
6. Read and accept the license agreement. Select *Next*.
7. Optionally, you can change the installation location. Select *Next*.
8. Optionally, change the *User Name*.

By default, the agent is installed using the currently running account. If you want FSSO to use another existing admin account, change the *User Name* using the format `DomainName \UserName`. For example if the account is `jsmith` and the domain is `example_corp` you would enter `example_corp\jsmith`.

9. In the *Password* field, enter the password for the account listed in the *User Name* field.
10. Select *Next*.
11. Enable as needed:
  - *Monitor user logon events and send the information to the FortiGate unit*
  - *Serve NTLM authentication requests coming from FortiGate*

By default, both methods are enabled. You can change these options after installation.

12. Select the access method to use for Windows Directory:
  - Select *Standard* to use Windows domain and username credentials.
  - Select *Advanced* if you will set up LDAP access to Windows Directory.

See [“Collector agent AD Access mode - Standard versus Advanced” on page 567](#).

13. Select *Next* and then select *Install*.
14. If you want to use DC Agent mode, ensure that *Launch DC Agent Install Wizard* is selected. This will start DC agent installation immediately after you select *Finish*.
15. Select *Finish*.



If you see an error such as Service Fortinet Single Sign On agent (service\_FSAE) failed to start, there are two possible reasons for this. Verify the user account you selected has sufficient privileges to run the FSSO service. Also verify the computer system you are attempting to install on is a supported operating system and version.

---

## DC agent installation

The FSSO\_Setup file contains both the Collector agent and DC Agent installers, but the DC Agent installer is also available separately as either a .exe or .msi file named DCAgent\_Setup.

### To install the DC Agent

1. If you have just installed the Collector agent, the FSSO - Install DC Agent wizard starts automatically. Otherwise, go to *Start > Programs > Fortinet > Fortinet Single Sign On Agent > Install DC Agent*.
2. Select *Next*.
3. Read and accept the license agreement. Select *Next*.
4. Optionally, you can change the installation location. Select *Next*.
5. Enter the *Collector agent IP address*.
6. If the Collector agent computer has multiple network interfaces, ensure that the one that is listed is on your network. The listed *Collector agent listening port* is the default. Only change this if the port is already used by another service.
7. Select *Next*.
8. Select the domains to monitor and select *Next*.
9. If any of your required domains are not listed, cancel the wizard and set up the proper trusted relationship with the domain controller. Then run the wizard again by going to *Start > Programs > Fortinet > Fortinet Single Sign On Agent > Install DC Agent*.
10. Optionally, select users that you do not want monitored. These users will not be able to authenticate to FortiGate units using FSSO. You can also do this later. See [“Configuring the FSSO Collector agent for Windows AD” on page 576](#).
11. Select *Next*.
12. Optionally, clear the check boxes of domain controllers on which you do not want to install the DC Agent.
13. Select the *Working Mode* as DC Agent Mode. While you can select Polling Mode here, in that situation you would not be installing a DC Agent. For more information, see [“DC Agent mode” on page 565](#) and [“Polling mode” on page 566](#).
14. Select *Next*.
15. Select *Yes* when the wizard requests that you reboot the computer.



If you reinstall the FSSO software on this computer, your FSSO configuration is replaced with default settings.

---

If you want to create a redundant configuration, repeat the procedure “[To install the Collector agent](#)” on page 572 on at least one other Windows AD server.



When you start to install a second Collector agent, cancel the Install Wizard dialog appears the second time. From the configuration GUI, the monitored domain controller list will show your domain controllers un-selected. Select the ones you wish to monitor with this Collector agent, and select *Apply*.

Before you can use FSSO, you need to configure it on both Windows AD and on the FortiGate units. “[Configuring FSSO on FortiGate units](#)” on page 589 will help you accomplish these two tasks.

### Installing FSSO without using an administrator account

Normally when installing services in Windows, it is best to use the Domain Admin account, as stated earlier. This ensures installation goes smoothly and uninterrupted, and when using the FSSO agent there will be no permissions issues. However, it is possible to install FSSO with a non-admin account in Windows 2003 or 2008 AD.



The following instructions for Windows 2003 are specific to the event log polling mode only. Do not use this procedure with other FSSO configurations.

#### Windows 2003

There are two methods in Windows 2003 AD for installing FSSO without an admin account — add the non-admin user to the security log list, and use a non-admin account with read-only permissions. A problem with the first method is that full rights (read, write, and clear) are provided to the event log. This can be a problem when audits require limited or no write access to logs. In those situations, the non-admin account with read-only permissions is the solution.

#### To add the non-admin user account to the Windows 2003 security log list

1. Go to *Default Domain Controller Security Settings > Security Settings > User Rights Assignment > Manage auditing and security log*.
2. Add the user account to this list.
3. Repeat these steps on every domain controller in Windows 2003 AD.
4. A reboot is required.

#### To use a non-admin account with read-only permissions to install FSSO on Windows 2003

The following procedure provides the user account specified with read only access to the Windows 2003 AD Domain Controller Security Event Log which allows FSSO to function.

1. Find out the SID of the account you intend to use.

Tools for this can be downloaded for free from <http://technet.microsoft.com/en-us/sysinternals/bb897417>.

2. Then create the permission string. For example:

```
(A; ; 0x1 ; ; ; S-1-5-21-4136056096-764329382-1249792191-1107)
```

A means Allow,

0x1 means Read, and

S-1-5-21-4136056096-764329382-1249792191-1107 is the SID.

3. Then, append it to the registry key

4. HKEY\_LOCAL\_MACHINE\System\CurrentControlSet\Services\Eventlog\Security\CustomSD.
5. Repeat these steps on every domain controller in Windows 2003 AD.
6. A reboot is required.

### Windows 2008

In Windows 2008 AD, if you do not want to use the Domain Admin account then the user account that starts the FSSO agent needs to be added to the Event Log Readers group.

When the user is added to the Event Log Readers group, that user is now allowed to have read only access to the event log and this is the minimal rights required for FSSO to work.

## Citrix TS agent installation

To install the Citrix TS agent, you must obtain the TSAgent\_Setup file from the [Fortinet Support web site](#). Perform the following installation procedure on the Citrix server.

### To install the FSSO TS agent

1. On the Citrix server, create an account with administrator privileges and a password that does not expire. See Citrix documentation for more information.
2. Log on to the account that you created in Step 1.
3. Double-click the TSAgent\_Setup installation file.  
The Fortinet SSO Terminal Server Agent Setup Wizard starts.
4. Select *Next*.
5. Read and accept the license agreement. Select *Next*.
6. Optionally, you can change the installation location. Select *Next*.
7. Verify that *This Host IP Address* is correct.
8. In the *FSSO Collector Agent List*, enter the IP address(es) of your Collector Agents.
9. Select *Next* and then select *Install*.  
The TS agent is installed.
10. Select *Finish*.

## Novell eDirectory agent installation

To install the eDirectory agent, you must obtain the FSSO\_Setup\_eDirectory file from the [Fortinet Support web site](#). Perform the following installation procedure on the computer that will run the eDirectory agent. This can be any server or domain controller that is part of your network. You will need to provide some setup information in step 7.

### To install the FSSO eDirectory agent

1. Create an account with administrator privileges and a password that does not expire. See Novell documentation for more information.
2. Log on to the account that you created in Step 1.
3. Double-click the FSSO\_Setup\_edirectory file to start the installation wizard.
4. Select *Next*.
5. Read and accept the license agreement. Select *Next*.
6. Optionally, change the installation location. Select *Next*.

7. Enter:

---

**eDirectory Server**

---

<b>Server Address</b>	Enter the IP address of the eDirectory server.
-----------------------	------------------------------------------------

---

<b>Use secure connection (SSL)</b>	Select to connect to the eDirectory server using SSL security.
------------------------------------	----------------------------------------------------------------

---

<b>Search Base DN</b>	Enter the base Distinguished Name for the user search.
-----------------------	--------------------------------------------------------

---

**eDirectory Authentication**

---

<b>Username</b>	Enter a username that has access to the eDirectory, using LDAP format.
-----------------	------------------------------------------------------------------------

---

<b>User password</b>	Enter the password.
----------------------	---------------------

---

8. Select *Next*.
9. Select *Install*. When the installation completes, select *Finish*.

## Updating FSSO agents on Windows AD

After FSSO is installed on your network, you may want to upgrade to a newer version. The following procedure helps ensure you have a trouble free upgrade. How you update FSSO depends on if you are using polling mode or DCagent mode.

For polling mode, since there are no DC agents you only need to upgrade the Collector. However in DCagent mode, each DC Agent must be updated as well.

### To update FSSO in DC Agent mode

1. Go to the system32 directory on all DC's and rename the `dcagent.dll` file to `dcagent.dll.old`.  
This ensures the when the upgrade is pushed to the DC it does not overwrite the old file. If there are any problems this makes it easy to revert to the old version.
2. Run the FSSO setup .exe file to update the collector. When this is completed, ignore any reboot message.
3. Go to *Programs > Fortinet > Fortinet Single Sign On Agent > Install DC Agent* and push the DC agent out to all servers. All DC's will now need to be rebooted so that the new DLL file is loaded.
4. After the reboot, go to all DC's and delete the `dcagent.dll.old` files.

## Configuring the FSSO Collector agent for Windows AD

On the FortiGate unit, security policies control access to network resources based on user groups. With Fortinet Single Sign On, this is also true but each FortiGate user group is associated with one or more Windows AD user groups. This is how Windows AD user groups get authenticated in the FortiGate security policy.

Fortinet Single Sign On sends information about Windows user logons to FortiGate units. If there are many users on your Windows AD domains, the large amount of information might affect the performance of the FortiGate units.



To avoid this problem, you can configure the Fortinet Single Sign On Collector agent to send logon information only for groups named in the FortiGate unit's security policies. See [“Configuring FortiGate group filters” on page 583](#).

On each server with a Collector agent, you will be

- [Configuring Windows AD server user groups](#)
- [Configuring Collector agent settings](#), including the domain controllers to be monitored
- [Configuring Directory Access settings](#)
- [Configuring the Ignore User List](#)
- [Configuring FortiGate group filters](#) for each FortiGate unit
- [Configuring FSSO ports](#)
- [Configuring alternate user IP address tracking](#)



In some environments where user IP addresses change frequently, it might be necessary to configure the alternate IP address tracking method. For more information, see [“Configuring alternate user IP address tracking” on page 585](#).

---

## Configuring Windows AD server user groups

FortiGate units control network resource access at the group level. All members of a user group have the same network access as defined in FortiGate security policies.

You can use existing Windows AD user groups for authentication to FortiGate units if you intend that all members within each group have the same network access privileges.

Otherwise, you need to create new user groups for this purpose.



If you change a user's group membership, the change does not take effect until the user logs off and then logs on again.

---



The FSSO Agent sends only Domain Local Security Group and Global Security Group information to FortiGate units. You cannot use Distribution group types for FortiGate access. No information is sent for empty groups.

---

Refer to Microsoft documentation for information about creating and managing Windows AD user groups.

## Configuring Collector agent settings

You need to configure which domain controllers the Collector agent will use and which domains to monitor for user logons. You can also alter default settings and settings you made during installation. These tasks are accomplished by configuring the FSSO Collector Agent, and selecting either Apply to enable the changes.

At any time to refresh the FSSO Agent settings, select Apply.

## To configure the Collector agent

1. From the Start menu, select *Programs > FortiNet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent*.

2. Enter the following information.

### Monitoring user logon events

By default, this is enabled to automatically authenticate users as they log on to the Windows domain. Disable the Monitor feature only if you have a large network where this feature will slow responses too much.

### Support NTLM authentication

By default, this is enabled to facilitate logon of users who are connected to a domain that does not have the FSSO DC Agent installed. Disable NTLM authentication only if your network does not support NTLM authentication for security or other reasons.

### Collector Agent Status

Shows RUNNING when Collector agent is active.

### Listening ports

You can change FSSO Collector Agent related port numbers if necessary.

#### FortiGate

TCP port for FortiGate units. Default 8000.

#### DC Agent

UDP port for DC Agents. Default 8002.

### Logging

#### Log level

Select the minimum severity level of logged messages.

#### Log file size limit (MB)

Enter the maximum size for the log file in MB.

#### View Log

View all Fortinet Single Sign On agent logs.

<b>Log logon events in separate logs</b>	<p>Record user login-related information separately from other logs. The information in this log includes</p> <ul style="list-style-type: none"> <li>• data received from DC agents</li> <li>• user logon/logoff information</li> <li>• workstation IP change information</li> <li>• data sent to FortiGate units</li> </ul>
<b>View Logon Events</b>	If <i>Log logon events in separate logs</i> is enabled, you can view user login-related information.
<b>Authentication</b>	
<b>Require authenticated connection from FortiGate</b>	Select to require the FortiGate unit to authenticate before connecting to the Collector agent.
<b>Password</b>	Enter the password that FortiGate units must use to authenticate. The maximum password length is 16 characters. The default password is “fortinetcanada”.
<b>Timers</b>	
<b>Workstation verify interval (minutes)</b>	<p>Enter the interval in minutes at which the Fortinet Single Sign On Collector agent connects to client computers to determine whether the user is still logged on. The default is every 5 minutes. The interval may be increased if your network has too much traffic.</p> <p><b>Note:</b> This verification process creates security log entries on the client computer.</p> <p>If ports 139 or 445 cannot be opened on your network, set the interval to 0 to prevent checking. See <a href="#">“Configuring FSSO ports” on page 584</a>.</p>
<b>Dead entry timeout interval</b>	<p>Enter the interval in minutes after which Fortinet Single Sign On Agent purges information for user logons that it cannot verify. The default is 480 minutes (8 hours).</p> <p>Dead entries usually occur because the computer is unreachable (such as in standby mode or disconnected) but the user has not logged off. A common reason for this is when users forget to logoff before leaving the office for the day.</p> <p>You can also prevent dead entry checking by setting the interval to 0.</p>

<b>IP address change verify interval</b>	<p>Fortinet Single Sign On Agent periodically checks the IP addresses of logged-in users and updates the FortiGate unit when user IP addresses change. IP address verification prevents users from being locked out if they change IP addresses, as may happen with DHCP assigned addresses.</p> <p>Enter the verification interval in seconds. The default is 60 seconds. You can enter 0 to prevent IP address checking if you use static IP addresses.</p> <p>This does not apply to users authenticated through NTLM.</p>
<b>Cache user group lookup result</b>	<p>Enable caching.</p> <p>Caching can reduce group lookups and increase performance.</p>
<b>Cache expire in (minutes)</b>	<p>Fortinet Single Sign On Agent caches group information for logged-in users.</p> <p>Enter the duration in minutes after which the cache entry expires. If you enter 0, the cache never expires.</p> <p>A long cache expire interval may result in more stale user group information. This can be an issue when a user's group information is changed.</p>
<b>Clear Group Cache</b>	<p>Clear group information of logged-in users.</p> <p>This affects all logged-in users, and may force them to re-login.</p>

3. You can select *Save&Close* now or leave the agent configuration window open to complete additional configuration in the following sections.



To view the version and build number information for your FSSO Collector Agent configuration, selecting the Fortinet icon in the upper left corner of the Collector agent Configuration screen and select *About Fortinet Single Sign On Agent configuration*.

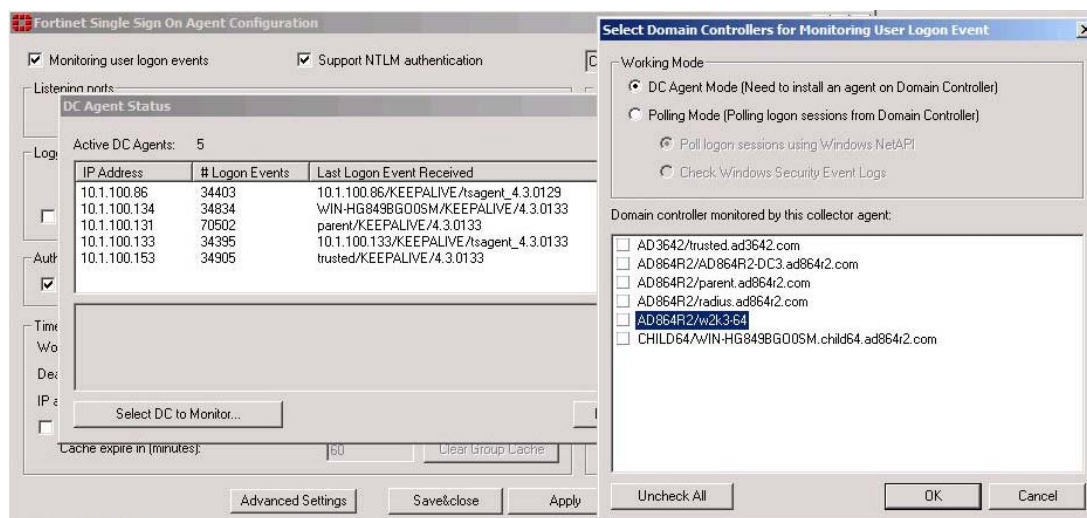
## Selecting Domain Controllers and working mode for monitoring

You can change which DC agents are monitored or change the working mode for logon event monitoring between DC agent mode and polling mode.

When polling mode is selected, it will poll port 445 of the domain controller every few seconds to see who is logged on.

1. From the Start menu select *Programs > Fortinet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent*.
2. In the *Common Tasks* section, select *Show Monitored DCs*.

### 3. Select *Select DC to Monitor*.



### 4. Choose the *Working Mode*.

- **DC Agent mode** — a Domain Controller agent monitors user logon events and passes the information to the Collector agent. This provides reliable user logon information, however you must install a DC agent on every domain controller in the domain.
- **Polling mode** — the Collector agent polls each domain controller for user logon information. Under heavy system load this might provide information less reliably. However installing a DC agent on each domain controller is not required in this mode.

You also need to choose the method used to retrieve logon information:

- Poll logon sessions using Windows NetAPI
- Check Windows Security Event Logs

For more information about these options, see [“Polling mode” on page 566](#).

5. In *Domain controller monitored by this collector agent*, select the collector agent that you installed.
6. Select *OK*. Select *Close*. Select *Save & Close*.

## Configuring Directory Access settings

The FSSO Collector Agent can access Windows Active Directory in one of two modes:

- **Standard** — the FSSO Collector Agent receives group information from the Collector agent in the *domain\user* format. This option is available on FortiOS 3.0 and later.
- **Advanced** — the FSSO Collector Agent obtains user group information using LDAP. The benefit of this method is that it is possible to nest groups within groups. This option is available on FortiOS 3.0 MR6 and later. The group information is in standard LDAP format.



If you change AD access mode, you must reconfigure your group filters to ensure that the group information is in the correct format.

### To configure Directory Access settings

1. From the Start menu select *Programs > Fortinet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent*.

2. In the *Common Tasks* section, select *Set Directory Access Information*.
3. The *Set Directory Access Information* dialog box opens.
4. From the *AD access mode* list, select either *Standard* or *Advanced*.
5. If you selected *Advanced AD* access mode, select *Advanced Setting* and configure the following settings and then select *OK*:

<b>AD server address</b>	Enter the address of your network's global catalog server.
<b>AD server port</b>	The default AD server port is 3268. This must match your server port.
<b>BaseDN</b>	Enter the Base distinguished name for the global catalog. This is the point in the tree that will be considered the starting point by default.
<b>Username</b>	If the global catalog accepts your Fortinet Single Sign On Agent agent's credentials, you can leave these fields blank.
<b>Password</b>	Otherwise, enter credentials for an account that can access the global catalog.

### BaseDN example

An example DN for Training Fortinet, Canada is `DN = ou=training, ou=canada, dc=fortinet, dc=com`. If you set the *BaseDN* to `ou=canada, dc=fortinet, dc=com` then when Fortinet Single Sign On Agent is looking up user credentials, it will only search the Canada organizational unit, instead of all the possible countries in the company. Its a short cut to entering less information and faster searches.

However, you may have problems if you narrow the *BaseDN* too much when you have international employees from the company visiting different offices. If someone from Fortinet Japan is visiting the Canada office in the example above, their account credentials will not be matched because they are in `DN = ou=japan, dc=fortinet, dc=com` instead of the *BaseDN* `ou=canada, dc=fortinet, dc=com`. The easy solution is to change the *BaseDN* to simply be `dc=fortinet, dc=com`. Then any search will check all the users in the company.

## Configuring the Ignore User List

The Ignore User List excludes users that do not authenticate to any FortiGate unit, such as system accounts. The logons of these users are not reported to FortiGate units. This reduces the amount of required resources on the FortiGate unit especially when logging logon events to memory.

### To configure the Ignore User List

1. From the Start menu select *Programs > Fortinet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent*.
2. In the *Common Tasks* section, select *Set Ignore User List*.  
The current list of ignored users is displayed. To view ignored usernames, expand each domain.
3. Do any of the following:
  - To remove a user from the list, select the check box beside the username and then select *Remove*. The user's login is no longer ignored.
  - To add users to be ignored, select *Add*, select the check box beside each required username, and then select *Add*.
4. Select *OK*.

## Configuring FortiGate group filters

FortiGate group filters actively control which user logon information is sent to each FortiGate unit. You need to configure the group filter list so that each FortiGate unit receives the correct user logon information for the user groups that are named in its security policies. These group filters help limit the traffic sent to the FortiGate unit, and help limit the logon events logged.

The maximum number of Windows AD user groups allowed on a FortiGate depends on the model. Low end models up to 300A support 256 Windows AD user groups, where mid and high end models support 1024 groups. This is per VDOM if VDOMs are enabled on the FortiGate unit.

You do not need to configure a group filter on the Collector agent if the FortiGate unit retrieves group information from Windows AD using LDAP. In that case, the Collector agent uses the list of groups you selected on the FortiGate unit as its group filter.

The filter list is initially empty. You need to configure filters for your FortiGate units using the Add function. At a minimum, create a default filter that applies to all FortiGate units without a defined filter.



If no filter is defined for a FortiGate unit and there is no default filter, the Collector agent sends all Windows AD group and user logon events to the FortiGate unit. While this normally is not a problem, limiting the amount of data sent to the FortiGate unit improves performance by reducing the amount of memory the unit uses to store the group list and resulting logs.

### To configure a FortiGate group filter

1. From the Start menu select *Programs > Fortinet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent*.
2. In the *Common Tasks* section, select *Set Group Filters*.

The FortiGate Filter List opens. It has the following columns:

<b>FortiGate SN</b>	The serial number of the FortiGate unit to which this filter applies.
<b>Description</b>	An optional description of the role of this FortiGate unit.
<b>Monitored Groups</b>	The Windows AD user groups that are relevant to the security policies on this FortiGate unit.
<b>Add</b>	Create a new filter.
<b>Edit</b>	Modify the filter selected in the list.
<b>Remove</b>	Remove the filter selected in the list.
<b>OK</b>	Save the filter list and exit.
<b>Cancel</b>	Cancel changes and exit.

3. Select *Add* to create a new filter. If you want to modify an existing filter, select it in the list and then select *Edit*.

4. Enter the following information and then select *OK*.

<b>Default filter</b>	Select to create the default filter. The default filter applies to any FortiGate unit that does not have a specific filter defined in the list.
<b>FortiGate Serial Number</b>	Enter the serial number of the FortiGate unit to which this filter applies. This field is not available if Default is selected.
<b>Description</b>	Enter a description of this FortiGate unit's role in your network. For example, you could list the resources accessed through this unit. This field is not available if Default is selected.
<b>Monitor the following groups</b>	The Collector agent sends to the FortiGate unit the user logon information for the Windows AD user groups in this list. Edit this list using the Add, Advanced and Remove buttons.
<b>Add</b>	<p>In the preceding single-line field, enter the Windows AD domain name and user group name, and then select Add. If you don't know the exact name, use the Advanced button instead.</p> <p>The format of the entry depends on the AD access mode (see <a href="#">“Configuring Directory Access settings” on page 581</a>):</p> <p>Standard: Domain\Group</p> <p>Advanced: cn=group, ou=corp, dc=domain</p>
<b>Advanced</b>	Select <i>Advanced</i> , select the user groups from the list, and then select <i>Add</i> .
<b>Remove</b>	Remove the user groups selected in the monitor list.

## Configuring FSSO ports

For FSSO to function properly a small number of TCP and UDP ports must be open through all firewalls on the network. These ports listed in this section assume the default FSSO ports are used.

### TCP ports for FSSO agent with client computers

Windows AD records when users log on but not when they log off. For best performance, Fortinet Single Sign On Agent monitors when users log off. To do this, Fortinet Single Sign On Agent needs read-only access to each client computer's registry over TCP port 139 or 445. Open at least one of these ports — ensure it is not blocked by firewalls.

If it is not feasible or acceptable to open TCP port 139 or 445, you can turn off Fortinet Single Sign On Agent logoff detection. To do this, set the Collector agent workstation verify interval to 0. The FSSO Collector Agent assumes that the logged on computer remains logged on for the duration of the Collector agent dead entry timeout interval — by default this is eight hours.

### Configuring ports on the Collector agent computer

On the computer where you install the Collector agent, you must make sure that the firewall does not block the listening ports for the FortiGate unit and the DC Agent. By default, these are TCP port 8000 and UDP port 8002. For more information about setting these ports, see [“Configuring Collector agent settings” on page 577](#).



## Configuring alternate user IP address tracking

In environments where user IP addresses change frequently, you can configure Fortinet Single Sign On Agent to use an alternate method to track user IP address changes. Using this method, Fortinet Single Sign On Agent responds more quickly to user IP address changes because it directly queries workstation IP addresses to match users and IP addresses.

This feature requires FSAE version 3.5.27 or later, Fortinet Single Sign On Agent any version, and FortiOS 3.0 MR7 or later.

### To configure alternate user IP address tracking

1. On the computer where the Collector agent is installed, go to *Start > Run*.
2. Enter `regedit` or `regedt32` and select *OK*.  
The Registry Editor opens.
3. Find the registry key `HKEY_LOCAL_MACHINE\SOFTWARE\Fortinet\FSAE\collectoragent`.
4. Set the `supportFSAEauth` value (dword) to `00000001`.  
If needed, create this new dword.
5. Close the Registry Editor.
6. From the Start menu select *Programs > Fortinet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent*.
7. Select *Apply*.  
The Fortinet Single Sign On Agent service restarts with the updated registry settings.

## Viewing FSSO component status

It is important to know the status of both your Collector agents and DC agents.

### Viewing Collector agent status

Use the *Show Service Status* to view your Collector agent information in the Status window. The Status window displays:

- the version of the software
- the status of the service
- the number of connected FortiGate units
- connected FortiGate information such as serial number, IP address, and time connected

### To view Collector agent status

1. From the Start menu select *Programs > Fortinet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent*.
2. In the *Common Tasks* section, select *Show Service Status*.

The Fortinet Single Sign On Collector agent Status window opens.

Optionally select *Get NTLM statistics* in the Status window to display NTLM information such as number of messages received, processed, failed, in the queue.

### Viewing DC agent status

Use the *Show Monitored DCs* to view the status of DC agents.

### To view domain controller agent status

1. From the Start menu select *Programs > Fortinet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent*.

2. In the *Common Tasks* section, select *Show Monitored DCs*.

For each DC Agent, the following information is displayed:

- IP address
- number of logon events received
- the last logon event
- when last logon was received.

To change which DC agents are monitored or change the working mode for logon event monitoring, select *Select DC to Monitor*.

## Configuring the FSSO TS agent for Citrix

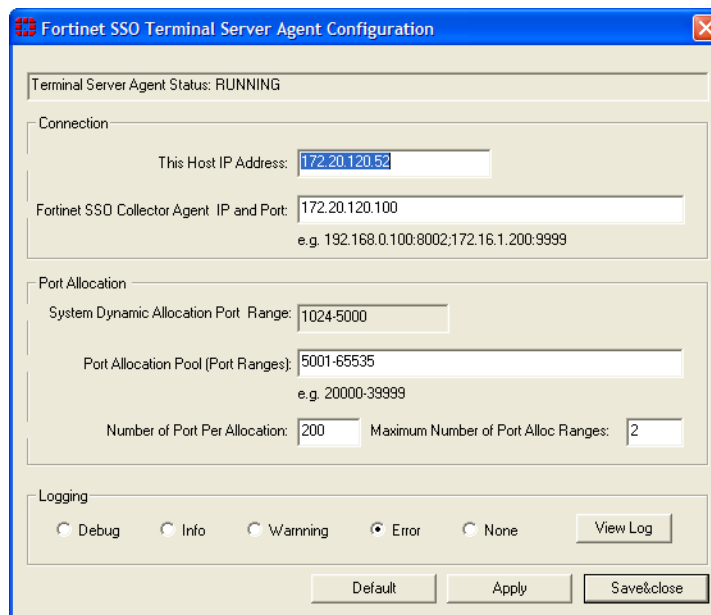
The FSSO TS agent works with the same FSSO Collector agent that is used for integration with Windows Active Directory. Install the Collector agent first. Follow the Collector agent installation procedure in “[Collector agent installation](#)” on page 572.

Configuration steps include:

- Install the Fortinet Citrix FSSO agent on the Citrix server.
- Install the Fortinet FSSO collector on a server on the network.
- Add the Citrix FSSO agent to the FortiGate Single-sign-On configuration.
- Add Citrix FSSO groups and users to an FSSO user group.
- Add an FSSO identity-based security policy that includes the Citrix FSSO user groups.

To change the TS agent configuration, select from the Start menu *Programs > Fortinet > Fortinet Single Sign-On Agent > TSAgent Config*. In addition to the host and Collector agent IP addresses that you set during installation, you can adjust port allocations for Citrix users. When a Citrix user logs on, the TS agent assigns that user a range of ports. By default each user has a range of 200 ports.

**Figure 124:**Configuring the TS agent



# Configuring the FSSO eDirectory agent for Novell eDirectory

You need to configure the eDirectory agent for it to communicate with eDirectory servers. You may have provided some of this information during installation.

This section includes:

- [Configuring the eDirectory agent](#)
- [Adding an eDirectory server](#)
- [Configuring a group filter](#)

## Configuring the eDirectory agent

You need to configure the eDirectory agent for it to communicate with eDirectory servers.

### To configure the eDirectory agent

1. From the Start menu select *Programs > Fortinet > eDirectory Agent > eDirectory Config Utility*.
2. The eDirectory Agent Configuration Utility dialog opens. Enter the following information and select *OK*.

---

#### eDirectory Authentication

**Username** Enter a username that has access to the eDirectory, using LDAP format.

**Password** Enter the password.

---

**Listening port** Enter the TCP port on which Fortinet Single Sign On Agent listens for connections from FortiGate units. The default is 8000. You can change the port if necessary.

---

**Refresh interval** Enter the interval in seconds between polls of the eDirectory server to check for new logons. The default is 30 seconds.

---

#### FortiGate Connection Authentication

**Require authenticated connection from FortiGate** Select to require the FortiGate unit to authenticate before connecting to the eDirectory Agent.

**Password** Enter the password that FortiGate units must use to authenticate. The maximum password length is 16 characters. The default password is "FortinetCanada".

---

**User logon Info Search Method** Select how the eDirectory agent accesses user logon information: *LDAP* or *Native* (Novell API). LDAP is the default.

If you select *Native*, you must also have the Novell Client installed on the PC.

---

#### Logging

**Log file size limit (MB)** Enter the maximum size for the log file in MB.

**View Log** View the current log file.

<b>Dump Session</b>	List the currently logged-on users in the log file. This can be useful for troubleshooting.
<b>Log level</b>	Select <i>Debug</i> , <i>Info</i> , <i>Warning</i> or <i>Error</i> as the minimum severity level of message to log or select <i>None</i> to disable logging.
<b>eDirectory Server List</b>	If you specified an eDirectory server during installation, it appears in this list.
<b>Add</b>	Add an eDirectory server. See .
<b>Delete</b>	Delete the selected eDirectory server.
<b>Edit</b>	Modify the settings for the selected server.
<b>Set Group Filters...</b>	Select the user groups whose user logons will be reported to the FortiGate unit. This is used only if user groups are not selected on the FortiGate unit.

## Adding an eDirectory server

Once the eDirectory agent is configured, you add one or more eDirectory servers.

### To add an eDirectory server

1. In the eDirectory Agent Configuration Utility dialog box (see the preceding procedure, “[Configuring the eDirectory agent](#)”), select *Add*.
2. The eDirectory Setup dialog box opens. Enter the following information and select OK:

<b>eDirectory Server Address</b>	Enter the IP address of the eDirectory server.
<b>Port</b>	If the eDirectory server does not use the default port 389, clear the <i>Default</i> check box and enter the port number.
<b>Use default credential</b>	Select to use the credentials specified in the eDirectory Configuration Utility. See “ <a href="#">Configuring the eDirectory agent</a> ” on page 587. Otherwise, leave the check box clear and enter a username and Password below.
<b>User name</b>	Enter a username that has access to the eDirectory, using LDAP format.
<b>User password</b>	Enter the password.
<b>Use secure connection (SSL)</b>	Select to connect to the eDirectory server using SSL security.
<b>Search Base DN</b>	Enter the base Distinguished Name for the user search.

## Configuring a group filter

The eDirectory agent sends user logon information to the FortiGate unit for all user groups unless you either configure an LDAP server entry for the eDirectory on the FortiGate unit and select the groups that you want to monitor or configure the group filter on the eDirectory agent.

If both the FortiGate LDAP configuration and the eDirectory agent group filter are present, the FortiGate user group selections are used.

#### To configure the group filter

1. From the Start menu select *Programs > Fortinet > eDirectory Agent > eDirectory Config Utility*.
2. Select *Set Group Filters*.
3. Do one of the following:
  - Enter group names, then select *Add*.
  - Select *Advanced*, select groups, and then select *Add*.
4. Select *OK*.

## Configuring FSSO on FortiGate units

To configure your FortiGate unit to operate with agent-based FSSO, you

- Configure any access to LDAP servers that might be necessary. Skip this step if you are using FSSO Standard mode. See [“Configuring LDAP server access” on page 589](#).
- Specify the Collector agent or Novell eDirectory agent that will provide user logon information. See [“Specifying your Collector agents or Novell eDirectory agents” on page 590](#).
- Add Active Directory user groups to FortiGate user groups. See [“Creating Fortinet Single Sign-On \(FSSO\) user groups” on page 591](#).
- Create security policies for FSSO-authenticated groups. See [“Creating security policies” on page 591](#).
- Optionally, specify a guest security policy to allow guest access. See [“Enabling guest access through FSSO security policies” on page 594](#).

## Configuring LDAP server access

LDAP access is required if your network has a Novell eDirectory agent or a Collector agent using Windows Advanced AD access mode. If you are using FSSO Standard mode, go to [“Specifying your Collector agents or Novell eDirectory agents” on page 590](#).

1. Go to *User & Device > Authentication > LDAP Servers* and select *Create New*.
2. Enter a *Name* to identify this server in FortiGate configurations.
3. Enter the *Server Name/IP* of the LDAP server.
4. Enter the *Distinguished Name*.
5. Set *Bind Type* to *Regular*.
6. In the *User DN* field, enter the administrative account name that you created for FSSO. For example, if the account is *FSSO\_Admin*, enter *“cn=FSSO\_Admin,cn=users”*.
7. Make sure that the *User DN* entry ends with a comma and append the string from the *Distinguished Name* field to the end of it.  
Example: *cn=FSSO\_Admin,cn=users,dc=office,dc=example,dc=com*
8. Enter the administrative account password in the *Password* field.
9. Select the *Test* button.  
A pop-up window near the top of the window should indicate *“Successful”*.
10. Select *OK*.

### To configure LDAP for FSSO - CLI example

```
config user ldap
 edit "ADserver"
 set server "10.11.101.160"
 set cnid "cn"
 set dn "cn=users,dc=office,dc=example,dc=com"
 set type regular
 set username
 "cn=administrator,cn=users,dc=office,dc=example,dc=com"
 set password set_a_secure_password
 next
end
```

## Specifying your Collector agents or Novell eDirectory agents

You need to configure the FortiGate unit to access at least one Collector agent or Novell eDirectory agent. You can specify up to five servers on which you have installed a Collector or eDirectory agent. The FortiGate unit accesses these servers in the order that they appear in the list. If a server becomes unavailable, the next one in the list is tried.

### To specify Collector agents - web-based manager

1. Go to *User & Device > Authentication > Single Sign-On* and select *Create New*.
2. In *Type*, select *Fortinet Single-Sign-On Agent*.
3. Enter a *Name* for the Windows AD server. This name appears in the list of Windows AD servers when you create user groups.
4. Enter the following information for each of up to five collector agents and select *OK*:

---

<b>Agent IP/Name</b>	Enter the IP address or the name of the server where this agent is installed. Maximum name length is 63 characters.  If the TCP port used for FSSO is not the default, 8000, you can change the setting in the CLI using the <code>config user fssso</code> command.  <a href="#">See "Configuring Collector agent settings" on page 577 .</a>
----------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

---

<b>Password</b>	Enter the password for the Collector agent or eDirectory agent. For the Collector agent, this is required only if you configured the agent to require authenticated access.
-----------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------

---

5. For Novell eDirectory or Windows AD with Collector agent in Advanced AD access mode
  - a. select the *LDAP Server* you configured previously. See ["Configuring LDAP server access" on page 589](#).
  - b. In *Users/Groups*, select the *Edit Users/Groups* tab and then select the users or groups that you want to monitor. Select the *View Users/Groups* tab to check your selection.
6. Select *OK*.

## To specify the FSSO Collector agent - CLI

In this example, the SSO server name is WinAD\_1 and the LDAP server is ADserver.

```
config user fssso
 edit WinAD_1
 set ldap-server ADserver
 set password ENC
 G7GQV7NEqilCM9jKmVmJJFVvhQ2+wtNEe9T0iYA5Sa+EgT2J8zhOrbkJFD
 r0RmY3c4LaoXdsoBczA1dONmcGfthTxxwGsigzGpbJdC71spFlQYtj
 set server 10.11.101.160
 set port 8000
 end
config user adgrp
 edit
```

## Creating Fortinet Single Sign-On (FSSO) user groups

You cannot use Windows or Novell groups directly in FortiGate security policies. You must create FortiGate user groups of the FSSO type and add Windows or Novell groups to them.

### To create a user group for FSSO authentication - web-based manager

1. Go to *User & Device > User > User Groups*.

2. Select *Create New*.

The New User Group dialog box opens.

3. In the *Name* box, enter a name for the group, FSSO\_Internet\_users for example.

4. In *Type*, select *Fortinet Single Sign-On (FSSO)*.

5. From the *Available Members* list, select the required FSSO groups.

Using the CTRL or SHIFT keys, you can select multiple groups.

6. Select the green right arrow button to move the selected groups to the *Members* list.

7. Select *OK*.

### To create the FSSO\_Internet-users user group - CLI

```
config user group
 edit FSSO_Internet_users
 set group-type fssso-service
 set member CN=Engineering,cn=users,dc=office,dc=example,dc=com
 CN=Sales,cn=users,dc=office,dc=example,dc=com
 end
```

## Creating security policies

Policies that require FSSO authentication are very similar to other security policies. Using identity-based policies, you can configure access that depends on the FSSO user group. This allows each FSSO user group to have its own level of access to its own group of services

In this situation, Example.com is a company that has its employees and authentication servers on an internal network. The FortiGate unit intercepts all traffic leaving the internal network and requires FSSO authentication to access network resources on the Internet. The following procedure configures the security policy for FSSO authentication. FSSO is installed and configured including the RADIUS server, FSSO Collector agent, and user groups on the FortiGate

For the following procedure, the internal interface is `port1` and the external interface connected to the Internet is `port2`. There is an address group for the internal network called `company_network`. The FSSO user group is called `fsso_group`, and the FSSO RADIUS server is `fsso_rad_server`.

**To configure an FSSO authentication security policy - web-based manager**

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Enter the following information.

<b>Policy Type</b>	Firewall
<b>Policy subtype</b>	User Identity
<b>Incoming Interface</b>	port1
<b>Source Address</b>	company_network
<b>Outgoing Interface</b>	port2
<b>Enable NAT</b>	Select

3. In *Configure Authentication Rules*, select *Create New*.
4. Enter

<b>Destination Address</b>	all
<b>Groups</b>	Select from the FSSO user groups that you created earlier.  FSSO_Guest_users is a default user group enabled when FSSO is configured. It allows guest users on the network who do not have FSSO account to still authenticate and have access to network resources. See <a href="#">“Enabling guest access through FSSO security policies” on page 594</a> .
<b>Schedule</b>	always
<b>Service</b>	HTTP, HTTPS, FTP, and Telnet
<b>Action</b>	ACCEPT
<b>Logging Options</b>	Select Log all Sessions. Logging FSSO logon events helps troubleshoot any FSSO related issues.
<b>UTM Security Profiles</b>	Enable AntiVirus, IPS, Web Filter, and Email Filter default profiles.

5. Select *OK*.  
A new line of information will appear in the identity-based policy table. The table lists the ID, user group or groups, the service or services, schedule, UTM, and logging selected for the rule. Use this display to verify your information was entered correctly.
6. Select *OK*.
7. Ensure the FSSO authentication policy is at the top of the list so it will be attempted to be matched before any other policy.



## To create a security policy for FSSO authentication - CLI

```
config firewall policy
 edit 0
 set srcintf internal
 set dstintf wan1
 set srcaddr company_network
 set dstaddr all
 set action accept
 set identity-based enable
 set nat enable
 config identity-based-policy
 edit 1
 set schedule any
 set groups company_network FSSO_guest_users
 set service HTTP HTTPS FTP TELNET
 end
 end
end
```

Here is an example of how this FSSO authentication policy is used. Example.com employee on the internal company network logs on to the internal network using their RADIUS username and password. When that user attempts to access the Internet, which requires FSSO authentication, the FortiGate authentication security policy intercepts the session, checks with the FSSO Collector agent to verify the user's identity and credentials, and then if everything is verified the user is allowed access to the Internet.

## Users belonging to multiple groups

Before FSSO 4.0 MR3, if a user belonged to multiple user groups, the first security policy to match any group that user belonged too was the only security policy applied. If that specific group did not have access to this protocol or resource where another group did, the user was still denied access. For example, `test_user` belongs to `group1` and `group2`. There are two FSSO authentication policies — one matches `group1` to authenticate FTP traffic and one matches `group2` to authenticate email traffic. The `group1` policy is at the top of the list of policies. If `test_user` wants to access an email server, the first policy encountered for a group `test_user` belongs to is the `group1` policy which does not allow email access and `test_user` is denied access. This is despite the next policy allowing access to email. If the order was reversed in this case, the traffic would be matched and the user's traffic would be allowed through the firewall. However if the policy order was reversed, FTP traffic would not be matched.

As of FSSO 4.0 MR3, if a user belongs to multiple groups multiple then attempts to match the group are attempted if applicable. Using the above example, when the attempt to match the `group1` policy is made and fails, the next policy with a group that `test_user` is a member of is attempted. In this case, the next policy is matched and access is granted to the email server.

When configuring this example the only difference between the policies is the services that are listed and the FSSO user group name.

Authenticating through multiple groups allows administrators to assign groups for specific services, and users who are members of each group have access to those services. For example there could be an FTP group, an email group, and a Telnet group.

## Enabling guest access through FSSO security policies

You can enable guest users to access FSSO security policies. Guests are users who are unknown to the Windows AD or Novell network and servers that do not logon to a Windows AD domain.

To enable guest access in your FSSO security policy, add an identity-based policy assigned to the built-in user group `FSSO_Guest_Users`. Specify the services, schedule and protection profile that apply to guest users — typically guests receive reduced access to a reduced set of services. See [“Creating security policies” on page 591](#).

## FortiOS FSSO log messages

There are two types of FortiOS log messages — firewall and event. FSSO-related log messages are generated from authentication events. These include user logon and log off events, and NTLM authentication events. These log messages are central to network accounting policies, and can also be useful in troubleshooting issues. For more information on firewall logging, see [“Enabling security logging” on page 522](#). For more information on logging, see the [FortiOS Handbook Log and Reporting chapter](#).

## Enabling authentication event logging

For the FortiGate unit to log events, that specific type of event must be enabled under logging.

When VDOMs are enabled certain options may not be available, such as CPU and memory usage events. You can enable event logs only when you are logged on to a VDOM; you cannot enable event logs globally.

To ensure you log all the events need, set the minimum log level to Notification or Information. Firewall logging requires Notification as a minimum. The closer to Debug level, the more information will be logged. While this extra information is useful, you must

### To enable event logging

1. Go to *Log&Report > Log Config > Log Settings*.
2. In *Event Logging*, select

---

<b>System activity event</b>	All system-related events, such as ping server failure and gateway status.
<b>User activity event</b>	All administration events, such as user logins, resets, and configuration updates.

---

Optionally you can enable any or all of the other logging event options.

3. Select *Apply*.

**Figure 125:**Authentication log messages

#	Date	Time	Level	Sub Type	ID	Action	Message	User Interface
1	2011-05-02	11:51:47	notice	auth	43011	authentication	User from 172.20.120.230 was timed out	test(172.20.120.230)
2	2011-05-02	11:43:31	notice	auth	43008	authentication	User test succeeded in authentication	HTTPS(172.20.120.230)
3	2011-05-02	11:43:23	notice	auth	43009	authentication	User techdoc failed in authentication	HTTPS(172.20.120.230)

Date	2011-05-02	Time	11:51:47
Level	notice	Sub Type	auth
ID	43011	Virtual Domain	root
Src	172.20.120.230	Dst	N/A
User	test	Group	testee
Policy ID	4	User Interface	test(172.20.120.230)
Action	authentication	Status	timed_out
Reason	Authentication timed out	Message	User from 172.20.120.230 was timed out

**Table 26:** List of FSSO related log messages

Message ID	Severity	Description
43008	Notification	Authentication was successful
43009	Notification	Authentication session failed
43010	Warning	Authentication locked out
43011	Notification	Authentication timed out
43012	Notification	FSSO authentication was successful
43013	Notification	FSSO authentication failed
43014	Notification	FSSO user logged on
43015	Notification	FSSO user logged off
43016	Notification	NTLM authentication was successful
43017	Notification	NTLM authentication failed

For more information on logging, see the [FortiOS Handbook Log and Reporting chapter](#).

## Testing FSSO

Once FSSO is configured, you can easily test to ensure your configuration is working as expected. For additional FSSO testing, see [“Troubleshooting FSSO” on page 596](#).

1. Logon to one of the stations on the FSSO domain, and access an Internet resource.
2. Connect to the CLI of the FortiGate unit, and if possible log the output.
3. Enter the following command:

```
diagnose debug authd fssolist
```

4. Check the output. If FSSO is functioning properly you will see something similar to the following:

```
----FSSO logons----
IP: 192.168.1.230 User: ADMINISTRATOR Groups: VLAD-AD/DOMAIN USERS
IP: 192.168.1.240 User: ADMINISTRATOR Groups: VLAD-AD/DOMAIN USERS
Total number of users logged on: 2
----end of FSSO logons----
```

The exact information will vary based on your installation.

5. Check the FortiGate event log, for FSSO-auth action or other FSSO related events with FSSO information in the message field. For a list of FSSO log message IDs, see [Table 26 on page 595](#).
6. To check server connectivity, run the following commands from the CLI:

```
FGT# diagnose debug enable
FGT# diagnose debug authd fssso server-status
FGT# Server Name Connection Status

SBS-2003 connected
```

## Troubleshooting FSSO

When installing, configuring, and working with FSSO some problems are quite common. A selection of these problems follows including explanations and solutions.

Some common Windows AD problems include:

- [General troubleshooting tips for FSSO](#)
- [User status “Not Verified” on the Collector agent](#)
- [After initial configuration, there is no connection to the Collector agent](#)
- [FortiGate performance is slow on a large network with many users](#)
- [Users from the Windows AD network are not able to access the network](#)
- [Users on a particular computer \(IP address\) can not access the network](#)
- [Guest users do not have access to network](#)
- [Can’t find the DCagent service](#)
- [User logon events not received by FSSO Collector agent](#)
- [User list from Windows AD is empty](#)
- [Mac OS X users can’t access external resources after waking from sleep mode](#)

## General troubleshooting tips for FSSO

The following tips are useful in many FSSO troubleshooting situations.

- To help locate the problem, configure a sniffer policy to capture FSSO logon messages along with other information.  
If FSSO is in use the log messages captured by a sniffer policy will include a user name if the IP address in the log message corresponds to the IP address of a user who has been authenticated with FSSO.
- Ensure all firewalls are allowing the FSSO required ports through.  
FSSO has a number of required ports that must be allowed through all firewalls or connections will fail. These include: ports 139, 389 (LDAP), 445, 636 (LDAP) 8000, and 8002.
- Ensure the Collector agent has at least 64kbps bandwidth to the FortiGate unit.  
If not the Collector agent does not have this amount of bandwidth, information FSSO information may not reach the FortiGate unit resulting in outages. The best solution is to configure traffic shaping between the FortiGate unit and the Collector agent to ensure that minimum bandwidth is always available.

## User status “Not Verified” on the Collector agent

When selecting “Show logon Users” in the Collector agent, some users may have their status set as “Not Verified”.

The Collector agent receives logon events for users from the DC agents, but Windows does not generate log out events. As such, the Collector agent needs to verify that the user is still logged on by checking the registry on that host.

If the Collector agent cannot connect to the host on ports 139 and 445 to perform this check, the host status is set to “Not Verified” and a log entry will be added to the Collector agent logs:

```
"01/01/2010 01:23:45 [1884] name_ip_match: failed to connect to
workstation: <Workstation Name> (192.168.1.1)"
```

### Solution

There are a few things that can cause the Collector agent not to be able to connect to the user's work station. Below is a list of the most common causes:

- Most commonly, a host firewall on the user's workstation or a router on the network prevents remote access on ports 139 and/or 445. Try opening the ports on the host firewall.
- If the remote registry service is not running on the user's workstation, the Collector agent will not be able to connect to the registry remotely. Make sure the remote registry service is running.
- This problem may also be caused by [a known MS upgrade issue](#).

Using `Regedit.exe`, edit “HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Control\SecurePipeServers”, set permissions for winreg and allow Local Service with R and W permissions.

## After initial configuration, there is no connection to the Collector agent

The Collector Agent has been configured but now cannot be contacted. This may be a regular connectivity problem. The section *Troubleshooting Connectivity* in the [FortiOS Handbook Troubleshooting chapter](#) will help locate and identify any network problems. Other solutions specific to FSSO are listed here.

## Solution

If there are no network problems that can be identified, try the following solutions.

- The Windows AD network must be configured before configuring the FortiGate unit. This includes the domain controller agents, and Collector agents.
- Ensure the DC agents point to the correct collector agent port and IP address.
- Ensure that TCP port 8000, and UDP port 8002 are not blocked.
- FSSO is very dependent on DNS, ensure the forward DNS zone has no stale records and after adding it to the domain if the DNS entry is not in the zone add it.
- An error in the DNI field on the FortiGate unit will prevent connections. Select the browse button next to the field to confirm it can connect correctly to the Windows AD server and return information. See
- If the secure check box is selected, ensure that LDAP v3 is being used since earlier LDAP does not support secure TLS connections.
- Ensure that the default LDAP ports are not being blocked on the network. These ports include port 389, and port 636. If you change the default ports, ensure both the FortiGate unit and the Windows AD server are using the same port numbers and that those ports are allowed through all firewalls on your network.
- If you are using FSSO in polling mode, ensure that port 445 is not blocked by firewalls.

## Collector Agent service freezing and shutting down

FSSO problem.

### Solution

- Reinstall FSSO.

## FortiGate performance is slow on a large network with many users

FSSO sends information about Windows user logons to FortiGate units. If there are many users on your Windows AD domains, the large amount of information might affect the performance of the FortiGate units. Logon tracking is logged to memory, and may reduce performance in extreme situations.

To avoid this problem, you can configure the Collector agent to send logon information only for groups named in the FortiGate unit's security policies. Also you can configure the Ignore User list on the FortiGate unit to avoid tracking unnecessary logons.

Also logging to memory can consume large amounts of FortiGate system memory. To lessen the memory used, change the logging from the default level of Information to a less frequent level such as Error or Warning. This results in less information being logged and frees system memory to improve overall FortiGate system performance. However, if you are trying to troubleshoot a problem one of the first things to do is to change the logging severity to Information or possibly even Debug to provide you with additional information while solving your problem.

## Solution

Add users to the Ignore User list. This is a best practice for admin accounts whose logon information will not be sent to the FortiGate unit. This is useful for automated accounts that may logon many times. Examples of accounts in this category include:

- IIS services
- AV
- other system accounts

For more information on configuring the Ignore users list, see [“Configuring the Ignore User List” on page 582](#).

## Users from the Windows AD network are not able to access the network

If nobody can access the network and your network has only one Collector agent, when it goes offline no users will have access. However if only some users can not access the network, it is likely that user group changes were made recently that are causing the problems.

### Solutions

- If there is only one Collector agent, configure additional Collector agents in the domain to act as backups. They will provide the redundancy required if the original collector goes offline. Remember to add them to the Fortinet Single Sign-On Agent entry under *User & Device > Authentication > Single Sign-On* on the web-based manager or `config user fssso` in the CLI. If the server and port for the new agent are not in the list, it will not be contacted.
- Ensure the Collector agent has at least 64kbps bandwidth to the FortiGate unit. If not, information FSSO information may not reach the FortiGate unit resulting in outages. The best solution is to configure traffic shaping between the FortiGate unit and the Collector agent to ensure that minimum bandwidth is always available.
- If some users can not connect, verify their Windows AD records to find groups in common, and investigate the state of those groups focusing on any recent changes. It may be a group or permission change is the reason.
- There may be a problem with the user list. See [“User list from Windows AD is empty” on page 600](#).

## Users on a particular computer (IP address) can not access the network

Windows AD Domain Controller agent gets the username and workstation where the logon attempt is coming from. If there are two computers with the same IP address and the same user trying to logon, it is possible for the authentication system to become confused and believe that the user on computer\_1 is actually trying to access computer\_2.

Windows AD does not track when a user logs out. It is possible that a user logs out on one computer, and immediately logs onto a second computer while the system still believes the user is logged on the original computer. While this is allowed, information that is intended for the session on one computer may mistakenly end up going to the other computer instead. The result would look similar to a hijacked session.

### Solutions

- Ensure each computer has separate IP addresses.
- Encourage users to logout on one machine before logging onto another machine.
- If multiple users have the same username, change the usernames to be unique.
- Shorten timeout timer to flush inactive sessions after a shorter time.

## Guest users do not have access to network

A group of guest users was created, but they don't have access.

### Solution

The group of the guest users was not included in a policy, so they do not fall under the guest account. To give them access, associate their group with a security policy.

Additionally, there is a default group called `FSSO_Guest_Users`. Ensure that group is part of an identity-based security policy to allow traffic.

## Can't find the DCagent service

The DCagent service can't be found in the list of regular windows services. This is because it has no associated Windows service.

Instead DCagent is really `dcagent.dll` and is located in the `Windows\system32` folder. This DLL file is loaded when windows boots up and it intercepts all logon events processed by the domain controller to send these events to the Collector agent (CA).

### Solution

#### To verify that the DCagent is installed properly

1. Check that `DCagent.dll` exists in `%windir%\system32` folder.
2. Check that the registry key exists:  
`[HKEY_LOCAL_MACHINE\SOFTWARE\Fortinet\FSAE\dcagent]`

If both exist, the DCagent is properly installed.

## User logon events not received by FSSO Collector agent

When a warning dialog is present on the screen on the Collector agent computer, the Collector agent will not receive any logon events. Once the dialog has been closed normal operation will resume.

If polling mode is enabled, it is possible the polling interval is too large. Use a shorter polling interval to ensure the collector agent is capturing all logon events.

If NetAPI polling mode is enabled, consider switching to Event log polling as it provides better accuracy.

## User list from Windows AD is empty

FSSO server is configured. I have received a list of windows AD groups. However, a user list is empty.

### Solution

There could be 2 problems:

In most cases, the FortiGate receives login information, but can't translate the Windows AD group into the protection profile. Make sure that all the required Windows AD groups are included in the FortiGate user groups and that all FortiGate user groups are included into the authentication security policy.

There may be a problem with AD FSSO service running on the Windows AD server.



### To ensure the problem is on windows side

1. Go to *Log&Report > Log Config*.
2. Enable firewall authentication event logging and debug level logging on the FortiGate.
3. Ask one or more users to log in into windows.
4. Check the FortiGate logs for the logon event from the Windows AD server.

If there is no new logon event entry in the logs, the problem is with Windows side. Use MS Windows AD documentation to troubleshoot the problem.

## Mac OS X users can't access external resources after waking from sleep mode

When client computers running Mac OS X (10.6.X and higher) wake up from sleep mode, the user must authenticate again to be able to access external resources. If the user does not re-authenticate, the user will maintain access to internal web sites, but will be unable to access any external resources.

This issue is caused by Mac OS X not providing sufficient information to the FSAE. This results in the FortiGate blocking access to the user because they cannot be authenticated.

### Solution

The security settings on client computer(s) must be configured to require that a username and password be entered when exiting sleep mode or screen saver. With this feature enabled in Mac OS X, the FortiGate will receive the authentication information it requires to authenticate the user and allow them access.

Note that if the user reverts their settings to disable the password requirement, this will cause the issue to reappear.

# SSO using RADIUS accounting records

A FortiGate unit can authenticate users transparently who have already authenticated on an external RADIUS server. Based on the user group to which the user belongs, the security policy applies the appropriate UTM profiles. RADIUS SSO is relatively simple because the FortiGate unit does not interact with the RADIUS server, it only monitors RADIUS accounting records that the server emits. These records include the user's IP address and user group.

After the initial set-up, changes to the user database, including changes to user group memberships, are made on the external RADIUS server, not on the FortiGate unit.

This section describes:

- [User's view of RADIUS SSO authentication](#)
- [Configuration Overview](#)
- [Configuring the RADIUS server](#)
- [Creating the FortiGate RADIUS SSO agent](#)
- [Defining local user groups for RADIUS SSO](#)
- [Creating security policies](#)
- [Example: webfiltering for student and teacher accounts](#)

## User's view of RADIUS SSO authentication

For the user, RADIUS SSO authentication is simple:

- The user connects to the RADIUS server and authenticates.
- The user attempts to connect to a network resource that is reached through a FortiGate unit. Authentication is required for access, but the user connects to the destination without being asked for logon credentials because the FortiGate unit knows that the user is already authenticated. FortiOS applies UTM features appropriate to the user groups that the user belongs to.

## Configuration Overview

The general steps to implement RADIUS Single Sign-On are:

1. If necessary, configure your RADIUS server. The user database needs to include user group information and the server needs to send accounting messages.
2. Create the FortiGate RADIUS SSO agent.
3. Define local user groups that map to RADIUS groups.
4. Create an identity-based security policy and create authentication rules as appropriate for the different user groups that are permitted access.

## Configuring the RADIUS server

You can configure FortiGate RSSO to work with most RADIUS-based accounting systems. In most cases, you only need to do the following to your RADIUS accounting system:

- Add a user group name field to customer accounts on the RADIUS server so that the name is added to the RADIUS Start record sent by the accounting system to the FortiOS unit. User group names do not need to be added for all users, only to the accounts of users who will use RSSO feature on the FortiGate unit.
- Configure your accounting system to send RADIUS Start records to the FortiOS unit. You can send the RADIUS Start records to any FortiGate network interface. If your FortiGate unit is operating with virtual domains (VDMs) enabled, the RADIUS Start records must be sent to a network interface in the management VDOM.

## Creating the FortiGate RADIUS SSO agent

Once you define a RADIUS SSO (RSSO) agent, the FortiGate unit will accept user logon information from any RADIUS server that has the same shared secret. You can create only one RSSO agent in each VDOM.

Before you create the RSSO agent, you need to allow RADIUS accounting information on the interface that connects to the RADIUS server.

### To enable RADIUS access on the interface - web-based manager

1. Go to *System > Network > Interfaces* and edit the interface to which the RADIUS server connected.
2. Select *Listen for RADIUS Accounting Messages*.
3. Select *OK*.

### To enable RADIUS access on the interface - CLI

In this example, the port2 interface is used.

```
config system interface
 edit port2
 append allowaccess radius-acct
 end
```

### To create a RADIUS SSO agent

1. Go to *User & Device > Authentication > Single Sign-On* and select *Create New*.
2. In *Type*, select *RADIUS Single-Sign-On Agent*.
3. Select *Use RADIUS Shared Secret* and enter the RADIUS server shared secret.
4. Select *Send RADIUS Responses*.
5. Select *OK*.

The Single Sign-On agent is named `RSSO_Agent`.

## To create a RADIUS SSO agent - CLI

In this example, the RADIUS server secret is “fortinet”.

```
config user radius
 edit RSSO_Agent
 set rspo enable
 set rspo-validate-request-secret enable
 set rspo-secret fortinet
 set rspo-radius-response enable
 end
```

When the RSSO agent is created in the web-based manager, it is automatically named RISSO\_Agent. You can use any name when creating the agent in the CLI.



## Selecting which RADIUS attributes are used for RSSO

For RADIUS SSO to work, FortiOS needs to know the user’s endpoint identifier (usually IP address) and RADIUS user group. There are default RADIUS attributes where FortiOS expects this information, but you can change these attributes in the `config user radius` CLI command.

**Table 27:** RSSO information and RADIUS attribute defaults

RSSO Information	RADIUS Attribute	CLI field
Endpoint identifier	Calling-Station-ID	rsso-endpoint-attribute
Endpoint block attribute	Called-Station-ID	rsso-endpoint-block-attribute
User group	Class	sso-attribute

The Endpoint block attribute can be used to block a user. If the attribute value is “Block”, FortiOS blocks all traffic from that user’s IP address. The RSSO fields are visible only when `rsso` is set to enable.

## Configuring logging for RSSO

In the `config user radius` CLI command, you can set the following flags in the `rsso-log-flags` field to determine which types of RSSO-related events are logged:

`accounting-event` — FortiOS did not find the expected information in a RADIUS record.

`accounting-stop-missed` — a user context entry expired without FortiOS receiving a RADIUS Stop message.

`context-missing` — FortiOS was not able to match a communication session with a user.

`endpoint-block` — FortiOS blocked a user because the RADIUS record’s endpoint block attribute had the value “Block”.

`profile-missing` — FortiOS cannot find a user group name in a RADIUS start message that matches the name of an RSSO user group in FortiOS.

`protocol-error` — A RADIUS protocol error occurred.

`radiusd-other` — Other events, described in the log message.

## Defining local user groups for RADIUS SSO

You cannot use RADIUS user groups directly in security policies. Instead, you create locally-defined user groups on the FortiGate unit and associate each of them with a RADIUS user group.

### To define local user groups for RADIUS SSO

1. Go to *User & Device > User > User Groups* and select *Create New*.
2. Enter a Name for the user group.
3. In *Type*, select *RADIUS Single Sign-On (RSSO)*.
4. In *RADIUS Attribute Value*, enter the name of the RADIUS user group this local user group represents.
5. Select *OK*.

### To define local user groups for RADIUS SSO

This example creates an RSSO user group called RSSO-1 that is associated with RADIUS user group “student”.

```
config user group
 edit RSSO-1
 set group-type rspo
 set sso-attribute-value student
 end
```

## Creating security policies

RADIUS SSO uses regular identity-based security policies. The RSSO user group you specify determines which users are permitted to use the policy. You can create multiple authentication rules so that various user groups can have different UTM features enabled, different permitted services, schedules, and so on.

### To create a security policy for RSSO - web-based manager

1. Go to *Policy > Policy > Policy*.
2. Select *Create New*.
3. Enter the following information.

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	User Identity
<b>Incoming Interface</b> <b>Source Address</b> <b>Outgoing Interface</b>	as needed
<b>Enable NAT</b>	Selected

4. In *Configure Authentication Rules*, select *Create New* and enter:

<b>Destination Address</b>	all
<b>Group(s)</b>	Select the user groups you created for RSSO. See <a href="#">“Defining local user groups for RADIUS SSO”</a> on page 605.
<b>User(s)</b>	not used
<b>Schedule Service</b>	as needed
<b>Action</b>	ACCEPT
<b>UTM Security Profiles</b>	Select UTM security profiles appropriate for the user group.

5. Select *OK*.
6. Repeat steps 4 and 5 for each user group that is allowed to use this security policy. Schedule, Service, and UTM profiles can be different for each group.
7. Select *OK*.

To ensure an RSSO-related policy is matched first, the policy should be placed higher in the security policy list than more general policies for the same interfaces.

#### To create a security policy for RSSO - CLI

In this example, an internal network to Internet policy enables web access for members of a student group and activates the appropriate UTM profiles.

```
config firewall policy
 edit 0
 set srcintf internal
 set dstintf wan1
 set srcaddr all
 set action accept
 set rso enable
 set identity-based enable
 set nat enable
 config identity-based-policy
 edit 1
 set schedule always
 set utm-status enable
 set groups "RSSO-student"
 set dstaddr "all"
 set service HTTP HTTPS
 set av-profile students
 set webfilter-profile students
 set spamfilter-profile students
 set dlp-sensor default
 set ips-sensor default
 set application-list students
 set profile-protocol-options "default"
 end
 end
 end
```

## Example: webfiltering for student and teacher accounts

The following example uses RADIUS SSO to apply web filtering to students, but not to teachers. Assume that the RADIUS server is already configured to send RADIUS Start and Stop records to the FortiGate unit. There are two RADIUS user groups, *students* and *teachers*, recorded in the default attribute *Class*. The workstations are connected to port1, port2 connects to the RADIUS server, and port3 connects to the Internet.

### Configure the student web filter profile

1. Go to *Security Profiles > Web Filter > Profiles* and select *Create New*.
2. Enter the following and select *Apply*.

<b>Name</b>	student
<b>Inspection Mode</b>	Proxy
<b>FortiGuard Categories</b>	Enable. Right-click the <i>Potentially Liable</i> category and select <i>Block</i> . Repeat for <i>Adult/Mature Content</i> and <i>Security Risk</i> .

### Enable RADIUS access on the port2 interface

1. Go to *System > Network > Interfaces* and edit the port2 interface.
2. Select *Listen for RADIUS Accounting Messages*.
3. Select *OK*.

### Create the RADIUS SSO agent

1. Go to *User & Device > Authentication > Single Sign-On* and select *Create New*.
2. In *Type*, select *RADIUS Single-Sign-On*.
3. Select *Use RADIUS Shared Secret* and enter the RADIUS server shared secret.
4. Select *Send RADIUS Responses*.
5. Select *OK*.

The Single Sign-On agent is named *RSSO\_Agent*.

### Define local user groups associated with the RADIUS SSO user groups

1. Go to *User & Device > User > User Groups* and select *Create New*.
2. Enter the following and select *OK*.

<b>Name</b>	RSSO-students
<b>Type</b>	RADIUS Single Sign-On (RSSO)
<b>RADIUS Attribute Value</b>	students

3. Select *Create New*, enter the following and select *OK*.

<b>Name</b>	RSSO-teachers
<b>Type</b>	RADIUS Single Sign-On (RSSO)
<b>RADIUS Attribute Value</b>	teachers

## Create a security policy for RSSO

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Enter

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	User Identity
<b>Incoming Interface</b>	port1
<b>Source Address</b>	all
<b>Outgoing Interface</b>	port3
<b>Enable NAT</b>	Selected

3. In *Configure Authentication Rules*, select *Create New* and enter:

<b>Destination Address</b>	all
<b>Group(s)</b>	RSSO-students
<b>Schedule</b>	always
<b>Service</b>	HTTP, HTTPS
<b>Action</b>	ACCEPT
<b>UTM Security Profiles</b>	Enable AntiVirus, Web Filter, IPS. In Web Filter, select the <i>student</i> profile.

4. Select *OK*.
5. In *Configure Authentication Rules*, select *Create New* and enter:

<b>Destination Address</b>	all
<b>Group(s)</b>	RSSO-teachers
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT
<b>UTM Security Profiles</b>	Enable AntiVirus and IPS.

6. Select *OK*.
7. Select *OK* to save the policy.
8. Repeat steps 4 and 5 for each user group that is allowed to use this security policy. Schedule, Service, and UTM profiles can be different for each group.
9. Select *OK*.



# Monitoring authenticated users

This section describes how to view lists of currently logged-in firewall and VPN users. It also describes how to disconnect users.

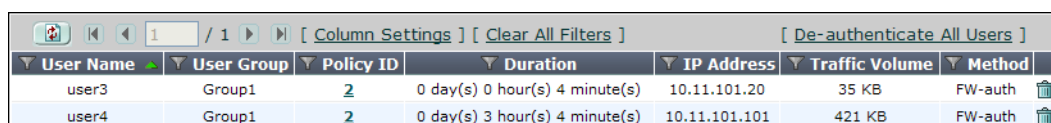
The following topics are included in this section:

- [Monitoring firewall users](#)
- [Monitoring SSL VPN users](#)
- [Monitoring IPsec VPN users](#)

## Monitoring firewall users

To monitor firewall users, go to *User & Device > Monitor > Firewall*.

**Figure 126:** Firewall users listed in monitor



User Name	User Group	Policy ID	Duration	IP Address	Traffic Volume	Method
user3	Group1	2	0 day(s) 0 hour(s) 4 minute(s)	10.11.101.20	35 KB	FW-auth
user4	Group1	2	0 day(s) 3 hour(s) 4 minute(s)	10.11.101.101	421 KB	FW-auth

You can de-authenticate a user by selecting the Delete icon for that entry.

You can filter the list of displayed users either by selecting the funnel icon for one of the column titles or selecting *Filter Settings*.

Select *Column Settings* to add or remove columns to the display, or rearrange the order of the columns displayed.

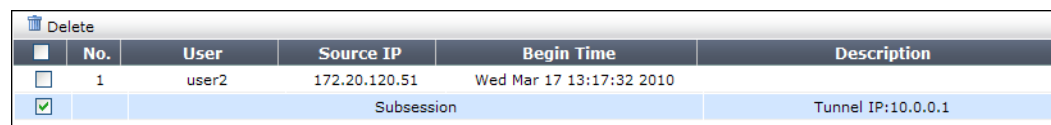
Optionally, you can select *De-authenticate all users*. Best practices dictate that this only be used in extreme cases since all users will momentarily lose their network resource connections.

## Monitoring SSL VPN users

You can monitor web-mode and tunnel-mode SSL VPN users by username and IP address.

To monitor SSL VPN users, go to *VPN > Monitor > SSL-VPN Monitor*. To disconnect a user, select the user and then select the *Delete* icon.

**Figure 127:** Monitoring SSL VPN users



No.	User	Source IP	Begin Time	Description
1	user2	172.20.120.51	Wed Mar 17 13:17:32 2010	
<input checked="" type="checkbox"/>			Subsession	Tunnel IP:10.0.0.1

The first line, listing the username and IP address, is present for a user with either a web-mode or tunnel-mode connection. The Subsession line is present only if the user has a tunnel mode connection. The *Description* column displays the virtual IP address assigned to the user's tunnel-mode connection.

For more information about SSL VPN, see the [FortiOS Handbook SSL VPN chapter](#).

### To monitor SSL VPN users - CLI

To list all of the SSL VPN sessions and their index numbers:

```
execute vpn sslvpn list
```

The output looks like this:

SSL-VPN Login Users:

Index	User	Auth	Type	Timeout	From	HTTPS in/out
0	user1	1		256	172.20.120.51	0/0

SSL-VPN sessions:

Index	User	Source IP	Tunnel/Dest IP
0	user2	172.20.120.51	10.0.0.1

You can use the Index value in the following commands to disconnect user sessions:

### To disconnect a tunnel-mode user

```
execute vpn sslvpn del-tunnel <index>
```

### To disconnect a web-mode user

```
execute vpn sslvpn del-web <index>
```

You can also disconnect multiple users:

### To disconnect all tunnel-mode SSL VPN users in this VDOM

```
execute vpn ssl del-all tunnel
```

### To disconnect all SSL VPN users in this VDOM

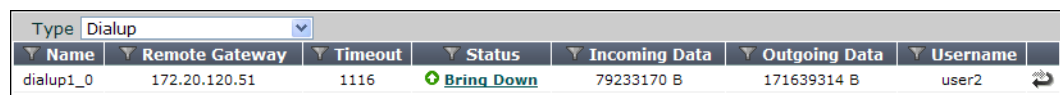
```
execute vpn ssl del-all
```

## Monitoring IPsec VPN users

To monitor IPsec VPN tunnels in the web-based manager, go to *VPN > Monitor > IPsec Monitor*. user names are available only for users who authenticate with XAuth.

You can close a tunnel by selecting its *Bring Down* link in the *Status* column.

**Figure 128:**Monitoring dialup VPN users



Name	Remote Gateway	Timeout	Status	Incoming Data	Outgoing Data	Username
dialup1_0	172.20.120.51	1116	<a href="#">Bring Down</a>	79233170 B	171639314 B	user2

For more information, see the [FortiOS Handbook IPsec chapter](#).

## Monitoring banned users

The Banned User list shows all IP addresses and interfaces blocked by NAC quarantine. The list also shows all IP addresses, authenticated users, senders, and interfaces blocked by Data Leak Prevention (DLP). The system administrator can selectively release users or interfaces from quarantine or configure quarantine to expire after a selected time period.

All sessions started by users or IP addresses on the Banned User list are blocked until the user or IP address is removed from the list. All sessions to an interface on the list are blocked until the interface is removed from the list.

You can configure NAC quarantine to add users or IP addresses to the Banned User list under the following conditions:

- **Users or IP addresses that originate attacks detected by IPS** - To quarantine users or IP addresses that originate attacks, enable and configure *Quarantine Attackers* in an IPS Sensor Filter.
- **IP addresses or interfaces that send viruses detected by virus scanning** - To quarantine IP addresses that send viruses or interfaces that accept traffic containing a virus, enable *Quarantine Virus Sender* in an antivirus profile.
- **Users or IP addresses that are banned or quarantined by Data Leak Prevention** -Set various options in a DLP sensor to add users or IP addresses to the Banned User list.

For more information, see [FortiOS Handbook UTM chapter](#).

Banned users are viewed from *User & Device > Monitor > Banned User*.

<b>Banned User page</b>	
Lists all banned users.	
<b>Page Controls</b>	Use to navigate through the list.
<b>Clear</b>	Removes all users and IP addresses from the Banned User list.
<b>#</b>	The position number of the user or IP address in the list.
<b>Ban key</b>	The Ban key.
<b>Application Protocol</b>	The protocol that was used by the user or IP address added to the Banned User list.
<b>Cause or rule</b>	The Fortinet function that caused the user or IP address to be added to the Banned User list. <i>Cause or rule</i> can be IPS, Antivirus, or Data Leak Prevention.
<b>Created</b>	The date and time the user or IP address was added to the Banned User list.
<b>Expires</b>	The date and time the user or IP address will be automatically removed from the Banned User list. If <i>Expires</i> is <i>Indefinite</i> , you must manually remove the user or host from the list.
<b>Delete</b>	Removes the selected user or IP address from the Banned User list.

## Monitoring IM users

User lists can be managed to allow or block certain users. Each user can be assigned a policy to allow or block activity for each IM protocol. Each IM function can be individually allowed or blocked providing the administrator the granularity to block the more bandwidth consuming features such as voice chat while still allowing text messaging. The IM user monitor list displays information about instant messaging users who are currently connected. The list can be filtered by protocol. After IM users connect through the firewall, the unit displays which users are

connected. You can analyze the list and decide which users to allow or block. A policy can be configured to handle unknown users.

Active IM users are viewed from *User & Device > Monitor > IM*.



IM users who are already logged on before changes are made to the IM user profile will not be affected until their next login. You cannot disconnect users who have already logged on by enabling logon blocking.

<b>IM page</b>	
Lists all active IM users that are currently active. This page allows you to view blocked users as well as users that are currently using a particular IM protocol, such as MSN.	
<b>Protocol</b>	Filter the list by selecting the protocol for which to display current users: AIM, ICQ, MSN, or Yahoo. All current users can also be displayed.
<b>#</b>	The position number of the IM user in the list.
<b>Protocol</b>	The protocol being used.
<b>User Name</b>	The name selected by the user when registering with an IM protocol. The same user name can be used for multiple IM protocols. Each user name/protocol pair appears separately in the list.
<b>Source IP</b>	The IP address where the user initiated the IM session from.
<b>Last Login</b>	The last time the current user used the protocol.
<b>Block</b>	Select to add the user name to the permanent black list. Each user name/protocol pair must be explicitly blocked by the administrator.

# Examples and Troubleshooting


This chapter provides an example of a FortiGate unit providing authenticated access to the Internet for both Windows network users and local users.

The following topics are included in this section:

- [Firewall authentication example](#)
- [LDAP Dial-in using member-attribute](#)
- [RADIUS SSO example](#)
- [Troubleshooting](#)

## Firewall authentication example

**Figure 129:**Example configuration



## Overview

In this example, there is a Windows network connected to Port 2 on the FortiGate unit and another LAN, Network\_1, connected to Port 3.

All Windows network users authenticate when they logon to their network. Members of the Engineering and Sales groups can access the Internet without entering their authentication credentials again. The example assumes that the Fortinet Single Sign On (FSSO) has already been installed and configured on the domain controller.

LAN users who belong to the Internet\_users group can access the Internet after entering their username and password to authenticate. This example shows only two users, User1 is authenticated by a password stored on the FortiGate unit, User2 is authenticated on an external authentication server. Both of these users are referred to as local users because the user account is created on the FortiGate unit.

## Creating a locally-authenticated user account

User1 is authenticated by a password stored on the FortiGate unit. It is very simple to create this type of account.

### To create a local user - web-based manager

1. Go to *User & Device > User > User Definition* and select *Create New*.
2. Enter the following information: username, Password.

<b>username</b>	User1
<b>Password</b>	hardtoguess

3. Select *OK*.

### To create a local user - CLI

```
config user local
 edit user1
 set type password
 set passwd hardtoguess
 end
```

## Creating a RADIUS-authenticated user account

To authenticate users using an external authentication server, you must first configure the FortiGate unit to access the server.

### To configure the remote authentication server - web-based manager

1. Go to *User & Device > Authentication > RADIUS Servers* and select *Create New*.
2. Enter the following information and select *OK*:

<b>Name</b>	OurRADIUSsrv
<b>Primary Server Name/IP</b>	10.11.101.15
<b>Primary Server Secret</b>	OurSecret
<b>Authentication Scheme</b>	Select <i>Use Default Authentication Scheme</i> .

### To configure the remote authentication server - CLI

```
config user radius
 edit OurRADIUSsrv
 set server 10.11.102.15
 set secret OurSecret
 set auth-type auto
 end
```

Creation of the user account is similar to the locally-authenticated account, except that you specify the RADIUS authentication server instead of the user's password.

### To configure a remote user - web-based manager

1. Go to *User & Device > User > User Definition* and select *Create New*.

2. Enter the following information and select *OK*:

<b>User Name</b>	User2
<b>Match user on RADIUS server</b>	Select this option and then select OurRADIUSsrv from the list.

#### To configure a remote user - CLI

```
config user local
 edit User2
 set name User2
 set type radius
 set radius-server OurRADIUSsrv
 end
```

## Creating user groups

There are two user groups: an FSSO user group for FSSO users and a firewall user group for other users. It is not possible to combine these two types of users in the same user group.

### Creating the FSSO user group

For this example, assume that FSSO has already been set up on the Windows network and that it uses Advanced mode, meaning that it uses LDAP to access user group information. You need to

- configure LDAP access to the Windows AD global catalog
- specify the collector agent that sends user logon information to the FortiGate unit
- select Windows user groups to monitor
- select and add the Engineering and Sales groups to an FSSO user group

#### To configure LDAP for FSSO - web-based manager

1. Go to *User & Device > Authentication > LDAP Servers* and select *Create New*.
2. Enter the following information:

<b>Name</b>	ADserver
<b>Server Name / IP</b>	10.11.101.160
<b>Distinguished Name</b>	dc=office,dc=example,dc=com
<b>Bind Type</b>	Regular
<b>User DN</b>	cn=FSSO_Admin,cn=users,dc=office,dc=example,dc=com
<b>Password</b>	set_a_secure_password

Leave other fields at their default values.

3. Select *OK*.

### To configure LDAP for FSSO - CLI

```
config user ldap
 edit "ADserver"
 set server "10.11.101.160"
 set dn "cn=users,dc=office,dc=example,dc=com"
 set type regular
 set username
 "cn=administrator,cn=users,dc=office,dc=example,dc=com"
 set password set_a_secure_password
 next
end
```

### To specify the collector agent for FSSO - web-based manager

1. Go to *User & Device > Authentication > Single Sign-On*.
2. Select *Fortinet Single Sign-On Agent*.
3. Enter the following information:

<b>Name</b>	WinGroups
<b>Primary Agent IP/Name</b>	10.11.101.160
<b>Password</b>	fortinet_canada
<b>LDAP Server</b>	ADserver

4. Select *Apply & Refresh*.

In a few minutes, the FortiGate unit downloads the list of user groups from the server.

### To specify the collector agent for FSSO - CLI

```
config user fsso
 edit "WinGroups"
 set ldap-server "ADserver"
 set password ENC
 G7GQV7NEqilCM9jKmVmJJFVvhQ2+wtNEe9T0iYA5Sa+EgT2J8zhOrbkJFD
 r0RmY3c4LaoXdsoBczA1dONmcGfthTxxwGsigzGpbJdC71spFlQYtj
 set server "10.11.101.160"
 end
```

### To create the FSSO\_Internet-users user group - web-based manager

1. Go to *User & Device > User > User Groups* and select *Create New*.
2. Enter the group name, *FSSO\_Internet\_users*.
3. Select *Fortinet Single Sign-On (FSSO)*.
4. In the *Available Members* list, select the *Engineering* and *Sales* groups and then select the right arrow button to move them to the *Members* list.
5. Select *OK*.



### To create the FSSO\_Internet-users user group - CLI

```
config user group
 edit FSSO_Internet_users
 set group-type fsso-service
 set member CN=Engineering,cn=users,dc=office,dc=example,dc=com
 CN=Sales,cn=users,dc=office,dc=example,dc=com
 end
```

### Creating the Firewall user group

The non-FSSO users need a user group too. In this example, only two users are shown, but additional members can be added easily.

#### To create the firewall user group - web-based manager

1. Go to *User & Device > User > User Groups* and select *Create New*.
2. Enter the following information and then select *OK*:

<b>Name</b>	Internet_users
<b>Type</b>	Firewall
<b>Members</b>	User1, User2

#### To create the firewall user group - CLI

```
config user group
 edit Internet_users
 set group-type firewall
 set member User1 User2
 end
```

### Defining policy addresses

Go to *Firewall Objects > Address > Addresses* and create the following addresses:

<b>Address Name</b>	Internal_net
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	10.11.102.0/24
<b>Interface</b>	Port 3

<b>Address Name</b>	Windows_net
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	10.11.101.0/24
<b>Interface</b>	Port 2

## Creating security policies

Two security policies are needed: one for firewall group who connect through port3 and one for FSSO group who connect through port2.

### To create a security policy for FSSO authentication - web-based manager

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Enter the following information:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	User Identity
<b>Incoming Interface</b>	Port2
<b>Source address</b>	Windows_net
<b>Outgoing Interface</b>	Port1
<b>Enable NAT</b>	Select

3. In Configuration *Authentication Rules*, select *Create New*.  
In the *New Authentication Rule* window, enter the following information, and then select OK:

<b>Destination Address</b>	all
<b>Group(s)</b>	FSSO_Internet_users
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>UTM Security Profiles</b>	Optionally, enable UTM profiles.

4. Select OK.

### To create a security policy for FSSO authentication - CLI

```
config firewall policy
 edit 0
 set srcintf port2
 set dstintf port1
 set srcaddr Windows_net
 set dstaddr all
 set action accept
 set identity-based enable
 set nat enable
 config identity-based-policy
 edit 1
 set schedule always
 set groups FSSO_Internet_users
 set service ANY
 end
 end
end
```

### To create a security policy for local user authentication - web-based manager

1. Go to *Policy > Policy* and select *Create New*.
2. Enter the following information:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	User Identity
<b>Incoming Interface</b>	Port3
<b>Source address</b>	Internal_net
<b>Outgoing Interface</b>	Port1
<b>Enable NAT</b>	Select

3. In Configuration *Authentication Rules*, select *Create New*.  
In the *New Authentication Rule* window, enter the following information, and then select OK:

<b>Destination Address</b>	all
<b>Group(s)</b>	Internet_users
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>UTM Security Profiles</b>	Optionally, enable UTM profiles.

4. Select OK.

### To create a security policy for local user authentication - CLI

```
config firewall policy
 edit 0
 set srcintf port3
 set dstintf port1
 set srcaddr internal_net
 set dstaddr all
 set action accept
 set identity-based enable
 set nat enable
 config identity-based-policy
 edit 1
 set schedule always
 set groups Internet_users
 set service ANY
 end
 end
 end
```

## LDAP Dial-in using member-attribute

In this example, users defined in MicroSoft Windows Active Directory (AD) are allowed to set up a VPN connection simply based on an attribute that is set to TRUE, instead of based on their user group. In AD the "Allow Dialin" property is activated in the user properties, and this sets the `msNPAllowDialin` attribute to "TRUE".

This same procedure can be used for other member attributes, as your system requires.

To accomplish this with a FortiGate unit, member-attribute must be set. This can only be accomplished through the CLI - the option is not available through the web-based manager.

Before configuring the FortiGate unit, ensure the AD server has the `msNPAllowDialin` attribute set to "TRUE" for the users in question. If not, those users will not be able to authenticate.

### To configure user LDAP member-attribute settings - CLI

```
config user ldap
 edit "ldap_server"
 set server "192.168.201.3"
 set cnid "sAMAccountName"
 set dn "DC=fortilabanz,DC=com,DC=au"
 set type regular
 set username "fortigate@sample.com"
 set password *****
 set member-attr "msNPAllowDialin"
 next
end
```

### To configure LDAP group settings - CLI

```
config user group
 edit "ldap_grp"
 set member "ldap"
 config match
 edit 1
 set server-name "ldap"
 set group-name "TRUE"
 next
 end
 next
end
```

Once these settings are in place, users that are a member of the `ldap` user group will be able to authenticate.

To ensure your settings are correct, here is the sample output from a `diag debug` command that shows the authentication process.

When the "Allow Dial-in" attribute is set to "TRUE" the following will likely be in the output:

```
get_member_of_groups-Get the memberOf groups.
get_member_of_groups- attr='msNPAllowDialin', found 1 values
get_member_of_groups-val[0]='TRUE'
fnbamd_ldap_get_result-Auth accepted
fnbamd_ldap_get_result-Going to DONE state res=0
fnbamd_auth_poll_ldap-Result for ldap svr 192.168.201.3 is SUCCESS
fnbamd_auth_poll_ldap-Passed group matching
```

If the attribute is not set but it is expected, the following will likely be in the output:

```
get_member_of_groups-Get the memberOf groups.
get_member_of_groups- attr='msNPAllowDialin', found 1 values
get_member_of_groups-val[0]='FALSE'
fnbamd_ldap_get_result-Auth accepted
fnbamd_ldap_get_result-Going to DONE state res=0
fnbamd_auth_poll_ldap-Result for ldap svr 192.168.201.3 is SUCCESS
fnbamd_auth_poll_ldap-Failed group matching
```

The only difference between these two outputs is the last line which is either passed or failed based on if the member-attribute is set to the expected value or not.

## RADIUS SSO example

A common RADIUS SSO topology involves a medium sized company network of users connecting to the Internet through the FortiGate unit, and authenticating with a RADIUS server. RADIUS SSO authentication was selected because it is fast and relatively easy to configure.

This section includes:

- [Assumptions](#)
- [Topology](#)
- [General configuration](#)
- [Configuring RADIUS](#)
- [Configuring FortiGate regular and RADIUS SSO security policies](#)
- [Testing](#)


### Assumptions

- VDOMs are not enabled
- The admin super\_admin administrator account will be used for all FortiGate unit configuration.
- Any other devices on the network do not affect the topology of this example, and therefore are not included.
- Anywhere settings are not described, they are assumed to be default values.
- A RADIUS server is installed on a server or FortiAuthenticator unit and uses default attributes.
- BGP is used for any dynamic routing.
- Authentication event logging under Log&Report has been configured.

## Topology

Example.com has an office with 20 users on the internal network. These users need access to the Internet to do their jobs. The office network is protected by a FortiGate-60C unit with access to the Internet through the wan1 interface, the user network on the internal interface, and all the servers are on the DMZ interface. This includes an Ubuntu Linux server running FreeRADIUS. For this example only two users will be configured — Pat Lee with an account name `plee`, or `plee@example.com`, and Kelly Green with an account name `kgreen`, or `kgreen@example.com`.

**Figure 130:**RADIUS SSO topology



## General configuration

1. [Configuring RADIUS with users, user group, and FortiGate information.](#)
2. [Configuring FortiGate interfaces](#)
3. [Configuring FortiGate regular and RADIUS SSO security policies](#)

## Configuring RADIUS

Configuring RADIUS includes configuring the RADIUS server such as FreeRADIUS, a radius client on user's computers, and configuring users in the system. For this example the two users will be Pat Lee, and Kelly Green. They belong to a group called `exampledotcom_employees`. When it is all configured, the RADIUS daemon needs to be started.

The users have a RADIUS client installed on their PCs that allows them to authenticate through the RADIUS server.

FreeRADIUS can be found on the [freeradius.org](http://freeradius.org) website. For any problems installing FreeRADIUS, see the FreeRADIUS documentation.

## Configuring FortiGate interfaces

Before configuring the RADIUS SSO security policy, configure FortiGate interfaces. This includes defining a DHCP server for the internal network as this type of network typically uses DHCP. The wan1 and dmz interfaces are assigned static IP addresses and do not need a DHCP server.

**Table 28:** FortiGate interfaces used in this example

Interface	Subnet	Act as DHCP Server	Devices
wan1	172.20.120.141	No	Internet Service Provider
dmz	10.11.101.100	No	Servers, including RADIUS server
internal	10.11.102.100	Yes: x.x.x.110-.250	Internal user network

**To configure FortiGate interfaces - web-based manager**

1. Go to *System > Network > Interfaces*.
2. Select wan1 to edit.
3. Enter the following information and select OK.

<b>Alias</b>	Internet
<b>Addressing Mode</b>	Manual
<b>IP/Network Mask</b>	172.20.120.141/255.255.255.0
<b>Administrative Access</b>	HTTPS, SSH
<b>Enable DHCP Server</b>	Not selected
<b>Comments</b>	Internet
<b>Administrative Status</b>	Up

4. Select dmz to edit.
5. Enter the following information and select OK.

<b>Alias</b>	Servers
<b>Addressing Mode</b>	Manual
<b>IP/Network Mask</b>	10.11.101.100/255.255.255.0
<b>Administrative Access</b>	HTTPS, SSH, PING, SNMP
<b>Enable DHCP Server</b>	Not selected
<b>Listen for RADIUS Accounting Messages</b>	Select
<b>Comments</b>	Servers
<b>Administrative Status</b>	Up

6. Select internal to edit.

7. Enter the following information and select OK.

<b>Alias</b>	Internal network
<b>Addressing Mode</b>	Manual
<b>IP/Network Mask</b>	10.11.102.100/255.255.255.0
<b>Administrative Access</b>	HTTPS, SSH, PING
<b>Enable DHCP Server</b>	Select
<b>Address Range</b>	10.11.102.110 - 10.11.102.250
<b>Netmask</b>	255.255.255.0
<b>Default Gateway</b>	Same as Interface IP
<b>DNS Server</b>	Same as System DNS
<b>Comments</b>	Internal network
<b>Administrative Status</b>	Up

## Configuring a RADIUS SSO Agent on the FortiGate unit

### To create a RADIUS SSO agent

1. Go to *User & Device > Authentication > Single Sign-On* and select *Create New*.
2. In *Type*, select *RADIUS Single-Sign-On Agent*.
3. Select *Use RADIUS Shared Secret* and enter the RADIUS server shared secret.
4. Select *Send RADIUS Responses*.
5. Select *OK*.

The Single Sign-On agent is named `RSSO_Agent`.

## Creating a RADIUS SSO user group

### To define a local user group for RADIUS SSO

1. Go to *User & Device > User > User Groups* and select *Create New*.
2. Enter a Name for the user group.
3. In *Type*, select *RADIUS Single Sign-On (RSSO)*.
4. In *RADIUS Attribute Value*, enter the name of the RADIUS user group this local user group represents.
5. Select *OK*.



## Configuring FortiGate regular and RADIUS SSO security policies

With the RADIUS server and FortiGate interfaces configured, security policies can be configured. This includes both RADIUS SSO and regular policies, as well as addresses and address groups. All policies require NAT to be enabled.

**Table 29:** security policies needed for RADIUS SSO

Seq. No.	From -> To	Type	Schedule	Description
1	internal -> wan1	RADIUS SSO	business hours	Authenticate outgoing user traffic.
2	internal -> wan1	regular	always	Allow essential network services and VoIP.
3	dmz -> wan1	regular	always	Allow servers to access Internet.
4	internal -> dmz	regular	always	Allow users to access servers.
5	any -> any	deny	always	Implicit policy denying all traffic that hasn't been matched



The RADIUS SSO policy must be placed at the top of the policy list so it is matched first. The only exception to this is if you have a policy to deny access to a list of banned users. In this case, that policy must go at the top so the RADIUS SSO does not mistakenly match a banned user or IP address.

This section includes:

- [Schedules, address groups, and services groups](#)
- [Configuring regular security policies](#)
- [Configuring RADIUS SSO security policy](#)

### Schedules, address groups, and services groups

This section lists the lists that need to be configured before security policies are created. Creating these lists is straight forward, so the essential information has been provided here but not step by step instructions. For more information on firewall related details, see

#### Schedules

Only one schedule needs to be configured — `business_hours`. This is a fairly standard Monday to Friday 8am to 5pm schedule, or whatever days and hours covers standard work hours at the company.

#### Address groups

The following address groups need to be configured before the security policies.

**Table 30:**

Address group Name	Interface	Address range included
--------------------	-----------	------------------------

**Table 30:**

internal_network	internal	10.11.102.110 to 10.11.102.250
company_servers	dmz	10.11.101.110 to 10.11.101.250

**Service groups**

The following service groups need to be configured before the security policies. Note that the services listed are suggestions and may include more or less as required.

**Table 31:**

Service group Name	Interface	Description of services to be included
essential_network_services	internal	Any network protocols required for normal network operation such as DNS, NTP, BGP.
essential_server_services	dmz	All the protocols required by the company servers such as BGP, HTTP, HTTPS, FTP, IMAP, POP3, SMTP, IKE, SQL, MYSQL, NTP, TRACEROUTE, SOCKS, and SNMP.
user_services	internal	Any protocols required by users HTTP, HTTP, FTP,

The following security policy configurations are basic and only include logging, and default AV and IPS.

**Configuring regular security policies**

Regular security policies allow or deny access for non-RADIUS SSO traffic. This is essential as there are network services—such as DNS, NTP, and FortiGuard—that require access to the Internet.

**To configure regular security policies - web-based manager**

1. Go to *Policy > Policy*, and select Create New.
2. Enter the following information, and select OK.

<b>Source Interface/Zone</b>	Internal
<b>Source Address</b>	internal_network
<b>Destination Interface/Zone</b>	wan1
<b>Destination Address</b>	all
<b>Schedule</b>	always
<b>Service</b>	essential_network_services
<b>Action</b>	ACCEPT
<b>Log Allowed Traffic</b>	enable
<b>Enable NAT</b>	enable
<b>UTM</b>	enable

<b>Enable Antivirus</b>	enable Default
<b>Enable IPS</b>	enable Default
<b>Enable VoIP</b>	enable Default
<b>Comments</b>	Essential network services

3. Select Create New, enter the following information, and select OK.

<b>Source Interface/Zone</b>	dmz
<b>Source Address</b>	company_servers
<b>Destination Interface/Zone</b>	wan1
<b>Destination Address</b>	all
<b>Schedule</b>	always
<b>Service</b>	essential_server_services
<b>Action</b>	ACCEPT
<b>Log Allowed Traffic</b>	enable
<b>Enable NAT</b>	enable
<b>UTM</b>	enable
<b>Enable Antivirus</b>	enable Default
<b>Enable IPS</b>	enable Default
<b>Comments</b>	Company servers accessing the Internet

4. Select Create New, enter the following information, and select OK.

<b>Source Interface/Zone</b>	Internal
<b>Source Address</b>	internal_network
<b>Destination Interface/Zone</b>	dmz
<b>Destination Address</b>	company_servers
<b>Schedule</b>	always
<b>Service</b>	all
<b>Action</b>	ACCEPT
<b>Log Allowed Traffic</b>	enable
<b>Enable NAT</b>	enable
<b>UTM</b>	enable

<b>Enable Antivirus</b>	enable Default
<b>Enable IPS</b>	enable Default
<b>Comments</b>	Access company servers

## Configuring RADIUS SSO security policy

The RADIUS SSO policy allows access for members of specific RADIUS groups.

### To configure RADIUS SSO security policy

1. Go to *Policy > Policy > Policy*.
2. Select *Create New*.
3. Enter the following information.

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	User Identity
<b>Incoming Interface</b>	Internal
<b>Source Address</b>	internal_network
<b>Outgoing Interface</b>	wan1
<b>Enable NAT</b>	Selected

4. In *Configure Authentication Rules*, select *Create New* and enter:

<b>Destination Address</b>	all
<b>Group(s)</b>	Select the user groups you created for RSSO.
<b>User(s)</b>	not used
<b>Schedule</b>	business_hours
<b>Service</b>	ALL
<b>Action</b>	ACCEPT
<b>UTM Security Profiles</b>	Enable AntiVirus, WebFilter, IPS, and Email Filter. In each case, select the default profile.

5. Select *OK*.
6. Repeat steps 4 and 5 for each user group that is allowed to use this security policy. Schedule, Service, and UTM profiles can be different for each group.
7. Select *OK*.

To ensure an RSSO-related policy is matched first, the policy should be placed higher in the security policy list than more general policies for the same interfaces.

## Testing

Once configured, a user only needs to logon to their PC using their RADIUS account. After that when they attempt to access an Internet website, the FortiGate unit will use their session

information to get their RADIUS information. Once the user is verified, they are allowed access to the website.


To test the configuration perform the following steps.

1. Have user plee logon to their PC, and try to access an Internet website.
2. The FortiGate unit will contact the RADIUS server for user plee's information.
3. Once confirmed, plee will have access to the website.

Each step generates log entries that enable you to verify that each step was successful.

If a step is unsuccessful, confirm your configuration is correct and see [“Troubleshooting dynamic profiles” on page 230](#).

**Figure 131:**RADIUS SSO test



## Troubleshooting

In the web-based manager, a good tool for troubleshooting is the packet counter column on the security policy page (*Policy > Policy*). This column displays the number of packets that have passed through this security policy. Its value when you are troubleshooting is that when you are testing your configuration (end to end connectivity, user authentication, policy use) watching the packet count for an increase confirms any other methods you may be using for troubleshooting. It provides the key of which policy is allowing the traffic, useful information if you expect a user to require authentication and it never happens. For more information about authentication security policies, see [“Authentication in security policies” on page 517](#).

This section addresses how to get more information from the CLI about users and user authentication attempts to help troubleshoot failed authentication attempts.

```
diag firewall iprope authuser
```

Shows the IP of where your computer is connected from. This is useful to confirm authorization and VPN settings.

```
diag firewall iprope resetauth
```

Clear all authorized users from the current list. Useful to force users to re-authenticate after system or group changes. However, this command may easily result in many users having to re-authenticate, so use carefully.

```
diag firewall auth list
```

List all the authorized users on this system.

```
diag rso query ip
```

```
diag rso query rso-key
```

Queries the RSSO database.

For more information on troubleshooting specific features, go to that section of this document. Most sections have troubleshooting information at the end of the section. In addition to that information, see the [FortiOS Handbook Troubleshooting chapter](#) for general troubleshooting information.

# Chapter 4 FortiOS Carrier

This FortiOS Handbook chapter contains the following sections:

[Overview of FortiOS Carrier features](#) provides an overview of the three major topics for FortiOS Carrier — Dynamic Profiles, MMS, and GTP.

[Carrier web-based manager settings](#) describes the web-based manager interface of FortiOS Carrier specific features.

[MMS Security features](#) describes FortiOS security features as they apply to MMS including MMS virus scanning, MMS file filtering, MMS content-based Antispam protection, and MMS DLP archiving.

[Message flood protection](#) describes setting thresholds to protect your MMS servers from receiving too many messages from the same sender.

[Duplicate message protection](#) describes setting thresholds to protect your MMS servers from receiving the same message from more than one sender.

[Configuring GTP on FortiOS Carrier](#) explains configuration of the more basic FortiOS Carrier GTP features.

[GTP message type filtering](#) explains this feature, and how to configure it on FortiOS Carrier.

[GTP identity filtering](#) explains this feature, and how to configure it on FortiOS Carrier.

[Troubleshooting](#) provides answer to common FortiOS Carrier GTP issues.

# Overview of FortiOS Carrier features

FortiOS Carrier specific features include Multimedia messaging service (MMS) protection, and GPRS Tunneling Protocol (GTP) protection.

All FortiGate units, carrier-enabled or not, are capable of handling Stream Control Transmission Protocol (SCTP) traffic, which is a protocol designed for and primarily used in Carrier networks.

This section includes:

- [Overview](#)
- [Registering FortiOS Carrier](#)
- [MMS background](#)
- [How FortiOS Carrier processes MMS messages](#)
- [MMS protection profiles](#)
- [Bypassing MMS protection profile filtering based on carrier endpoints](#)
- [Applying MMS protection profiles to MMS traffic](#)
- [GTP basic concepts](#)
- [Parts of a GTPv1 network](#)
- [GPRS network common interfaces](#)
- [Packet flow through the GPRS network](#)
- [SCTP](#)

## Overview

FortiOS Carrier provides all the features found on FortiGate units plus added features specific to carrier networks: MMS and GTP.

## MMS

MMS is a standard for sending messages that include multimedia content between mobile phones. MMS is also popular as a method of delivering news and entertainment content including videos, pictures, and text. Carrier networks include four different MMS types of messages — MM1, MM3, MM4, and MM7. See [“MMS background” on page 633](#).

## GTP

The GPRS Tunneling Protocol (GTP) runs on GPRS carrier networks. GPRS is a GSM packet radio standard. It provides more efficient usage of the radio interface so that mobile devices can share the same radio channel. FortiOS supports GTPv1 release 7.15.0 and GTPv1 release 8.12.0.

GPRS provides direct connections to the Internet (TCP/IP) and X.25 networks for point-to-point services (connection-less/connection oriented) and point-to-multipoint services (broadcast).

GPRS currently supports data rates from 9.6 kbps to more than 100 kbps, and it is best suited for burst forms of traffic. GPRS involves both radio and wired components. The mobile phone sends the message to a base station unit (radio based) that converts the message from radio to



wired, and sends the message to the carrier network and eventually the Internet (wired carrier network). See “GTP basic concepts” on page 645.

## Registering FortiOS Carrier

In FortiOS 5.0, the Carrier registration process has changed. Devices are no longer shipped with Carrier functions built-in, so that administrators will have more leeway in configuring Carrier functions on existing devices.

A license is purchased from Fortinet, and is delivered by mail as a scratch card. The contained registration code is entered into the CLI according to the instructions on the card. The FortiGate will then factory reset itself into Carrier mode.

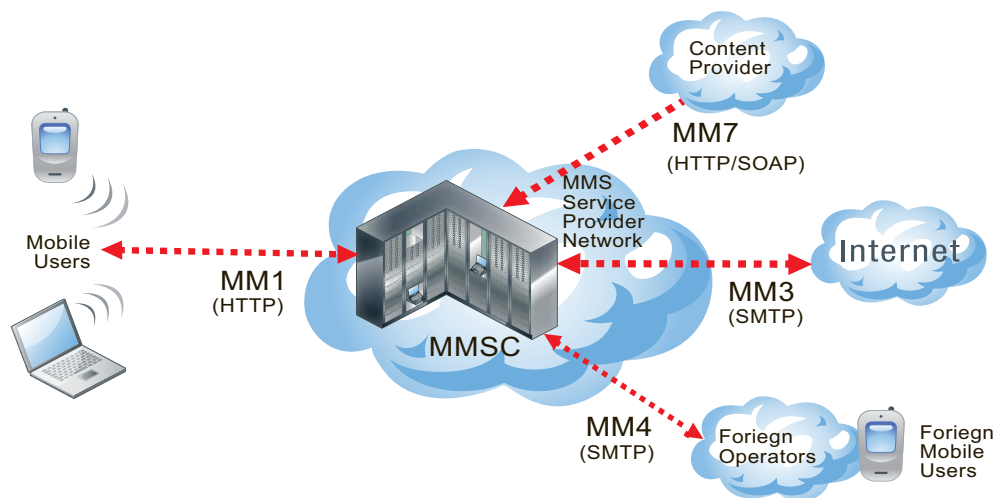
Only certain models support Carrier mode. Contact Fortinet Support for more information and firmware images.

## MMS background

MMS is a common method for mobile users to send and receive multimedia content. A Carrier network supports MMS across its network. This makes up the MMS Service Provider Network (MSPN).

Messages can be sent or received between the MMSC and a number of other services including the Internet, content providers, or other carriers. Each of these different service connections uses different MMS formats including MM1 and MM7 messages (essentially HTTP format), and MM3 and MM4 messages (SMTP formatted). These different formats reflect the different purposes and content for each type of MMS message.

**Figure 132:**MMS content interfaces



## MMS content interfaces

MMS messages are sent from devices and servers to other devices and servers using MMS content interfaces

There are eight interfaces defined for the MMS standard, referred to as MM1 through MM8. The most important of these interfaces for the transfer of data is the MM1 interface, as this defines how mobile users communicate from the mobile network to the Multimedia Message Service

Center (MMSC). MMS content to be monitored and controlled comes from these mobile users and is going to the provider network.

Other MMS content interfaces that connect a service provider network to other external sources can pose threats as well. MM3 handles communication between the Internet and the MMSC and is a possible source of viruses and other content problems from the Internet. MM4 handles communication between different content provider MMSCs. Filtering MM4 content protects the service provider network from content sent from foreign service providers and their subscribers. Finally MM7 is used for communication between content providers and the MMSC. Filtering MM3 content can also keep harmful content off of the service provider network.

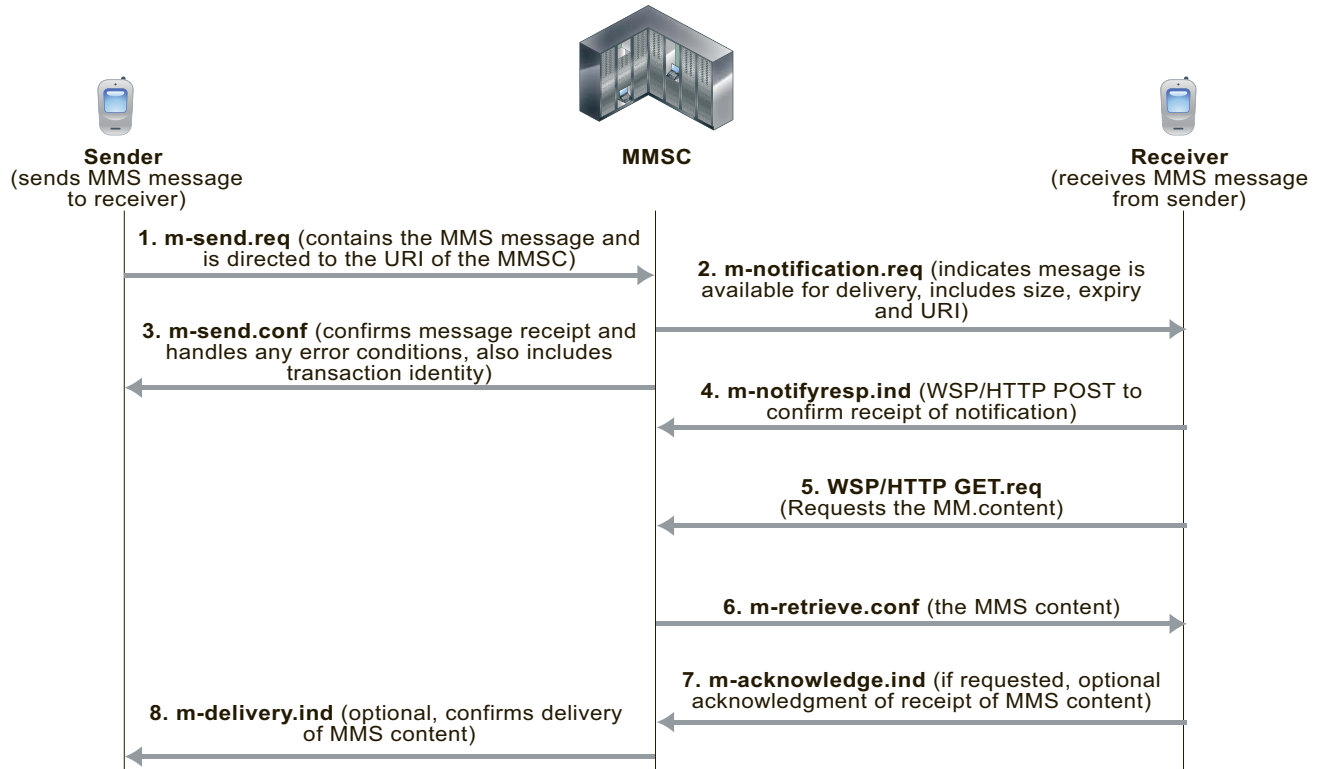
**Table 32:** MMS content interfaces

Type	Transaction	Similar to
<b>MM 1</b>	Handset to MMSC	HTTP
<b>MM 3</b>	Between MMSC and Internet	SMTP
<b>MM 4</b>	Between Operator MMSCs	SMTP
<b>MM 7</b>	Content Providers to MMSC	HTTP and SOAP

## How MMS content interfaces are applied

As shown in [Figure 133](#), the sender's mobile device encodes the MMS content in a form similar to MIME email message (MMS MIME content formats are defined by the MMS Message Encapsulation specification). The encoded message is then forwarded to the service provider's MMSC. Communication between the sending device and the MMSC uses the MM1 content interface. The MM1 content interface establishes a connection and sends an MM1 send request (`m-send.req`) message that contains the MMS message. The MMSC processes this request and sends back an MM1 send confirmation (`m-send.conf`) HTTP response indicating the status of the message — accepted or an error occurred, for example.

**Figure 133:**MM1 transactions between senders and receivers and the MMSC

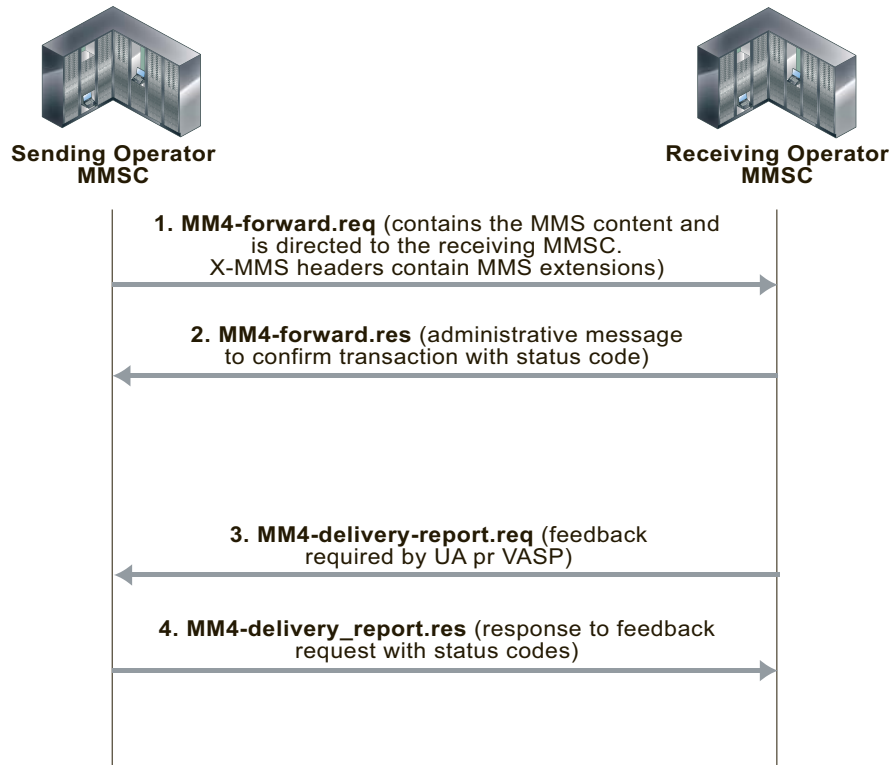


If the recipient is on another carrier, the MMSC forwards the message to the recipient's carrier. This forwarding uses the MM4 content interface for forwarding content between operator MMSCs (see [Figure 134](#)).

Before the MMSC can forward the message to the final recipient, it must first determine if the receiver's handset can receive MMS messages using the MM1 content interface. If the recipient can use the MM1 content interface, the content is extracted and sent to a temporary storage server with an HTTP front-end.

To retrieve the message, the receiver's handset establishes a connection with the MMSC. An HTTP get request is then sent from the recipient to the MMSC. This message contains the URL where the content of the message is stored. The MMSC responds with a retrieve confirmation (`m-retrieve.conf`) HTTP response that contains the message.

**Figure 134:**MM4 messages sent between operator MMSCs



This causes the receiver's handset to retrieve the content from the embedded URL. Several messages are exchanged to indicate status of the delivery attempt. Before delivering content, some MMSCs also include a content adaptation service that attempts to modify the multimedia content into a format suitable for the recipient's handset.

If the receiver's handset is not MM1 capable, the message can be delivered to a web based service and the receiver can view the content from a normal Internet browser. The URL for the content can be sent to the receiver in an SMS text message. Using this method, non-MM1 capable recipients can still receive MMS content.

The method for determining whether a handset is MMS capable is not specified by the standards. A database is usually maintained by the operator, and in it each mobile phone number is marked as being associated with a legacy handset or not. It can be a bit hit and miss since customers can change their handset at will and this database is not usually updated dynamically.


Email and web-based gateways from MMSC to the Internet use the MM3 content interface. On the receiving side, the content servers can typically receive service requests both from WAP and normal HTTP browsers, so delivery via the web is simple. For sending from external sources to handsets, most carriers allow MIME encoded message to be sent to the receiver's phone number with a special domain.

## How FortiOS Carrier processes MMS messages

MMS messages can be vectors for propagating undesirable content such as spam and viruses. FortiOS Carrier can scan MMS messages sent using the MM1, MM3, MM4, and MM7 content interfaces. You can configure FortiOS Carrier to scan MMS messages for spam and viruses by configuring and adding MMS protection profiles and adding the MMS protection profiles to security policies. You can also use MMS protection profiles to apply content blocking, carrier

endpoint filtering, MMS address translation, sending MMS notifications, DLP archiving of MMS messages, and logging of MMS message activity.

**Figure 135:**FortiOS Carrier MMS processing



FortiOS Carrier can send MMS messages to senders informing those senders that their devices are infected. FortiOS Carrier can also send MMS notifications to administrators to inform them of suspicious activity on their networks.

For message floods and duplicate messages, FortiOS Carrier does not send notifications to message senders but does send notifications to administrators and sends messages to sender handsets to complete MM1 and MM4 sessions.

Where MMS messaging uses the TCP/IP set of protocols, SMS text messaging uses the Signaling System Number 7 (SS7) set of protocols, which is not supported by FortiOS.

## FortiOS Carrier and MMS content scanning

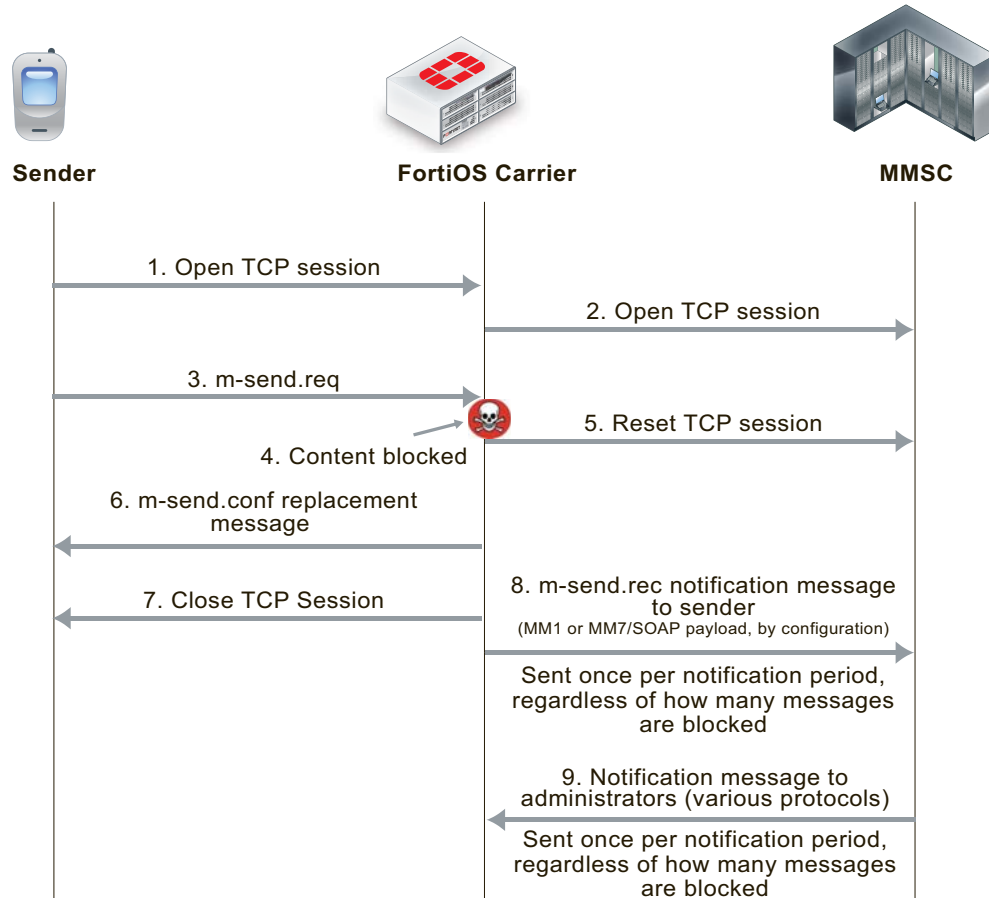
The following section applies to MMS content scanning, including virus scanning, file filtering, content spam filtering, carrier endpoint filtering, and MMS content checksum filtering.

### MM1 Content Scanning

During MM1 content scanning a message is first transmitted from the sender, establishing a connection with the MMSC. FortiOS Carrier intercepts this connection and acts as the endpoint. FortiOS Carrier then establishes its own connection to the MMSC. Once connected, the client transmits its `m-send.req` HTTP post request to FortiOS Carrier which scans it according to the MMS protection profile settings. If the content is clean, the message is forwarded to the MMSC. The MMSC returns `m-send.conf` HTTP response through FortiOS Carrier to the sender.

If FortiOS Carrier blocks the message (for example because a virus was found, see [Figure 136](#)), FortiOS Carrier resets the connection to the MMSC and sends `m-send.conf` HTTP response back to the sender. The response message can be customized using replacement messages. FortiOS Carrier then terminates the connection. Sending back an `m-send.conf` message prevents the sender from trying to send the message again.

**Figure 136:**MM1 MMS scanning of message sent by sender (blocking m.send.req messages)



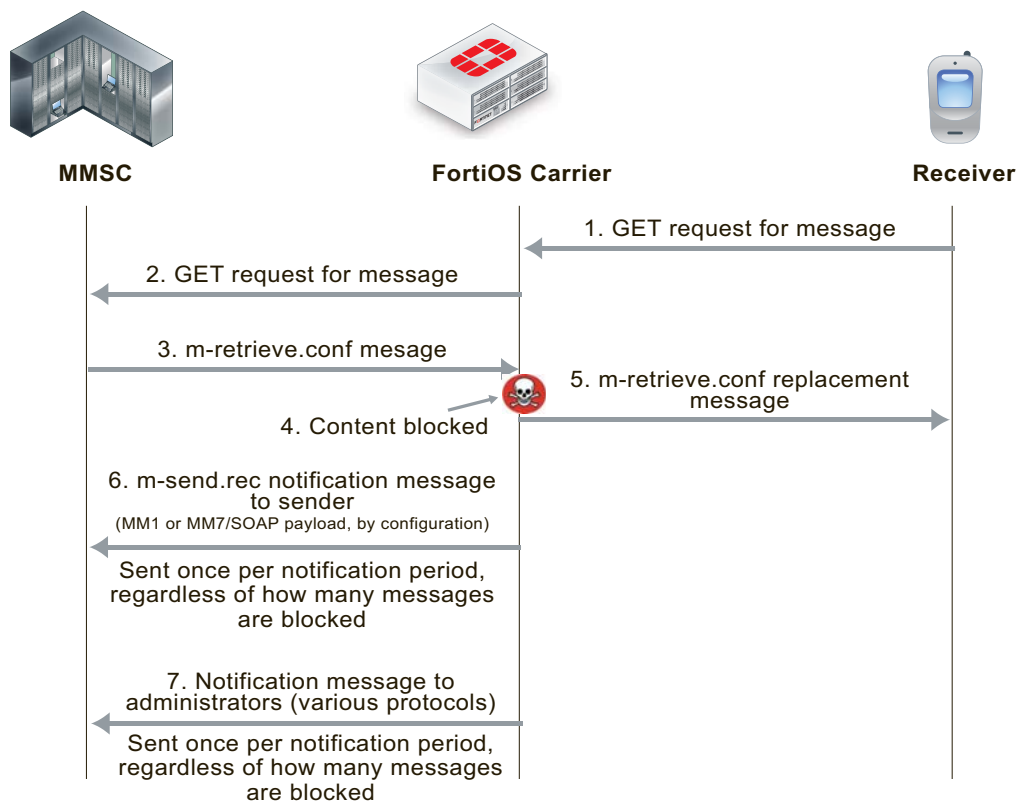
FortiOS Carrier also sends `m-send.rec` notifications messages to the MMSC that are then forwarded to the sender to notify them of blocked messages.

### Filtering message retrieval

FortiOS Carrier intercepts the connection to the MMSC, and the `m-retrieve.conf` HTTP response from the MMSC is scanned according to the MMS content scanning settings. If the content is clean, the response is forwarded back to the client. If the content is blocked, FortiOS Carrier drops the connection to the MMSC. It then builds an `m-retrieve.conf` message from the associated replacement message and transmits this back to the client.

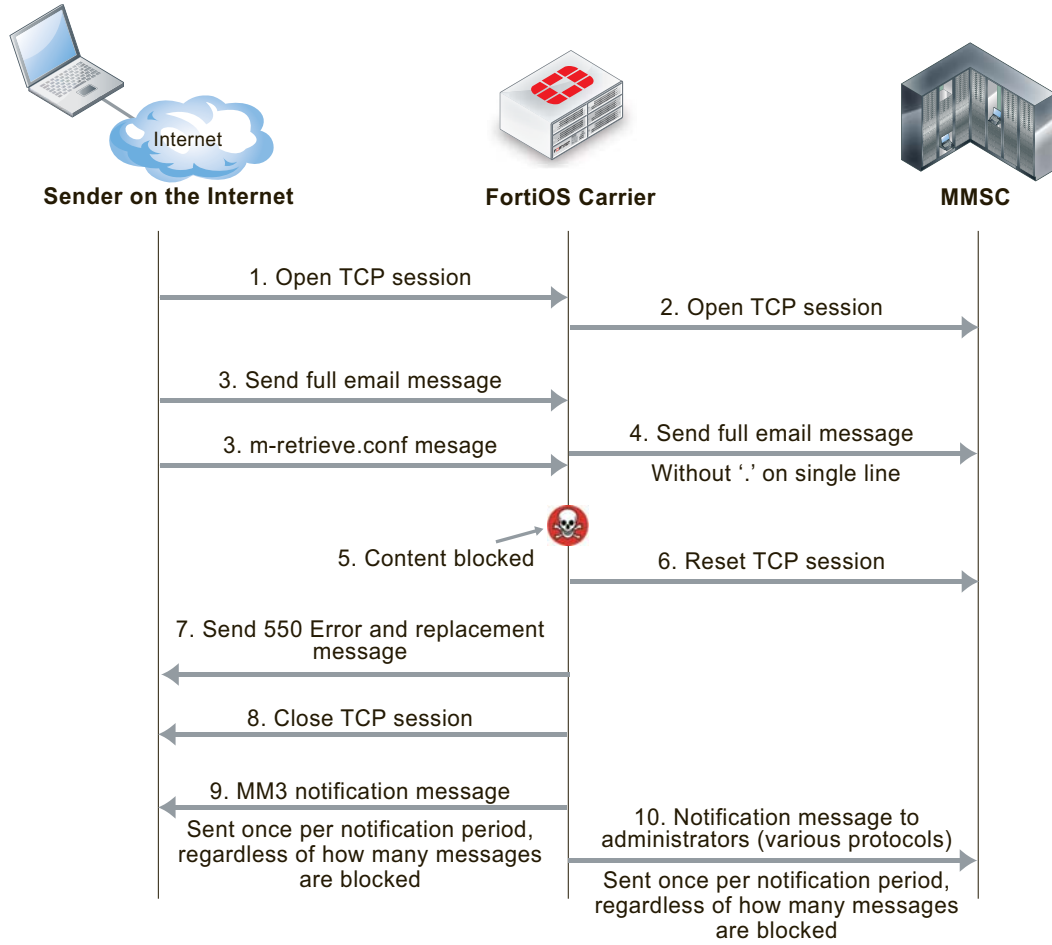
FortiOS Carrier also sends `m-send.rec` notifications messages to the MMSC that are then forwarded to the receiver to notify them of blocked messages.

**Figure 137:MM1 MMS scanning of messages received by receiver (blocking m.retrieve.conf messages)**



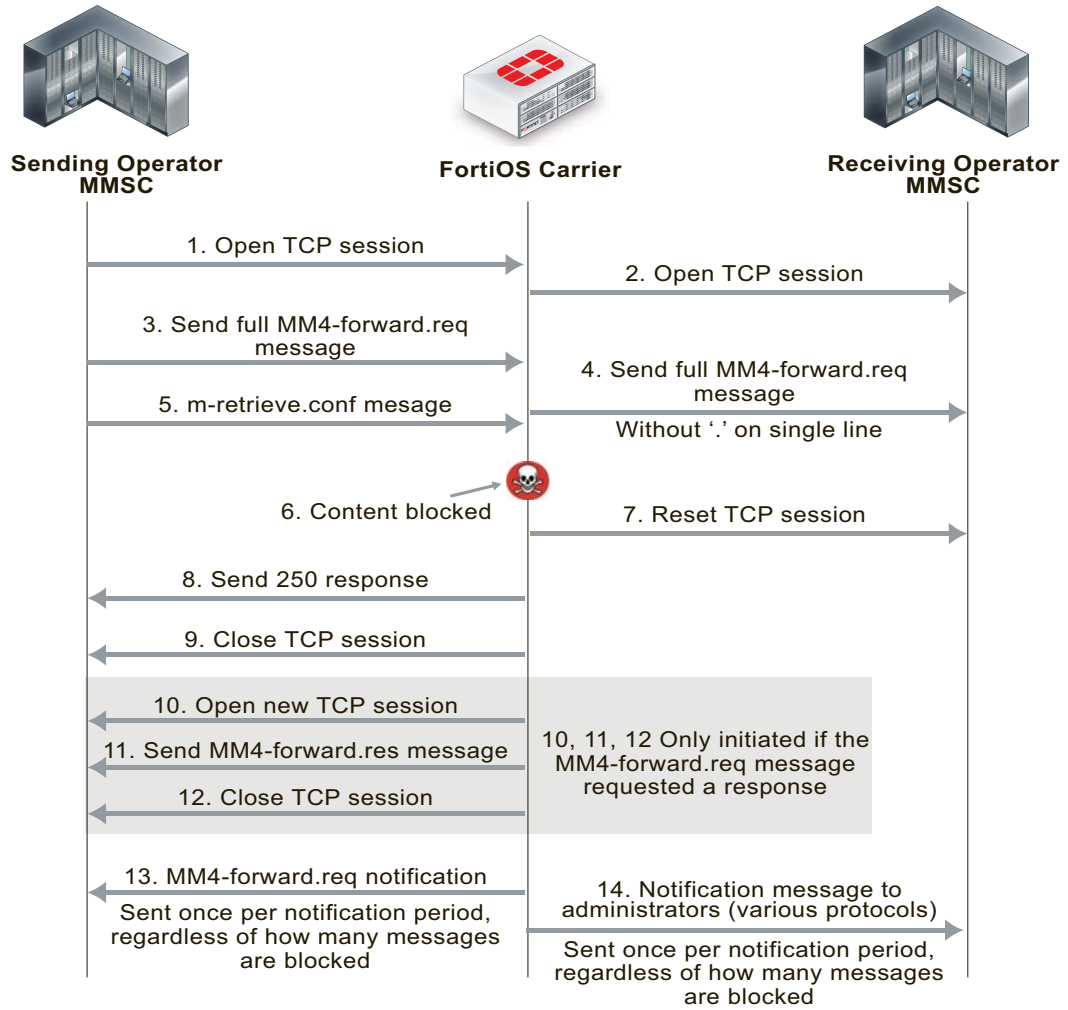
Filtering MM3 and MM4 messages works in a similar way to MM1 (see [Figure 138](#) and [Figure 139](#)). FortiOS Carrier intercepts connections to the MMSC, and scans messages as configured. When messages are blocked, FortiOS Carrier closes sessions as required, sends confirmation messages to the sender, notifies administrators, and notifies senders and receivers of messages.

**Figure 138:MM3 MMS scanning of messages sent from a sender on the Internet to an MMSC**

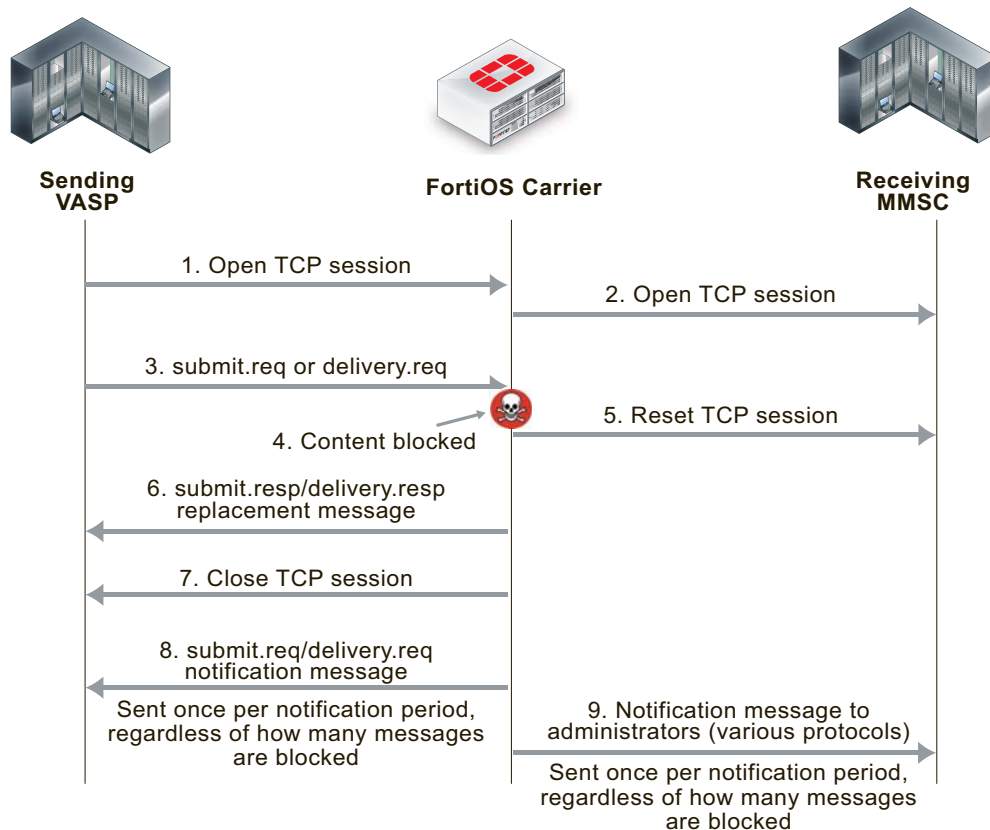




**Figure 139:MM4 MMS scanning of messages sent between operator MMSCs**



**Figure 140:MM7 MMS scanning of messages sent between a VASP and an MMSC**



## FortiOS Carrier and MMS duplicate messages and message floods

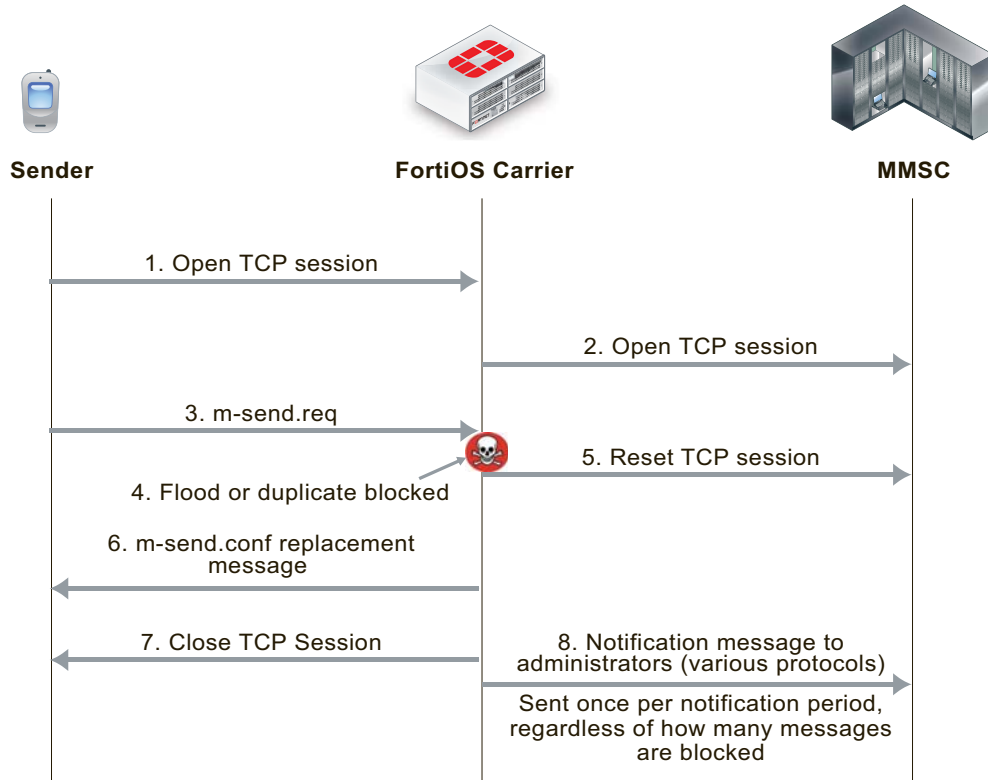
FortiOS Carrier detects duplicate messages and message floods for the MM1 and MM4 interfaces. How FortiOS Carrier detects and responds to duplicate messages and message floods is different from how FortiOS Carrier detects and responds to viruses and other MMS scanning protection measures.

For message floods and duplicate messages, the sender does not receive notifications about floods or duplicate messages, as if the sender is an attacker they can gain useful information about flood and duplicate thresholds. Plus, duplicate messages and message floods are usually a result of a large amount of messaging activity and filtering of these messages is designed to reduce the amount of unwanted messaging traffic. Adding to the traffic by sending notifications to senders and receivers could result in an increase in message traffic.

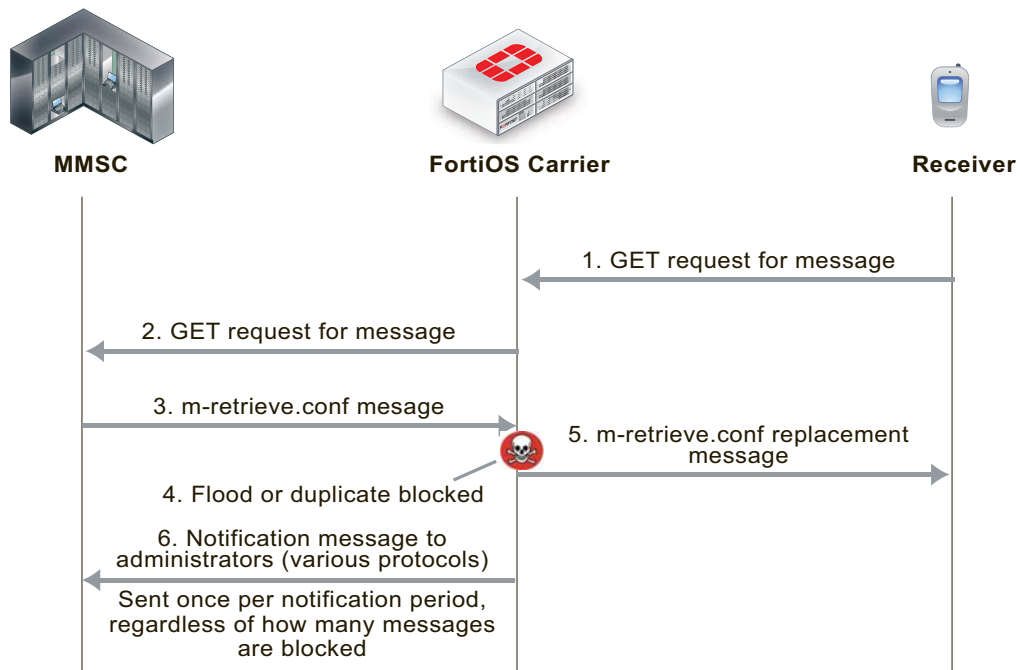
You can create up to three thresholds for detecting duplicate messages and message floods. For each threshold you can configure the FortiOS Carrier unit to respond by logging the activity, archiving or quarantining the messages, notifying administrators of the activity, and by blocking the messages. In many cases you may only want to configure blocking for higher activity thresholds, and to just monitor and send administrator notifications at lower activity thresholds.

When a block threshold is reached for MM1 messages, FortiOS Carrier sends `m-send.conf` or `m-retrieve.conf` messages to the originator of the activity. These messages are sent to end the MM1 sessions, otherwise the originator would continue to re-send the blocked message. When a block threshold is reached for MM4, FortiOS Carrier sends a `MM4-forward.res` message to close the MM4 session. An MM4 message is sent only if initiated by the originating `MM4-forward.req` message.

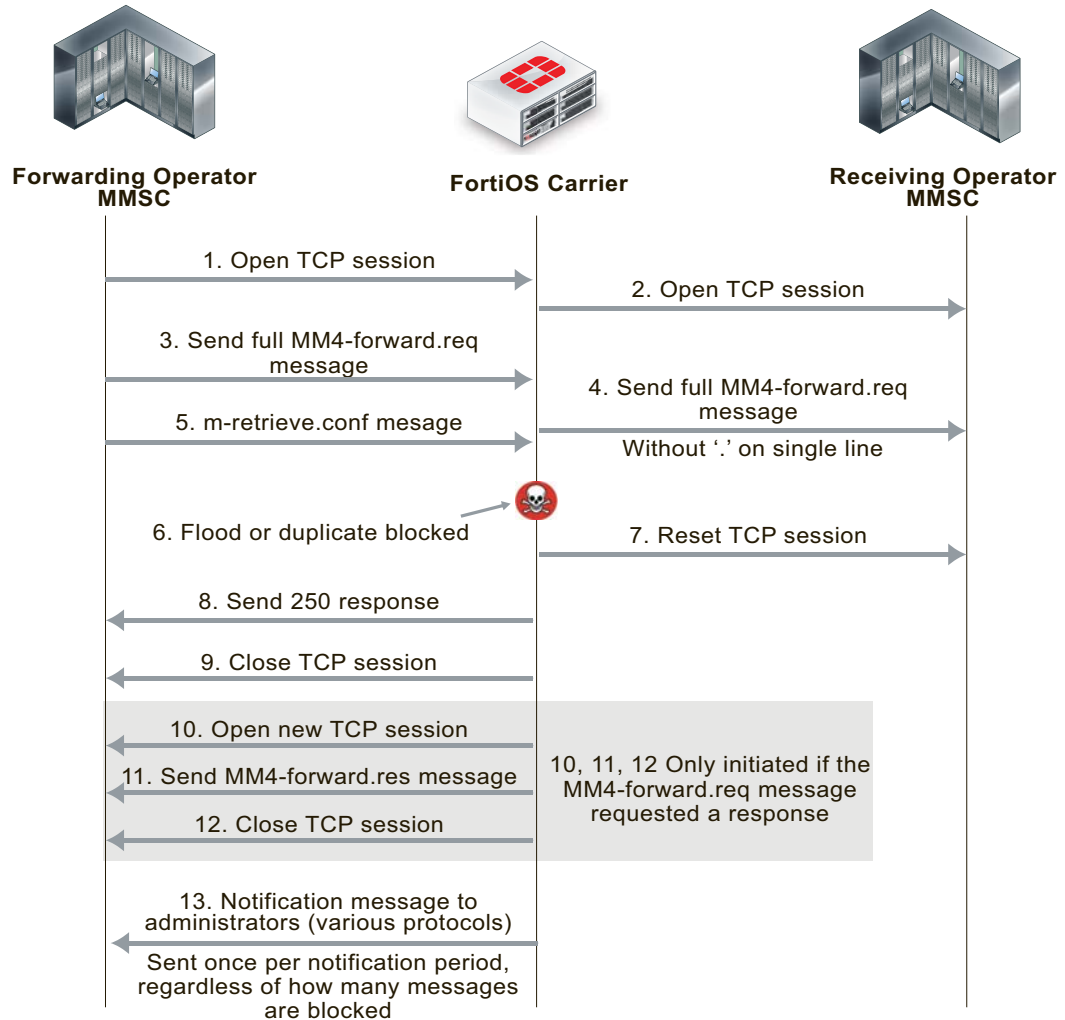
**Figure 141:**MM1 message flood and duplicate message blocking of sent messages



**Figure 142:**MM1 message flood and duplicate message blocking of received messages



**Figure 143:**MM4 message flood and duplicate message blocking



## MMS protection profiles

An MMS protection profile is a group of settings that you can apply to an MMS session matched by a security policy.

MMS protection profiles are easy to configure and can be used by more than one security policy. You can configure a single MMS protection profile for the different traffic types handled by a set of security policies that require identical protection levels and types. This eliminates the need to repeatedly configure those same MMS protection profile settings for each individual security policy.

For example, while traffic between trusted and untrusted networks might need strict protection, traffic between trusted internal addresses might need only moderate protection. You would configure two separate MMS protection profiles to provide the different levels of protection: one for traffic between trusted networks, and one for traffic between trusted and untrusted networks.

Once you have configured the MMS Protection Profile, you need to add it to a security policy to apply the profile to MMS traffic.

See [“MMS Security features” on page 701](#).

## Bypassing MMS protection profile filtering based on carrier endpoints

You can use carrier endpoint filtering to exempt MMS sessions from MMS protection profile filtering. Carrier endpoint filtering matches carrier endpoints in MMS sessions with carrier endpoint patterns. If you add a carrier endpoint pattern to a filter list and set the action to exempt from all scanning, all messages from matching carrier endpoints bypass MMS protection profile filtering. See [“Bypassing message flood protection based on user’s carrier endpoints” on page 727](#).

## Applying MMS protection profiles to MMS traffic

To apply an MMS protection profile you must first create the MMS protection profile and then add the MMS protection profile to a security policy by enabling the Carrier security profile. The MMS protection profile then applies itself to the traffic accepted by that security policy.

MMS protection profiles can contain settings relevant to many different services. Each security policy uses the subset of the MMS protection profile settings that apply to the sessions accepted by the security policy. In this way, you might define just one MMS protection profile that can be used by many security policies, each policy using a different or overlapping subset of the MMS protection profile.

### To add an MMS protection profile to a security policy

1. Go to *Security Profiles > Carrier > MMS Profile*.
2. Select *Create New* to add an MMS protection profile.
3. Configure as needed, and save.
4. Go to *Policy*.
5. Select *Create New* to add a security policy, or select an existing policy and *Edit* to add the MMS profile.
6. Configure the security policy as required.
7. Enable *MMS Profile*, and select the MMS profile to add to the security policy.
8. Select *OK*.

## GTP basic concepts

GPRS currently supports data rates from 9.6kbps to more than 100 kbps, and is best suited for burst forms of traffic. GPRS involves both radio and wired components. The mobile phone sends the message to a base station unit (radio based), and the base station unit sends the message to the carrier network and eventually the Internet (wired carrier network).

The network system then either sends the message back to a base station and to the destination mobile unit, or forwards the message to the proper carrier’s network where it gets routed to the mobile unit.

This section includes:

- [PDP Context](#)
- [GPRS security](#)

### PDP Context

The packet data protocol (PDP) context is a connection between a mobile station and the end address that goes through the SGSN and GGSN. It includes identifying information about the

mobile customer used by each server or device to properly forward the call data to the next hop in the carrier network, typically using a GTP tunnel between the SGSN and GGSN.

When a mobile customer has an active voice or data connection open, both the SGSN and GGSN have the PDP context information for that customer and session.

When a mobile phone attempts to communicate with an address on an external packet network, either an IP or X.25 address, the mobile station that phone is connected to opens a PDP context through the SGSN and GGSN to the end address. Before any traffic is sent, the PDP context must first be activated.

The information included in the PDP context includes the customer's IP address, the IMSI number of the mobile handset, and the tunnel endpoint ID for both the SGSN and GGSN. The ID is a unique number, much like a session ID on a TCP/IP firewall. All this information ensures a uniquely identifiable connection is made.

Since one mobile device may have multiple connections open at one time, such as data connections to different Internet services and voice connections to different locations, there may be more than one PDP context with the same IP address making the extra identifying information required.

The endpoint that the mobile phone is connecting to only knows about the GGSN — the rest of the GPRS connection is masked by the GGSN.

Along the PDP context path, communication is accomplished in using three different protocols.

- The connection between the Mobile Station and SGSN uses the SM protocol.
- Between SGSN and GGSN GTP is used.
- Between GGSN and the endpoint either IP or X.25 is used.

FortiOS Carrier is concerned with the SGSN to GGSN part of the PDP context — the part that uses GTP.

For more about PDP context, see [“Tunnel Management Messages” on page 745](#).

## Creating a PDP context

While FortiOS Carrier is concerned mostly with the SGSN to GGSN part of the PDP Context, knowing the steps involved in creating a PDP context helps understand the role each device, protocol, and message type plays.

Both mobile stations and GGSNs can create PDP contexts.

### A Mobile Station creates a PDP context

1. The Mobile Station (MS) sends a `PDP activation request` message to the SGSN including the MS PDP address, and APN.
2. Optionally, security functions may be performed to authenticate the MS.
3. The SGSN determines the GGSN address by using the APN identifier.
4. The SGSN creates a downlink GTP tunnel to send IP packets between the GGSN and SGSN.
5. The GGSN creates an entry in its PDP context table to deliver IP packets between the SGSN and the external packet switching network.
6. The GGSN creates an uplink GTP tunnel to route IP-PDU from SGSN to GGSN.
7. The GGSN then sends back to the SGSN the result of the PDP context creation and if necessary the MS PDP address.
8. The SGSN sends an `Activate PDP context accept` message to the MS by returning negotiated the PDP context information and if necessary the MS PDP address.
9. Now traffic can pass from the MS to the external network endpoint.

### A GGSN creates a PDP context

1. The network receives an IP packet from an external network.
2. The GGSN checks if the PDP Context has already been created.
3. If not, the GGSN sends a `PDU notification request` to the SGSN in order to initiate a PDP context activation.
4. The GGSN retrieves the IP address of the appropriate SGSN address by interrogating the HLR from the IMSI identifier of the MS.
5. The SGSN sends to the MS a request to activate the indicated PDP context.
6. The PDP context activation procedure follows the one initiated by the MS. See [“A Mobile Station creates a PDP context” on page 646](#).
7. When the PDP context is activated, the IP packet can be sent from the GGSN to the MS.

### Terminating a PDP context

A PDP context remains open until it is terminated. To terminate the PDP context an MS sends a `Deactivate PDP context` message to the SGSN, which then sends a `Delete PDP Context` message to the GGSN. When the SGSN receives a PDP context deletion acknowledgment from the GGSN, the SGSN confirms to the MS the PDP context deactivation. The PDP can be terminated by the SGSN or GGSN as well with a slight variation of the order of the messages passed.

When the PDP Context is terminated, the tunnel it was using is deleted as well. If this is not completed in a timely manner, it is possible for someone else to start using the tunnel before it is deleted. This hijacking will result in the original customer being overbilled for the extra usage. Anti-overbilling helps prevent this. See [“Configuring Anti-overbilling in FortiOS Carrier” on page 742](#).

## GPRS security

The GPRS network has some built-in security in the form of GPRS authentication. However this is minimal, and is not sufficient for carrier network security needs. A GTP firewall, such as FortiOS Carrier, is required to secure the Gi, Gn, and Gp interfaces.

### GPRS authentication

GPRS authentication is handled by the SGSN to prevent unauthorized GPRS calls from reaching the GSM network beyond the SGSN (the base station system, and mobile station). Authentication is accomplished using some of the customer's information with a random number and uses two algorithms to create ciphers that then allow authentication for that customer.

User identity confidentiality ensures that customer information stays between the mobile station and the SGSN — no identifying information goes past the SGSN. Past that point other numbers are used to identify the customer and their connection on the network.

Periodically the SGSN may request identity information from the mobile station to compare to what is on record, using the IMEI number.

Call confidentiality is achieved through the use of a cipher, similar to the GPRS authentication described earlier. The cipher is applied between the mobile station and the SGSN. Essentially a cipher mask is XOR'd with each outgoing frame, and the receiving side XORs with its own cipher to result in the original frame and data.

## Parts of a GTPv1 network

A sample GTP network consists of the end handset sender, the sender's mobile station, the carrier's network including the SGSN and GGSN, the receiver's mobile station, and the receiver handset.

When a handset moves from one mobile station and SGSN to another, the handset's connection to the Internet is preserved because the tunnel the handset has to the Internet using GTP tracks the user's location and information. For example, the handset could move from one cell to another, or between countries.

The parts of a GPRS network can be separated into the following groups according to the roles of the devices:

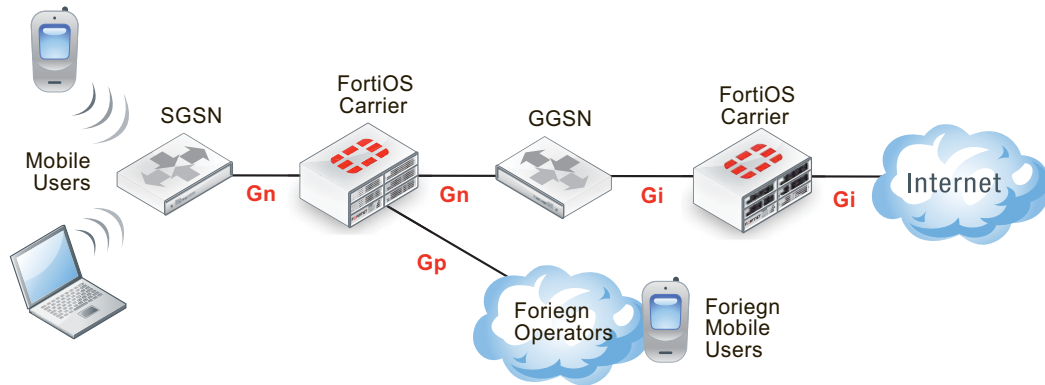
- Radio access to the GPRS network is accomplished by mobile phones and mobile stations (MS). See [“Radio access” on page 649](#).
- Transport the GPRS packets across the GPRS network is accomplished by SGSNs and GGSNs, both local and remote, by delivering packets to the external services. See [“Transport” on page 649](#).
- Billing and records are handled by CDF, CFR, HLR, and VLR devices. See [“Billing and records” on page 652](#).

GPRS networks also rely on access points and PDP contexts as central parts of the communication structure. These are not actual devices, but they are still critical .

For more information on APN, see [“Access Point Number \(APN\)” on page 753](#). For more information on PDP Context, see [“PDP Context” on page 645](#).

These devices, their roles, neighboring devices, the interfaces and protocols they use are outlined in the following table. These devices and their connections can be viewed in the [“Packet flow through the GPRS network” on page 654](#).

**Figure 144:**Carrier network showing the interfaces used (GTPv1)



**Table 33:**Devices on a GTPv1 network

Device role	Neighboring Devices	Interfaces used	Protocols used
Mobile Users	Mobile Stations (MS)	Radio Access Technology (RAT)	
Mobile Stations (MS)	Mobile Users, SGSN	Gb	IP, Frame Relay



**Table 33:** Devices on a GTPv1 network

SGSN (local)	<b>MS, SGSN (local or remote), GGSN (local and remote), CDR, CFR, HLR, VLR</b>	Ga, Gb, Gn, Gp, Gz	IP, Frame Relay, GTP, GTP'
SGSN (remote)	<b>SGSN (local)</b>	Gn	GTP
GGSN (local)	SGSN (local or remote), GGSN (local and remote), CDR, CFR, HLR, VLR	Ga, Gi, Gn, Gp, Gz	IP, GTP, GTP'
GGSN (remote)	SGSN (local), WAP gateway, Internet, other external services	Gi, Gp	IP, GTPv1
CDR, CFR	SGSN (local), GGSN (local)	Ga, Gz	GTP'
HLR, VLR	SGSN (local), GGSN (local)	Ga, Gz	GTP'

## Radio access

For a mobile phone to access the GPRS core network, it must first connect to a mobile station. This is a cellular tower that is connected to the carrier network.

How the mobile phone connects to the mobile station (MS) is determined by what Radio Access Technologies (RATs) are supported by the MS.

## Transport

Transport protocols move data along the carrier network between radio access and the Internet or other carrier networks.

FortiOS Carrier should be present where information enters the Carrier network, to ensure the information entering is correct and not malicious. This means a Carrier-enabled FortiGate unit intercepts the data coming from the SGSN or foreign networks destined for the SSGN or GGSN onto the network, and after the GGSN as the data is leaving the network.

## GTP

GPRS Tunnelling Protocol (GTP) is a group of IP-based communications protocols used to carry General Packet Radio Service (GPRS) within Global System for Mobile Communications (GSM) and Universal Mobile Telecommunications System (UMTS) networks. It allows carriers to transport actual cellular packets over their network via tunneling. This tunneling allows users to move between SGSNs and still maintain connection to the the Internet through the GGSN.

GTP has three versions version 0, 1, and 2. GTP1 and GTP2 are supported by FortiOS Carrier. The only GTP commands that are common to all forms of GTP are the echo request/response commands that allow GSNs to verify up to once every 60 seconds that neighboring GSNs are alive.

## GTPv0

There have been three versions of GTP to date. The original version of GTP (version 0) has the following differences from version GTPv1.

- the tunnel identification is not random
- there are options for transporting X.25
- the fixed port number 3386 is used for all functions, not just charging
- optionally TCP is allowed as a transport instead of UDP
- not all message types are supported in version 0

## GTPv1

On a GPRS network, Packet Data Protocol (PDP) context is a data structure used by both the Serving GPRS Support Node (SGSN) and the Gateway GPRS Support Node (GGSN). The PDP context contains the subscribers information including their access point, IP address, IMSI number, and their tunnel endpoint ID for each of the SGSN and GGSN.

The Serving GPRS Support Node (SGSN) is responsible for the delivery of data packets from and to the mobile stations within its geographical service area. Its tasks include packet routing and transfer, mobility management (attach/detach and location management), logical link management, and authentication and charging functions. The location register of the SGSN stores location information (e.g., current cell, current VLR) and user profiles (e.g., IMSI, address(es) used in the packet data network) of all GPRS users registered with this SGSN.

## GTPv1-C

GTPv1-C refers to the control layer of the GPRS Transmission network. This part of the protocol deals with network related traffic.

FortiOS Carrier handles GTPv1-C in GTPv1 by using the Tunnel Endpoint Identifier (TEID), IP address and a Network layer Service Access Point Identifier (NSAPI), sometimes called the application identifier, as an integer value that is part of the PDP context header information used to identify a unique PDP context in a mobile station, and SGSN.

For more information on GTPv1-C, see [“GTP-C messages” on page 745](#).

## GTPv1-U

GTPv1-U is defined in 3GPP TS 29.281 and refers to the user layer of the GPRS Tunneling network. This part of the protocol deals with user related traffic, user tunnels, and user administration issues.

A GTPv1-U tunnel is identified by a TEID, an IP address, and a UDP port number. This information uniquely identifies the limb of a GTPv1 PDP context. The IP address and the UDP port number define a UDP/IP path, a connectionless path between two endpoints (i.e. SGSN or GGSN). The TEID identifies the tunnel endpoint in the receiving GTPv1-U protocol entity; it allows for the multiplexing and demultiplexing of GTP tunnels on a UDP/IP path between a given GSN-GSN pair. For more information on GTPv1-U, see [“GTP-U messages” on page 746](#).

The GTP core network consists of one or more SGSNs and GGSNs.

## GGSN

The Gateway GPRS Support Node (GGSN) connects the GPRS network on one side via the SGSN to outside networks such as the Internet. These outside networks are called packet data networks (PDNs). The GGSN acts as an edge router between the two different networks — the GGSN forwards incoming packets from the external PDN to the addressed SGSN and the

GGSN also forwards outgoing packets to the external PDN. the GGSN also converts the packets from the GPRS packets with SGSN to the external packets, such as IP or X.25.


## SGSN

The Serving GPRS Support Node (SGSN) connects the GPRS network to GTPv1 compatible mobile stations, and mobile units (such as UTRAN and ETRAN) on one side and to the gateway node (GGSN), which leads to external networks, on the other side. Each SGSN has a geographical area, and mobile phones in that area connect to the GPRS network through this SGSN. The SGSN also maintains a location register that contains customer’s location and user profiles until they connect through a different SGSN at which time the customer information is moved to the new SGSN. This information is used for packet routing and transfer, mobility management also known as location management, logical link management, and authentication and billing functions.

## GTPv2

GTPv2, defined in 3GPP TS 29.274, is dramatically different from GTPv1, defined in 3GPP TS 29.060. Where in GTPv1 the tunnel is between the SGSN and the GGSN, in GTPv2 The SGSN is between the MME and the LTE Serving Gateway (S-GW), beyond which is the PDN gateway (P-GW). Even tunnel management messages have changed significantly.

**Figure 145:**Network diagram for GTPv2



**Table 34:** Device roles on a GTPv2 network

Device role	Neighboring Devices	Interfaces used	Protocols used
Mobile Users	Mobile Stations (MS)	Radio Access Technology (RAT)	--
GTPv1 Mobile Stations (MS)	Mobile Users, SGSN	Gb	IP, Frame Relay
GTPv2 Mobile Stations (MS)	Mobile Users, MME	???	IP, Frame Relay
SGSN (local)	<b>GTPv1 MS, SGSN, S-GW</b>	???	IP, Frame Relay, GTPv1, GTP'

**Table 34:** Device roles on a GTPv2 network

S-GW	SGSN, MME, P-GW	???	IP, GTPv2, GTP'
P-GW	S-GW, Internet, other external services	???	IP, GTPv2

### GTPv2-C

GTPv2-C is the control layer messaging for GTPv2. It is used by LTE mobile stations, SGSN units for backwards compatibility, and SGWs that are the gateway to other networks. The messaging is very different from GTPv1. GTPv2-C is required to communicate with the Mobility Management Entity (MME) to create, change and delete EPS bearers when handover events happen, and to create Forwarding tunnels. The protocol is also used to communicate with the Serving Gateway (SGW) which has the S-GW and PDN-GW interfaces, and the Serving GPRS Support Node (SGSN).

### MME

MME essentially fills the role of the SGSN in a GTPv1 network — it is how the mobile stations gain access to the Carrier network. GTPv2 supports different mobile stations than GTPv1, so MME handles the GTPv2 MSes and SGSN handles the GTPv1 MSes

## Billing and records

A major part of the GPRS network is devoted to billing. Customer billing requires enough information to identify the customer, and then billing specific information such as connection locations and times, as well as amount of data transferred. A modified form of GTP called GTP' is used for billing. The home location records and visitor location records store information about customers that is critical to billing.

### GTP' (GTP prime)

GTP is used to handle tunnels of user traffic between SGSNs and GGSNs. However for billing purposes, other devices that are not supported by GTP are required. GTP' (GTP prime) is a modified form of GTP and is used to communicate with these devices such as the Charging Data Function (CDF) that communicates billing information to the Charging Gateway Function (CGF). In most cases, GTP' transports user records from many individual network elements, such as the GGSNs, to a centralised computer which then delivers the charging data more conveniently to the network operator's billing center, often through the CGF. The core network sends charging information to the CGF, typically including PDP context activation times and the quantity of data which the end user has transferred.

GTP' is used by the Ga and Gz interfaces to transfer billing information. GTP' uses registered UDP/TCP port 3386. GTP' defines a different header, additional messages, field values, as well as a synchronisation protocol to avoid losing or duplicating CDRs on CGF or SGSN/GGSN failure. Transferred CDRs are encoded in ASN.1.

### HLR

The Home Location Register (HLR) is a central database that contains details of each mobile phone subscriber that is authorized to use the GSM core network. There can be several logical, and physical, HLRs per public land mobile network (PLMN), though one international mobile subscriber identity (IMSI)/MSISDN pair can be associated with only one logical HLR (which can span several physical nodes) at a time. The HLRs store details of every SIM card issued by the mobile phone operator. Each SIM has a unique identifier called an IMSI which is the primary key to each HLR record.

## VLR

The Visitor Location Register (VLR) is a database which stores information about all the mobile devices that are currently under the jurisdiction of the Mobile Switching Center which it serves. Of all the information the VLR stores about each Mobile Station, the most important is the current Location Area Identity (LAI). This information is vital in the call setup process.

Whenever an MSC detects a new MS in its network, in addition to creating a new record in the VLR, it also updates the HLR of the mobile subscriber, informing it of the new location of that MS.

For more information on GTP', see [“GTP-U and Charging Management Messages” on page 746](#).

## GPRS network common interfaces

There are interfaces for each connection on the GPRS network. An interface is an established standard form of communication between two devices. Consider a TCP/IP network. In addition to the transport protocol (TCP) there are other protocols on that network that describe how devices can expect communications to be organized, just like GPRS interfaces.

### Interfaces between devices on the network

There are a series of interfaces that define how different devices on the carrier network communicate with each other. These interfaces are called Ga to Gz, and each one defines how a specific pair of devices will communicate. For example Gb is the interface between the base station and the SGSN, and Gn is one possible interface between the SGSN and GGSN.

The SGSN and GGSN keep track of the CDR information and forward it to the Charging Data Function (CDF) using the Gr interface between the SGSN and home location register (HLR), Gs interface between the SGSN and MSC (VLR), Gx interface between the GGSN and the Charging Rules Function (CRF), Gy between the GGSN and online charging system (OCS), and finally Gz which is the off-line (CDR-based) charging interface between the GSN and the CG that uses GTP'.

Each of these interfaces on the GPRS network has a name in the format of G<sub>x</sub> where x is a letter of the alphabet that determines what part of the network the interface is used in. It is common for network diagrams of GPRS networks to include the interface name on connections between devices. See [“Packet flow through the GPRS network” on page 654](#).



The Carrier-enabled FortiGate unit only provides protection on the Gn, Gp, and Gi interfaces.

**Table 35:** GPRS network interfaces, their roles, and billing

Name	Device connections that use this interface	Traffic Protocol used	Its role or how it affects billing
<b>Ga</b>	CDR and GSN (SGSNs and GGSNs)	GTP' - GTP modified to include CDR role	CDR have the accounting records, that are compiled in the GSN and then sent to the Charging Gateway (CG)

**Table 35:** GPRS network interfaces, their roles, and billing

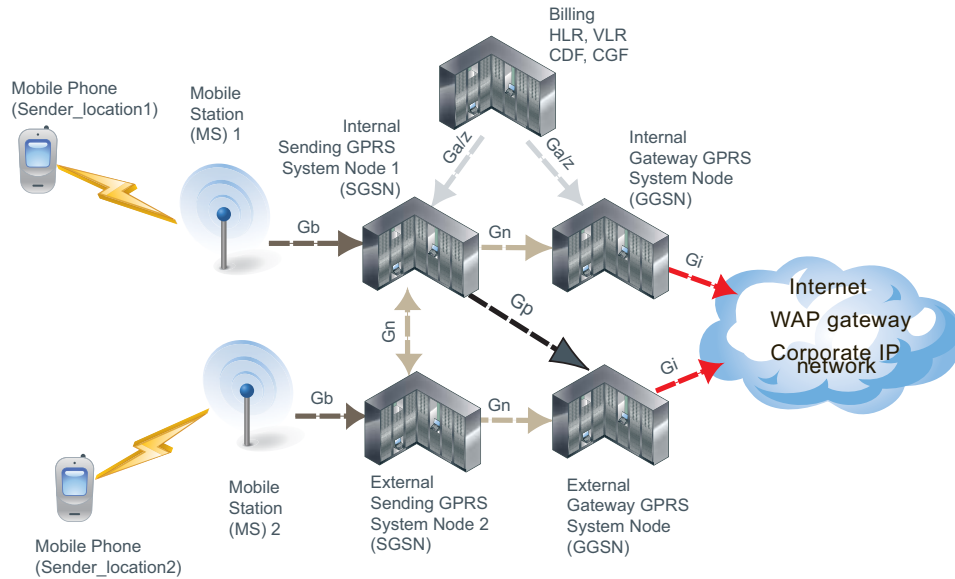
<b>Gb</b>	MS and SGSN	Frame Relay or IP	When an IP address moves to a new MS, the old MS may continue to use and bill that IP address.
<b>Gi</b>	GGSN and public data networks (PDNs)	IP based	This is the connection to the Internet. If the GTP tunnel is deleted without notifying the Gi interface, the connection may remain open incurring additional charges. FortiOS Carrier adds this interface to a firewall. See <a href="#">“Anti-overbilling with FortiOS Carrier”</a> on page 742.
<b>Gn</b>	SGSN and external SGSNs and internal GGSNs	GTP	When the GTP tunnel is deleted, need to inform other interfaces immediately to prevent misuse of connections remaining open. FortiOS Carrier adds this interface to a firewall.
<b>Gp</b>	Internal SGSN and external GGSNs	GTP	
<b>Gz</b>	GSN (SGSN and GGSN) and the charging gateway (CG)	GTP'	Used for the offline charging interface. Ga is used for online charging.

Corporate customers may have a direct connection to the Gi interface for higher security. The Gi interface is normally an IP network, though a tunnelling protocol such as GRE or IPsec may be used instead.

## Packet flow through the GPRS network

To better understand the GPRS network, we will follow the path data takes for a normal connection. For this example a call placed from a mobile phone involves accessing services on the Internet.

**Figure 146:**Sample GPRS network topology



1. A mobile phone places a call using a mobile station (MS). This connection between the mobile phone and the MS is a radio connection using one of the radio access technologies. See [“Radio Access Technology \(RAT\) type”](#) on page 753.
2. The MS connects to a GPRS System Node (GSN) specifically a Sending GSN. This connection uses the Gb interface and typically uses IP address or Frame Relay.
3. The SGSN checks the mobile phone information located in the home location register (HLR) or visitor location register (VLR) to ensure there is subscriber information for that phone. If this mobile phone is from another network, the SGSN uses the VLR and updates its home carrier’s information with its current location and information. This connection involves the Ga or Gz interfaces, and uses the GTP’ protocol for communication.
4. The SGSN checks to make sure the phone did not transfer this connection from a different MS. If it did, the connection has already been established (along with the billing) and is handed off to this SGSN. If the call is being handed over from another SGSN, it will use the Gn interface between the two SGSNs.
5. The SGSN sends GTP messages to the local external Gateway GSN (GGSN) to create a GTP tunnel for this PDP context to access the Internet. It is possible that a remote GGSN has access to a service, such as a WAP gateway, that the local GGSN is missing. In this situation, the local SGSN uses the Gp interface to connect to the remote GGSN. Both the Gn and Gp interfaces use GTP.
6. The both the local and remote GGSNs connect to external services outside the GPRS network. These services can include a WAP gateway, a corporate IP network directly connected to the GPRS network, or the Internet. The connection from the GGSN to the external services uses the Gi interface.

## SCTP

As of FortiOS version 5.0, the FortiGate can now natively handle SCTP (Stream Control Transport Protocol) traffic, as an alternative to TCP and UDP for use in Carrier networks. The FortiGate handles SCTP as if it would any other traffic.

## Overview

SCTP is a connection-oriented transport protocol that overcomes some of the limitations of both TCP and UDP that prevent reliable transfer of data over IP-based networks (such as those used by telephony systems and carrier networks). The 'Stream' in SCTP refers to the sequence of user messages or packets that are considered at the same time to be individual objects and also treated as a whole by networked systems. SCTP is less vulnerable to congestion and flooding due to more advanced error handling and flood protection built into the protocol.

**Table 36:** SCTP features as compared to TCP and UDP

Feature	SCTP	TCP	UDP
State required at each endpoint	yes	yes	no
Reliable data transfer	yes	yes	no
Congestion control and avoidance	yes	yes	no
Message boundary conservation	yes	no	yes
Path MTU discovery and message fragmentation	yes	yes	no
Message bundling	yes	yes	no
Multi-homed hosts support	yes	no	no
Multi-stream support	yes	no	no
Unordered data delivery	yes	no	yes
Security cookie against SYN flood attack	yes	no	no
Built-in heartbeat (reachability check)	yes	no	N/A

All of these features are built into the design of the Protocol, and the structure of SCTP packets and networks. The FortiGate unit interprets the traffic and provides the necessary support for maintenance and verification features, but the features are not FortiGate specific. These features are documented in greater detail below.

### State required at each endpoint

Constant back and forth acknowledgement and content verification messages are sent between all SCTP peer endpoints, and all endpoints' state machine actions must be synchronized for traffic to flow.

### Reliable data transfer

SCTP places data and control information (eg. source, destination, verification) into separate messages, both sharing the same header in the same SCTP packet. This allows for constant verification of the contained data at both ends and along the path, preventing data loss or fragmentation. As well, data is not sent in an interruptible stream as in TCP. See [Message bundling](#) below for more information.

### Congestion control and avoidance

Built-in, constantly updating path detection and monitoring automatically redirect packets along alternate paths in case of traffic congestion or inaccessible destinations. For deliberate/malicious congestion control, see [Security cookie against SYN flood attack](#) below.



## Message boundary conservation

SCTP is designed in such a way that no matter how messages are divided, redirected, or fragmented, the message boundaries will be maintained within the packets, and all messages cannot be appended without tripping verification mechanisms.

## Path MTU discovery and message fragmentation

SCTP is capable of Path Maximum Transmission Unit discovery, as outlined in RFC4821. Two specific alterations have been made to how SCTP handles MTU. First, that endpoints will have separate MTU estimates for each possible multi-homed endpoint. Second, that bundled message fragments (as explained below) will be directed based on MTU calculations, so that retransmissions (if necessary) will be sent without delay to alternate addresses.

## Message bundling

SCTP is a message-oriented protocol, which means that despite being a streaming data protocol, it transports a sequence of specific messages, rather than transporting a stream of bytes (like TCP). Since some data transmissions are small enough to not require a complete message's worth of content, so multiple pieces of content will be transmitted simultaneously within the messages.

## Multi-homed hosts support

SCTP supports multi-homing, which is a network structure in which one or multiple sources/destinations has more than one IP address. SCTP can adapt to multi-homing scenarios and redirect traffic to alternate IP addresses in case of failure.

## Multi-stream support

Due to the message bundling feature allowing for multiple pieces of content to be sent in messages at once, SCTP can 'multi-stream' content, by deliberately dividing it among messages at a fixed rate, so that multiple types of content (eg. both images and text) can be loaded at once, at the same pace.

## Unordered data delivery

With control messages in every packet to provide verification of any packet's data and its place in the stream, the data being transmitted can actually arrive in any order, and verify that all has arrived or that some is missing.

## Security cookie against SYN flood attack

Since every packet contains verification of its place in the stream, it makes it easy for the protocol to detect when redundant, corrupted or malicious packets flood the path, and they are automatically dropped when necessary.

## Built-in heartbeat (reachability check)

Endpoints automatically send specific control chunks among the other SCTP packet information to peer endpoints, to determine the reachability of the destination. Heartbeat acknowledgement packets are returned if the destination is available.

## SCTP Firewall

FortiGate stateful firewalls will protect and inspect SCTP traffic, according to RFC4960. SCTP over IPsec VPN is also supported. The FortiGate device is inserted as a router between SCTP endpoints. It checks SCTP Syntax for the following information:

- Source and destination port
- Verification Tag
- Chunk type, chunk flags, chunk length
- Sequence of chunk types
- Associations

The firewall also oversees and maintains several SCTP security mechanisms:

- SCTP four-way handshake
- SCTP heartbeat
- NAT over SCTP

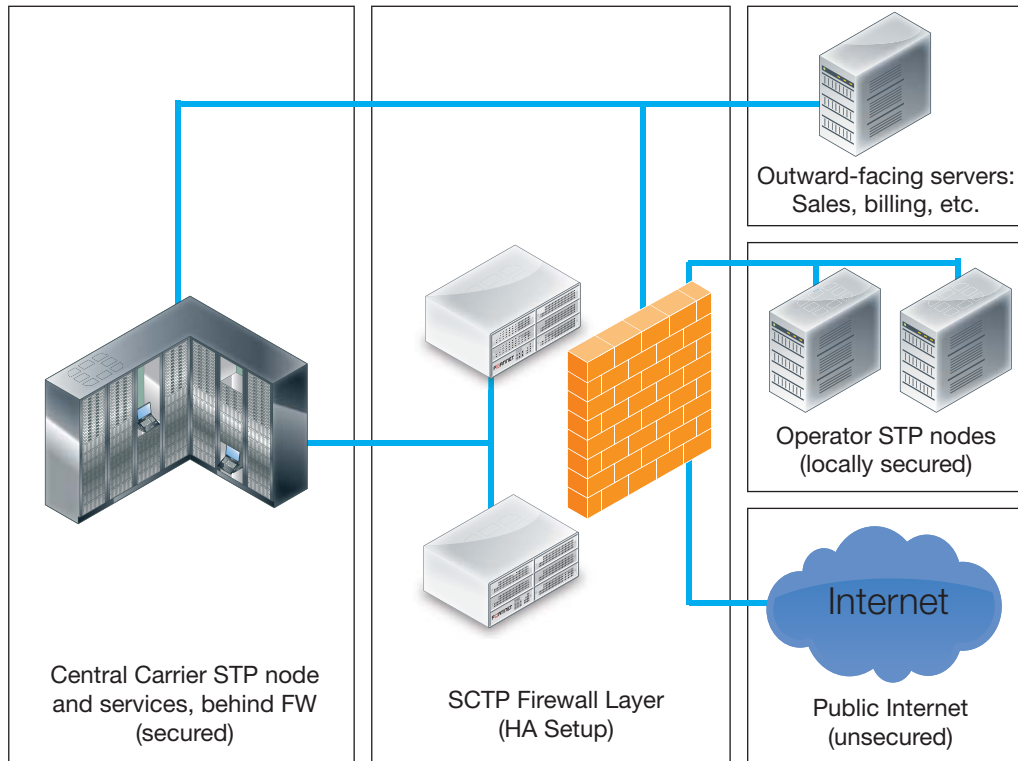
The firewall has IPS DoS protection against known threats to SCTP traffic, including INIT/ACK flood attacks, and SCTP fuzzing.

## SCTP example scenario

An ideal SCTP configuration for a Carrier serving multiple operators/service providers involves a unified Firewall, securing all incoming and outgoing traffic over the Carrier network, whether it be standard web traffic, GTP or other carrier traffic, or corporate traffic for the Carrier company.

One best practice method to provide a unified firewall with built-in redundancy is to make use of multiple FortiGate units, connected in a High Availability cluster. Also, there are additional methods that can be applied to ease the complexity of managing multiple services, functions, and traffic types across multiple devices.

**Figure 147:**Sample SCTP Network Topology

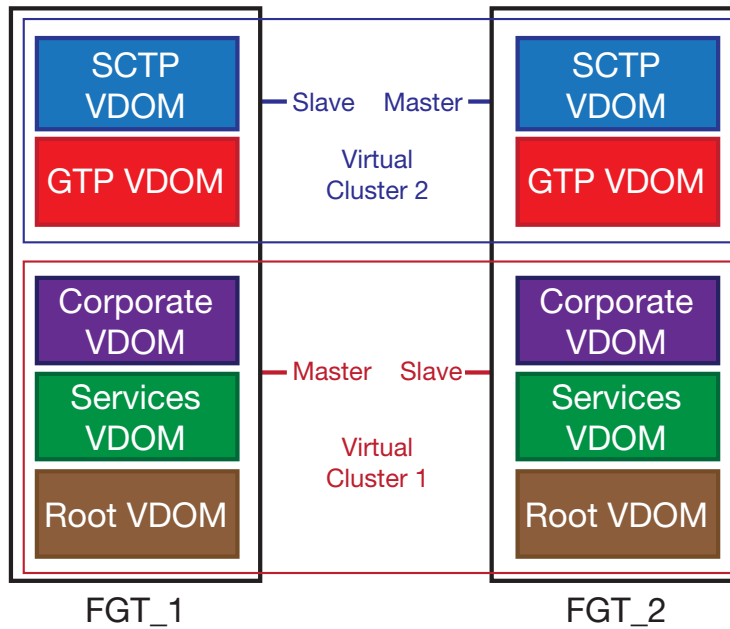


In this example, the firewall layer is configured with two FortiGate devices to act as an HA cluster, providing automatic load balancing and failover detection for the main firewall.

The two devices together make up the firewall, through which all traffic passes. Virtual Domains are created within the FortiGate units, distributing services and traffic into individual VDOMs, allowing them to be monitored and secured individually, to help mitigate possible threats to Carrier networks that target specific services. Individual departments or administrators can manage specific VDOMs, or the FortiGates can be collectively managed centrally by network administrators.

The VDOMs are distributed as below:

**Figure 148:**VDOM distribution between SCTP Firewall Layer FortiGate units



One FortiGate handles basic FortiGate services and non-Carrier traffic. Configuring virtual clustering across the two FortiGates allows one to mirror its VDOMs across to the other unit.

The second FortiGate can then primarily provide Carrier-specific services and handle SCTP, Gi and GTP traffic, using the first FortiGate as the slave unit in a second virtual cluster.

# Carrier web-based manager settings

The Carrier menu provides settings for configuring FortiOS Carrier features within the Security Profiles menu. These features include MMS and GTP profiles.

In *Security Profiles > Carrier*, you can configure profiles and settings for MMS and GTP. In the Carrier menu, you can configure an MMS profile and then apply it to a security policy. You can also configure GTP profiles and apply those to security policies as well.

This topic includes the following:

- [MMS profiles](#)
- [MMS Content Checksum](#)
- [Notification List](#)
- [Message Flood](#)
- [Duplicate Message](#)
- [Carrier Endpoint Filter Lists](#)
- [GTP Profile](#)

## MMS profiles

Since MMS profiles can be used by more than one security policy, you can configure one profile for the traffic types handled by a set of security policies requiring identical protection levels and types, rather than repeatedly configuring those same profile settings for each individual security policy.



If the security policy requires authentication, do not select the MMS profile in the security policy. This type of profile is specific to the authenticating user group. For details on configuring the profile associated with the user group, see User Groups in the [User Authentication guide](#).

For example, while traffic between trusted and untrusted networks might need strict protection, traffic between trusted internal addresses might need moderate protection. To provide the different levels of protection, you might configure two separate protection profiles: one for traffic between trusted networks, and one for traffic between trusted and untrusted networks.

Once you have configured the MMS profile, you can then apply the profile to MMS traffic by applying it to a security policy.

MMS profiles can contain settings relevant to many different services. Each security policy uses the subset of the MMS profile settings that apply to the sessions accepted by the security policy. In this way, you might define just one MMS profile that can be used by many security policies, each policy using a different or overlapping subset of the MMS profile.

The MMS Profile page contains options for each of the following:

- MMS scanning
- MMS Bulk Email Filtering Detection
- MMS Address Translation
- MMS Notifications
- DLP Archive
- Logging

## MMS profile configuration settings

The following are MMS profile configuration settings in *Security Profiles > Carrier > MMS Profile*.

<p><b>MMS Profile page</b></p> <p>Lists each individual MMS profile that you created. On this page, you can edit, delete or create an MMS profile.</p>	
<b>Create New</b>	Creates a new MMS profile. When you select <i>Create New</i> , you are automatically redirected to the New MMS Profile page.
<b>Edit</b>	Modifies settings within an MMS profile. When you select <i>Edit</i> , you are automatically redirected to the Edit MMS Profile.
<b>Delete</b>	Removes an MMS profile from the list on the MMS Profile page. <p>To remove multiple MMS profiles from within the list, on the MMS Profile page, in each of the rows of the profiles you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all MMS profiles from the list, on the MMS Profile page, select the check box in the check box column, and then select <i>Delete</i>.</p>
<b>Name</b>	The name of the MMS profile.
<b>Ref.</b>	Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page ( <i>Security Profiles &gt; Antivirus &gt; Profiles</i> ), 1 appears in <i>Ref.</i> . <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> <li>• <b>View the list page for these objects</b> – automatically redirects you to the list page where the object is referenced at.</li> <li>• <b>Edit this object</b> – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page.</li> <li>• <b>View the details for this object</b> – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.</li> </ul>
<p><b>New MMS Profile page</b></p> <p>Provides settings for configuring an MMS profile. This page also provides settings for configuring DLP archives and logging.</p>	
<b>Profile Name</b>	Enter a name for the profile.

<b>Comments</b>	Enter a description about the profile. This is optional.
<b>MMS Scanning</b>	Configure MMS Scanning options. See <a href="#">“MMS scanning options” on page 663.</a>
<b>MMS Bulk Email Filtering Detection</b>	Configure MMS Bulk Email options. See <a href="#">“MMS bulk email filtering options” on page 665.</a>
<b>MMS Address Translation</b>	Configure MMS Address Translation options. See <a href="#">“MMS Address Translation options” on page 668</a>
<b>MMS Notifications</b>	Configure MMS Notification options. <a href="#">“MMS Notifications” on page 669.</a>
<b>DLP Archive</b>	Configure DLP archive option. See <a href="#">“DLP Archive options” on page 672</a>
<b>Logging</b>	Configure logging options. See <a href="#">“Logging” on page 673.</a>

### MMS scanning options

You can configure MMS scanning protection profile options to apply virus scanning, file filtering, content filtering, carrier endpoint blocking, and other scanning to MMS messages transmitted using the MM1, MM3, MM4 and MM7 protocols.

The following are the MMS Scanning options that are available within an MMS profile. You can create an MMS profile in *Security Profiles > Carrier > MMS Profile* or edit an existing one. You must expand MMS Scanning to access the following options.

<b>MMS Scanning section of the New MMS Profile page</b>	
<b>Monitor Only</b>	Select to cause the unit to record log messages when MMS scanning options find a virus, match a file name, or match content using any of the other MMS scanning options. Select this option to be able to report on viruses and other problems in MMS traffic without affecting users.  <b>Tip:</b> Select <i>Remove Blocked</i> if you want the unit to actually remove content intercepted by MMS scanning options.
<b>Virus Scan</b>	Select to scan attachments in MMS traffic for viruses.  Since MM1 and MM7 use HTTP, the oversize limits for HTTP and the HTTP antivirus port configuration also applies to MM1 and MM7 scanning.  MM3 and MM4 use SMTP and the oversize limits for SMTP and the SMTP antivirus port configuration also applies to MM3 and MM4 scanning.
<b>Scan MM1 message retrieval</b>	Select to scan message retrievals that use MM1. If you enable <i>Virus Scan</i> for all MMS interfaces, messages are also scanned while being sent. In this case, you can disable MM1 message retrieval scanning to improve performance.

<b>Quarantine</b>	Select to quarantine the selected MMS traffic
<b>Remove Blocked</b>	Select to remove blocked content from each protocol and replace it with the replacement message.  Select <i>Constant</i> if the unit is to preserve the length of the message when removing blocked content, as may occur when billing is affected by the length of the message.  <b>Tip:</b> If you only want to monitor blocked content, select <i>Monitor Only</i> .
<b>Content Filter</b>	Select to filter messages based on matching the content of the message with the words or patterns in the selected web content filter list.  For information about adding a web content filter list, see the <a href="#">FortiGate CLI Reference</a> .
<b>Carrier Endpoint Block</b>	Select to add <i>Carrier Endpoint Filtering</i> in this MMS profile. Select the carrier endpoint filter list to apply it to the profile. For information about carrier endpoint filtering, see “ <a href="#">Carrier Endpoint Filter Lists</a> ” on page 681.
<b>MMS Content Checksum</b>	Select to add MMS Content Checksum in this MMS profile. Select the MMS content checksum list to apply it to the profile.
<b>Pass Fragmented Messages</b>	Select to pass fragmented MM3 and MM4 messages. Fragmented MMS messages cannot be scanned for viruses. If you do not select these options, fragmented MM3 and MM4 message are blocked.
<b>Comfort Clients</b>	Select client comforting for MM1 and MM7 sessions.  Since MM1 and MM7 messages use HTTP, MM1 and MM7 client comforting operates like HTTP client comforting.
<b>Comfort Servers</b>	Select server comforting for each protocol.  Similar to client comforting, you can use server comforting to prevent server connection timeouts that can occur while waiting for the unit to buffer and scan large POST requests from slow clients.
<b>Interval (1-900 seconds)</b>	Enter the time in seconds before client and server comforting starts after the download has begun, and the time between sending subsequent data.
<b>Amount (1-10240 bytes)</b>	The number of bytes sent by client or server comforting at each interval.



<b>Oversized MMS Message</b>	<p>Select <i>Block</i> or <i>Pass</i> for files and email messages exceeding configured thresholds for each protocol.</p> <p>The oversize threshold refers to the final size of the message, including attachments, after encoding by the client. Clients can use a variety of encoding types; some result in larger file sizes than the original attachment. As a result, a file may be blocked or logged as oversized even if the attachment is several megabytes smaller than the oversize threshold.</p>
<b>Threshold (1KB - 800 MB)</b>	<p>Enter the oversized file threshold and select KB or MB. If a file is larger than the threshold the file is passed or blocked depending on the <i>Oversized MMS Message</i> setting. The web-based manager displays the allowed threshold range. The threshold maximum is 10% of the unit's RAM.</p>

### MMS bulk email filtering options

You can use the MMS bulk email filtering options to detect and filter MM1 and MM4 message floods and duplicate messages. You can configure three thresholds that define a flood of message activity and three thresholds that define excessive duplicate messages. The configuration of each threshold includes the response actions for the threshold.

The configurable thresholds for each of the flood and duplicate sensors and must be enabled in sequence. For example, you can enable Flood Threshold 1 and Flood Threshold 2, but you cannot disable Flood Threshold 1 and enable Flood Threshold 2.

You can also add MSISDN to the bulk email filtering configuration and select a subset of the bulk email filtering options to applied to these individual MSISDNs.

You must first select MM1 and/or MM4 to detect excessive message duplicates. If excessive message duplicates are detected, the unit will perform the *Duplicate Message Action* for the specified duration.

You can configure three duplicate message thresholds and enable them with separate values and actions. They are labeled Duplicate Threshold 1 through 3 and must be enabled in sequence. For example, you can enable Duplicate Threshold 1 and Duplicate Threshold 2, but you cannot disable Duplicate Threshold 1 and enable Duplicate Threshold 2.

When traffic accepted by a security policy that contains an MMS profile with duplicate message configured receives MM1 or MM4 duplicate messages that match a threshold configured in the MMS protection profile, the unit performs the duplicate message action configured for the matching threshold.

You can configure three message flood thresholds and enable them with separate values and actions. They are labeled Flood Threshold 1 through 3 and must be enabled in sequence. For example, you can enable Flood Threshold 1 and Flood Threshold 2, but you cannot disable Flood Threshold 1 and enable Flood Threshold 2.

When traffic accepted by a security policy that contains an MMS protection profile with message flooding configured experiences MM1 or MM4 message flooding that matches a threshold configured in the MMS profile, the unit performs the message flood action configured for the matching threshold.

<p><b>MMS Bulk Email Filtering Detection section of the New MMS Profile page</b>  <b>This section of the New MMS Profile page contains numerous sections where you can configure specific settings for flood threshold, duplicate threshold and recipient MSISDNs.</b></p>	
<p><b>Message Flood section of the new MMS Profile page</b>  The message flood settings for each flood threshold. Expand each to configure settings for a threshold.</p>	
<b>Flood Threshold 1</b>	Expand to reveal the flood threshold settings for Flood Threshold 1. The settings for Flood Threshold 1 are the same for Flood Threshold 2 and 3.
<b>Enable</b>	Select to apply Flood Threshold 1 to the MSISDN exception.
<b>Message Flood Window</b>	Enter the period of time during which a message flood will be detected if the <i>Message Flood Limit</i> is exceeded. The message flood window can be 1 to 2880 minutes (48 hours).
<b>Message Flood Limit</b>	Enter the number of messages which signifies a message flood if exceeded within the <i>Message Flood Window</i> .
<b>Message Flood Block Time</b>	Enter the amount of time during which the unit performs the <i>Message Flood Action</i> after a message flood is detected.
<b>Message Flood Action</b>	Select one or more actions that the unit is to perform when a message flood is detected.
<b>Flood Threshold 2 Flood Threshold 3</b>	Expand to configure settings for Flood Threshold 2 or 3 respectively.
<p><b>Duplicate Message section of the new MMS Profile page</b>  <b>The duplicate message threshold settings. Expand each to configure settings for a threshold.</b></p>	
<b>MM1 Retrieve Duplicate Enable</b>	Select to scan MM1 <code>mm1-retr</code> messages for duplicates. By default, <code>mm1-retr</code> messages are not scanned for duplicates as they may often be the same without necessarily being bulk or spam.
<b>Enable</b>	Select to enable the selected duplicate message threshold and to make the rest of the options available for configuration.
<b>Duplicate Message Window</b>	Enter the period of time during which excessive message duplicates will be detected if the Duplicate message Limit it exceeded. The duplicate message window can be 1 to 2880 minutes (48 hours).
<b>Duplicate Message Limit</b>	Enter the number of messages which signifies excessive message duplicates if exceeded within the Duplicate Message Window.

	<b>Duplicate Message Block Time</b>	Enter the amount of time during which the unit will perform the Duplicate Message Action after a message flood is detected.
	<b>Duplicate Message Action</b>	Select one or more actions that the unit is to perform when excessive message duplication is detected.
	<b>Duplicate Threshold 2 Duplicate Threshold 3</b>	Expand to configure settings for Duplicate Threshold 2 or 3 respectively.
<b>Recipient MSISDN section of the New MMS Profile page</b>		
<p>The recipient Mobile Subscriber Integrated Services Digital Network Number (MSISDN) settings for each recipient MSISDN. When you select <i>Create New</i>, you are automatically redirected to the New MSISDN page.</p> <p>You need to save the profile before you can add MSISDNs.</p>		
	<b>Recipient MSISDN</b>	The recipient MSISDN.
	<b>Flood Threshold 1</b>	Check to enable Flood Threshold 1 settings for this MSISDN.
	<b>Flood Threshold 2</b>	ICheck to enable Flood Threshold 2 settings for this MSISDN..
	<b>Flood Threshold 3</b>	ICheck to enable Flood Threshold 3 settings for this MSISDN..
	<b>Duplicate Threshold 1</b>	Check to enable Duplicate Threshold 1 settings for this MSISDN..
	<b>Duplicate Threshold 2</b>	Check to enable Duplicate Threshold 2 settings for this MSISDN..
	<b>Duplicate Threshold 3</b>	Check to enable Duplicate Threshold 3 settings for this MSISDN..
	<b>Edit</b>	Modifies the settings of a Recipient MSISDN in the Recipient MSISDN list. When you select <i>Edit</i> , you are automatically redirected to the New MSISDN page.
	<b>Delete</b>	Removes a Recipient MSISDN in the Recipient MSISDN list within the Recipient MSISDN section of the page.
<b>New MSISDN page</b>		
	<b>Create New</b>	Creates a new Recipient MSISDN. When you select <i>Create New</i> , you are automatically redirected to the New MSISDN page.
	<b>Recipient MSISDN</b>	Enter a name for the recipient MSISDN.
	<b>Flood Threshold 1</b>	Select to apply Flood Threshold 1 to the MSISDN exception.
	<b>Flood Threshold 2</b>	Select to apply Flood Threshold 2 to the MSISDN exception.

<b>Flood Threshold 3</b>	Select to apply Flood Threshold 3 to the MSISDN exception.
<b>Duplicate Threshold 1</b>	Select to apply Duplicate Threshold 1 to the MSISDN exception.
<b>Duplicate Threshold 2</b>	Select to apply Duplicate Threshold 2 to the MSISDN exception.
<b>Duplicate Threshold 3</b>	Select to apply Duplicate Threshold 3 to the MSISDN exception.

### MMS Address Translation options

The sender's carrier endpoint is used to provide logging and reporting details to the mobile operator and to identify the sender of infected content.

When MMS messages are transmitted, the *From* field may or may not contain the sender's address. When the address is not included, the sender information will not be present in the logs and the unit will not be able to notify the user if the message is blocked unless the sender's address is made available elsewhere in the request.

The unit can extract the sender's address from an extended HTTP header field in the HTTP request. This field must be added to the HTTP request before it is received by the unit. If this field is present, it will be used instead of the sender's address in the MMS message for logging and notification. If this header field is present when a message is retrieved, it will be used instead of the *To* address in the message. If this header field is not present the content of the *To* header field is used instead.

Alternatively, the unit can extract the sender's address from a cookie.

You can configure MMS address translation to extract the sender's carrier endpoint so that it can be added to log and notification messages. You can configure MMS address translation settings to extract carrier endpoints from HTTP header fields or from cookies. You can also configure MMS address translation to add an endpoint prefix to the extracted carrier endpoints. For more information, see Dynamic Profiles and Endpoints in the [User Authentication guide](#).

<b>MMS Address Translation section of the New MMS Profile page</b>	
<b>Sender Address Source</b>	Select to extract the sender's address from the <i>HTTP Header Field</i> or a <i>Cookie</i> . You must also specify the identifier that contains the carrier endpoint.

<p><b>Sender Address Identifier</b></p>	<p>Enter the sender address identifier that includes the carrier endpoint. The default identifier is <code>x-up-calling-line-id</code>.</p> <p>If the <i>Sender Address Source</i> is <i>HTTP Header Field</i>, the address and its identifier in the HTTP request header takes the format:</p> <pre>&lt;Sender Address Identifier&gt;: &lt;MSISDN_value&gt;</pre> <p>Where the <code>&lt;MSISDN_value&gt;</code> is the carrier endpoint. For example, the HTTP header might contain:</p> <pre>x-up-calling-line-id: 6044301297</pre> <p>where <code>x-up-calling-line-id</code> would be the Sender Address Identifier.</p> <p>If the <i>Sender Address Source</i> is <i>Cookie</i>, the address and its identifier in the HTTP request header's <code>Cookie</code> field takes the format of attribute-value pairs:</p> <pre>Cookie: id=&lt;cookie-id&gt;;        &lt;Sender Address Identifier&gt;=&lt;MSISDN Value&gt;</pre> <p>For example, the HTTP request headers might contain:</p> <pre>Cookie: id=0123jf!a;x-up-calling-line-id=6044301297</pre> <p>where <code>x-up-calling-line-id</code> would be the <i>Sender Address Identifier</i>.</p>
<p><b>Convert Sender Address From / To HEX</b></p>	<p>Select to convert the sender address from ASCII to hexadecimal or from hexadecimal to ASCII. This is required by some applications.</p>
<p><b>Add Carrier Endpoint Prefix for Logging / Notification</b></p>	<p>Select the following to enable adding endpoint prefixes for logging and notification.</p>
	<p><b>Enable</b></p> <p>Select to enable adding the country code to the extracted carrier endpoint, such as the MSISDN, for logging and notification purposes. You can limit the number length for the test numbers used for internal monitoring without a country code.</p>
	<p><b>Prefix</b></p> <p>Enter a carrier endpoint prefix to be added to all carrier endpoints. Use the prefix to add extra information to the carrier endpoint in the log entry.</p>
	<p><b>Minimum Length</b></p> <p>Enter the minimum length of the country code information being added. If this and Maximum Length are set to zero (0), length is not limited.</p>
	<p><b>Maximum Length</b></p> <p>Enter the maximum length of the country code information being added. If this and Minimum Length are set to zero (0), length is not limited.</p>

### MMS Notifications

MMS notifications are messages that a unit sends when an MMS profile matches content in an MM1, MM3, MM4 or MM7 session. For example, the MMS profile detects a virus or uses

content blocking to block a web page, text message or email. You can send notifications to the sender of the message using same protocol and the addressing headers in the original message. You can also configure MMS notifications to send notification messages to another destination (such as a system administrator) using the MM1, MM3, MM4 or MM7 protocol.

You need to enable one or more *Notification Types* or you can add an *Antivirus Notification List* to enable sending notifications,.

You can also use MMS notifications options to configure how often notifications are sent. The unit sends notification messages immediately for the first event, then at a configurable interval if events continue to occur. If the interval does not coincide with the window of time during which notices may be sent, the unit waits to send the notice in the next available window. Subsequent notices contain a count of the number of events that have occurred since the previous notification.

There are separate notifications for each notification type, including virus events. Virus event notifications include the virus name. Up to three viruses are tracked for each user at a time. If a fourth virus is found, one of the existing tracked viruses is removed from the list.

The notifications are MM1 *m-send-req* messages sent from the unit directly to the MMSC for delivery to the client. The host name of the MMSC, the URL to which *m-send-req* messages are sent, and the port must be specified.

<b>MMS Notification section of the New MMS Profile page</b>	
<b>Antivirus Notification List</b>	<p>Optionally select an antivirus notification list to select a list of virus names to send notifications for. The unit sends a notification message whenever a virus name or prefix in the antivirus notification list matches the name of a virus detected in a session scanned by the MMS protection profile. Select <i>Disabled</i> if you do not want to use a notification list. To create an antivirus notification list, see <a href="#">“Notification List” on page 675</a>.</p> <p>Instead of selecting a notification list you can configure the <i>Virus Scan Notification Type</i> to send notifications for all viruses.</p>
<b>Message Protocol</b>	<p>In each column, select the protocol used to send notification messages. You can use a different protocol to send the notification message than the protocol on which the violation was sent. The MMS Notifications options change depending on the message protocol that you select.</p> <p>If you select a different message protocol, you must also enter the User Domain. If selecting MM7 you must also enter the Message Type.</p>
<b>Message Type</b>	<p>Select the MM7 message type to use if sending notifications using MM7. Options include <i>deliver.REQ</i> and <i>submit.REQ</i></p>

<b>Detect Server Details</b>	<p>Select to use the information in the headers of the original message to set the address of the notification message. If you do not select this option, you can enter the required addressing information manually.</p> <p>You cannot select <i>Detect Server Details</i> if you are sending notification messages using a different message protocol.</p> <p>If you select <i>Detect Server Details</i>, you cannot change the <i>Port</i> where the notification is being sent.</p>
<b>Hostname</b>	<p>Enter the FQDN or the IP address of the server where the notifications will be sent.</p>
<b>URL</b>	<p>Enter the URL of the server. For example if the notification is going to <code>www.example.com/home/alerts</code>, the URL is <code>/home/alerts</code>.</p> <p>This option is available only when <i>Message Protocol</i> is <i>mm1</i> or <i>mm7</i>.</p>
<b>Port</b>	<p>Enter the port number of the server.</p> <p>You cannot change the <i>Port</i> if <i>Detect Server Details</i> is enabled.</p>
<b>Username</b>	<p>Enter the user name required for sending messages using this server (optional).</p> <p>This option is available only when <i>Message Protocol</i> is <i>mm7</i>.</p>
<b>Password</b>	<p>Enter the password required for sending messages using this server (optional).</p> <p>This option is available only when <i>Message Protocol</i> is <i>mm7</i>.</p>
<b>VASP ID</b>	<p>Enter the value-added-service-provider (VASP) ID to be used when sending a notification message. If a VAS is not offered by the mobile provider, it is offered by a third party or a VAS provider or content provider (CP).</p> <p>This option is available only when <i>Message Protocol</i> is <i>mm7</i>.</p>
<b>VAS ID</b>	<p>Enter the value-added-service (VAS) ID to be used when sending a notification message. A VAS is generally any service beyond voice calls and fax.</p> <p>This option is available only when <i>Message Protocol</i> is <i>mm7</i>.</p>
<b>All Notification Types</b>	<p>In each column, select notification for all MMS event types for that MMS protocol, then enter the amount of time and select the time unit for notice intervals.</p> <p>Alternatively, expand <i>All Notification Types</i>, and then select notification for individual MMS event types for each MMS protocol. Then enter the amount of time and select the time unit for notice intervals.</p> <p>Not all event types are available for all MMS protocols.</p>

<b>Content Filter</b>	In each column, select to notify when messages are blocked by the content filter, then enter the amount of time and select the time unit for notice intervals.
<b>File Block</b>	In each column, select to notify when messages are blocked by file block, then enter the amount of time and select the time unit for notice intervals.
<b>Carrier Endpoint Block</b>	In each column, select to notify when messages are blocked, then enter the amount of time and select the time unit for notice intervals.
<b>Flood</b>	In each column, select to notify when message flood events occur, then enter the amount of time and select the time unit for notice intervals.
<b>Duplicate</b>	In each column, select to notify when duplicate message events occur, then enter the amount of time and select the time unit for notice intervals.
<b>MMS Content Checksum</b>	In each column, select to notify when the content within an MMS message is scanned and banned because of the checksum value that was matched.
<b>Virus Scan</b>	In each column, select to notify when the content within an MMS message is scanned for viruses.
<b>Notifications Per Second Limit</b>	For each MMS protocol, enter the number of notifications to send per second. If you enter zero (0), the notification rate is not limited.
<b>Day of Week</b>	For each MMS protocol, select the days of the week the unit is allowed to send notifications.
<b>Window Start Time</b>	For each MMS protocol, select the time of day to begin the message alert window. By default, the message window starts at 00:00. You can change this if you want to start the message window later in the day. When configured, notification outside this window will not be sent.
<b>Window Duration</b>	For each MMS protocol, select the time of day at which to end the message alert window. By default, the message window ends at 00:24. You can change this if you want to end the message window earlier in the day.  When configured, notification outside this window will not be sent

### DLP Archive options

Select DLP archive options to archive MM1, MM3, MM4, and MM7 sessions. In addition to the MMS profile's DLP archive options, you can:

- Archive MM1 and MM7 message floods
- Archive MM1 and MM7 duplicate messages
- Select *DLP archiving* for carrier endpoint patterns in a *Carrier Endpoint List* and select the *Carrier Endpoint Block* option in the *MMS Scanning* section of an MMS Profile



The unit only allows one sixteenth of its memory for transferring content archive files. For example, for units with 128MB RAM, only 8MB of memory is used when transferring content archive files. Best practices dictate to not enable full content archiving if antivirus scanning is also configured because of these memory constraints.

#### DLP Archive section of the New MMS Profile page

<b>Display DLP meta-information on the system dashboard</b>	Select each required protocol to display the content archive summary in the Log and Archive Statistics dashboard widget on the System Dashboard.
<b>Archive to FortiAnalyzer/FortiGuard</b>	<p>Select the type of archiving that you want for the protocol (MM1, MM3, MM4, and MM7). You can choose from Full, Summary or None.</p> <p><b>None</b> – Do not send content archives.</p> <p><b>Summary</b> – Send content archive metadata only. Includes information such as date and time, source and destination, request and response size, and scan result.</p> <p><b>Full</b> – Send content archive both metadata and copies of files or messages.</p> <p>In some cases, FortiOS Carrier may not archive content, or may make only a partial content archive, regardless of your selected option. This behavior varies by prerequisites for each protocol.</p> <p>This option is available only if a FortiAnalyzer unit or FortiGuard Analysis and Management Service is configured.</p>

## Logging

You can enable logging in an MMS profile to write event log messages when the MMS profile options that you have enabled perform an action. For example, if you enable MMS antivirus protection, you could also use the MMS profile logging options to write an event log message every time a virus is detected.

You must first configure how the unit stores log messages so that you can then record these logs messages. For more information, see the [FortiOS Handbook Logging and Reporting chapter](#).

#### Logging section of the New MMS Profile page

<b>MMS-Antivirus</b>		If antivirus settings are enabled for this MMS profile, select the following options to record <i>Antivirus Log</i> messages.
	<b>Viruses</b>	Record a log message when this MMS profile detects a virus.
	<b>Blocked Files</b>	Record a log message when antivirus file filtering enabled in this MMS profile blocks a file.

	<b>Intercepted Files</b>	Record a log message when this MMS profile intercepts a file.
	<b>Oversized Files/Emails</b>	Record a log message when this MMS profile encounters an oversized file or email message. Oversized files and email messages cannot be scanned for viruses.
<b>MMS Scanning</b>		If MMS scanning settings are enabled for this MMS profile, select the following options to record <i>Email Filter Log</i> messages.
	<b>Notification Messages</b>	Select to log the number of MMS notification messages sent.
	<b>Bulk Messages</b>	Select to log MMS Bulk AntiSpam events. You must also select which protocols to write log messages for in the MMS bulk email filtering part of the MMS profile. For more information, see <a href="#">“MMS bulk email filtering options” on page 665</a> .
	<b>Carrier Endpoint Filter Block</b>	Select to log MMS carrier endpoint filter events, such as MSISDN filtering.
	<b>MMS Content Checksum</b>	Select to log MMS content checksum activity.
	<b>Content Block</b>	Select to log content blocking events.

## MMS Content Checksum

The MMS Content Checksum menu allows you to configure content checksum lists.

Configure MMS content checksum lists in *Security Profiles > Carrier > MMS Content Checksum* using the following table.

---

### MMS Content Checksum page

Lists each individual content checksum list that you created. On this page, you can edit, delete or create a content checksum list.

<b>Create New</b>	Creates a new MMS content checksum list. When you select <i>Create New</i> , you are automatically redirected to the New List. This page provides a name field and comment field. You must enter a name to go to MMS Content Checksum Settings page.
<b>Edit</b>	Modifies settings to a MMS content checksum. When you select <i>Edit</i> , you are automatically redirected to the MMS Content Checksum Settings page.

<b>Delete</b>	<p>Removes an MMS content checksum from the page.</p> <p>To remove multiple content checksum lists from within the list, on the MMS Content Checksum page, in each of the rows of the content checksum lists you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all content checksum lists from list, on the MMS Content Checksum page, select the check box in the check box column and then select <i>Delete</i>.</p>
<b>Name</b>	The name of the MMS content checksum list that you created.
<b># Entries</b>	The number of checksums that are included in the content checksum list.
<b>MMS Profiles</b>	The MMS profile or profiles that have the MMS content checksum list applied. For example if two different MMS profiles use this content checksum list, they will both be listed here.
<b>Comments</b>	A description given to the MMS content checksum.
<b>Ref.</b>	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (<i>Security Profiles &gt; AntiVirus &gt; Profiles</i>), 1 appears in <i>Ref.</i> .</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> <li>• <b>View the list page for these objects</b> – automatically redirects you to the list page where the object is referenced at.</li> <li>• <b>Edit this object</b> – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page.</li> <li>• <b>View the details for this object</b> – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.</li> </ul>

## Notification List

The Notification List menu allows you to configure a list of viruses. This virus list provides a list for scanning viruses in MMS messages. You can use one virus list in multiple MMS profiles, and configure multiple virus lists.

## Notification list configuration settings

The following are notification list configuration settings in *Security Profiles > Carrier > Notification List*.

<b>Notification List page</b> Lists all the notification lists that you created. On this page you can edit, delete or create a new notification list.	
<b>Create New</b>	Creates a new notification list. When you select <i>Create New</i> , you are automatically redirected to the New List page. You must enter a name to go to the Notification List Settings page.
<b>Edit</b>	Modifies settings within the notification list. When you select <i>Edit</i> , you are automatically redirected to the Notification List Settings page.
<b>Delete</b>	Removes a notification list from the list on the Notification List page.  To remove multiple notification lists from within the list, on the Notification List page, in each of the rows of the notification lists you want removed, select the check box and then select <i>Delete</i> .  To remove all notification lists from the list, on the Notification List page, select the check box in the check box column and then select <i>Delete</i> .
<b>Name</b>	The name of the MMS content checksum list that you created.
<b># Entries</b>	The number of checksums that are included in that content checksum list.
<b>MMS Profiles</b>	The MMS profile or profiles that are associated with
<b>Comments</b>	A description given to the MMS notification list.

<b>Ref.</b>	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (<i>Security Profiles &gt; Antivirus &gt; Profiles</i>), 1 appears in <i>Ref.</i> .</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> <li>• <b>View the list page for these objects</b> – automatically redirects you to the list page where the object is referenced at.</li> <li>• <b>Edit this object</b> – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page.</li> <li>• <b>View the details for this object</b> – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.</li> </ul>
<p><b>Notification List Settings page</b></p> <p>Provides settings for configuring a notification list, which is a list of viruses and is used for scanning viruses in MMS messages. This list is called the Antivirus Notification List in an MMS profile.</p>	
<b>Name</b>	If editing the name of a notification list, enter the new name in this field. You must select <i>OK</i> to save the change.
<b>Comments</b>	If you want to enter a comment, enter the comment in the field. You must select <i>OK</i> to save the change.
<b>Create New</b>	Creates a notification entry in the list. When you select <i>Create New</i> , you are automatically redirected to the New Entry page.
<b>Edit</b>	Modifies settings within a notification list. When you select <i>Edit</i> , you are automatically redirected to the Edit Entry page.
<b>Delete</b>	<p>Removes a notification entry from the list on the page.</p> <p>To remove multiple notification entries from within the list, on the Notification List Settings page, in each of the rows of the entries you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all notification entries from the list, on the Notification List Settings page, select the check box in the check box column and then select <i>Delete</i>.</p>
<b>Enable</b>	Enables a notification entry that is disabled.
<b>Disable</b>	Disables a notification entry so that it is not active and available for use, but it is not deleted.

<b>Remove All Entries</b>	Removes all notification entries that are listed on the Notification List Settings page.
<b>Enable</b>	Displays whether or not the checksum is enabled.
<b>Virus Name/Profile</b>	The name of the virus that was added to the list.
<b>Entry Type</b>	The type of match that will be used to match the virus stated in the notification list to the actual virus that is found.
<b>New Entry page</b>	
<b>Virus Name/Profile</b>	Enter the virus name.
<b>Entry Type</b>	Select the type of match that will be used to match the virus stated in the notification list to the actual virus that is found.
<b>Enable</b>	Select to enable the virus in the list.

## Message Flood

The convenience offered by MM1 and MM4 messaging can be abused by users sending spam or attempting to overload the network with an excess of messages. MMS flood prevention can help prevent this type of abuse. A message flood occurs when a single subscriber sends a volume of messages that exceed the flood threshold that you set. The threshold defines the maximum number of messages allowed, the period during which the subscriber sent messages are considered, and the length of time the sender is restricted from sending messages after a flood is detected. For example, for the first threshold you may determine that any subscriber who sends more than 100 MM1 messages in an hour (60 minutes) will have all outgoing messages blocked for 30 minutes.

<b>Action</b>	<b>Description</b>
<b>Log</b>	Add a log entry indicating that a message flood has occurred. You must also enable logging for <i>MMS Scanning &gt; Bulk Messages</i> in the <i>Logging</i> section of the MMS protection profile.
<b>DLP Archive</b>	Save the first message to exceed the flood threshold, or all the messages that exceed the flood threshold, in the DLP archive. DLP archiving flood messages may not always produce useful results. Since different messages can be causing the flood, reviewing the archived messages may not be a good indication of what is causing the problem since the messages could be completely random.
<b>All messages</b>	All the messages that exceed the flood threshold will be saved in the DLP archive.
<b>First message only</b>	Save only the first message to exceed the flood threshold in the DLP archive. Other messages in the flood are not saved. For message floods this may not produce much useful information since a legitimate message could trigger the flood threshold.

Action	Description
<b>Intercept</b>	Messages that exceed the flood threshold are passed to the recipients, but if quarantine is enabled for intercepted messages, a copy of each message will also be quarantined for later examination. If the quarantine of intercepted messages is disabled, the <i>Intercept</i> action has no effect.
<b>Block</b>	Messages that exceed the flood threshold are blocked and will not be delivered to the message recipients. If quarantine is enabled for blocked messages, a copy of each message will be quarantined for later examination.
<b>Alert Notification</b>	If the flood threshold is exceeded, the Carrier-enabled FortiGate unit will send an MMS flood notification message.  In the web-based manager when <i>Alert Notification</i> is selected it displays the fields to configure the notification.

Flood protection for MM1 messages prevents your subscribers from sending too many messages to your MMSC. Configuring flood protection for MM4 messages prevents another service provider from sending too many messages from the same subscriber to your MMSC.

### Message flood configuration settings

The following are message flood configuration settings in *Security Profiles > Carrier > Message Flood*.

<b>Message Flood page</b> Lists the large amount of messages that are being sent to you, from outside sources.	
<b>Delete</b>	Removes messages from the list.  To remove multiple messages from within the list, on the Message Flood page, in each row of the messages you want removed, select the check box and then select <i>Delete</i> .  To remove all messages from the list, on the Message Flood page, select the check box in the check box column and then select <i>Delete</i> .
<b>Remove All Entries</b>	Removes all messages from the list.
<b>Protocol</b>	The protocol used.
<b>MMS Profile</b>	The MMS profile that is used.
<b>Sender</b>	The sender's email address.
<b>Level</b>	The level of severity of the message.
<b>Count</b>	The count column can be up or down and these settings can be turned off by selecting beside the column's name.
<b>Window Size (minutes)</b>	The time in minutes.

<b>Timer (minutes:seconds)</b>	The time in seconds and in minutes. The timer column can be up or down and these settings turned off by selecting beside the column's name.
<b>Page Controls</b>	Use to navigate through the list.

## Duplicate Message

Duplicate message protection for MM1 messages prevents multiple subscribers from sending duplicate messages to your MMSC. Duplicate message protection for MM4 messages prevents another service provider from sending duplicate messages from the same subscriber to your MMSC.

The unit keeps track of the sent messages. If the same message appears more often than the threshold value that you have configured, action is taken. Possible actions are logging the duplicate messages, blocking or intercepting them, archiving, and sending an alert to inform an administrator that duplicate messages are occurring.

### Duplicate message configuration settings

View duplicate messages in *Security Profiles > Carrier > Duplicate Message*.

<b>Duplicate Message page</b> Lists duplicates of messages that were sent to you.	
<b>Delete</b>	Removes a message from the list.  To remove multiple duplicate messages from within the list, on the Message Flood page, in each row of the messages you want removed, select the check box and then select <i>Delete</i> .  To remove all duplicate messages from the list, on the Message Flood page, select the check box in the check box column and then select <i>Delete</i> .
<b>Page Controls</b>	Use to navigate through the list.
<b>Remove All Entries</b>	Removes all duplicate messages from the list.
<b>Protocol</b>	Either MM1 or MM4
<b>Profile</b>	The MMS profile that logs the detection.
<b>Checksum</b>	The checksum of the MMS message.
<b>Status</b>	Either flagged or blank. Flagged means that the actions defined in the MMS profile are taken. For more information, see <a href="#">“MMS bulk email filtering options” on page 665</a> .
<b>Count</b>	Displays the number of messages in the last window of time.



<b>Window Size (minutes)</b>	The period of time during which a message flood will be detected if the Message Flood Limit is exceeded.
<b>Timer (minutes:seconds)</b>	Either the time left in the window if the message is unflagged, or the time until the message will be unflagged if it is already flagged.

## Carrier Endpoint Filter Lists

A carrier endpoint filter list contains carrier endpoint patterns. A pattern can match one carrier endpoint or can use wildcards or regular expressions to match multiple carrier endpoints. For each pattern, you select the action that the unit takes on a message when the pattern matches a carrier endpoint in the message. Actions include blocking the message, exempting the message from MMS scanning, and exempting the message from all scanning. You can also configure the pattern to intercept the message and content archive the message to a FortiAnalyzer unit.

### Carrier endpoint filter lists configuration settings

The following are Carrier endpoint filter list configuration settings in *Security Profiles > Carrier > Carrier Endpoint Filter Lists*.

<b>Carrier Endpoint Filter Lists page</b>	
Lists all the endpoint filters that you created. On this page, you can edit, delete or create a new endpoint filter list.	
<b>Create New</b>	Creates a new endpoint filter list. When you select <i>Create New</i> , you are automatically redirected to the New List page. You must enter a name to go to the Carrier Endpoint Filter Lists Settings page.
<b>Edit</b>	Modifies settings within an endpoint filter list in the list.
<b>Delete</b>	Removes an endpoint filter in the list.  To remove multiple endpoint filter lists from within the list, on the Carrier Endpoint Filter List page, in each of the rows of the endpoint filter lists you want removed, select the check box and then select <i>Delete</i> .  To remove all endpoint filter lists from the list, on the Carrier Endpoint Filter List page, select the check box in the check box column and then select <i>Delete</i> .
<b>Name</b>	The name of the endpoint filter.
<b># Entries</b>	The number of carrier endpoint patterns in each carrier endpoint filter list.
<b>MMS Profiles</b>	The MMS profile that the carrier endpoint filter list is added to.
<b>Comments</b>	A description about the endpoint filter.

<b>Ref.</b>	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (<i>Security Profiles &gt; Antivirus &gt; Profiles</i>), 1 appears in <i>Ref.</i> .</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> <li>• <b>View the list page for these objects</b> – automatically redirects you to the list page where the object is referenced at.</li> <li>• <b>Edit this object</b> – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page.</li> <li>• <b>View the details for this object</b> – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.</li> </ul>
-------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

**Carrier Endpoint Filter Lists Settings page**  
Provides settings for configuring an endpoint filter.

<b>Name</b>	The name you entered on the New List page, after selecting <i>Create New</i> on the Carrier Endpoint Filter page.
<b>Comments</b>	A description about the endpoint filter. You can add one here if you did not enter one on the New List page.
<b>Create New</b>	Creates a new endpoint filter list. When you select <i>Create New</i> , you are automatically redirected to the New Entry page.
<b>Edit</b>	Select to modify the settings of a pattern in the list.
<b>Delete</b>	Select to remove a pattern in the list.
<b>Enable</b>	Enables a disabled pattern in the list.
<b>Disable</b>	Disables a pattern in the list.
<b>Remove All Entries</b>	Removes all patterns in the list on the Carrier Endpoint Filter Lists Settings page.
<b>Enable</b>	Indicates whether or not the pattern is enabled.

<b>Pattern</b>	Enter or change the pattern that FortiOS Carrier uses to match with carrier endpoints. The pattern can be a single carrier endpoint or consist of wildcards or Perl regular expressions that will match more than one carrier endpoint. Set <i>Pattern Type</i> to correspond to the pattern that you want to use.
<b>Action</b>	Select the action taken by FortiOS Carrier for messages from a carrier endpoint that matches the carrier endpoint pattern:
<b>Pattern Type</b>	The type of pattern chosen.
<b>New Entry page</b>	
<b>Pattern</b>	Enter or change the pattern that FortiOS Carrier uses to match with carrier endpoints. The pattern can be a single carrier endpoint or consist of wildcards or Perl regular expressions that will match more than one carrier endpoint. Set <i>Pattern Type</i> to correspond to the pattern that you want to use.
<b>Action(s)</b>	Select the action taken by FortiOS Carrier for messages from a carrier endpoint that matches the carrier endpoint pattern:
	<b>Content Archive</b>
	MMS messages from the carrier endpoint are delivered, the message content is DLP archived according to MMS DLP archive settings. Content archiving is also called DLP archiving.
	<b>Intercept</b>
	MMS messages from the carrier endpoint are delivered. Based on the quarantine configuration, attached files may be removed and quarantined.
<b>Pattern Type</b>	Select a pattern type as one of Single Carrier Endpoint, Wildcard or Regular Expression.  Wildcard and Regular Expression will match multiple patterns where Single Carrier Endpoint matches only one.
<b>Enable</b>	Select to enable this carrier endpoint filter pattern.

## GTP Profile

You can configure multiple GTP profiles within the GTP menu. GTP profiles concern GTP activity flowing through the unit. These GTP profiles are then applied to a security policy.

### GTP profile configuration settings

The following are GTP profile configuration settings in *Security Profiles > Carrier > GTP Profile*.

<b>GTP Profile page</b>	
Lists each GTP profile that you have created. On this page, you can edit, delete or create a new GTP profile.	
<b>Create New</b>	Creates a new GTP profile. When you select <i>Create New</i> , you are automatically redirected to the New page.
<b>Edit</b>	Modifies settings within a GTP profile in the list. When you select <i>Edit</i> , you are automatically redirected to Edit page.
<b>Delete</b>	Removes a GTP profile from the list.  To remove multiple GTP profiles from within the list, on the GTP Profile page, in each of the rows of the profiles you want removed, select the check box and then select <i>Delete</i> .  To remove all GTP profiles from within the list, on the GTP Profile page, select the check box in the check box column and then select <i>Delete</i> .
<b>Name</b>	The name of the GTP profile.
<b>Ref.</b>	Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page ( <i>Security Profiles &gt; Antivirus &gt; Profiles</i> ), 1 appears in <i>Ref.</i> .  To view the location of the referenced object, select the number in <i>Ref.</i> , and the Object Usage window appears displaying the various locations of the referenced object.  To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:  <b>View the list page for these objects</b> – automatically redirects you to the list page where the object is referenced at.  <b>Edit this object</b> – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page.  <b>View the details for this object</b> – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy’s settings appear within the table.
<b>New GTP Profile page</b> <b>Provides settings for configuring a GTP profile.</b>	
<b>Name</b>	Enter a name for the GTP profile.
<b>General Settings</b>	Configure general options for the GTP profile. See <a href="#">“General settings options” on page 685</a> .
<b>Message Type Filtering</b>	Configure filtering for messages. See <a href="#">“Message type filtering options” on page 687</a> .

<b>APN Filtering</b>	Configure filtering options for APN. See <a href="#">“APN filtering options” on page 687</a> .
<b>IMSI Filtering</b>	Configure filtering options for IMSI. See <a href="#">“Basic filtering options” on page 689</a> .
<b>Advanced Filtering</b>	Configure advanced filtering options. See <a href="#">“Advanced filtering options” on page 690</a> .
<b>IE removal policy</b>	Configure IE removal policy options. See <a href="#">“Information Element (IE) removal policy options” on page 694</a> .
<b>Encapsulated IP Traffic Filtering</b>	Configure filtering options for encapsulated IP traffic. See <a href="#">“Encapsulated IP traffic filtering options” on page 694</a> .
<b>Encapsulated Non-IP End User Address Filtering</b>	Configure filtering options for encapsulated non-IP end user addresses. See <a href="#">“Encapsulated non-IP end user traffic filtering options” on page 695</a> .
<b>Protocol Anomaly</b>	Configure protocol anomaly options. See <a href="#">“Protocol Anomaly prevention options” on page 696</a> .
<b>Anti-Overbilling</b>	Configure anti-overbilling options. See <a href="#">“Anti-Overbilling options” on page 697</a> .
<b>Log</b>	Configure log options. See <a href="#">“Log options” on page 697</a> .

### General settings options

The following are mostly house keeping options that appear in the General Settings area of the GTP configuration page.

---

#### General Settings section of the New GTP Profile page

---

**Sequence Number Validation** Enable to check that packets are not duplicated or out of order. GTP packets contain a Sequence Number field.

This number tells the receiving GGSN the order of the packets it is receiving. Normally the GGSN compares this sequence number in the packets with its own sequence counter — if the two do not match, the packet is dropped. This sequence number validation can be off-loaded to the FortiOS Carrier freeing up resources on the GGSN.

---

**GTP-in-GTP** Select *Allow* to enable GTP packets to be allowed to contain GTP packets, or a GTP tunnel inside another GTP tunnel.

To block all GTP-in-GTP packets, select *Deny*.

---

**Minimum Message Length** Enter the shortest possible message length in bytes. Normally this is controlled by the protocol, and will vary for different message types. If a packet is smaller than this limit, it is discarded as it is likely malformed and a potential security risk.

The default minimum message length is 0 bytes.

---

---

<b>Maximum Message Length</b>	<p>Enter the maximum allowed length of a GTP packet in bytes.</p> <p>A GTP packet contains three headers and corresponding parts GTP, UDP, and IP. If a packet is larger than the maximum transmission unit (MTU) size, it is fragmented to be delivered in multiple packets. This is inefficient, resource intensive, and may cause problems with some applications.</p> <p>By default the maximum message length is 1452 bytes.</p>
<b>Tunnel Limit</b>	<p>Enter the maximum number of tunnels allowed open at one time. For additional GTP tunnels to be opened, existing tunnels must first be closed.</p> <p>This feature can help prevent a form of denial of service attack on your network. This attack involves opening more tunnels than the network can handle and consuming all the network resources doing so. By limiting the number of tunnels at any one time, this form of attack will be avoided.</p> <p>The tunnel limiting applies to the Handover Group, and Authorized SGSNs and GGSNs.</p>
<b>Tunnel Timeout</b>	<p>Enter the maximum number of seconds that a GTP tunnel is allowed to remain active. After the timeout the unit deletes GTP tunnels that have stopped processing data. A GTP tunnel may hang for various reasons. For example, during the GTP tunnel tear-down stage, the "delete pdap context response" message may get lost. By setting a timeout value, you can configure the FortiOS Carrier firewall to remove the hanging tunnels.</p> <p>The default is 86400 seconds, or 24 hours.</p>
<b>Control plane message rate limit</b>	<p>Enter the number of packets per second to limit the traffic rate to protect the GSNs from possible Denial of Service (DoS) attacks. The default limit of 0 does not limit the message rate.</p> <p>GTP DoS attacks can include:</p> <ul style="list-style-type: none"><li>• <b>Border gateway bandwidth saturation:</b> A malicious operator can connect to your GRX and generate high traffic towards your Border Gateway to consume all the bandwidth.</li><li>• <b>GTP flood:</b> A GSN can be flooded by illegitimate traffic</li></ul>
<b>Handover Group</b>	<p>Select the allowed list of IP addresses allowed to take over a GTP session when the mobile device moves locations.</p> <p>Handover is a fundamental feature of GPRS/UMTS, which enables subscribers to seamlessly move from one area of coverage to another with no interruption of active sessions. Session hijacking can come from the SGSN or the GGSN, where a fraudulent GSN can intercept another GSN and redirect traffic to it. This can be exploited to hijack GTP tunnels or cause a denial of service.</p> <p>When the handover group is defined it acts like a whitelist with an implicit default deny at the end — the GTP address must be in the group or the GTP message will be blocked. This stops handover requests from untrusted GSNs.</p>

---

---

**Authorized SGSNs** Use *Authorized SGSNs* to only allow authorized SGSNs to send packets through the unit and to block unauthorized SGSNs. Go to *Firewall Objects > Address > Addresses* and add the IP addresses of the authorized SGSNs to a firewall address or address group. Then set *Authorized SGSNs* to this firewall address or address group.

You can use *Authorized SGSNs* to allow packets from SGSNs that have a roaming agreement with your organization.

---

**Authorized GGSNs** Use *Authorized GGSNs* to only allow authorized GGSNs to send packets through the unit and to block unauthorized GGSNs. Go to *Firewall Objects > Address > Addresses* and add the IP addresses of the authorized GGSNs to a firewall address or address group. Then set *Authorized GGSNs* to this firewall address or address group.

You can use *Authorized GGSNs* to allow packets from SGSNs that have a roaming agreement with your organization.

---

### Message type filtering options

On the *New GTP Profile* page, you can select to allow or deny the different types of GTP messages, which is referred to as message type filtering. You must expand the *Message Type Filtering* section to access the settings.

The messages types include Path Management, Tunnel Management, Location Management, Mobility Management, MBMS, and GTP-U and Charging Management messages.



For enhanced security, Fortinet best practices dictate that you set *Unknown Message Action* to deny. This will block all unknown GTP message types, some of which may be malicious.

To configure message type filter options, expand *Message Type Filtering* in the GTP profile.

### APN filtering options

An Access Point Name (APN) is an Information Element (IE) included in the header of a GTP packet. It provides information on how to reach a network.

An APN has the following format:

```
<network_id>[.mnc<mnc_int>.mcc<mcc_int>.gprs]
```

Where:

- `<network_id>` is a network identifier or name that identifies the name of a network, for example, `example.com` or `internet`.
- `[.mnc<mnc_int>.mcc<mcc_int>.gprs]` is the optional operator identifier that uniquely identifies the operator's PLMN, for example `mnc123.mcc456.gprs`.

Combining these two examples results in a complete APN of `internet.mnc123.mcc456.gprs`.

By default, the unit permits all APNs. However, you can configure APN filtering to restrict roaming subscribers' access to external networks.

APN filtering applies only to the GTP *create pdp request* messages. The unit inspects GTP packets for both APN and selected modes. If both parameters match and APN filter entry, the unit applies the filter to the traffic.

Additionally, the unit can filter GTP packets based on the combination of an IMSI prefix and an APN. For more information, see [“Basic filtering options” on page 689](#).



You cannot add an APN when creating a new profile.

APN Filtering section on the New GTP Profile page	
<b>Enable APN Filter</b>	Select to enable APN filtering.
<b>Default APN Action</b>	Select the default action for APN filtering. If you select <i>Allow</i> , all sessions are allowed except those blocked by individual APN filters. If you select <i>Deny</i> , all sessions are blocked except those allowed by individual APN filters.
<b>Value</b>	The APN to be filtered.
<b>Mode</b>	The type of mode chosen that indicates where the APN originated and whether the Home Location Register (HLR) has verified the user subscription:
<b>Action</b>	The type of action that will be taken.
<b>Edit</b>	Modifies the settings within the filter. When you select <i>Edit</i> , the Edit window appears, which allows you to modify the settings of the APN.
<b>Delete</b>	Removes the APN from the list within the table, in the APN Filtering section.
<b>Add APN</b>	Adds a new APN filter to the list. When you select <i>Add APN</i> , the New window appears, which allows you to configure the APN settings.
New APN page	
<b>Value</b>	Enter an APN to be filtered. You can include wild cards to match multiple APNs. For example, the value <i>internet*</i> would match all APNs that begin with <i>internet</i> .
<b>Mode</b>	Select one or more of the available modes to indicate where the APN originated and whether the Home Location Register (HLR) has verified the user subscription.
	<b>Mobile Station provided</b>
	MS-provided APN, subscription not verified, indicates that the mobile station (MS) provided the APN and that the HLR did not verify the user's subscription to the network.
	<b>Network provided</b>
	Network-provided APN, subscription not verified, indicates that the network provided a default APN because the MS did not specify one, and that the HLR did not verify the user's subscription to the network.



<b>Subscription Verified</b>	MS or Network-provided APN, subscription verified, indicates that the MS or the network provided the APN and that the HLR verified the user's subscription to the network
<b>Action</b>	Select <i>Allow</i> or <i>Deny</i> .

### Basic filtering options


The International Mobile Station Identity (IMSI) is used by a GPRS Support Node (GSN) to identify a mobile station. Three elements make up every IMSI:

- the mobile country code (MCC)
- the mobile network code (MNC)
- the mobile subscriber identification number (MSIN).

The subscriber's home network—the public land mobile network (PLMN)—is identified by the IMSI prefix, formed by combining the MCC and MNC.

By default, the unit allows all IMSIs. You can add IMSI prefixes to deny GTP traffic coming from non-roaming partners. Any GTP packets with IMSI prefixes not matching the prefixes you set will be dropped. GTP *Create pdp* request messages are filtered and only IMSI prefixes matching the ones you set are permitted. Each GTP profile can have up to 1000 IMSI prefixes set.

An IMSI prefix and an APN can be used together to filter GTP packets if you set an IMSI filter entry with a non-empty APN.



You cannot add an IMSI when creating a new profile. You must add it after the profile has been created and you are editing the profile.

<b>IMSI Filtering section of the New GTP Profile page</b>	
<b>Enable IMSI Filter</b>	Select to enable IMSI filtering.
<b>Default IMSI Action</b>	Select the default action for IMSI filtering. If you select <i>Allow</i> , all sessions are allowed except those blocked by individual IMSI filters. If you select <i>Deny</i> , all sessions are blocked except those allowed by individual IMSI filters.
<b>APN</b>	The APN that is part of the IMSI that will be filtered.
<b>MCC-MNC</b>	The MCC-MNC part of the IMSI that will be filtered.
<b>Mode</b>	The type of mode that indicates where the APN originated and whether the Home Location Register (HLR) has verified the user subscription.
<b>Action</b>	The type of action that will be taken.
<b>Edit</b>	Modifies settings to an IMSI filter. When you select <i>Edit</i> , the Edit window appears, which allows you to modify the IMSI filter's settings.
<b>Delete</b>	Removes an IMSI filter from within the table, in the IMSI Filtering section.

<b>Add IMSI</b>	Adds a new IMSI filter to the list. When you select <i>Add IMSI</i> , the New window appears, which allows you to configure IMSI filter settings.
<b>New IMSI page</b>	
<b>APN</b>	Enter the APN part of the IMSI to be filtered.
<b>MCC-MNC</b>	Enter the MCC-MCC part of the IMSI to be filtered.
<b>Mode</b>	Select one or more of the available modes to indicate where the APN originated and whether the Home Location Register (HLR) has verified the user subscription.
<b>Mobile Station provided</b>	MS-provided APN, subscription not verified, indicates that the mobile station (MS) provided the APN and that the HLR did not verify the user's subscription to the network.
<b>Network provided</b>	Network-provided APN, subscription not verified, indicates that the network provided a default APN because the MS did not specify one, and that the HLR did not verify the user's subscription to the network.
<b>Subscription Verified</b>	MS or Network-provided APN, subscription verified, indicates that the MS or the network provided the APN and that the HLR verified the user's subscription to the network.
<b>Action</b>	Select <i>Allow</i> or <i>Deny</i> .

### Advanced filtering options

The FortiOS Carrier firewall supports advanced filtering against the attributes RAT, RAI, ULI, APN restriction, and IMEI-SV in GTP to block specific harmful GPRS traffic and GPRS roaming traffic. The following table shows some of the GTP context requests and responses that the firewall supports.

**Table 37:** Attributes supported by FortiCarrier firewalls

	<b>GTP Create PDP Context Request</b>	<b>GTP Create PDP Context Response</b>	<b>GTP Update PDP Context Request</b>	<b>GTP Update PDP Context Response</b>
<b>APN</b>	yes	yes	-	
<b>APN Restriction</b>	yes	-	-	yes
<b>IMEI-SV</b>	yes	-	-	-
<b>IMSI</b>	yes	-	yes	-
<b>RAI</b>	yes	-	yes	-

**Table 37:** Attributes supported by FortiCarrier firewalls

	GTP Create PDP Context Request	GTP Create PDP Context Response	GTP Update PDP Context Request	GTP Update PDP Context Response
APN	yes	yes	-	
RAT	yes	-	yes	-
ULI	yes	-	yes	-

When editing a GTP profile, select *Advanced Filtering* > *Add* to create and add a rule. When the rule matches traffic it will either allow or deny that traffic as selected in the rule.

<b>Advanced Filtering section on New GTP Profile page</b>	
<b>Enable</b>	Select to enable advanced filtering.
<b>Default Action</b>	Select the default action for advanced filtering. If you select <i>Allow</i> , all sessions are allowed except those blocked by individual advanced filters. If you select <i>Deny</i> , all sessions are blocked except those allowed by individual advanced filters.
<b>Messages</b>	The messages, for example, Create PDP Context Request.
<b>APN Restriction</b>	The APN restriction.
<b>RAT Type</b>	The RAT types associated with that filter.
<b>ULI</b>	The ULI pattern.
<b>RAI</b>	The RAI pattern.
<b>IMEI</b>	The IMEI pattern.
<b>Action</b>	The action that will be taken.
<b>Edit</b>	Modifies the filter's settings. When you select <i>Edit</i> , the Edit window appears, which allows you to modify the filter's settings.
<b>Delete</b>	Removes a filter from the list.
<b>Add</b>	Adds a filter to the list. When you select <i>Add</i> , the New window appears, which allows you to configure settings for messages, APN, IMSI, MSISDN, RAT type, ULI, RAI, IMEI patterns as well as the type of action.
<b>New Filtering page</b>	
<b>Messages</b>	The PDP content messages this profile will match.
<b>Create PDP Context Request</b>	Select to allow create PDP context requests.


<b>Create PDP Context Response</b>	Select to allow create PDP context responses.
<b>Update PDP Context Request</b>	Select to allow update PDP context requests.
<b>Update PDP Context Response</b>	Select to allow update PDP context responses.
<b>APN</b>	Enter the APN.
<b>APN Mode</b>	Select an APN mode as one or more of <ul style="list-style-type: none"> <li>• Mobile Station provided</li> <li>• Network provided</li> <li>• Subscription provided</li> </ul> This field is only available when an APN has been entered.
<b>Mobile Station provided</b>	MS-provided PAN, subscription not verified, indicates that the mobile station (MS) provided the APN and that the HLR did not verify the user's subscription to the network.
<b>Network provided</b>	Network-provided APN, subscription not verified, indicates that the network provided a default APN because the MS did not specify one, and that the HLR did no verify the user's subscription to the network.
<b>Subscription verified</b>	MS or Network-provided APN, subscription verified, indicates that the MS or the network provided the APN and that the HLR verified the user's subscription to the network.
<b>APN Restriction</b>	Select the type of restriction that you want. You can choose all of the types, or one of the types. You cannot choose multiple types. Types include: <ul style="list-style-type: none"> <li>• all</li> <li>• Public-1</li> <li>• Public-2</li> <li>• Private-1</li> <li>• Private-2</li> </ul>
<b>IMSI</b>	Enter the IMSI.
<b>MSISDN</b>	Enter the MSISDN.

<b>RAT Type</b>	<p>Optionally select the RAT type as any combination of the following:</p> <ul style="list-style-type: none"> <li>• Any</li> <li>• UTRAN</li> <li>• GERAN</li> <li>• Wifi</li> <li>• GAN</li> <li>• HSPA</li> </ul> <p>Some RAT types are GTPv1 specific.</p>
<b>ULI pattern</b>	Enter the ULI pattern.
<b>RAI pattern</b>	Enter the RAI pattern.
<b>IMEI pattern</b>	Enter the IMEI pattern.
<b>Action</b>	Select either <i>Allow</i> or <i>Deny</i> .

### Adding an advanced filtering rule

When adding a rule, use the following formats:

- Prefix, for example, range 31\* for MCC matches MCC from 310 to 319.
- Range, for example, range 310-319 for MCC matches MCC from 310 to 319.
- Mobile Country Code (MCC) consists of three digits. The MCC identifies the country of domicile of the mobile subscriber.
- Mobile Network Code (MNC) consists of two or three digits for GSM/UMTS applications. The MNC identifies the home PLMN of the mobile subscriber. The length of the MNC (two or three digits) depends on the value of the MCC. Best practices dictate not to mix two and three digit MNC codes within a single MCC area.
- Location Area Code (LAC) is a fixed length code (of 2 octets) identifying a location area within a PLMN. This part of the location area identification can be coded using a full hexadecimal representation except for the following reserved hexadecimal values: 0000 and FFFE. These reserved values are used in some special cases when no valid LAI exists in the MS (see 3GPP TS 24.008, 3GPP TS 31.102 and 3GPP TS 51.011).
- Routing Area Code (RAC) of a fixed length code (of 1 octet) identifies a routing area within a location.
- CI or SAC of a fixed length of 2 octets can be coded using a full hexadecimal expression.
- Type Allocation Code (TAC) has a length of 8 digits.
- Serial Number (SNR) is an individual serial number identifying each equipment within each TAC. SNR has a length of 6 digits.
- Software Version Number (SVN) identifies the software version number of the mobile equipment. SVN has a length of 2 digits.



You cannot add an advanced filtering rule when creating a new profile. You must add it after the profile has been created and you are editing the profile.

## Information Element (IE) removal policy options

In some roaming scenarios, the unit is installed on the border of the PLMN and the GRX. In this configuration, the unit supports information element (IE) removal policies to remove any combination of R6 IEs (RAT, RAI, ULI, IMEI-SV and APN restrictions) from the types of messages described in [“Advanced filtering options” on page 690](#), prior to forwarding the messages to the HGGSN (proxy mode).

<b>IE removal policy section of the New GTP Profile page</b>	
<b>Enable</b>	Select to enable this option.
<b>SGSN address of message IE</b>	The firewall address or address group that contains the SGSN addresses.
<b>IEs to be removed</b>	The IE types that will be removed. These include APN Restriction, RAT, RAI, ULI, and IMEI.
<b>Add</b>	Adds an IE removal policy. When you select <i>Add</i> , the New window appears, which allows you to configure the IE policy.
<b>Edit</b>	Modifies settings from within the IE removal policy. When you select <i>Edit</i> , the Edit window appears, which allows you to modify the settings within the policy.
<b>Delete</b>	Removes the IE removal policy from the list.
<b>New IE policy page</b>	
<b>SGSN address</b>	Select a firewall address or address group that contains SGSN addresses.
<b>IEs to be removed</b>	Select one or more IE types to be removed. These include APN Restriction, RAT, RAI, ULI, and IMEI.

## Encapsulated IP traffic filtering options

You can use encapsulated IP traffic filtering to filter GTP sessions based on information contained in the data stream. to control data flows within your infrastructure. You can configure IP filtering rules to filter encapsulated IP traffic from mobile stations by identifying the source and destination policies. For more information, see [“When to use encapsulated IP traffic filtering” on page 740](#).

Expand *Encapsulated IP Traffic Filtering* in the GTP profile to reveal the options.

---

### Encapsulated IP Traffic Filtering section of the New GTP Profile page

<b>Enable IP Filter</b>	Select to enable encapsulated IP traffic filtering options.
<b>Default IP Action</b>	Select the default action for encapsulated IP traffic filtering. If you select <i>Allow</i> , all sessions are allowed except those blocked by individual encapsulated IP traffic filters. If you select <i>Deny</i> , all sessions are blocked except those allowed by individual encapsulated IP traffic filters.

<b>Source</b>	Select a source IP address from the configured firewall IP address or address group lists. Any encapsulated traffic originating from this IP address will be a match if the destination also matches.
<b>Destination</b>	Select a destination IP address from the configured firewall IP address or address group lists. Any encapsulated traffic being sent to this IP address will be a match if the destination also matches.
<b>Action</b>	The type of action that will be taken. Select to Allow or Deny encapsulated traffic between this source and Destination.
<b>Edit</b>	Modifies the source, destination or action settings.
<b>Add IP Policy</b>	Adds a new encapsulated IP traffic filter. When you select <i>Add IP Policy</i> , the New window appears which allows you to configure IP policy settings.
<b>New (window)</b>	
<b>Source</b>	Select the source firewall address or address group.
<b>Destination</b>	Select the destination firewall address or address group.
<b>Action</b>	Select <i>Allow</i> or <i>Deny</i> .

### Encapsulated non-IP end user traffic filtering options

Depending on the installed environment, it may be beneficial to detect GTP packets that encapsulate non-IP based protocols. You can configure the FortiOS Carrier firewall to permit a list of acceptable protocols, with all other protocols denied.

The encoded protocol is determined in the PDP Type Organization and PDP Type Number fields within the End User Address Information Element. The PDP Type Organization is a 4-bit field that determines if the protocol is part of the ETSI or IETF organizations. Values are zero and one, respectively. The PDP Type field is one byte long. Both GTP specifications list only PPP, with a PDP Type value of one, as a valid ETSI protocol. PDP Types for the IETF values are determined in the "Assigned PPP DLL Protocol Numbers" sections of RFC1700. The PDP types are compressed, meaning that the most significant byte is skipped, limiting the protocols listed from 0x00 to 0xFF.

<b>Encapsulated Non-IP End User Address Filtering section of the New GTP Profile page</b>	
<b>Enable Non-IP Filter</b>	Select to enable encapsulated non-IP traffic filtering.
<b>Default Non-IP Action</b>	Select the default action for encapsulated non-IP traffic filtering. If you select <i>Allow</i> , all sessions are allowed except those blocked by individual encapsulated non-IP traffic filters. If you select <i>Deny</i> , all sessions are blocked except those allowed by individual encapsulated non-IP traffic filters.
<b>Type</b>	The type chosen, <i>AESTI</i> or <i>IETF</i> .
<b>Start Protocol</b>	The beginning protocol port number range.

<b>End Protocol</b>	The end of the protocol port number range.
<b>Action</b>	The type of action that will be taken.
<b>Edit</b>	Modify a non-IP filter's settings in the list. When you select <i>Edit</i> , the Edit window appears, which allows you to modify the Non-IP policy settings.
<b>Delete</b>	Remove a non-IP policy from the list.
<b>Add Non-IP Policy</b>	Add a new encapsulated non-IP traffic filter. When you select <i>Add Non-IP Policy</i> , you are automatically redirected to the New page.
<b>New (window)</b>	
<b>Type</b>	Select <i>AESTI</i> or <i>IETF</i> .
<b>Start Protocol End Protocol</b>	Select a start and end protocol from the list of protocols in RFC 1700. Allowed range includes 0 to 255 (0x00 to 0xff). Some common protocols include: <ul style="list-style-type: none"> <li>• 33 (0x0021) Internet Protocol</li> <li>• 35 (0x0023) OSI Network Layer</li> <li>• 63 (0x003f) NETBIOS Framing</li> <li>• 65 (0x0041) Cisco Systems</li> <li>• 79 (0x004f) IP6 Header Compression</li> <li>• 83 (0x0053) Encryption</li> </ul>
<b>Action</b>	Select <i>Allow</i> or <i>Deny</i> .

### Protocol Anomaly prevention options

Use protocol anomaly detection options to detect or deny protocol anomalies according to GTP standards and tunnel state. Protocol anomaly attacks involve malformed or corrupt packets that typically fall outside of the protocol specifications. Packets cannot pass through if they fail the sanity check.

<b>Protocol Anomaly section of the New GTP Profile page</b>	
<b>Invalid Reserved Field</b>	GTP version 0 (GSM 09.60) headers specify a number of fields that are marked as "Spare" and contain all ones (1). GTP packets that have different values in these fields are flagged as anomalies. GTP version 1 (GSM 29.060) makes better use of the header space and only has one, 1-bit, reserved field. In the first octet of the GTP version1 header, bit 4 is set to zero.
<b>Reserved IE</b>	Both versions of GTP allow up to 255 different Information Elements (IE). However, a number of Information Elements values are undefined or reserved. Packets with reserved or undefined values will be filtered.
<b>Miss Mandatory IE</b>	GTP packets with missing mandatory Information Elements (IE) will not be passed to the GGSN.



<b>Out of State Message</b>	<p>The GTP protocol requires a certain level of state to be kept by both the GGSN and SGSN. Some message types can only be sent when in a specific GTP state. Packets that do not make sense in the current state are filtered or rejected</p> <p>Both versions of GTP allow up to 255 different message types. However, a number of message type values are undefined or reserved.</p> <p>Best practices dictate that packets with reserved or undefined values will be filtered.</p>
<b>Out of State IE</b>	GTP Packets with out of order Information Elements are discarded.
<b>Spoofed Source Address</b>	<p>The End User Address Information Element in the PDP Context Create &amp; Response messages contain the address that the mobile station (MS) will use on the remote network. If the MS does not have an address, the SGSN will set the End User Address field to zero when sending the initial PDP Context Create message. The PDP Context Response packet from the GGSN will then contain an address to be assigned to the MS. In environments where static addresses are allowed, the MS will relay its address to the SGSN, which will include the address in the PDP Context Create Message. As the MS address is negotiated within the PDP Context creation handshake, any packets originating from the MS that contain a different source address are detected and dropped.</p>

### Anti-Overbilling options

You can configure the FortiOS Carrier firewall to prevent overbilling subscribers for traffic over the. To enable anti-overbilling, you must configure both the Gn/Gp firewall and the Gi firewall.

Expand *Anti-Overbilling* in the GTP profile to reveal these settings.

<b>Anti-Overbilling section of the New GTP Profile page</b>	
<b>Gi Firewall IP Address</b>	The IP address of the unit's interface configured as a Gi gateway.
<b>Port</b>	The SG security port number. The default port number is port 21123. Change this number if your system uses a different SG port.
<b>Interface</b>	Select the unit interface configured as a Gi gateway.
<b>Security Context ID</b>	Enter the security context ID. This ID must match the ID entered on the server Gi firewall. The default security context ID is 696.

### Log options

All the GTP logs are treated as a subtype of the event logs. To enable GTP logging, you must:

- configure the GTP log settings in a GTP profile
- enable GTP logging when you configure log and report settings.

### To enable GTP logging after a GTP profile has been configured

1. Go to *Log & Report > Log Config > Log Settings*.
2. Select *Event Logging*, and select *GTP service event*.
3. Select *Apply*.

<b>Log section of the New GTP Profile page</b>	
<b>Log Frequency</b>	<p>Enter the number of messages to drop between logged messages.</p> <p>An overflow of log messages can sometimes occur when logging rate-limited GTP packets exceed their defined threshold. To conserve resources on the syslog server and the Carrier-enabled FortiGate unit, you can specify that some log messages are dropped. For example, if you want only every twentieth message to be logged, set a logging frequency of 20. This way, 20 messages are skipped and the next logged.</p> <p>Acceptable frequency values range from 0 to 2147483674. When set to '0', no messages are skipped.</p>
<b>Forwarded Log</b>	Select to log forwarded GTP packets.
<b>Denied Log</b>	Select to log GTP packets denied or blocked by this GTP profile.
<b>Rate Limited Log</b>	Select to log rate-limited GTP packets.
<b>State Invalid Log</b>	Select to log GTP packets that have failed stateful inspection.
<b>Tunnel Limit Log</b>	Select to log packets dropped because the maximum limit of GTP tunnels for the destination GSN is reached.

<b>Extension Log</b>	<p>Select to log extended information about GTP packets. When enabled, this additional information will be included in log entries:</p> <ul style="list-style-type: none"> <li>• IMSI</li> <li>• MSISDN</li> <li>• APN</li> <li>• Selection Mode</li> <li>• SGSN address for signaling</li> <li>• SGSN address for user data</li> <li>• GGSN address for signaling</li> <li>• GGSN address for user data</li> </ul>
<b>Traffic count Log</b>	<p>Select to log the total number of control and user data messages received from and forwarded to the GGSNs and SGSNs that the unit protects.</p> <p>The unit can report the total number of user data and control messages received from and forwarded to the GGSNs and SGSNs it protects. Alternately, the total size of the user data and control messages can be reported in bytes. The unit differentiates between traffic carried by each GTP tunnel, and also between GTP-User and GTP-Control messages.</p> <p>The number of messages or the number of bytes of data received from and forwarded to the SGSN or GGSN are totaled and logged if a tunnel is deleted.</p> <p>When a tunnel is deleted, the log entry contains:</p> <ul style="list-style-type: none"> <li>• Timestamp</li> <li>• Interface name (if applicable)</li> <li>• SGSN IP address</li> <li>• GGSN IP address</li> <li>• TID</li> <li>• Tunnel duration time in seconds</li> <li>• Number of messages sent to the SGSN</li> <li>• Number of messages sent to the GGSN</li> </ul>

### Specifying logging types

You can configure the unit to log GTP packets based on their status with GTP traffic logging.

The status of a GTP packet can be any of the following 5 states:

- **Forwarded** - a packet that the unit transmits because the GTP policy allows it
- **Prohibited** - a packet that the unit drops because the GTP policy denies it
- **Rate-limited** - a packet that the unit drops because it exceeds the maximum rate limit of the destination GSN
- **State-invalid** - a packet that the unit drops because it failed stateful inspection
- **Tunnel-limited** - a packet that the unit drops because the maximum limit of GTP tunnels for the destination GSN is reached.

The following information is contained in each log entry:

- Timestamp
- Source IP address
- Destination IP address
- Tunnel Identifier (TID) or Tunnel Endpoint Identifier (TEID)
- Message type
- Packet status: forwarded, prohibited, state-invalid, rate-limited, or tunnel-limited
- Virtual domain ID or name
- Reason to be denied if applicable.

# MMS Security features

FortiOS Carrier includes all the Security features of FortiOS with extra features specific to MMS carrier networks.

This section includes:

- [Why scan MMS messages for viruses and malware?](#)
- [MMS virus scanning](#)
- [MMS file filtering](#)
- [MMS content-based Antispam protection](#)
- [MMS DLP archiving](#)

## Why scan MMS messages for viruses and malware?

The requirement for scanning MM1 content comes from the fact that MMS is an increasingly popular technique for propagating malware between mobile devices. See [“MMS virus scanning” on page 702](#).

### Example: COMMWARRIOR

This is a virus for Series 60 type cell phones, such as Nokia, operating Symbian OS version 6 [or higher]. The object of the virus is to spread to other phones using Bluetooth and MMS as transport avenues. The targets are selected from the contact list of the infected phone and also sought via Bluetooth searching for other Bluetooth-enabled devices (phones, printers, gaming devices etc.) in the proximity of the infected phone.

This virus is more than a proof of concept - it has proven successfully its ability to migrate from a zoo collection to being in-the-wild. Currently, this virus is being reported in over 18 different countries around Europe, Asia and North America.

When the virus first infects a cell phone, a prompt is displayed asking the recipient if they want to install “Caribe”. Symptoms of an infected phone may include rapid battery power loss due to constant efforts by the virus to spread to other phones via a Bluetooth seek-and-connect outreach.

The following variants among others are currently scanned by the FortiOS Carrier devices, in addition to more signatures that cover all known threats.

- **SymbOS/COMWAR.V10B!WORM**
- Aliases: SymbOS.Commwarrior.B, SymbOS/Commwar.B, SymbOS/Commwar.B!wm, SymbOS/Commwar.B-net, SymbOS/Commwarrior.b!sis, SymbOS/Comwar.B,

SymbOS/Comwar.B!wm, SymbOS/Comwar.B-wm, SYMBOS\_COMWAR.B,  
SymbOS/Comwar.1.0.B!wormSYMBOS/COMWAR.V10B.SP!WORM [spanish version]

- First Discovered In The Wild: July 04, 2007
- Impact Level: 1
- Virus Class: Worm
- Virus Name Size: 23,320
- **SymbOS/Commwar.A!worm**
- Aliases: Commwarrior-A, SymbOS.Commwarrior.A [NAV], SymbOS/Commwar.A-net, SymbOS/Commwar\_ezboot.A-ne, SymbOS/Comwar.A, SymbOS/Comwar.A-wm, SYMBOS\_COMWAR.A [Trend]
- First Discovered In The Wild: May 16 2005
- Impact Level: 1
- Virus Class: Worm
- Virus Name Size: 27,936
- SymbOS/Commwarriie.C-wm
- Aliases: None
- First Discovered In The Wild: Oct 17 2005
- Impact Level: 1
- Virus Class: File Virus
- Virus Name Size: None

For the latest list of threats Fortinet devices detect, go to the [FortiGuard Center Resource Library's Mobile index](#).

## MMS virus scanning

You can use MMS virus scanning to scan content contained within MMS messages for viruses. FortiOS Carrier virus scanning can be applied to the MM1, MM3, MM4, and MM7 interfaces to detect and remove content containing viruses at many points in an MMS network. Perhaps the most useful interface to apply virus scanning would be the MM1 interface to block viruses sent by mobile users before they get into the service provider network.

To go to MMS virus scanning, go to *Security Profiles > Carrier > MMS Profile*, select an existing or create a new profile, and expand *MMS Scanning*. See “[MMS scanning options](#)” on page 663.

This section includes:

- [MMS virus monitoring](#)
- [MMS virus scanning blocks messages \(not just attachments\)](#)
- [Scanning MM1 retrieval messages](#)
- [Configuring MMS virus scanning](#)
- [Removing or replacing blocked messages](#)
- [Carrier Endpoint Block](#)
- [MMS Content Checksum](#)
- [Passing or blocking fragmented messages](#)
- [Client comforting](#)
- [Server comforting](#)
- [Handling oversized MMS messages](#)

## MMS virus monitoring

To enable MMS virus monitoring, expand *MMS Scanning* and enable *Monitor only* for the selected MMS types.

This feature causes the FortiOS Carrier unit to record log messages when MMS scanning options find a virus, match a file name, or match content using any of the other MMS scanning options. Selecting this option enables reporting on viruses and other problems in MMS traffic without affecting users.

## MMS virus scanning blocks messages (not just attachments)

To enable MMS virus scanning, expand *MMS Scanning* and enable *Virus Scan* for the selected MMS types.

Because MM1 and MM7 use HTTP, the oversize limits for HTTP and the HTTP antivirus port configurations also apply to MM1 and MM7 scanning. See

MM3 and MM4 use SMTP and the oversize limits for SMTP and the SMTP antivirus port configurations also apply to MM3 and MM4 scanning.

The message contents will be scanned for viruses, matched against the file extension blocking lists and scanned for banned words. All these items will be configured via the standard GUI interfaces available for the other protocols and will be controlled at the protection profile level with new options specifically for the MM1 messages.

The FortiOS Carrier unit extracts the sender's Mobile Subscriber Integrated Services Digital Network Number (MSISDN) from the HTTP headers if available. The `POST` payload will be sent to the scanunits which will parse the MMS content and scan each message data section. If any part of the data is to be blocked, the proxy will be informed, the connection to the MMSC will be reset and the Carrier-enabled FortiGate unit will return an `HTTP 200 OK` message with an `m-send-conf` payload to the client to prevent a retry. Finally the appropriate logging, alert, and replacement message events will be triggered.

For client notification, the `x-mms-response-status` and `x-mms-response-text` fields can also be customized as required.

## Scanning MM1 retrieval messages

To scan MM1 retrieval messages, expand *MMS Scanning* and select *Scan MM1 message retrieval*.

Select to scan message retrievals that use MM1. If you enable *Virus Scan* for all MMS interfaces, messages are also scanned while being sent. In this case, you can disable MM1 message retrieval scanning to improve performance.

## Configuring MMS virus scanning

To configure MMS virus scanning, expand *MMS Scanning* and enable *Virus Scan*.

Once applied to a security policy, the MMS protection profile will then perform virus scans on all traffic accepted by that policy.

## Removing or replacing blocked messages

To remove blocked messages, expand *MMS Scanning* and select *Remove Blocked* for the selected MMS types.

Select *Remove Blocked* remove blocked content from each protocol and replace it with the replacement message. If FortiOS Carrier is to preserve the length of the message when

removing blocked content, as may occur when billing is affected by the length of the message, select *Constant*.

If you only want to monitor blocked content, select *Monitor Only*.

## Carrier Endpoint Block

A carrier endpoint defines a specific client on the carrier network. Typically the client IP address is used to identify the client, however on a carrier network this may be impractical when the client is using a mobile device. Other identifying information such as the MSISDN number is used instead.

This information can be used to block a specific endpoint on the network. Reasons for blocking may include clients whose accounts have expired, clients from another carrier, clients who have sent malicious content (phishing, exploits, viruses, etc), or other violations of terms of use.

### Enabling carrier endpoint blocking

To enable carrier endpoint blocking you first need to create a carrier endpoint filter list, and then enable it.

#### To enable carrier endpoint blocking - web-based manager

1. Create a carrier endpoint filter list. See
2. Go to *Security Profiles > Carrier > MMS Profile*.
3. Select *Create New*, or select an existing profile to edit and select *Edit*.
4. Expand MMS Scanning.
5. Select one or more types of MMS messaging to enable endpoint blocking on.
6. Select the carrier endpoint filter list to use in matching the endpoints to be blocked.



In MMS Profile, endpoints can only be blocked.

### Create a carrier endpoint filter list

A carrier endpoint filter list contains one or more carrier endpoints to match. When used in MMS scanning entries in the filter list that are matched are blocked.

You can configure multiple filter lists for different purposes and groups of clients, such as blocking clients, clients with different levels of service agreements, and clients from other carriers. See [“Carrier endpoint filter lists configuration settings” on page 681](#).

#### To create a carrier endpoint filter list - web-based manager

1. Go to *Security Profiles > Carrier > Carrier Endpoint Filter Lists*.
2. Select *Create New*.
3. Enter a descriptive name for the filter list, such as `blocked_clients` or `CountryX_clients`, and select OK.
4. Select *Create New* to add one or more entries to the list.
5. Select OK to return to display the list of filter lists.

### Configuring endpoint filter list entries

For each single endpoint or group of endpoints have part of their identifying information in common, you create an entry in the endpoint filter list.



For example a `blocked_clients` filter list may include entries for single endpoints added as each one needs to be blocked and a group of clients from a country that does not allow certain services.

### To configure an endpoint filter list entry - web-based manager

1. Select *Create New*.
2. Enter the following information and select OK.

<b>Name</b>	Name of endpoint filter list. Select this name in an MMS protection profile.
<b>Comments</b>	Optional description of the endpoint filter list.
<b>Check/Uncheck All</b>	Select the check box to enable all endpoint patterns in the MMS filter list.  Clear the check box to disable all entries on the MMS filter list.  You can also select or clear individual check boxes to enable or disable individual endpoint patterns.
<b>Pattern</b>	The pattern that FortiOS Carrier uses to match with endpoints. The pattern can be a single endpoint or consist of wildcards or Perl regular expressions that will match more than one endpoint. For more on wildcard and regular expressions, see <i>Using wildcards and Perl regular expressions</i> in the <a href="#">UTM handbook</a> chapter.
<b>Action</b>	Select the action taken by FortiOS Carrier for messages from a carrier endpoint that matches the endpoint pattern:  None - No action is taken.  Block - MMS messages from the endpoint are not delivered and FortiOS Carrier records a log message.  <b>Exempt from mass MMS</b> - MMS messages from the endpoint are delivered and are exempt from mass MMS filtering. Mass MMS filtering is configured in MMS protection profiles and is also called MMS Bulk Email Filtering and includes MMS message flood protection and MMS duplicate message detection. A valid use of mass MMS would be when a service provider notifies customers of a system-wide event such as a shutdown.  <b>Exempt from all scanning</b> - MMS messages from the endpoint are delivered and are exempt from all MMS protection profile scanning.
<b>Content Archive</b>	MMS messages from the endpoint are delivered, the message content is DLP archived according to MMS DLP archive settings. Content archiving is also called DLP archiving.
<b>Intercept</b>	MMS messages from the endpoint are delivered. Based on the quarantine configuration, attached files may be removed and quarantined.

<b>Pattern Type</b>	The pattern type: <i>Wildcard</i> , <i>Regular Expression</i> , or <i>Single Endpoint</i> .
<b>Enable</b>	Select to enable this endpoint filter pattern.

## Blocking network access based on endpoints

You can use endpoint IP filtering to block traffic from source IP addresses associated with endpoints. You can also configure FortiOS Carrier to record log messages whenever endpoint IP filtering blocks traffic. Endpoint IP filtering blocks traffic at the IP level, before the traffic is accepted by a security policy.

To configure endpoint IP filtering, go to *Security Profiles > Carrier > IP Filter* and add endpoints to the IP filter list. For each endpoint you can enable or disable both blocking traffic and logging blocked traffic.



You cannot add endpoint patterns to the endpoint IP filter list. You must enter complete and specific endpoints that are valid for your network.



The only action available is block. You cannot use endpoint IP filtering to exempt endpoints from IP filtering or to content archive or quarantine communication sessions.

FortiOS Carrier looks in the current user context list for the endpoints in the IP filter list and extracts the source IP addresses for these endpoints. Then any communication session with a source IP address that matches one of these IP addresses is blocked at the IP level, before the communication session is accepted by a security policy.

FortiOS Carrier dynamically updates the list of IP addresses to block as the user context list changes. Only these updated IP addresses are blocked by endpoint IP filtering.

For information about the carrier endpoints and the user context list, including how entries are added to and removed from this list, see [For more information on carrier endpoints, see the FortiOS Handbook User Authentication chapter.](#)

## MMS Content Checksum

The MMS content checksum feature attempts to match checksums of known malicious MMS messages, and on a successful match it will be blocked. The checksums are applied to each part of the message—attached files and message body have separate checksums. These checksums are created with CRC-32, the same method as FortiAnalyzer checksums.

For example, if an MMS message contains a browser exploit in the message body, you can add the checksum for that message body to the list, and future occurrences of that exact message will be blocked. Content will be replaced by the content checksum block notification replacement message for that type of MMS message, and if it is enabled the event will be logged.

One possible implementation would to configure all .sis files to be intercepted. When one is found to be infected or malicious it would be added to the MMS content checksum list.

To use this feature a list of one or more malicious checksums must be created and then the feature is enabled using that list. For a detailed list of options, see [“MMS Content Checksum” on page 674.](#)

### To configure an MMS content checksum list

1. Go to *Security Profiles > Carrier > MMS Content Checksum*.
2. Select *Create New*.
3. Enter a name for the list of checksums, and select OK.  
You are taken to the edit screen for that new list.
4. Select *Create New* to add a checksum.
5. Enter the *Name* and *Checksum*, and select OK.  
The checksum is added to the list.

To add more checksums to the list, repeat steps 4 and 5.

To remove a checksum from the list you can either delete the checksum or simply disable it and leave it in the list.

To enable MMS content checksums, expand *MMS Scanning* and select *MMS Content Checksum* for the selected MMS types. Select the checksum list to match.

## Passing or blocking fragmented messages

Select to pass fragmented MM3 and MM4 messages. Fragmented MMS messages cannot be scanned for viruses. If you do not select these options, fragmented MM3 and MM4 message are blocked.

The *Interval* is the time in seconds before client comforting starts after the download has begun, and the time between sending subsequent data.

The *Amount* is the number of bytes sent by client or server comforting at each interval.

## Client comforting

In general, client comforting is available for MM1 and MM7 messaging and provides a visual display of progress for web page loading or HTTP or FTP file downloads. Client comforting does this by sending the first few packets of the file or web page being downloaded to the client at configured time intervals so that the client is not aware that the download has been delayed. The client is the web browser or FTP client. Without client comforting, clients and their users have no indication that the download has started until the Carrier-enabled FortiGate unit has completely buffered and scanned the download. During this delay users may cancel or repeatedly retry the transfer, thinking it has failed.

The appearance of a client comforting message (for example, a progress bar) is client-dependent. In some instances, there will be no visual client comforting cue.

During client comforting, if the file being downloaded is found to be infected, then the Carrier-enabled FortiGate unit caches the URL and drops the connection. The client does not receive any notification of what happened because the download to the client had already started. Instead the download stops, and the user is left with a partially downloaded file.

If the user tries to download the same file again within a short period of time, then the cached URL is matched and the download is blocked. The client receives the Infection cache message replacement message as a notification that the download has been blocked. The number of URLs in the cache is limited by the size of the cache.



Client comforting can send unscanned (and therefore potentially infected) content to the client. Only enable client comforting if you are prepared to accept this risk. Keeping the client comforting interval high and the amount low will reduce the amount of potentially infected data that is downloaded.

## MM1 and MM7 client comforting steps

Since MM1 and MM7 messages use HTTP, MM1 and MM7 client comforting operates like HTTP client comforting.

The following steps show how client comforting works for a download of a 1 Mbyte file with the client comforting interval set to 20 seconds and the client comforting amount set to 512 bytes.

1. The client requests the file.
2. The Carrier-enabled FortiGate unit buffers the file from the server. The connection is slow, so after 20 seconds about one half of the file has been buffered.
3. The Carrier-enabled FortiGate unit continues buffering the file from the server, and also sends 512 bytes to the client.
4. After 20 more seconds, the FortiGate unit sends the next 512 bytes of the buffered file to the client.
5. When the file has been completely buffered, the client has received the following amount of data:  
$$ca * (T/ci) \text{ bytes} == 512 * (40/20) == 512 * 2 == 1024 \text{ bytes,}$$
where  $ca$  is the client comforting amount,  $T$  is the buffering time and  $ci$  is the client comforting interval.
6. If the file does not contain a virus, the Carrier-enabled FortiGate unit sends the rest of the file to the client. If the file is infected, the FortiGate closes the data connection but cannot send a message to the client.

## Server comforting

Server comforting can be selected for each protocol.

Similar to client comforting, you can use server comforting to prevent server connection timeouts that can occur while waiting for FortiOS Carrier to buffer and scan large `POST` requests from slow clients.

The *Interval* is the time in seconds before client and server comforting starts after the download has begun, and the time between sending subsequent data.

The *Amount* is the number of bytes sent by client or server comforting at each interval.

## Handling oversized MMS messages

Select *Block* or *Pass* for files and email messages exceeding configured thresholds for each protocol.

The oversize threshold refers to the final size of the message, including attachments, after encoding by the client. Clients can use a variety of encoding types; some result in larger file sizes than the original attachment. As a result, a file may be blocked or logged as oversized even if the attachment is several megabytes smaller than the oversize threshold.

## MM1 sample messages

```
Internet Protocol, Src Addr: 10.128.206.202 (10.128.206.202), Dst Addr:
10.129.192.190 (10.129.192.190)
Transmission Control Protocol, Src Port: 34322 (34322), Dst Port: http
(80), Seq: 1, Ack: 1, Len: 1380
 Source port: 34322 (34322)
 Destination port: http (80)
 Header length: 20 bytes
 Flags: 0x0010 (ACK)
```

Window size: 24840  
Checksum: 0x63c1 (correct)

## HTTP proxy

```
Hypertext Transfer Protocol
POST / HTTP/1.1\r\n
 Request Method: POST
 Request URI: /
 Request Version: HTTP/1.1
Host: 10.129.192.190\r\n
Accept: */*,
application/vnd.wap.sic,application/vnd.wap.mms-message,text/x-hdml,image/mng,image/x-mng,video/mng,video/x-mng,image/bmp\r\n
Accept-Charset: utf-8,*\r\n
Accept-Language: en\r\n
Content-Length: 25902\r\n
Content-Type: application/vnd.wap.mms-message\r\n
User-Agent: Nokia7650/1.0 SymbianOS/6.1 Series60/0.9
Profile/MIDP-1.0 Configuration/CLDC-1.0 UP.Link/6.2.1\r\n
x-up-devcap-charset: utf-8\r\n
x-up-devcap-max-pdu: 102400\r\n
x-up-uplink: magh-ip.mi.vas.omnitel.it\r\n
x-wap-profile: "http://nds.nokia.com/uaprof/N7650r200.xml"\r\n
x-up-subno: 1046428312-826\r\n
x-up-calling-line-id: 393475171234\r\n
x-up-forwarded-for: 10.211.4.12\r\n
x-forwarded-for: 10.211.4.12\r\n
Via: 1.1 magh-ip.mi.vas.omnitel.it\r\n
\r\n
```

## Scan engine

```
MMS Message Encapsulation, Type: m-send-req
X-Mms-Message-Type: m-send-req (0x80)
X-Mms-Transaction-ID: 1458481935
X-Mms-MMS-Version: 1.0
From: <insert address>
To: 3475171234/TYPE=PLMN
X-Mms-Message-Class: Personal (0x80)
X-Mms-Expiry: 21600.000000000 seconds
X-Mms-Priority: Normal (0x81)
X-Mms-Delivery-Report: No (0x81)
X-Mms-Read-Report: No (0x81)
Content-Type: application/vnd.wap.multipart.related;
start=<1822989907>; type=application/smil
 Start: <1822989907>
 Type: application/smil
Data (Post)
 Multipart body
 Part: 1, content-type: text/plain
 Content-Type: text/plain; charset=iso-10646-ucs-2;
name=Ciao.txt
 Charset: iso-10646-ucs-2
 Name: Ciao.txt
Headers
 Content-Location: Ciao.txt
```

```
Line-based text data: text/plain
\377\376C\000i\000a\000o\000
[Unreassembled Packet: MMSE]
```

## MMS file filtering

Use MMS file filtering to apply antivirus file filtering to MMS traffic. Select a file filter list to apply. To configure MMS file filtering, go to *Security Profiles > Carrier > MMS Profile*, select an existing or create a new profile, and expand *MMS Scanning*.

Configure the FortiGate file filter to block files by:

- **File pattern:** Files can be blocked by name, extension, or any other pattern. File pattern blocking provides the flexibility to block potentially harmful content.
- File pattern entries are not case sensitive. For example, adding `*.exe` to the file pattern list also blocks any files ending in `.EXE`.
- In addition to the built-in patterns, you can specify more file patterns to block.
- **File type:** Files can be blocked by type, without relying on the file name to indicate what type of files they are. When blocking by file type, the FortiGate unit analyzes the file and determines the file type regardless of the file name.

For standard operation, you can choose to disable file filter in the protection profile, and enable it temporarily to block specific threats as they occur.

The FortiGate unit can take either of these actions toward files that match a configured file pattern or type:

- **Allow:** the file is allowed to pass.
- **Block:** the file is blocked and a replacement messages will be sent to the user. If both file filter and virus scan are enabled, the Carrier-enabled FortiGate unit blocks files that match the enabled file filter and does not scan these files for viruses.
- **Intercept:** the file will be archived to the local hard disk or the FortiAnalyzer unit.

The Carrier-enabled FortiGate unit also writes a message to the virus log and sends an alert email message if configured to do so.

Files are compared to the enabled file patterns and then the file types from top to bottom. If a file does not match any specified patterns or types, it is passed along to antivirus scanning (if enabled). In effect, files are passed if not explicitly blocked.

Using the allow action, this behavior can be reversed with all files being blocked unless explicitly passed. Simply enter all the file patterns or types to be passed with the allow attribute. At the end of the list, add an all-inclusive wildcard (`*.*`) with a block action. Allowed files continue to antivirus scanning (if enabled) while files not matching any allowed patterns are blocked by the wildcard at the end.

## Built-in patterns and supported file types

The FortiGate unit is preconfigured with a default list of file patterns:

- executable files (\*.bat, \*.com, and \*.exe)
- compressed or archive files (\*.gz, \*.rar, \*.tar, \*.tgz, and \*.zip)
- dynamic link libraries (\*.dll)
- HTML application (\*.hta)
- Microsoft Office files (\*.doc, \*.ppt, \*.xl?)
- Microsoft Works files (\*.wps)
- Visual Basic files (\*.vb?)
- screen saver files (\*.scr)
- program information files (\*.pif)
- control panel files (\*.cpl)

The FortiGate unit can take actions against the following file types:

**Table 38:** Supported file types

arj	activemime	aspack	base64	bat	binhex	bzip	bzip2
cab	class	cod	elf	exe	fsg	gzip	hlp
hta	html	jad	javascript	lzh	mime	msc	msoffic e
petite	prc	rar	sis	tar	upx	uue	zip
unknown	ignored						



The “unknown” type is any file type that is not listed in the table. The “ignored” type is the traffic the FortiGate unit typically does not scan. This includes primarily streaming audio and video.

### Filtering based on file name

There are filenames that are known to be associated with malware such as viruses and trojans. There are filenames you may associate with other undesirable content in addition to malware. In these situations you want to select specific filenames to filter.

You do not have to match the entire filename. For example if you wanted to block all files with the word `trojan` in them you could use wildcards to accomplish this - `*trojan*` . This allows you to select the entire filename, part of the filename, or just the file type to match.

The following procedure creates a filter list called `filterExampleFiles` that filters two files called `exampleTrojanFile.abc` and `*trojan*.def` . When completed, this file filter list can be included in an MMS profile.

#### To create a file filter based on file name - web-based manager

1. Go to *Security Profiles > Data Leak Prevention > File Filter*.
2. Select *Create New*, to create a new file filtering list.
3. Name the list `filterExampleFiles`.

4. Select *Create New* to add a filter to the list.
5. Select *File Name Pattern* for *Filter Type*.
6. Enter `exampleTrojanFile.abc` .
7. Enter *Block* for the *Action*.
8. Select *Enable*, and *OK*.
9. Select *Create New* to add a filter to the list.
10. Select *File Name Pattern* for *Filter Type*.
11. Enter `*trojan*.def` .
12. Enter *Block* for the *Action*.
13. Select *Enable*, and *OK*.

### Filtering based on file type

When filtering files, it is often useful to filter based on the file type. When malware finds a file type that allows them access to a system, the filename will change but the file type will remain the same. Even for preventing applications that are not malware but simply undesirable, filtering based on file type is often the easiest method.

Simply matching the file type, .zip for example, may not be as accurate a method as using the built-in patterns. If users see that .zip attachments are blocked, they may simply rename the file so the filters will allow it through. Checking against patterns can help prevent this bypassing.

There are two possible methods available to filter based on file type. If the file type is one of the built-in patterns, you can use them - for example blocking PalmOS files on your network since Palm devices are not supported. Otherwise, you can simply use wildcards to match the file type.

The following example will filter all batch files (.bat).

#### To filter files based on file type using file name pattern - web-based manager

1. Go to *Security Profiles > Data Leak Prevention > File Filter*.
2. Select *Create New* and name the list `blockedFileTypes`.
3. Select *Create New* to add files to the list.
4. Select *File name pattern* for *Filter Type*.
5. Enter `*.bat` for *Pattern*.
6. Select an *Action* of *Block*.
7. Select *Enable* and *OK*.
8. At the file filter list, select *OK*.

The file filter is now available to be used in an MMS profile, and will block all .bat files that MMS profile matches.

#### To filter files based on file type using file type - web-based manager

1. Go to *Security Profiles > Data Leak Prevention > File Filter*.
2. Select *Create New* and name the list `blockedFileTypes`.
3. Select *Create New* to add files to the list.
4. Select *File Type* for *Filter Type*.
5. Select *Batch File (bat)* for *File Type*.
6. Select an *Action* of *Block*.
7. Select *Enable* and *OK*.
8. At the file filter list, select *OK*.



The file filter is now available to be used in an MMS profile, and will block all batch files (that use .bat file extension) that the MMS profile matches.

## MMS file filtering blocks messages (not just attachments)

When MMS file filtering finds a matching file in an MMS message, the entire message is blocked. This action is more secure, and can reduce the amount of processing required for that message. For example if one MMS message includes three files, and the first one is blocked then the other files won't be scanned or attempted to be matched because the whole message is already being blocked.

## Configuring MMS file filtering

To apply MMS file filtering you must begin with a file filter list. You can create your own list or use the built-in patterns list. You then must create the file filter, then add the file filter list to an MMS profile, and then add the MMS profile to a security policy. The MMS profile, and the corresponding file filter then applies to the traffic accepted by the security policy.

The following procedure creates a file filtering list called `MMS_file_filter` that is used in the MMS profile called `filtering_profile`. The filter will be applied to all MMS message types.

### To apply MMS file filtering

1. Go to *Security Profiles > Data Leak Prevention > File Filter* and create a file filter list called `MMS_file_filter`.
2. Select *Create New* to add entries to filter specific file types.
3. Select *OK*.
4. Go to *Security Profiles > Carrier > MMS Profile* and create a new profile called `filtering_profile`.
5. Expand *MMS Scanning*, and select all the MMS message types.
6. Select `MMS_file_filter` from the *Option* drop down menu.
7. Set other settings as required in the MMS profile.
8. Select *OK*.
9. Go to *Policy*, and select *Create New*.
10. Within the new policy select *Security Profiles > MMS Profile*, and select `filtering_profile`.
11. Configure the security policy as required.
12. Select *OK*.

## Sender notifications and logging

In most cases you will notify the sender that they are causing problems on the network — either by sending malware content, flooding the network, or some other unwanted activity. The notification assumes the sender is unaware of their activity and will stop or correct it when notified.

However, senders who are notified may use this information to circumvent administration's precautions. For example if flood notification is set to 1000 messages per minute, a notified user may simply reduce their message to 990 messages per minute if this flood is intentional. For this reason, not all problems include sender notifications.

There are two methods of notifying senders:

- [MMS notifications](#)
- [Replacement messages](#)

And three details to consider for logging and notifying administrators:

- [Logging and reporting](#)
- [MMS logging options](#)
- [SNMP](#)

## MMS notifications

MMS notifications enable you to customize notifications for many different situations and differently for all the supported MMS message protocols — MM1, MM3, MM4, and MM7.

MMS notification types include:

- Content Filter
- File Block
- Carrier Endpoint Block
- Flood
- Duplicate
- MMS Content Checksum
- Virus Scan

*Day of Week*, *Window start time* and *Window Duration* define what days and what time of day alert notifications will be sent. This allows you to control what alerts are sent on weekends. It also lets you control when to start sending notifications each day. This can be useful if system maintenance is performed at the same time each night — you might want to start alert notifications after maintenance has completed. Another reason to limit the time alert messages are sent could be to limit message traffic to business hours.

**Figure 149:**Notifications screen for FortiOS Carrier MMS Profile

	MM1	MM3	MM4	MM7
Message Protocol	mm1	mm3	mm4	mm7
Message Type				deliver.REQ
Detect Server Details	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Hostname				
URL	/			/
Port	80	25	25	80
Username				
Password				
VASP ID				
VAS ID				
All Notification Types	<input type="checkbox"/> 24 hour(s)	<input type="checkbox"/> 24 hour(s)	<input type="checkbox"/> 24 hour(s)	<input type="checkbox"/> 24 hour(s)
Notifications Per Second Limit	0	0	0	0
Day of Week	Sun Mon Tue Wed Thu Fri Sat <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/>	Sun Mon Tue Wed Thu Fri Sat <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/>	Sun Mon Tue Wed Thu Fri Sat <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/>	Sun Mon Tue Wed Thu Fri Sat <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/>
Window Start Time	00 : 00	00 : 00	00 : 00	00 : 00
Window Duration	24 : 00	24 : 00	24 : 00	24 : 00

For MMS Notification options, see [“MMS Notifications” on page 669](#).

## Replacement messages

FortiGate units send replacement messages when messages or content is blocked, quarantined, or otherwise diverted from the receiver. In its place a message is sent to notify the receiver what happened.

With FortiOS Carrier MMS replacement messages, send and receive message types are supported separately and receive their own custom replacement messages. This allows the network to potentially notify both the sender and receiver of the problem.

For example the replacement message *MM1 send-req file block message* is sent to the device that sent one or more files that were banned. The default message that is sent is `This device has sent %%NUM_MSG%% messages containing banned files in the last %%DURATION%% hours`. The two variables are replaced by the appropriate values.

Replacement messages are not as detailed or specific as MMS notifications, but they are also not as complicated to configure. They are also useful when content has been removed from an MMS message that was still delivered.

For more information on replacement messages, see [“MMS Replacement messages” on page 115](#).

## Logging and reporting

With each virus infection, or file block, a syslog message is generated. The format of this syslog message is similar to:

```
2005-09-22 19:15:47 deviceid=FGT5001ABCDEF1234 logid=0211060ABC
 type=virus subtype=infected level=warning src=10.1.2.3
 dst=10.2.3.4 srcintf=port1 dstintf=port2 service=mm1
 status=blocked from="<sending MSISDN>" to="<receiving MSISDN>"
 file="eicar.com.txt" virus="EICAR_TEST_FILE" msg="The file
 eicar.com.txt is infected with EICAR_TEST_FILE. ref
 http://www.fortinet.com/VirusEncyclopedia/search/encyclopediaSearch.do?method=quickSearchDirectly&virusName=EICAR_TEST_FILE"
```

Note that the *from* and *to* fields are samples and not real values.

## MMS logging options

You can enable logging in an MMS protection profile to write event log messages when the MMS protection profile options that you have enabled perform an action. For example, if you enable MMS antivirus protection, you could also use the MMS protection profile logging options to write an event log message every time a virus is detected.

To record these log messages you must first configure how the FortiOS Carrier unit stores log messages.

To configure MMS content archiving, go to *Security Profiles > Carrier > MMS Profile*. Select *Create New* or select the *Edit* icon beside an existing profile. Expand *MMS Bulk AntiSpam Detection > Logging*. Complete the fields as described in the following table and select *OK*. For more a detailed list of options, see [“Logging” on page 673](#).

## SNMP

A simple SNMP trap will be generated to inform the operators' alerting system that a virus has been detected. This SNMP trap could contain the sending and receiving MSISDN, however the initial solution would reflect the current behavior, i.e. only the fact that a virus has been detected will be communicated.

## MMS content-based Antispam protection

Expand *MMS Scanning* and select *Content Filter* in an MMS protection profile to create content filter black/white lists that block or allow MMS messages based on the content of the message.

### Overview

A school computer lab may block age-inappropriate content. A place of business may block unproductive content. A public access internet cafe may block offensive and graphic content. Each installation has its own requirements for what content needs to be blocked, and in what language.

FortiOS Carrier provides the ability to create custom local dictionaries, black lists, and white lists in multiple languages enables you to protect your customers from malicious content around the world.

Content-based protection includes:

- [Configurable dictionary](#)
- [Black listing](#)
- [White listing](#)

### Configurable dictionary

You can create a dictionary of configurable terms and phrases using the CLI. The text of MMS messages will be searched for these terms and phrases. Add content filter lists that contain content that you want to match in MMS messages. For every match found, a score is added. If enough matches are found to set the total score above the configured threshold, the MMS message is blocked.

You can add words, phrases, wild cards and Perl regular expressions to create content patterns that match content in MMS messages. For more on wildcard and regular expressions, see *Using wildcards and Perl regular expressions* in the [UTM](#) handbook chapter.

For each pattern you can select *Block* or *Exempt*.

- *Block* adds an antispam black list pattern. A match with a block pattern blocks a message depending on the score of the pattern and the content filter threshold.
- *Exempt* adds an antispam white list pattern. A match with an exempt pattern allows the message to proceed through the FortiOS Carrier unit, even if other content patterns in the same content filter list would block it.

If a pattern contains a single word, the FortiOS Carrier unit searches for the word in MMS messages. If the pattern contains a phrase, the FortiOS Carrier unit searches for all of the words in the phrase. If the pattern contains a phrase in quotation marks, the FortiOS Carrier unit searches for the whole phrase.

You can create patterns with Simplified Chinese, Traditional Chinese, Cyrillic, French, Japanese, Korean, Spanish, Thai, or Western character sets.

### Black listing

Black listing is the practice of banning entries on the list. For example if an IP address continuously sends viruses, it may be added to the black list. That means any computers that consult that list will not communicate with that IP address.

Sometimes computers or devices can be added to black lists for a temporary problem, such as a virus that is removed when notified. However, as a rule short of contacting the administrator in person to manually be removed from the black list, users have to wait and they generally will be removed after a period without problem.

## White listing

White listing is the practice of adding all critical IP addresses to a list, such as company email and web servers. Then if those servers become infected and start sending spam or viruses, those servers are not blocked. This allows the critical traffic through, even if there might be some malicious traffic as well. Blocking all traffic from your company servers would halt company productivity.

## Scores and thresholds

Each content pattern includes a score. When a MMS message is matched with a pattern the score is recorded. If a message matches more than one pattern or matches the same pattern more than once, the score for the message increases. When the total score for a message equals or exceeds the threshold the message is blocked.

The default score for a content filter list entry is 10 and the default threshold is 10. This means that by default a message is blocked by a single match. You can change the scores and threshold so that messages can only be blocked if there are multiple matches. For example, you may only want to block messages that contain the phrase “example” if it appears twice. To do this, add the “example” pattern, set action to block and score to 5. Keep the threshold at 10. If “example” is found twice or more in a message the score adds up 10 (or more) and the message is blocked.

## Configuring content-based antispam protection

### To apply content-based antispam protection - CLI

```
config webfilter content
 edit <filter_table_number>
 set name <filter_table_name>
 config entries
 edit <phrase or regexp you want to block>
 set action {block | exempt}
 set lang <phrase language>
 set pattern-type {wildcard | regexp}
 set score <phrase score>
 set status {enable | disable}
 end
 end
 end
```

## Configuring sender notifications

When someone on the MMS network sends an MMS message that is blocked, in most cases you will notify the sender. Typically an administrator is notified in addition to the sender so action can be taken if required.

There are two types of sender notifications available in FortiOS Carrier:

- [MMS notifications](#)
- [Replacement messages](#)

### MMS notifications

MMS notifications to senders are configured in *Security Profiles > Carrier > MMS Profile*, under MMS Notifications.

In this section you can configure up to four different notification recipients for any combination of MM1/3/4/7 protocol MMS messages. Also for MM7 messages the message type can be `submit.REQ` or `deliver.REQ`.

Useful settings include:

- delay in message based on notification type
- limit on notifications per second to prevent a flood
- schedules for notifications
- log in details for MM7 messages.

For more information on MMS notifications, see [“Notifying message flood senders and receivers” on page 725](#) and [“MMS Notifications” on page 669](#).

## Replacement messages

Replacement messages are features common to both FortiOS and FortiOS Carrier, however FortiOS Carrier has additional messages for the MMS traffic.

While each MMS protocol has its own different replacement messages, the one common to all MMS protocols is the *MMS blocked content replacement message*. This is the message that the receiver of the message sees when their content is blocked.

For more information on replacement messages, see [“MMS Replacement messages” on page 115](#).

## MMS DLP archiving

You can use DLP archiving to collect and view historical logs that have been archived to a FortiAnalyzer unit or the FortiGuard Analysis and Management service. DLP archiving is available for FortiAnalyzer when you add a FortiAnalyzer unit to the FortiOS Carrier configuration. The FortiGuard Analysis and Management server becomes available when you subscribe to the FortiGuard Analysis and Management Service.

You can configure full DLP archiving and summary DLP archiving. Full DLP archiving includes all content, for example, full email DLP archiving includes complete email messages and attachments. Summary DLP archiving includes just the meta data about the content, for example, email message summary records include only the email header.

You can archive MM1, MM3, MM4, and MM7 content:

### Configuring MMS DLP archiving

Select DLP archive options to archive MM1, MM3, MM4, and MM7 sessions. For each protocol you can archive just session metadata (*Summary*), or metadata and a copy of the associated file or message (*Full*).

In addition to MMS protection profile DLP archive options you can:

- Archive MM1 and MM7 message floods
- Archive MM1 and MM7 duplicate messages
- Select *DLP archiving* for carrier endpoint patterns in a *Carrier Endpoint List* and select the *Carrier Endpoint Block* option in the *MMS Scanning* section of an MMS Protection Profile

FortiOS Carrier only allows one sixteenth of its memory for transferring content archive files. For example, for Carrier-enabled FortiGate units with 128MB RAM, only 8MB of memory is used when transferring content archive files. Best practices dictate to not enable full content archiving if antivirus scanning is also configured because of these memory constraints.

### To configure MMS DLP archiving - web-based manager

1. Go to *Security Profiles > Carrier > MMS Profile*.
2. Select *Create New* or select the *Edit* icon beside an existing profile.
3. Expand *MMS Bulk AntiSpam Detection > Content Archive*.
4. Complete the fields as described in “[DLP Archive options](#)” on page 672.
5. Select *OK*.

## Viewing DLP archives

You can view DLP archives from the Carrier-enabled FortiGate unit web-based manager. Archives are historical logs that are stored on a log device that supports archiving, such as a FortiAnalyzer unit.

These logs are accessed from either *Log & Report > DLP Archive* or if you subscribed to the FortiCloud service, you can view log archives from there.

The *DLP Archive* menu is only visible if one of the following is true.

- You have configured the FortiGate unit for remote logging and archiving to a FortiAnalyzer unit.
- You have subscribed to FortiCloud.

The following tabs are available when you are viewing DLP archives for one of these protocols.

- *E-mail* to view POP3, IMAP, SMTP, POP3S, IMAPS, SMTPS, and spam email archives.
- *Web* to view HTTP and HTTPS archives.
- *FTP* to view FTP archives.
- *IM* to view AIM, ICQ, MSN, and Yahoo! archives.
- *MMS* to view MMS archives.
- *VoIP* to view session control (SIP, SIMPLE and SCCP) archives.

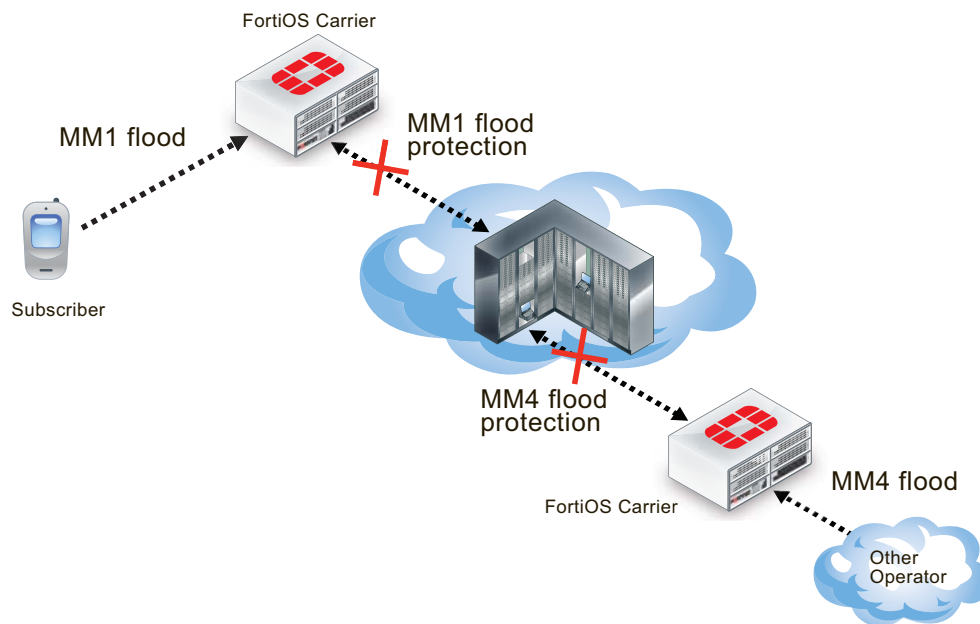
If you need to view log archives in Raw format, select *Raw* beside the *Column Settings* icon.

# Message flood protection

The convenience offered by MM1 and MM4 messaging can be abused by users sending spam or attempting to overload the network with an excess of messages. MMS flood prevention can help prevent this type of abuse.

Flood protection for MM1 messages prevents your subscribers from sending too many messages to your MMSC. Configuring flood protection for MM4 messages prevents another service provider from sending too many messages from the same subscriber to your MMSC.

**Figure 150:**MM1 and MM4 flood protection



The FortiOS Carrier unit keeps track of the number of messages each subscriber sends for the length of time you specify. If the number of messages a subscriber sends exceeds the threshold, a configured action is taken. Possible actions are logging the flood, blocking or intercepting messages in the flood, archiving the flood messages, and sending an alert message to inform the administrator that the flood is occurring.

You can create three different thresholds to take different levels of action at different levels of activity.

With this highly configurable system, you can prevent subscribers from sending more messages than you determine is acceptable, or monitor anyone who exceeds the thresholds.


## Setting message flood thresholds

A message flood occurs when a single subscriber sends a volume of messages that exceeds the flood threshold you set. The threshold defines the maximum number of messages allowed, the period during which the subscriber sent messages are considered, and the length of time the sender is restricted from sending messages after a flood is detected.

If a subscriber exceeds the message flood threshold and is blocked from sending more messages, any further attempts to send messages will re-start the block period. You must also



enable logging for *MMS Scanning > Bulk Messages* in the Logging section of the MMS protection profile.



A subscriber is still able to receive messages while they are blocked from sending messages.

## Example

For example, for the first threshold you may determine that any subscriber who sends more than 100 MM1 messages in an hour (60 minutes) will have all messages blocked for half an hour (30 minutes).

Using this example, if the subscriber exceeds the flood threshold, they are blocked from sending message for 30 minutes. If the subscriber tries to send any message after 15 minutes, the message will be blocked and the block period will be reset again to 30 minutes. The block period must expire with no attempts to send a message. Only then will the subscriber be allowed to send more messages.

### To configure MM1 message flood threshold - web-based manager

1. Go to *Security Profiles > Carrier > MMS Profile*.
2. Select *Create New*.
3. Enter `MM1 flood` for *Profile Name*.
4. Expand *MMS Bulk Email Filtering Detection*.
5. Enter the following information, and select *OK*.

MM1 (first column)	
<b>Enable</b>	Enable
<b>Message Flood Window</b>	60 minutes
<b>Message Flood Limit</b>	100
<b>Message Flood Block Time</b>	30 minutes
<b>Message Flood Action</b>	Block

### To configure MM1 message flood threshold - CLI

```
config firewall mms-profile
 edit profile_name
 config flood mm1
 set status1 enable
 set window1 60
 set limit1 100
 set action1 block
 set block-time1 30
 end
 end
```

The threshold values that you set for your network will depend on factors such as how busy your network is and the kinds of problems that your network and your subscribers encounter. For example, if your network is not too busy you may want to set message flood thresholds

relatively high so that only an exceptional situation will exceed a flood threshold. Then you can use log messages and archived MMS messages to determine what caused the flood.

If your subscribers are experiencing problems with viruses that send excessive amounts of messages, you may want to set thresholds lower and enable blocking to catch problems as quickly as possible and block access to keep the problem from spreading.

## Flood actions

When the Carrier-enabled FortiGate unit detects a message flood, it can take any combination of the five actions that you can configure for the flood threshold. For detailed options, see [“Message Flood” on page 678](#).

## Notifying administrators of floods

You can configure alert notifications for message floods by selecting the Alert Notification message flood action.

The FortiOS Carrier unit sends alert notifications to administrators using the MM1, MM3, MM4, or MM7 content interface. To send an alert notification you must configure addresses and other settings required for the content interface.

For example, to send notifications using the MM1 content interface you must configure a source MSISDN, hostname, URL, and port to which to send the notification. You can also configure schedules for when to send the notifications.

Finally you can add multiple MSISDN numbers to the MMS protection profile and set which flood thresholds to send to each MSISDN.

## Example – three flood threshold levels with different actions for each threshold

You can set up to three threshold levels to take different actions at different levels of activity.

The first example threshold records log messages when a subscriber’s handset displays erratic behavior by sending multiple messages using MM1 at a relatively low threshold. The erratic behavior could indicate a problem with the subscriber’s handset. For example, you may have determined for your network that if a subscriber sends more the 45 messages in 30 minutes that you want to record log messages as a possible indication or erratic behavior.

From the web-based manager in an MMS profile set message *Flood Threshold 1* to:

<b>Enable</b>	Selected
<b>Message Flood Window</b>	30 minutes
<b>Message Flood Limit</b>	45
<b>Message Flood Action</b>	Log

From the CLI:

```
config firewall mms-profile
 edit profile_name
 config flood mm1
```

```

 set status1 enable
 set window1 30
 set limit1 45
 set action1 log
 end
end

```

Set a second higher threshold to take additional actions when a subscriber sends more than 100 messages in 30 minutes. Set the actions for this threshold to log the flood, archive the message that triggered the second threshold, and block the sender for 15 minutes.

From the web-based manager in an MMS profile set message *Flood Threshold 2* to:

<b>Enable</b>	Selected
<b>Message Flood Window</b>	30 minutes
<b>Message Flood Limit</b>	100
<b>Message Block Time</b>	15 minutes
<b>Message Flood Action</b>	Log, DLP archive First message only, Block

From the CLI:

```

config firewall mms-profile
 edit profile_name
 config flood mml
 set status2 enable
 set window2 30
 set limit2 100
 set action2 block log archive-first
 set block-time2 15
 end
 end
end

```

Set the third and highest threshold to block the subscriber for an extended period and send an administrator alert if the subscriber sends more than 200 messages in 30 minutes. Set the actions for this threshold to block the sender for four hours (240 minutes), log the flood, archive the message that triggered the third threshold, and send an alert to the administrator.

From the web-based manager in an MMS profile set message *Flood Threshold 3* to:

<b>Enable</b>	Selected
<b>Message Flood Window</b>	30 minutes
<b>Message Flood Limit</b>	200
<b>Message Block Time</b>	240 minutes
<b>Message Flood Action</b>	Log, Block, Alert Notification

Because you have selected the *Alert Notification* action you must also configure alert notification settings. For this example, the source MSISDN is 5551234—telephone number 555-1234. When administrators receive MMS messages from this MSISDN they can assume a message flood has been detected.

In this example, alert notifications are sent by the FortiOS Carrier unit to the MMSC using MM1. The host name of the MMSC is `mmscexample`, the MMSC URL is `/`, and the port used by the MMSC is 80. In this example, the alert notification window starts at 8:00am and extends for eight hours on weekdays (Monday-Friday) and the minimum interval between message flood notifications is two hours.

<b>Source MSISDN</b>	5551234
<b>Message Protocol</b>	MM1
<b>Hostname</b>	mmscexample
<b>URL</b>	/
<b>Port</b>	80
<b>Notifications Per Second Limit</b>	0
<b>Window Start Time</b>	8:00
<b>Window Duration</b>	8:00
<b>Day of Week</b>	Mon, Tue, Wed, Thu, Fri, Sat
<b>Interval</b>	2 hours

From the CLI:

```

config firewall mms-profile
 edit profile_name
 config notification alert-flood-1
 set alert-src-msisdn 5551234
 set set msg-protocol mm1
 set mmsc-hostname mmscexample
 set mmsc-url /
 set mmsc-port 80
 set rate-limit 0
 set tod-window-start 8:00
 set tod-window-duration 8:00
 set days-allowed monday tuesday wednesday thursday friday
 set alert-int 2
 set alert-int-mode hours
 end
 end

```

You must also add the MSISDNs of the administrators to be notified of the message flood. In this example, the administrator flood threshold 3 alert notifications are sent to one administrator with MSISDN 5554321.

To add administrator's MSISDNs for flood threshold 3 from the web-based manager when configuring a protection profile, select *MMS Bulk Email Filtering Detection > Recipient MSISDN > Create New*.

<b>MSISDN</b>	5554321
<b>Flood Level 3</b>	Select

From the CLI:

```
config firewall mms-profile
 edit profile_name
 config notif-msisdn
 edit 5554321
 set threshold flood-thresh-3
 end
 end
 end
```

## Notifying message flood senders and receivers

The FortiOS Carrier unit does not send notifications to the sender or receiver that cause a message flood. If the sender or receiver is an attacker and is explicitly informed that they have exceeded a message threshold, the attacker may try to determine the exact threshold value by trial and error and then find a way around flood protection. For this reason, no notification is set to the sender or receiver.

However, FortiOS Carrier does have replacement messages for sending reply confirmations to MM1 senders and receivers and for MM4 senders for blocked messages identified as message floods. For information about how FortiOS Carrier responds when message flood detection blocks a message, see [“FortiOS Carrier and MMS duplicate messages and message floods” on page 642](#).

### Responses to MM1 senders and receivers

When the FortiOS Carrier unit identifies an MM1 message sent by a sender to an MMSC as a flood message and blocks it, the FortiOS Carrier unit returns a message submission confirmation (`m-send.conf`) to the sender — otherwise the sender’s handset would keep retrying the message. The `m-send.conf` message is sent only when the MM1 message flood action is set to Block. For other message flood actions the message is actually delivered to the MMSC and the MMSC sends the `m-send.conf` message.

You can customize the `m-send.conf` message by editing the *MM1 send-conf flood message* MM1 replacement message (from the CLI the `mm1-send-conf-flood` replacement message). You can customize the response status and message text for this message. The default response status is “Content not accepted”. To hide the fact that FortiOS Carrier is responding to a flood, you can change the response status to “Success”. The default message text informs the sender that the message was blocked. You could change this to something more generic.

For example, the following command sets the submission confirmation response status to “Success” and changes the message text to “Message Sent OK”:

```
config system replacemsg mm1 mm1-send-conf-flood
 set rsp-status ok
 set rsp-text "Message Sent OK"
end
```

When the FortiOS Carrier unit identifies an MM1 message received by a receiver from an MMSC as a flood message and blocks it, the FortiOS Carrier unit returns a message retrieval confirmation (`m-retrieve.conf`) to the sender (otherwise the sender’s handset would keep retrying the message). The `m-retrieve.conf` message is sent only when the MM1 message flood action is set to Block. For other message flood actions the message is actually delivered to the receiver, so the MMSC sends the `m-retrieve.conf` message.

You can customize the `m-retrieve.conf` message by editing the *MM1 retrieve-conf flood message* MM1 replacement message (from the CLI the `mm1-retr-conf-flood` replacement message). You can customize the class, subject, and message text for this message.

For example, you could use the following command make the response more generic:

```
config system replacemsg mml mml-retr-conf-flood
 set subject "Message blocked"
 set message "Message temporarily blocked by carrier"
end
```

## Forward responses for MM4 message floods

When the FortiOS Carrier unit identifies an MM4 message as a flood message and blocks it, the FortiOS Carrier unit returns a message forward response (MM4\_forward.res) to the forwarding MMSC (otherwise the forwarding MMSC would keep retrying the message). The MM4\_forward.res message is sent only when the MM4 message flood action is set to Block and the MM4-forward.req message requested a response. For more information, see [“FortiOS Carrier and MMS duplicate messages and message floods” on page 642](#).

You can customize the MM4\_forward.res message by editing the *MM4 flood message* MM4 replacement message (from the CLI the `mm4-flood` replacement message). You can customize the response status and message text for this message. The default response status is “Content not accepted” (`err-content-not-accept`). To hide the fact that the FortiOS Carrier unit is responding to a flood, you can change the response status to “Success”. The default message text informs the sender that the message was blocked. You could change this to something more generic.

For example, the following command sets the submission confirmation response status to “Success” and changes the message text to “Message Sent OK” for the MM4 message forward response

```
config system replacemsg mm4 mm4-flood
 set rsp-status ok
 set rsp-text "Message Forwarded OK"
end
```

## Viewing DLP archived messages

If *DLP Archive* is a selected message flood action, the messages that exceed the threshold are saved to the MMS DLP archive. The default behavior is to save all of the offending messages, but you can configure the DLP archive setting to save only the first message that exceeds the threshold. This still provides a sample of the offending messages without requiring as much storage.

### To select only the first message in a flood for DLP archiving - web-based manager

1. Go to *Security Profiles > Carrier > MMS Profile*.
2. Edit an existing MMS Profile.
3. Expand the *MMS Bulk Email Filtering Detection* section, the *Message Flood* subsection, and the desired *Flood Threshold* subsection.
4. Next to *DLP Archive*, select *First message only* from the dropdown menu.
5. Select *OK*.

## Order of operations: flood checking before duplicate checking

Although duplicate checking involves only examination and comparison of message contents and not the sender or recipient, and flood checking involves only totalling the number of

messages sent by each subscriber regardless of the message content, there are times when a selection of messages exceed both flood and duplicate thresholds.

The Carrier-enabled FortiGate unit checks for message floods before checking for duplicate messages. Flood checking is less resource-intensive and if the flood threshold invokes a *Block* action, the blocked messages are stopped before duplicate checking occurs. This saves both time and FortiOS Carrier system resources.



The duplicate scanner will only scan content. It will not scan headers. Content must be exactly the same. If there is any difference at all in the content, it will not be considered a duplicate.

## Bypassing message flood protection based on user's carrier endpoints

You can use carrier endpoint filtering to exempt MMS sessions from message flood protection. Carrier endpoint filtering matches carrier endpoints in MMS sessions with carrier endpoint patterns.

If you add a carrier endpoint pattern to a filter list and set the action to exempt from mass MMS, all messages from matching carrier endpoints bypass message flood protection. This allows legitimate bulk messages, such as system outage notifications, to be delivered without triggering message flood protection.

For more information on carrier endpoints, see the [FortiOS Handbook User Authentication](#) chapter.

## Configuring message flood detection

To have the Carrier-enabled FortiGate unit check for message floods, you must first configure the flood threshold in an MMS profile, select the MMS profile in a security policy. All the traffic examined by the security policy will be checked for message floods according to the threshold values you set in the MMS profile.

### Configure the MMS profile - web-based manager

1. Go to *Firewall Objects > MMS Profile*.
2. If you are editing an MMS profile, select the *Edit* icon of the MMS profile.  
If you are create a new MMS profile, select *Create New* and enter a profile name.
3. Expand *MMS Bulk Email Filtering Detection*.
4. Expand *Message Flood*.
5. Expand *Flood Threshold 1*.
6. Select the *Enable* check box for MM1 messages, MM4 messages, or both.
7. In the *Message Flood Window* field, enter the length of time the Carrier-enabled FortiGate unit will keep track of the number of messages each subscriber sends.  
If the Carrier-enabled FortiGate unit detects the quantity of messages specified in the *Message Flood Limit* sent during the number of minutes specified in the *Message Flood Window*, a message flood is in progress.
8. In the *Message Flood Limit* field, enter the number of messages required to trigger the flood.
9. In the *Message Flood Block Time* field, enter the length of time a user will be blocked from sending messages after causing the message flood.
10. Select the message flood actions the Carrier-enabled FortiGate unit will take when the message flood is detected.

11. Select OK.

#### **Configure the security policy - web-based manager**

1. Go to *Policy*.
2. Select the *Edit* icon of the security policy that controls the traffic in which you want to detect message floods.
3. Select the *MMS Profile* check box to enable the use of a protection profile.
4. Select the MMS protection profile from the list.
5. Select OK.

## **Sending administrator alert notifications**

When message floods are detected, the Carrier-enabled FortiGate unit can be configured to notify you immediately with an MMS message. Enable this feature by selecting Alert Notification in the message flood action. Each message flood threshold can be configured separately.

This section includes:

- [Configuring how and when to send alert notifications](#)
- [Configuring who to send alert notifications to](#)

### **Configuring how and when to send alert notifications**

You can configure different alert notifications for MM1 and MM4 message floods. You can configure the FortiOS Carrier unit to send these alert notifications using the MM1, MM3, MM4, or MM7 content interface. Each of these content interfaces requires alert notification settings that the FortiOS Carrier unit uses to communicate with a server using the selected content interface.

For the MM1 content interface you require:

- The hostname of the server
- The URL of the server (usually “/”)
- The server port (usually 80)

For the MM3 and MM4 content interfaces you require:

- The hostname of the server
- The server port (usually 80)
- The server user domain



For the MM7 content interface you require:

- The message type
- *submit.REQ* to send a notification message to the sender in the form of a submit request. The message goes from a VAS application to the MMSC.
- *deliver.REQ* to send a notification message to the sender in the form of a deliver request. The message goes from the MMSC to a VAS application.
- The hostname of the server
- The URL of the server (usually “/”)
- The server port (usually 80)
- A user name and password to connect to the server
- The value-added-service-provider (VASP) ID
- The value-added-service (VAS) ID

For more information, see [“MMS notifications” on page 714](#).

### **To configure administrator alert notifications - web-based manager**

1. Go to *Firewall Objects > MMS Profile* and edit or add a new MMS protection profile.
2. Expand *MMS Bulk Email Filtering Detection*.  
There are three message flood thresholds.
3. Expand the threshold that you want to configure alert notification for.
4. For *Message Flood Action*, select the *Alert Notification* check box. Alert notification options appear.
5. For the *Source MSISDN*, enter the MSISDN from which the alert notification message will be sent.
6. Select the Message Protocol the alert notification will use: *MM1*, *MM3*, *MM4*, or *MM7*.
7. Add the information required by FortiOS Carrier to send messages using the selected message protocol:
8. For *Notifications Per Second Limit*, enter the number of notifications to send per second.  
Use this setting to reduce control the number of notifications sent by the FortiOS Carrier unit. If you enter zero (0), the notification rate is not limited.
9. If required, change *Window Start Time* and *Window Duration* configure when the FortiOS Carrier unit sends alert notifications.  
By default, notifications are sent at any time of the day. You can change the Window Start Time if you want to delay sending alert messages. You can also reduce the Window Duration if you want to stop sending alert notifications earlier.  
For example, you might not want FortiOS Carrier sending notifications except during business hours. In this case the Window Start Time could be 9:00 and the Window Duration could be 8:00 hours.  
You can set different alert notifications for each message threshold. For example, you could limit the message window for lower thresholds and set it to 24 hours for higher thresholds. This way administrators will only receive alert notifications outside of business hours for higher thresholds.
10. For *Day of Week*, select the days of the week to send notifications.  
For example, you may only want to send alert notifications on weekends for higher thresholds.

11. In the *Interval field*, enter the maximum frequency that alert notification messages will be sent, in minutes or hours.

All alerts occurring during the interval will be included in a single alert notification message to reduce the number of alert messages that are sent.

## Configuring who to send alert notifications to

In each MMS protection profile you add a list of recipient MSISDNs. For each of these MSISDNs you select the message flood threshold that triggers sending notifications to this MSISDN.

### To configure the alert notification recipients - web-based manager

1. Go to *Firewall Objects > MMS Profile*.
2. Select the *Edit* icon of the MMS profile in which you want to configure the alert notification recipients.
3. Expand *MMS Bulk Email Filtering Detection*.
4. Expand *Recipient MSISDN*.
5. Select *Create New*.
6. In the *New MSISDN* window, enter the MSISDN to use for flood threshold alert notification.
7. Select the duplicate thresholds at which to send alert notifications to the MSISDN.



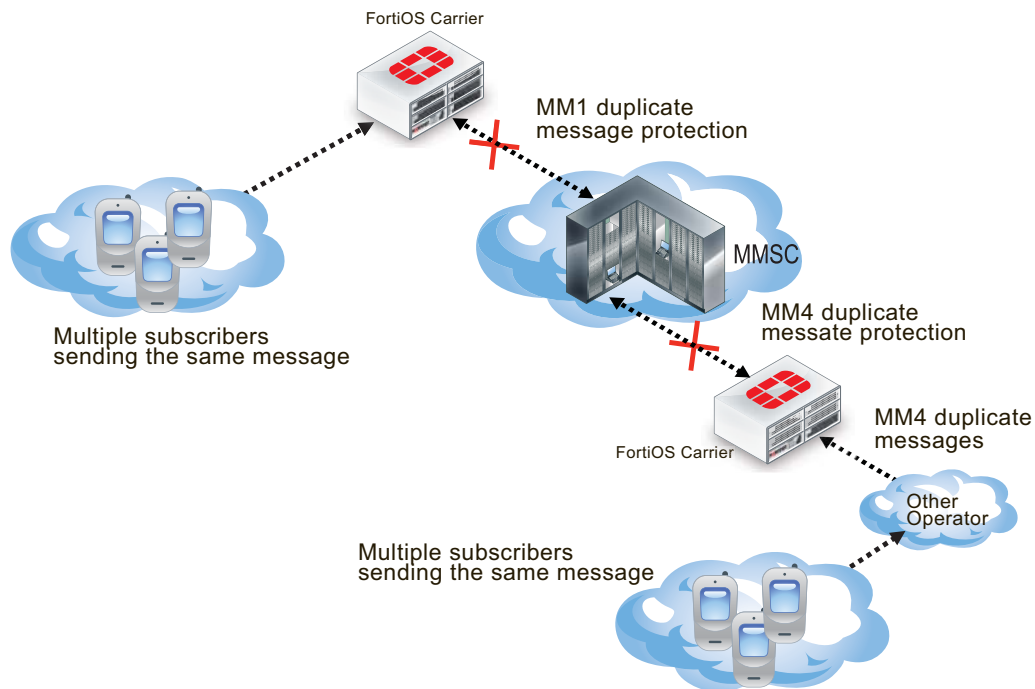
For the flood threshold to be able to send an alert notification to the MSISDN, the alert notification action must be enabled and configured within the flood threshold.

# Duplicate message protection

The convenience offered by MM1 and MM4 messaging can be abused by users sending spam or other unwanted messages. Often, the same message will be sent by multiple subscribers. The message can be spam, viral marketing, or worm-generated messages. MMS duplicate prevention can help prevent this type of abuse by keeping track of the messages being sent.

Duplicate message protection for MM1 messages prevents multiple subscribers from sending duplicate messages to your MMSC. Duplicate message protection for MM4 messages prevents another service provider from sending duplicate messages from the same subscriber to your MMSC. This can help prevent a potential flood that would otherwise become widespread between carriers.

**Figure 151:**MM1 and MM4 duplicate message protection



The FortiOS Carrier unit keeps track of the sent messages. If the same message appears more often than the threshold value you configure, then action is taken. Possible actions are logging the duplicates, blocking or intercepting duplicate messages, archiving the duplicate messages, and sending an alert to inform an administrator that duplicates are occurring.

With this highly configurable system, you can prevent the transmission of duplicate messages when there are more than you determine is acceptable.

For detailed configuration options, see [“Duplicate Message” on page 680](#).

## Using message fingerprints to identify duplicate messages

The Carrier-enabled FortiGate unit detects duplicates by keeping a record of all the messages travelling on the network and comparing new messages to those that have already been sent.

Rather than save the messages, the FortiOS carrier creates a checksum using the message body and subject. This serves as a fingerprint to identify the message. If another message with the same message body and subject appears, the fingerprint will also be the same and the Carrier-enabled FortiGate unit will recognize it as a duplicate.

By creating and saving message fingerprints instead of saving the messages, the Carrier-enabled FortiGate unit can save resources and time.

## Messages from any sender to any recipient

Duplicate message detection will detect duplicate messages regardless of the sender or recipient. To do this, message fingerprints are generated using only the message body and subject. The sender, recipient, and other header information is not included.

If multiple messages appear with the same subject and message body, the Carrier-enabled FortiGate unit will recognize them as being the same.

## Setting duplicate message thresholds

The FortiOS Carrier recognizes all duplicate messages, but it will take action when it detects a volume of duplicate messages that exceed the duplicate threshold you set. The threshold defines the maximum number of duplicate messages allowed, the period during which the messages are considered, and the length of time the duplicate message can not be sent by anyone.

For example, you may determine that once a duplicate message is sent more than 300 times in an hour, any attempt to send the same duplicate message will be blocked for 30 minutes.

If a particular duplicate message exceeds the duplicate message threshold and is blocked, any further attempts to send the same message will re-start the block period.

Using the example above, if the duplicate message count exceeds the duplicate threshold, any attempt to send a copy of the duplicate message will be blocked for 30 minutes. If a subscriber tries to send a copy of the message after waiting 15 minutes, the message will be blocked and the block period will be reset to 30 minutes. The block period must expire with no attempts to send a duplicate message. Only then will a subscriber be allowed to send the message.

Non-duplicate messages will not reset the block period.

## Duplicate message actions

When the Carrier-enabled FortiGate unit detects that a duplicate message has exceeded duplicate threshold, it can take any combination of the five actions you configure for the duplicate threshold.

Action	Description
Log	Add a log entry indicating that a duplicate message event has occurred. You must also enable logging for <i>MMS Scanning &gt; Bulk Messages</i> in the <i>Logging</i> section of the MMS protection profile.
DLP Archive	

<b>All messages</b>	Save all the messages that exceed the duplicate threshold in the DLP archive.
<b>First message only</b>	Save the first message to exceed the duplicate threshold in the DLP archive. Subsequent messages that exceed the duplicate threshold will not be saved.
<b>Intercept</b>	Messages that exceed the duplicate threshold are passed to the recipients, but if quarantine is enabled for intercepted messages, a copy of each message is also quarantined for later examination. If the quarantine of intercepted messages is disabled, the <i>Intercept</i> action has no effect.
<b>Block</b>	Messages that exceed the duplicate threshold are blocked and will not be delivered to the message recipients. If quarantine is enabled for blocked messages, a copy of each blocked message is quarantined for later examination.
<b>Alert Notification</b>	If the duplicate threshold is exceeded, the Carrier-enabled FortiGate unit will send an MMS duplicate message notification message.

## Notifying duplicate message senders and receivers

The FortiOS Carrier unit does not send notifications to the sender or receiver of duplicate messages. If the sender or receiver is an attacker and is explicitly informed that they have exceeded a message threshold, the attacker may try to determine the exact threshold value by trial and error and then find a way around duplicate message protection. For this reason, no notification is set to the sender or receiver.

However, the FortiOS Carrier unit does have replacement messages for sending reply confirmations to MM1 senders and receivers and for MM4 senders for blocked messages identified as duplicate messages. For information about how FortiOS Carrier responds when message flood detection blocks a message, see [“FortiOS Carrier and MMS duplicate messages and message floods” on page 642](#).

## Responses to MM1 senders and receivers

When the FortiOS Carrier unit identifies an MM1 message sent by a sender to an MMSC as a duplicate message and blocks it, the FortiOS Carrier unit returns a message submission confirmation (m-send.conf) to the sender (otherwise the sender’s handset would keep retrying the message). The m-send.conf message is sent only when the MM1 duplicate message action is set to Block. For other duplicate message actions the message is actually delivered to the MMSC and the MMSC sends the m-send.conf message.

You can customize the m-send.conf message by editing the *MM1 send-conf duplicate message* MM1 replacement message (from the CLI the `mm1-send-conf-dupe` replacement message). You can customize the response status and message text for this message. The default response status is “Content not accepted”. To hide the fact that the FortiOS Carrier unit is responding to a duplicate message, you can change the response status to “Success”. The default message text informs the sender that the message was blocked. You could change this to something more generic.

For example, the following command sets the submission confirmation response status to “Success” and changes the message text to “Message Sent OK”:

```
config system replacemsg mm1 mm1-send-conf-dupe
```

```
set rsp-status ok
set rsp-text "Message Sent OK"
end
```

When the FortiOS Carrier unit identifies an MM1 message received by a receiver from an MMSC as a duplicate message and blocks it, the FortiOS Carrier unit returns a message retrieval confirmation (m-retrieve.conf) to the sender (otherwise the sender's handset would keep retrying). The m-retrieve.conf message is sent only when the `MM1duplicate` message action is set to Block. For other message flood actions the message is actually received by the receiver, so the MMSC sends the m-retrieve.conf message.

You can customize the m-retrieve.conf message by editing the *MM1 retrieve-conf duplicate message* MM1 replacement message (from the CLI the `mm1-retr-conf-dupe` replacement message). You can customize the class, subject, and message text for this message.

For example, you could use the following command make the response more generic:

```
config system replacemsg mm1 mm1-retr-conf-dupe
 set subject "Message blocked"
 set message "Message temporarily blocked by carrier"
end
```

## Forward responses for duplicate MM4 messages

When the FortiOS Carrier unit identifies an MM4 message as a duplicate message and blocks it, the FortiOS Carrier unit returns a message forward response (MM4\_forward.res) to the forwarding MMSC (otherwise the forwarding MMSC would keep retrying the message). The MM4\_forward.res message is sent only when the MM4 duplicate message action is set to Block and the MM4-forward.req message requested a response. For more information, see [“FortiOS Carrier and MMS duplicate messages and message floods” on page 642](#).

You can customize the MM4\_forward.res message by editing the *MM4 duplicate message* MM4 replacement message (from the CLI the `mm4-dupe` replacement message). You can customize the response status and message text for this message. The default response status is “Content not accepted” (`err-content-not-accept`). To hide the fact that the FortiOS Carrier unit is responding to a duplicate message, you can change the response status to “Success”. The default message text informs the sender that the message was blocked. You could change this to something more generic.

For example, the following command sets the submission confirmation response status to “Success” and changes the message text to “Message Forwarded OK”:

```
config system replacemsg mm4 mm4-dupe
 set rsp-status ok
 set rsp-text "Message Forwarded OK"
end
```

## Viewing DLP archived messages

If *DLP Archive* is a selected duplicate message action, the messages that exceed the threshold are saved to the MMS DLP archive. The default behavior is to save all of the offending messages but you can configure the DLP archive setting to save only the first message that exceeds the threshold. See [“Viewing DLP archived messages” on page 726](#).

## Order of operations: flood checking before duplicate checking

Although duplicate checking involves only examination and comparison of message contents and not the sender or recipient, and flood checking involves only totalling the number of messages sent by each subscriber regardless of the message content, there are times when a selection of messages exceed both flood and duplicate thresholds.

The Carrier-enabled FortiGate unit checks for message floods before checking for duplicate messages. Flood checking is less resource-intensive and if the flood threshold invokes a *Block* action, the blocked messages are stopped before duplicate checking occurs. This saves both time and FortiOS Carrier system resources.

## Bypassing duplicate message detection based on user's carrier endpoints

You can use carrier endpoint filtering to exempt MMS sessions from duplicate message detection. Carrier endpoint filtering matches carrier endpoints in MMS sessions with carrier endpoint patterns. If you add a carrier endpoint pattern to a filter list and set the action to exempt from mass MMS, all messages from matching carrier endpoints bypass duplicate message detection. For more information about endpoints, see [FortiOS Handbook User Authentication chapter](#).

## Configuring duplicate message detection

To have the Carrier-enabled FortiGate unit check for duplicate messages, configure the duplicate threshold in an MMS profile, and select the MMS profile in a security policy.

All traffic matching the security policy will be checked for duplicate messages according to the settings in the MMS profile.



The duplicate scanner will only scan content. It will not scan headers. Content must be exactly the same. If there is any difference at all in the content, it will not be considered a duplicate.

The modular nature of the profiles allows you great flexibility in how you configure the scanning options. MMS profiles can be used in any number of policies, with different GTP profiles.

In a complex configuration, there may be many security policies, each with a different MMS profile. For a simpler network, you may have many security policies all using the same MMS profile.

## Sending administrator alert notifications

When duplicate messages are detected, the Carrier-enabled FortiGate unit can be configured to notify you immediately with an MMS message. Enable this feature by selecting Alert Notification in the duplicate message action. Each duplicate message threshold can be configured separately.

This section includes:

- [Configuring how and when to send alert notifications](#)
- [Configuring who to send alert notifications to](#)

## Configuring how and when to send alert notifications

You can configure different alert notifications for MM1 and MM4 duplicate messages. You can configure the FortiOS Carrier unit to send these alert notifications using the MM1, MM3, MM4, or MM7 content interface. Each of these content interfaces requires alert notification settings that the FortiOS Carrier unit uses to communicate with a server using the selected content interface.

For the MM1 content interface you require:

- The hostname of the server
- The URL of the server (usually “/”)
- The server port (usually 80)

For the MM3 and MM4 content interfaces you require:

- The hostname of the server
- The server port (usually 80)
- The server user domain

For the MM7 content interface you require:

- The message type
- *submit.REQ* to send a notification message to the sender in the form of a submit request. The message goes from a VAS application to the MMSC.
- *deliver.REQ* to send a notification message to the sender in the form of a deliver request. The message goes from the MMSC to a VAS application.
- The hostname of the server
- The URL of the server (usually “/”)
- The server port (usually 80)
- A user name and password to connect to the server
- The value-added-service-provider (VASP) ID
- The value-added-service (VAS) ID

### To configure administrator alert notifications - web-based manager

1. Go to *Security Profiles > Carrier > MMS Profile* and edit or add a new MMS protection profile.
2. Expand *MMS Bulk Email Filtering Detection*.  
There are three duplicate message thresholds.
3. Expand the threshold that you want to configure alert notification for.
4. For *Duplicate Message Action*, select the *Alert Notification* check box. Alert notification options appear.
5. For the *Source MSISDN*, enter the MSISDN from which the alert notification message will be sent.
6. Select the Message Protocol the alert notification will use: *MM1*, *MM3*, *MM4*, or *MM7*.
7. Add the information required by FortiOS Carrier to send messages using the selected message protocol:
8. For *Notifications Per Second Limit*, enter the number of notifications to send per second.  
Use this setting to reduce control the number of notifications sent by the FortiOS Carrier unit. If you enter zero (0), the notification rate is not limited.



9. If required, change *Window Start Time* and *Window Duration* configure when the FortiOS Carrier unit sends alert notifications.

By default, notifications are sent at any time of the day. You can change the Window Start Time if you want to delay sending alert messages. You can also reduce the Window Duration if you want to stop sending alert notifications earlier.

For example, you might not want FortiOS Carrier sending notifications except during business hours. In this case the Window Start Time could be 9:00 and the Window Duration could be 8:00 hours.

You can set different alert notifications for each message threshold. For example, you could limit the message window for lower thresholds and set it to 24 hours for higher thresholds. This way administrators will only receive alert notifications outside of business hours for higher thresholds.

10. For *Day of Week*, select the days of the week to send notifications.

For example, you may only want to send alert notifications on weekends for higher thresholds.

11. In the *Interval field*, enter the maximum frequency that alert notification messages will be sent, in minutes or hours.

All alerts occurring during the interval will be included in a single alert notification message to reduce the number of alert messages that are sent.

## Configuring who to send alert notifications to

In each MMS protection profile you add a list of recipient MSISDNs. For each of these MSISDNs you select the duplicate threshold that triggers sending notifications to this MSISDN.

### To configure the alert notification recipients - web-based manager

1. Go to *Security Profiles > Carrier > MMS Profile*.
2. Select the *Edit* icon of the MMS profile in which you want to configure the alert notification recipients.
3. Expand *MMS Bulk Email Filtering Detection*.
4. Expand *Recipient MSISDN*.
5. Select *Create New*.
6. In the *New MSISDN* window, enter the MSISDN to use for duplicate threshold alert notification.

Select the duplicate thresholds at which to send alert notifications to the MSISDN.



For the duplicate threshold to be able to send an alert notification to the MSISDN, the duplicate message threshold alert notification action must be enabled and configured.

# Configuring GTP on FortiOS Carrier

Configuring GTP support on FortiOS Carrier involves configuring a number of areas of features. Some features require longer explanations, and have their own chapters. The other features are addressed here.

This section includes:

- [GTP support on the Carrier-enabled FortiGate unit](#)
- [Configuring General Settings on the Carrier-enabled FortiGate unit](#)
- [Configuring Encapsulated Filtering in FortiOS Carrier](#)
- [Configuring the Protocol Anomaly feature in FortiOS Carrier](#)
- [Configuring Anti-overbilling in FortiOS Carrier](#)
- [Logging events on the Carrier-enabled FortiGate unit](#)

## GTP support on the Carrier-enabled FortiGate unit

The FortiCarrier unit needs to have access to all traffic entering and exiting the carrier network for scanning, filtering, and logging purposes. This promotes one of two configurations — hub and spoke, or bookend.

A hub and spoke configuration with the Carrier-enabled FortiGate unit at the hub and the other GPRS devices on the spokes is possible for smaller networks where a lower bandwidth allows you to divide one unit into multiple virtual domains to fill multiple roles on the carrier network. It can be difficult with a single FortiOS Carrier as the hub to ensure all possible entry points to the carrier network are properly protected from potential attacks such as [“Relayed network attacks” on page 741](#).

A bookend configuration uses two Carrier-enabled FortiGate units to protect the carrier network between them with high bandwidth traffic. One unit handles traffic from mobile stations, SGSNs, and foreign carriers. The other handles GGSN and data network traffic. Together they ensure the network is secure.

The Carrier-enabled FortiGate unit can access all traffic on the network. It can also verify traffic between devices, and verify that the proper GPRS interface is being used. For example there is no reason for a Gn interface to be used to communicate with a mobile station — the mobile station will not know what to do with the data — so that traffic is blocked.



When you are configuring your Carrier-enabled FortiGate unit's GTP profile, you must first configure the APN. It is critical to GTP communications — no traffic will flow without the APN.

The Carrier-enabled FortiGate unit does more than just forward and route GTP packets over the network. It also performs:

- [Packet sanity checking](#)
- [GTP stateful inspection](#)
- [Protocol anomaly detection and prevention](#)
- [HA](#)
- [Virtual domain support](#)

## Packet sanity checking

The FortiOS Carrier firewall checks the following items to determine if a packet conforms to the UDP and GTP standards:

- GTP release version number — must be 0, 1, or 2
- Settings of predefined bits
- Protocol type
- UDP packet length

If the packet in question does not confirm to the standards, the FortiOS Carrier firewall drops the packet, so that the malformed or forged traffic will not be processed.

## GTP stateful inspection

Apart from the static inspection (checking the packet header), the FortiOS Carrier firewall performs stateful inspection.

Stateful inspection provides enhanced security by keeping track of communications sessions and packets over a period of time. Both incoming and outgoing packets are examined. Outgoing packets that request specific types of incoming packets are tracked; only those incoming packets constituting a proper response are allowed through the firewall.

The FortiOS Carrier firewall can also index the GTP tunnels to keep track of them.

Using the enhanced Carrier traffic policy, the FortiOS Carrier firewall can block unwanted encapsulated traffic in GTP tunnels, such as infrastructure attacks. Infrastructure attacks involve attempts by an attacker to connect to restricted machines, such as GSN devices, network management systems, or mobile stations. If these attempts to connect are detected, they are to be flagged immediately by the firewall .

## Protocol anomaly detection and prevention

The FortiOS Carrier firewall detects and optionally drops protocol anomalies according to GTP standards and specific tunnel states. Protocol anomaly attacks involve malformed or corrupt packets that typically fall outside of protocol specifications. These packets are not seen on a production network. Protocol anomaly attacks exploit poor programming practices when decoding packets, and are typically used to maliciously impair system performance or elevate privileges.

FortiOS Carrier also detects IP address spoofing inside GTP data channel.

See [“Configuring the Protocol Anomaly feature in FortiOS Carrier” on page 742.](#)

## HA

FortiOS Carrier active-passive HA provides failover protection for the GTP tunnels. This means that an active-passive cluster can provide FortiOS Carrier firewall services even when one of the cluster units encounters a problem that would result in complete loss of connectivity for a stand-alone FortiOS Carrier firewall. This failover protection provides a backup mechanism that can be used to reduce the risk of unexpected downtime, especially for mission-critical environments.

FortiOS HA synchs TCP sessions by default, but UDP sessions are not synchronized by default. However synchronizing a session is only part of the solution if the goal is to continue GTP processing on a synchronized session after a HA switch. For that to be successful we also need to synch the GTP tunnel state. So, once the master completes tunnel setup then the GTP tunnel is synchronized to the slave.

GTP traffic will only flow without interruption on a HA switch if bidirectional GTP policies have been configured: an internal (GTP server) to external (all) UDP port GTP policy, and an external (all) to internal (GTP server) UDP port GTP policy. If either policy is missing then traffic may be interrupted until traffic flows in the opposite direction.

For more information on HA in FortiOS, see [“High Availability for FortiOS 5.0” on page 1116](#).

## Virtual domain support

FortiOS Carrier is suited to both large and smaller carriers. A single Carrier-enabled FortiGate unit can serve either one large carrier, or several smaller ones through virtual domains. As with any FortiGate unit, Carrier-enabled units have the ability to split their resources into multiple virtual units. This allows smaller carriers to use just the resources that they need without wasting the extra. For more information on virtual domains in FortiOS, see [“Virtual Domains” on page 2332](#).

## Configuring General Settings on the Carrier-enabled FortiGate unit

To configure the GTP General Settings, go to *Security Profiles > Carrier > GTP Profile*, and edit a GTP profile. Expand *General Settings* to configure settings. See [“General settings options” on page 685](#).

## Configuring Encapsulated Filtering in FortiOS Carrier

Encapsulated traffic on the GPRS network can come in a number of forms as it includes traffic that is “wrapped up” in another protocol. This detail is important for firewalls because it requires “unwrapping” to properly scan the data inside. If encapsulated packets are treated as regular packets, that inside layer will never be scanned and may allow malicious data into your network.

On Carrier-enabled FortiGate units, GTP related encapsulated filtering falls under encapsulated IP traffic filtering, and encapsulated non-IP end user address filtering.

### Configuring Encapsulated IP Traffic Filtering

Generally there are a very limited number of IP addresses that are allowed to encapsulate GPRS traffic. For example GTP tunnels are a valid type of encapsulation when used properly. This is the GTP tunnel which uses the Gp or Gn interfaces between SGSNs and GGSNs. However, a GTP tunnel within a GTP tunnel is not accessible — FortiOS Carrier will either block or forward the traffic, but is not able to open it for inspection.

The ability to filter GTP sessions is based on information contained in the data stream and provides operators with a powerful mechanism to control data flows within their infrastructure. You can also configure IP filtering rules to filter encapsulated IP traffic from Mobile Stations.

To configure the Encapsulated IP Traffic Filtering, go to *Security Profiles > Carrier > GTP Profile*, and edit a GTP profile. Expand *Encapsulated IP Traffic Filtering* to configure settings. See [“Encapsulated IP traffic filtering options” on page 694](#).

### When to use encapsulated IP traffic filtering

The following are the typical cases that need encapsulated IP traffic filtering:

#### Mobile station IP pools

In a well-designed network, best practices dictate that the mobile station address pool is to be completely separate from the GPRS network infrastructure range of addresses. Encapsulated

IP packets originating from a mobile station will not contain source or destination addresses that fall within the address range of GPRS infrastructures. In addition, traffic originating from the users handset will not have destination/source IP addresses that fall within any Network Management System (NMS) or Charging Gateway (CG) networks.

### Communication between mobile stations

Mobile stations on the same GPRS network are not able to communicate with other mobile stations. Best practices dictate that packets containing both source and destination addresses within the mobile station's range of addresses are to be dropped.

### Direct mobile device or internet attacks

It may be possible for attackers to wrap attack traffic in GTP protocols and submit the resulting GTP traffic directly to a GPRS network element from their mobile stations or a node on the Internet. It is possible that the receiving SGSN or GGSN would then strip off the GTP header and attempt to route the underlying attack. This underlying attack could have any destination address and would probably have a source address spoofed as if it were valid from that PLMN.



You cannot add an IE removal policy when you are creating a new profile.

### Relayed network attacks

Depending on the destination the attack could be directly routed, such as to another node of the PLMN, or rewrapped in GTP for transmission to any destination on the Internet outside the PLMN depending on the routing table of the GSN enlisted as the unwitting relay.

The relayed attack could have any source or destination addresses and could be any of numerous IP network attacks, such as an attack to hijack a PDP context, or a direct attack against a management interface of a GSN or other device within the PLMN. Best practices dictate that any IP traffic originating on the Internet or from an MS with a destination address within the PLMN is to be filtered.

## Configuring Encapsulated Non-IP End User Address Filtering

Much of the traffic on the GPRS network is in the form of IP traffic. However some parts of the network do not use IP based addressing, so the Carrier-enabled FortiGate unit is unable to perform Encapsulated IP Traffic Filtering.

Depending on the installed environment, it may be beneficial to detect GTP packets that encapsulate non-IP based protocols. You can configure the FortiOS Carrier firewall to permit a list of acceptable protocols, with all other protocols denied.

The encoded protocol is determined in the PDP Type Organization and PDP Type Number fields within the End User Address Information Element. The PDP Type Organization is a 4-bit field that determines if the protocol is part of the ETSI or IETF organizations. Values are zero and one, respectively. The PDP Type field is one byte long. Both GTP specifications only list PPP, with a PDP Type value of one, as a valid ETSI protocol. PDP Types for the IETF values are determined in the "Assigned PPP DLL Protocol Numbers" sections of RFC 1700. The PDP types are compressed, meaning that the most significant byte is skipped, limiting the protocols listed from 0x00 to 0xFF.

To configure the Encapsulated Non-IP End User Address Filtering, go to *Security Profiles > Carrier > GTP Profile*, and edit a GTP profile. Expand *Encapsulated Non-IP End User Address Filtering* to configure settings. See [“Encapsulated non-IP end user traffic filtering options” on page 695](#).

## Configuring the Protocol Anomaly feature in FortiOS Carrier

When anomalies do happen, it is possible for the anomaly to interrupt network traffic or consume network resources — if precautions are not taken. Anomalies can be generated by accident or maliciously, but both methods can have the same results — degrading the performance of the carrier network, or worse.

To configure GTP protocol anomalies, go to *Security Profiles > Carrier > GTP Profile*, and edit a GTP profile. Expand the *Protocol Anomaly* option. See [“Protocol Anomaly prevention options” on page 696](#).

The following are some examples:

- The GTP header specifies the length of the packet excluding the mandatory GTP header. In GTP version 0 (GSM 09.60), the mandatory GTP header size is 20 bytes, whereas GTP version 1 (GSM 29.060) specifies that the minimum length of the GTP header is 8 bytes. The GTP packet is composed of the header, followed by Information Elements typically presented in a Type-Length-Value format. It is possible for an attacker to create a GTP packet with a GTP header field length that is incompatible with the length of the necessary information elements.
- The same concepts are true for GTP version 2 headers even though there are different fields in them.
- It is similarly possible for an attacker to create a packet with an invalid IE length. Invalid lengths may cause protocol stacks to allocate incorrect amounts of memory, and thereby cause crashes or buffer overflows.

By default the FortiOS Carrier firewall detects these problems, as well as other protocol anomalies, and drops the packets. All protocol anomaly options are set to *Deny* by default. However, you can change the policy to allow them.

## Configuring Anti-overbilling in FortiOS Carrier

This section includes:

- [Overbilling in GPRS networks](#)
- [Anti-overbilling with FortiOS Carrier](#)

### Overbilling in GPRS networks

GPRS overbilling attacks can be prevented with a properly configured Carrier-enabled FortiGate unit.

Overbilling can occur when a subscriber returns his IP address to the IP pool. Before the billing server closes it, the subscriber's session is still open and vulnerable. If an attacker takes control of the subscriber's IP address, he can send or receive data and the subscriber will be billed for the traffic.

Overbilling can also occur when an available IP address is reassigned to a new mobile station (MS). Subsequent traffic by the previous MS may be forwarded to the new MS. The new MS would then be billed for traffic it did not initiate.

### Anti-overbilling with FortiOS Carrier

The Carrier-enabled FortiGate unit can be configured to assist with anti-overbilling measures. These measures ensure that the customer is only billed for connection time and data transfer that they actually use.

Anti-overbilling on the Carrier-enabled FortiGate unit involves:

- the administrator configuring the overbilling settings in the GTP profile to notify the Gi firewall when a GTP tunnel is deleted
- the unit clearing the sessions when the Gi firewall receives a notification from the Gn/Gp firewall about a GTP tunnel being deleted This way, the Gi firewall prevents overbilling by blocking traffic initiated by other users.

The three locations to configure anti-overbilling options include:

- *System > Network > Interface > Gi Gatekeeper* — edit an interface, and enable to monitor Gi anti-overbilling traffic on this interface
- *System > Admin > Settings > Gi Gatekeeper Settings* — set the context ID and port that anti-overbilling will take place on.
- *Security Profiles > Carrier > GTP Profile > Anti-Overbilling* — the IP address, port, interface and context ID to use for anti-overbilling measures.

For detailed options, see “[Anti-Overbilling options](#)” on page 697.

## Logging events on the Carrier-enabled FortiGate unit

Logging on the Carrier-enabled FortiGate unit is just like logging on any other FortiOS unit. The only difference with FortiOS Carrier is that there are a few additional events that you can log beyond the regular ones. These additional events are covered here. For more information on other logging issues, see the [Logging and Reporting Guide](#) and [FortiOS CLI Reference](#).

To enable FortiOS Carrier logging, go to *Log&Report > Event Log*, and ensure *GTP service event* is enabled. Once this option is selected, the logging options under *Security Profiles > Carrier > GTP Profile* will be active.

To change FortiOS Carrier specific logging event settings, go to *Security Profiles > Carrier > GTP Profile* and edit a GTP profile. Expand the *Log* section to change the settings. For detailed options, see “[Log options](#)” on page 697.

The following information is contained in each log entry:

<b>Timestamp</b>	The time and date when the log entry was recorded
<b>Source IP address</b>	The sender’s IP address.
<b>Destination IP address</b>	The receiver’s IP address. The sender-receiver pair includes a mobile phone on the GPRS local network, and a device on a network external to the GPRS network, such as the Internet.
<b>Tunnel Identifier (TID) Tunnel Endpoint Identifier (TEID)</b>	An identifier for the start and endpoints of a GTP tunnel. This information uniquely defines all tunnels. It is important for billing information based on the length of time the tunnel was active and how much data passed over the tunnel.
<b>Message type</b>	For available message types, see “ <a href="#">Common message types on carrier networks</a> ” on page 745.
<b>Packet status</b>	What action was performed on the packet. This field matches the logging options while you are configuring GTP logging. See “ <a href="#">Anti-overbilling with FortiOS Carrier</a> ” on page 742.  The status can be one of forwarded, prohibited, state-invalid, rate-limited, or tunnel-limited

<b>Virtual domain ID or name</b>	A Carrier-enabled FortiGate unit can be divided into multiple virtual units, each being a complete and self-contained virtual FortiCarrier unit. This field indicates which virtual domain (VDOM) was responsible for the log entry. If VDOMs are not enabled on your unit, this field will be <code>root</code> .
<b>Reason to be denied if applicable</b>	If the packet that generated this log entry was denied or blocked, this field will include what part of FortiOS denied or blocked that packet. Such as firewall, antivirus, webfilter, or spamfilter.

An example of the above log message format is for a Tunnel deleted log entry. When a tunnel is deleted, the log entry contains the following information:

- Timestamp
- Interface name (if applicable)
- SGSN IP address (source IP)
- GGSN IP address (destination IP)
- Tunnel ID
- Tunnel duration time in seconds
- Number of messages sent to the SGSN
- Number of messages sent to the GGSN



# GTP message type filtering

FortiOS Carrier supports message filtering in GTP by the type of message.

This section includes:

- [Common message types on carrier networks](#)
- [Configuring message type filtering in FortiOS Carrier](#)

## Common message types on carrier networks

Carrier networks include many types of messages — some concern the network itself, others are content moving across the network, and still others deal with handshaking, billing, or other administration based issues.

GTP contains two major parts GTP for the control plane (GTP-C) and GTP for user data tunnelling (GTP-U). Outside of those areas there are only unknown message types.

### GTP-C messages

GTP-C contains the networking layer messages. These address routing, versioning, and other similar low level issues.

When a subscriber requests a Packet Data Protocol (PDP) context, the SGSN will send a create PDP context request GTP-C message to the GGSN giving details of the subscriber's request. The GGSN will then respond with a create PDP context response GTP-C message which will either give details of the PDP context actually activated or will indicate a failure and give a reason for that failure. This is a UDP message on port 212.

GTP-C message types include Path Management Messages, Location Management Messages, and Mobility Management Messages.

#### Path Management Messages

Path management is used by one GSN to detect if another GSN is alive, or if it has restarted after a failure.

The path management procedure checks if a given GSN is alive or has been restarted after a failure. In case of SGSN restart, all MM and PDP contexts are deleted in the SGSN, since the associated data is stored in a volatile memory. In the case of GGSN restart, all PDP contexts are deleted in the GGSN.

#### Tunnel Management Messages

The tunnel management procedures are used to create, update, and delete GTP tunnels in order to route IP PDUs between an MS and an external PDN via the GSNs.

The PDP context contains the subscriber's session information when the subscriber has an active session. When a mobile wants to use GPRS, it must first attach and then activate a PDP context. This allocates a PDP context data structure in the SGSN that the subscriber is currently visiting and the GGSN serving the subscriber's access point.

Tunnel management procedures are defined to create, update, and delete tunnels within the GPRS backbone network. A GTP tunnel is used to deliver packets between an SGSN and a

GGSN. A GTP tunnel is identified in each GSN node by a TEID, an IP address, and a UDP port number.

### Location Management Messages

The location-management procedure is performed during the network-requested PDP context activation procedure if the GGSN does not have an SS7 MAP interface (i.e., Gc interface). It is used to transfer location messages between the GGSN and a GTP-MAP protocol-converting GSN in the GPRS backbone network.

Location management subprocedures are used between a GGSN that does not support an SS7 MAP interface (i.e., Gc interface) and a GTP-MAP protocol-converting GSN. This GSN supports both Gn and Gc interfaces and is able to perform a protocol converting between GTP and MAP.

### Mobility Management Messages

The MM procedures are used by a new SGSN in order to retrieve the IMSI and the authentication information or MM and PDP context information in an old SGSN. They are performed during the GPRS attach and the inter-SGSN routing update procedures.

The MM procedures are used between SGSNs at the GPRS-attach and inter-SGSN routing update procedures. An identity procedure has been defined to retrieve the IMSI and the authentication information in an old SGSN. This procedure may be performed at the GPRS attach. A recovery procedure enables information related to MM and PDP contexts in an old SGSN to be retrieved. This procedure is started by a new SGSN during an inter-SGSN RA update procedure.

## GTP-U messages

GTP-U is focused on user related issues including tunneling, and billing. GTP-U message types include MBMS messages, and GTP-U and Charging Management Messages

### MBMS messages

Multimedia Broadcast and Multicast Services (MBMS) have recently begun to be offered over GSM and UMTS networks on UTRAN and GERAN radio access technologies. MBMS is mainly used for mobile TV, using up to four GSM timeslots for one MBMS connection. One MBMS packet flow is replicated by GGSN, SGSN and RNCs.

MBMS is split into the MBMS Bearer Service and the MBMS User Service. The MBMS User Service is basically the MBMS Service Layer and offers a Streaming- and a Download Delivery Method. The Streaming Delivery method can be used for continuous transmissions like Mobile TV services. The Download Method is intended for "Download and Play" services.

### GTP-U and Charging Management Messages

SGSNs and GGSNs listen for GTP-U messages on UDP port 2152.

GTP' (GTP prime) is used for billing messages. It uses the common GTP messages (GTP Version Not Supported, Echo Request and Echo Response) and adds additional messages related to billing procedures.

## Unknown Action messages

If the system doesn't know what type of message it is, it falls into this category. This is an important category of message because malformed messages may appear and need to be handled with security in mind.



Fortinet best practices dictate that you set *Unknown Action messages* to deny for security reasons.

## Configuring message type filtering in FortiOS Carrier

GPRS Tunnelling Protocol (GTP) is a group of IP-based communications protocols used to carry General Packet Radio Service (GPRS) traffic within Global System for Mobile Communications (GSM) and Universal Mobile Telecommunications System (UMTS) networks. It allows carriers to transport actual cellular packets over their network via tunneling.

In the CLI, there is a keyword for each type of GTP message for both message filtering, and for message rate limiting.



GTP message rate limiting is only accessible from the CLI using the command `configure firewall gtp`.

### To configure GTP message type filtering - web-based manager

1. Go to *Security Profiles > Carrier > GTP Profile*.
2. Select *Create New*.
3. Enter a name for this profile such as `msg_type_filtering`.
4. Select *Message Type Filtering* to expand it.
5. For each type of message in the list, select Allow or Deny. All messages are set to Allow by default.



Fortinet best practices dictate that the unknown message action should be set to Deny for security reasons as this will block malformed messages.

6. Optionally select and configure any other GTP features for this profile, such as logging.
7. Select OK to save the profile.
8. Apply the `msg_type_filtering` profile a security policy configured for GTP tunnel traffic.

### To configure GTP message filtering and block Unknown Message Action messages- CLI

```
config firewall gtp
 edit msg_type_filtering
 config message-filter
 set unknown-message-action deny
 next
 end
end
```

## Message Type Fields

Each of the following message types can be allowed or denied by your Carrier-enabled FortiGate unit depending on your carrier network and GTP traffic.

The message types include:

- [Unknown Message Action](#)
- [Path Management Messages](#)
- [Tunnel Management Messages](#)
- [Location Management Messages](#)
- [Mobility Management Messages](#)
- [MBMS messages](#)
- [GTP-U and Charging Management Messages](#)

### Unknown Message Action

Set this message type to deny.

Many attempts to hack into a carrier network will result in this unknown message type and therefore it is denied for security reasons.

### Path Management Messages

Message Type	Used by	Description
Echo Request/Response	GTP-C, GTP-U, GTP'	Echo Request is sent on a path to another GSN to determine if the other node is alive. Echo Response is the reply.
Version not Supported	GTP-C, GTP-U, GTP'	There are multiple versions of GTP. Both devices communicating must use the same version of GTP, or this message will be the response.
Support Extension Headers Notification		Extensions are optional parts that a device can choose to support or not. If a device includes these extensions, it must include headers for the extensions to sure ensure proper formatting.

### Tunnel Management Messages

Message Type	Used by	Description
Create PDP Context Request/ Response	GTP-C	Sent from an SGSN to a GGSN node as part of a GPRS PDP Context Activation procedure or the Network-Requested PDP Context Activation procedure. A valid request initiates the creation of a tunnel.
Update PDP Context Request/ Response	GTP-C	Used when PDP Context information changes, such as when a mobile device changes location.
Delete PDP Context Request/ Response	GTP-C	Used to terminate a PDP Context, and confirm the context has been deleted.

Create AA PDP Context Request/ Response	GTP-C	Sent as part of the GPRS Anonymous Access PDP Context Activation. It is used to create a tunnel between a context in the SGSN and a context in the GGSN.
Delete AA PDP Context Request/ Response	GTP-C	Sent as part of the GPRS PDP Anonymous Access Context Deactivation procedure to deactivate an activated PDP Context. It contains Cause and Private Extension Information Elements
Error Indication	GTP-U	Sent to the GGSN when a tunnel PDU is received for the following conditions: <ul style="list-style-type: none"> <li>– No PDP context exists</li> <li>– PDP context is inactive</li> <li>– No MM context exists</li> <li>– GGSN deletes its PDP context when the message is received.</li> </ul>
PDU Notification Request/ Response/ Reject Request/ Reject Response	GTP-C	When receiving a Tunneled PDU (T-PDU), the GGSN checks if a PDP context is established for the given PDP address. If no PDP context has been established, the GGSN may initiate the Network-requested PDP Context Activation procedure by sending a PDU Notification Request to the SGSN.  <b>Reject Request</b> - Sent when the PDP context requested by the GGSN cannot be established.

## Location Management Messages

Message Type	Used By	Description
Send Routing Information for GPRS Request/ Response	GTP-C	Sent by the GGSN to obtain location information for the MS. This message type contains the IMSI of the MS and Private Extension.
Failure Report Request/ Response	GTP-C	Sent by the GGSN to the HLR when a PDU reject message is received.  The GGSN requests the HLR to set the flag and add the GGSN to the list of nodes to report to when activity from the subscriber that owns the PDP address is detected.  The message contains the subscriber IMSI and Private Extension
Note MS GPRS Present Request/ Response	GTP-C	When the HLR receives a message from a mobile with MDFG set, it clears the MDFG and sends the Note MS Present message to all GGSN's in the subscriber's list.  This message type contains subscriber IMSI, GSN Address and Private Extension

## Mobility Management Messages

Message Type	Used By	Description
Identification Request/Response	GTP-C	Sent by the new SGSN to the old SGSN to request the IMSI for a MS when a GPRS Attach is done with a P-TMSI and the MS has changed SGSNs since the GPRS Detach was done.
SGSN context Request/ Response/ Acknowledge	GTP-C	Sent by the new SGSN to the old SGSN to request the MM and PDP Contexts for the MS.
Forward Relocation Request/ Response/ Complete/ Complete Acknowledge	GTP-C	Indicates mobile activation/deactivation within a Routing Area. This prevents paging of a mobile that is not active (visited VLR rejects calls from the HLR or applies Call Forwarding). Note that the mobile station does not maintain an attach/detach state.  SRNS contexts contain for each concerned RAB the sequence numbers of the GTP-PDUs next to be transmitted in uplink and downlink directions.
Relocation Cancel Request/ Response	GTP-C	Send to cancel the relocation of a connection.
Forward SRNS Context/ Context Acknowledge	GTP-C	This procedure may be used to trigger the transfer of SRNS contexts from RNC to CN (PS domain) in case of inter system forward handover.
RAN Information Relay	GTP-C	Forward the Routing Area Network (RAN) information.  A Routing Area (RA) is a subset of a GSM Location Area (LA). A RA is served by only one SGSN. Ensures that regular radio contact is maintained by the mobile

## MBMS messages

Message Type	Used By	Description
MBMS Notification Request/ Response/ Reject Request/ Reject Response	GTP-C	Notification of the radio access devices.
Create MBMS Context Request/ Response	GTP-C	Request to create an active MBMS context. The context will be pending until the response is received.  Once active, the MBMS context allows the MS to receive data from a specific MBMS source

Update MBMS Context Request/Response	GTP-C	
Delete MBMS Context Request/Response	GTP-C	Request to deactivate the MBMS context. When the response is received, the MBMS context will be inactive.

### GTP-U and Charging Management Messages

Message Type	Used By	Description
G-PDU	GTP-C, GTP-U	GPRS Packet data unit delivery message.
Node Alive Request/Response	GTP-C, GTP-U	Used to inform rest of network when a node starts service.
Redirection Request/Response	GTP-C, GTP-U	Used to divert the flow of CDRs from the CDFs to another CGF when the sender is being removed, or they are used when the CGF has lost its connection to a downstream system.
Data Record Transfer Request/Response	GTP-C, GTP-U	Used to reliably transport CDRs from the point of generation (SGSN/GGSN) to non-volatile storage in the CGF

# GTP identity filtering

FortiOS Carrier supports a number of filtering methods based on subscriber identity such as APN filtering, IMSI filtering, and advanced filtering.

This section includes:

- [IMSI on carrier networks](#)
- [Other identity and location based information elements](#)
- [Configuring APN filtering in FortiOS Carrier](#)
- [Configuring IMSI filtering in FortiOS Carrier](#)
- [Configuring advanced filtering in FortiOS Carrier](#)


## IMSI on carrier networks

The International Mobile Subscriber Identity (IMSI) number is central to identifying users on a carrier network. It is a unique number that is assigned to a cell phone or mobile device to identify it on the GSM or UTM network.

Typical the IMSI number is stored on the SIM card of the mobile device and is sent to the network as required.

An IMSI number is 15 digits long, and includes the Mobile Country Code (MCC), Mobile Network Code (MNC), and Mobile Station Identification Number (MSIN).

**Figure 152:**IMSI codes



The Home Network Identity (HNI) is made up of the MCC and MNC. The HNI is used to fully identify a user's home network. This is important because some large countries have more than one country code for a single carrier. For example a customer with a mobile carrier on the East Coast of the United States would have a different MCC than a customer on the West Coast with the same carrier because even through the MNC would be the same the MCC would be different — the United States uses MCCs 310 to 316 due to its size.

If an IMSI number is not from the local carrier's network, IMSI analysis is performed to resolve the number into a Global Title which is used to access the user's information remotely on their home carrier's network for things like billing and international roaming.

See [“Configuring IMSI filtering in FortiOS Carrier” on page 756](#).

## Other identity and location based information elements

IMSI focuses on the user, their location, and carrier network. There are other numbers used to identify different user related Information Elements (IE).



These identity and location based elements include:

- [Access Point Number \(APN\)](#)
- [Mobile Subscriber Integrated Services Digital Network \(MSISDN\)](#)
- [Radio Access Technology \(RAT\) type](#)
- [User Location Information \(ULI\)](#)
- [Routing Area Identifier \(RAI\)](#)
- [International Mobile Equipment Identity \(IMEI\)](#)

### Access Point Number (APN)

The Access Point Number (APN) is used in GPRS networks to identify an IP packet data network that a user wants to communicate with. The Network Identifier describes the network and optionally the service on that network that the GGSN is connected to. The APN also includes the MCC and MCN, which together locate the network the GGSN belongs to. An example of an APN in the Barbados using Digicel as the carrier that is connecting to the Internet is `internet.mcc342.mnc750.gprs`.

When you are configuring your Carrier-enabled FortiGate unit's GTP profiles, you must first configure the APN. It is critical to GTP communications and without it no traffic will flow.

The access point can then be used in a DNS query to a private DNS network. This process (called APN resolution) gives the IP address of the GGSN which serves the access point. At this point a PDP context can be activated. See [“Configuring APN filtering in FortiOS Carrier” on page 755](#).

### Mobile Subscriber Integrated Services Digital Network (MSISDN)

This is a 15-digit number that, along with the IMSI, uniquely identifies a mobile user. Normally this number includes a 2-digit country code, a 3-digit national destination code, and a 10-digit subscriber number or the phone number of the mobile device, and because of that may change over time if the user changes their phone number. The MSISDN number follows the ITU-T E.164 numbering plan.

### Radio Access Technology (RAT) type

The RAT type represents the radio technology used by the mobile device. This can be useful in determining what services or content can be sent to a specific mobile device. FortiOS Carrier supports:

- **UMTS Terrestrial Radio Access Network (UTRAN)**, commonly referred to as 3G, routes many types of traffic including IP traffic. This is one of the faster types.
- **GSM EDGE Radio Access Network (GERAN)** is a key part of the GSM network which routes both phone calls and data.
- **Wireless LAN (WLAN)** is used but not as widely as the other types. It is possible for the mobile device to move from one WLAN to another such as from an internal WLAN to a commercial hot spot.
- **Generic Access Network (GAN)** can also be called unlicensed mobile access (UMA). It routes voice, data, and SIP over IP networks. GAN is commonly used for mobile devices that have a dual-mode and can hand-off between GSM and WLANs.
- **High Speed Packet Access (HSPA)** includes two other protocols High Speed Downlink and Uplink Packet Access protocols (HSDPA and HSUPA respectively). It improves on the older WCDMA protocols by better using the radio bandwidth between the mobile device and the radio tower. This results in an increased data transfer rate for the user.

RAT type is part of advanced filtering configuration. See [“Configuring advanced filtering in FortiOS Carrier” on page 757](#).

## User Location Information (ULI)

Gives Cell Global Identity/Service Area Identity (CGI/SAI) of where the mobile station is currently located. The ULI and the RAI are commonly used together to identify the location of the mobile device.

ULI is part of advanced filtering configuration. See [“Configuring advanced filtering in FortiOS Carrier” on page 757](#).

## Routing Area Identifier (RAI)

Routing Areas (RAs) divide the carrier network and each has its own identifier (RAI). When a mobile device moves from one routing area to another, the connection is handled by a different part of the network. There are normally multiple cells in a routing area. There is only one SSGN per routing area. The RAI and ULI are commonly used to determine a user’s location.

RAI is part of advanced filtering configuration. See [“Configuring advanced filtering in FortiOS Carrier” on page 757](#).

## International Mobile Equipment Identity (IMEI)

IMEI is a unique 15-digit number used to identify mobile devices on mobile networks. It is very much like the MAC address of a TCP/IP network card for a computer. It can be used to prevent network access by a stolen phone — the carrier knows the mobile phone’s IMEI, and when it is reported stolen that IMEI is blocked from accessing the carrier network no matter if it has the same SIM card as before or not. It is important to note that the IMEI stays with the mobile phone or device where the other information is either location based or stored on the removable SIM card.

IMEI type is part of advanced filtering configuration. See [“Configuring advanced filtering in FortiOS Carrier” on page 757](#).

## When to use APN, IMSI, or advanced filtering

At first glance APN, IMSI, and advanced filtering have parts in common. For example two can filter on APN, and another two can filter on IMSI. The difficulty is knowing when to use which type of filtering.

**Figure 153:**Identity filtering comparison

Filtering type	Filter on the following data:	When to use this type of filtering
APN	APN	Filter based on GTP tunnel start or destination
IMSI	IMSI, MCC-MNC	Filter based on subscriber information
Advanced	PDP context, APN, IMSI, MSISDN, RAT type, ULI, RAI, IMEI	When you want to filter based on: user phone number (MSISDN) what wireless technology the user employed to get on the network (RAT type) user location (ULI and RAI) handset ID, such as for stolen phones (IMEI)

APN filtering is very specific — the only identifying information that is used to filter is the APN itself. This will always be present in GTP tunnel traffic, so all GTP traffic can be filtered using this value. See [“Configuring APN filtering in FortiOS Carrier” on page 755](#).

IMSI filtering can use a combination of the APN and MCC-MNC numbers. The MCC and MNC are part of the APN, however filtering on MCC-MNC separately allows you to filter based on country and carrier instead of just the destination of the GTP Tunnel. See [“Configuring IMSI filtering in FortiOS Carrier” on page 756](#).

Advanced filtering can go into much deeper detail covering PDP contexts, MSISDN, IMEI, and more not to mention APN, and IMSI as well. If you can’t find the information in APN or IMSI that you need to filter on, then use Advanced filtering. See [“Configuring advanced filtering in FortiOS Carrier” on page 757](#).

## Configuring APN filtering in FortiOS Carrier

To configure APN filtering go to *Security Profiles > Carrier > GTP Profile*. Select a profile or create a new one, and expand *APN filtering*.



When you are configuring your Carrier-enabled FortiGate unit’s GTP profiles, you must first configure the APN. It is critical to GTP communications and without it no traffic will flow.

For more information on APN, see [“Access Point Number \(APN\)” on page 753](#).

<b>Enable APN Filter</b>	Select to enable filtering based on APN value.
<b>Default APN Action</b>	Select either Allow or Deny for all APNs that are not found in the list. The default is Allow.
<b>Value</b>	Displays the APN value for this entry. Partial matches are allowed using wildcard. For example *.mcc333.mcn111.gprs would match all APNs from country 333 and carrier 111 on the gprs network.

<b>Mode</b>	Select one or more of the methods used to obtain APN values.  Mobile Station provided - The APN comes from the mobile station where the mobile device connected. This is the point of entry into the carrier network for the user's connection.  Network provided - The APN comes from the carrier network.  Subscription Verified - The user's subscription has been verified for this APN. This is the most secure option.
<b>Action</b>	One of allow or deny to allow or block traffic associated with this APN.
<b>Delete icon</b>	Select to remove this APN entry from the list.
<b>Edit icon</b>	Select to change the information for this APN entry.
<b>Add APN</b>	Select to add an APN to the list. Not active while creating GTP profile, only when editing an existing GTP profile.  Save all changes before adding APNs. A warning to this effect will be displayed when you select the <i>Add APN</i> button.

The Add APN button is not activated until you save the new GTP profile. When you edit that GTP profile, you will be able to add new APNs.

## Configuring IMSI filtering in FortiOS Carrier

In many ways the IMSI on a GPRS network is similar to an IP address on a TCP/IP network. Different parts of the number provide different pieces of information. This concept is used in IMSI filtering on FortiOS Carrier.

To configure IMSI filtering go to *Security Profiles > Carrier > GTP Profile* and expand *IMSI filtering*.

While both the APN and MCC-MCN fields are optional, without using one of these fields the IMSI entry will not be useful as there is no information for the filter to match.

<b>Enable IMSI Filter</b>	Select to turn on IMSI filtering.
<b>Default IMSI Action</b>	Select Allow or Deny. This action will be applied to all IMSI numbers except as indicated in the IMSI list that is displayed.  The default value is Allow.
<b>APN</b>	The Access Point Number (APN) to filter on.  This field is optional.
<b>MCC-MNC</b>	The Mobile Country Code (MCC) and Mobile Network Code (MNC) to filter on. Together these numbers uniquely identify the carrier and network of the GGSN being used.  This field is optional.

<b>Mode</b>	<p>Select the source of the IMSI information as one or more of the following:</p> <p><b>Mobile Station provided</b> - the IMSI number comes from the mobile station the mobile device is connecting to.</p> <p><b>Network provided</b> - the IMSI number comes from the GPRS network which could be a number of sources such as the SGSN, or HLR.</p> <p><b>Subscription Verified</b> - the IMSI number comes from the user's home network which has verified the information.</p> <p>While Subscription Verified is the most secure option, it may not always be available. Selecting all three options will ensure the most complete coverage.</p>
<b>Action</b>	Select the action to take when this IMSI information is encountered. Select one of Allow or Deny.
<b>Delete Icon</b>	Select the delete icon to remove this IMSI entry.
<b>Edit Icon</b>	Select the edit icon to change information for this IMSI entry.
<b>Add IMSI</b>	<p>Select to add an IMSI to the list. Not active while creating GTP profile, only when editing an existing GTP profile.</p> <p>Save all changes before adding IMSIs. A warning to this effect will be displayed when you select the <i>Add IMSI</i> button.</p>

Also see [“Basic filtering options”](#) on page 689.

## Configuring advanced filtering in FortiOS Carrier

Compared to ADN or IMSI filtering, advanced filtering is well named. Advanced filtering can be viewed as a catch-all filtering option — if ADN or IMSI filtering doesn't do what you want, then advanced filtering will. The advanced filtering can use more information elements to provide considerably more granularity for your filtering.

<b>Enable</b>	Select to turn on advanced filtering.
<b>Default Action</b>	Select Allow or Deny as the default action to take when traffic does not match an entry in the advanced filter list .

<b>Messages</b>	<p>Optionally select one or more types of messages this filter applies to:</p> <p>Create PDP Context Request, Create PDP Context Response, Update PDP Context Request, or Update PDP Context Response.</p> <p>Selecting <i>Create PDP Context Response</i> or <i>Update PDP Context Response</i> limits RAT type to only GAN and HSPA, and disables the APN, APN Mode, IMSI, MSISDN, ULI, RAI, and IMEI fields.</p> <p>To select <i>Update PDP Context Request</i>, APN Restriction must be set to <i>all</i>. Selecting <i>Update PDP Context Request</i> disables the APN, MSISDN, and IMEI fields.</p> <p>if all message types are selected, only the RAT Types of GAN and HSPA are available to select.</p>
<b>APN Restriction</b>	<p>APN Restriction either allows all APNs or restricts the APNs to one of four categories — Public-1, Public-2, Private-1, or Private-2. This can also be combined with a specific APN or partial APN as well as specifying the APN mode. See <a href="#">“Access Point Number (APN)” on page 753</a>.</p>
<b>RAT Type</b>	<p>Select one or more of the Radio Access Technology Types listed. These fields control how a user accesses the carrier’s network. You can select one or more of UTRAN, GERAN, WLAN, GAN, HSPA, or any. See <a href="#">“Radio Access Technology (RAT) type” on page 753</a>.</p>
<b>ULI</b>	<p>The user location identifier. Often the ULI is used with the RAI to locate a user geographically on the carrier’s network.</p> <p>The ULI is disabled when <i>Create PDP Context Response</i> or <i>Update PDP Context Response</i> messages are selected.</p> <p>See <a href="#">“User Location Information (ULI)” on page 754</a>.</p>
<b>RAI</b>	<p>The router area identifier. There is only one SGSN per routing area on a carrier network. This is often used with ULI to locate a user geographically on a carrier network.</p> <p>The RAI is disabled when <i>Create PDP Context Response</i> or <i>Update PDP Context Response</i> messages are selected.</p> <p>See <a href="#">“Routing Area Identifier (RAI)” on page 754</a>.</p>
<b>IMEI</b>	<p>The International Mobile Equipment Identity. The IMEI uniquely identifies mobile hardware, and can be used to block stolen equipment.</p> <p>The IMEI is only available when <i>Create PDP Context Request</i> or no messages are selected.</p> <p>See <a href="#">“International Mobile Equipment Identity (IMEI)” on page 754</a></p>
<b>Action</b>	<p>Select Allow or Deny as the action when this filter matches traffic.</p> <p>The default is Allow.</p>
<b>Delete Icon</b>	<p>Select to delete this entry from the list.</p>

<b>Edit Icon</b>	Select to edit this entry.
<b>Add</b>	Select to add an advanced filter to the list. Not active while creating GTP profile, only when editing an existing GTP profile.  Save all changes before adding advanced filters. A warning to this effect will be displayed when you select the <i>Add</i> button.

Also see [“Advanced filtering options”](#) on page 690.

# Troubleshooting

This section offers troubleshooting options for Carrier-related issues.

This section includes:

- [FortiOS Carrier diagnose commands](#)
- [Applying IPS signatures to IP packets within GTP-U tunnels](#)
- [GTP packets are not moving along your network](#)

## FortiOS Carrier diagnose commands

This section includes diagnose commands specific to FortiOS Carrier features such as GTP.

### GTP related diagnose commands

This CLI command allows you to gain information on GTP packets, logs, statistics, and other information.

```
diag firewall gtp <command>
```

<b>apn list &lt;gtp_profile&gt;</b>	The APN list entries in the specified GTP profile
<b>auth-ggsns show &lt;gtp_profile&gt;</b>	The authorized GGSNs entries for the specified GTP profile. Any GGSNs not on this list will not be recognized.
<b>auth-sgsns show &lt;gtp_profile&gt;</b>	The authorized SGSNs list entries for the specified GTP profile. Any SGSNs not on this list will not be recognized.
<b>handover-grp show &lt;gtp_profile&gt;</b>	The handover group showing the range of allowed handover group IP addresses. The handover group acts like a whitelist of allowed GTP addresses with a default deny at the end – if the GTP address is not on the list, it is denied.
<b>ie-remove-policy list &lt;gtp_profile&gt;</b>	List of IE policies in the IE removal policy for this GTP profile. The information displayed includes the message count for this policy, the length of the SGSN, the list of IEs, and list of SGSN IP addresses.
<b>imsi list &lt;gtp_profile&gt;</b>	IMSI filter entries for this GTP profile. The information displayed includes the message count for this filter, length of the IMSI, the length of the APN and IMSI, and of course the IMSI and APN values.
<b>invalid-sgsns-to-long list &lt;gtp_profile&gt;</b>	List of SGSNs that do not match the filter criteria. These SGSNs will be logged.
<b>ip-policy list &lt;gtp_profile&gt;</b>	List the IP policies including message count for each policy, the action to take, the source and destination IP addresses or ranges, and masks.
<b>noip-policy &lt;gtp_profile&gt;</b>	List the non-IP policies including the message count, which mode, the action to take, and the start and end protocols to be used by decimal number.



<b>path {list   flush}</b>	Select list or flush. List the GTP related paths in FortiOS Carrier memory. Flush the GTP related paths from memory.
<b>policy list &lt;gtp_policy&gt;</b>	The GTP advanced filter policy information for this GTP profile. The information displayed for each entry includes a count for messages matching this filter, a hexadecimal mask of which message types to match, the associated flags, action to take on a match, APN selection mode, MSISDN, RAT types, RAI, ULI, and IMEI.
<b>profile list</b>	Displays information about the configured GTP profiles. You will not be able to see the bulk of the information if you do not log the output to a file.
<b>runtime-stat flush</b>	Select to flush the GTP runtime statistics from memory.
<b>stat</b>	Display the GTP runtime statistics — details on current GTP activity. This information includes how many tunnels are active, how many GTP profiles exist, how many IMSI filter entries, how many APN filter entries, advanced policy filter entries, IE remove policy filter entries, IP policy filter entries, clashes, and dropped packets.
<b>tunnel {list   flush}</b>	Select one of list or flush. List lists all the GTP tunnels currently active. Flush clears the list of active GTP tunnels.

## Applying IPS signatures to IP packets within GTP-U tunnels

GTP-U (GTP user data tunnelling) tunnels carry user data packets, signalling messages and error information. GTP-U uses UDP port 2152. Carrier-enabled FortiGate units can apply IPS intrusion protection and detection to GTP-U user data sessions.

To apply IPS to GTP-U user data sessions, add an IPS Sensor to a profile and add the profile to a security policy that accepts GTP-U tunnels. The security policy Service field must be set to GTP or ANY to accept GTP-U packets.

The Carrier-enabled FortiGate unit intercepts packets with destination port 2152, removes the GTP header and handles the packets as regular IP packets. Applying an IPS sensor to the IP packets, the Carrier-enabled FortiGate unit can log attacks and pass or drop packets depending on the configuration of the sensor.

If the packet is GTP-in-GTP, or a nested tunnel, the packets are passed or blocked without being inspected.

### To apply an IPS sensor to GTP-U tunnels

1. Go to *Security Profiles > Intrusion Protection > IPS Sensors* and select *Create New* to add an IPS Sensor.
2. Configure the IPS Sensor to detect attacks and log, drop, or pass attack packets.  
See the Intrusion Protection section of the [FortiOS UTM Guide](#).
3. Go to *Policy > Policy* and apply the IPS sensor to the security policy.
4. Go to *Policy > Policy* and select *Create New* to add a security policy or select a security policy.

5. Configure the security policy to accept GTP traffic.  
In the security policy configure the source and destination settings to match the GTP traffic. Service to GTP or ANY so that the security policy accepts GTP traffic.
6. Select the GTP profile within the security policy.
7. Configure any other required security policy settings.
8. Select *OK* to save the security policy.

## GTP packets are not moving along your network

When GTP packets are not getting to their destination, this could be caused by any one of a number of issues. General troubleshooting principals apply here.

The following sections provide some suggestions on how to troubleshoot this issue:

- [Attempt to identify the section of your network with the problem](#)
- [Ensure you have an APN configured](#)
- [Check the logs and adjust their settings if required](#)
- [Check the routing table](#)
- [Perform a sniffer trace](#)
- [Generate specific packets to test the network](#)

### Attempt to identify the section of your network with the problem

The first step is to determine how widespread this problem is. Does it affect the whole GPRS network, or just one or two devices?

If the entire network is has this problem, the solution is likely a more general one such as ensuring the security policies allow GTP traffic to pass, the GTP profile specifies SSGNs and GSGNs, or ensuring the GTP general settings are not overly limiting.

If one part of the network is affected, the problem is more likely centered around configurations with those network devices specified such as the handover group, or authorized SGSNs/GGSNs. It is also possible that small portions of the network may have hardware related issues such as cabling or faulty hardware. This section does not address those issues, and assumes hardware is not the problem.

The handover group is a whitelist of GTP addresses allowed to handle GTP messages. If a device's address is not on this list, it will be denied.

### Ensure you have an APN configured

When you configure your GTP profile, ensure you first configure the APN. Without it, there will be no flow of traffic. The APN is used in nearly all GTP communications and without it, the Carrier-enabled FortiGate unit doesn't have the information it needs.

### Check the logs and adjust their settings if required

During normal operation, the log settings will show any problems on the network but may not provide the level of details required to fully troubleshoot the problem. The reason for this is that the level of detail required for troubleshooting would quickly overwhelm the daily logs without any real benefit.

GTP related events in the event log will have message IDs in the range 41216 to 41222. For more information on GTP log messages, see the [Log Message Reference](#). For more information on logging in general, see the [Logging and Reporting handbook chapter](#).

Once there is a problem to troubleshoot, check the logs to trace the traffic patterns and narrow down the possible sources of the problem. There may be enough detail for you to locate and fix the problem without changing the log settings.



Remember to set any changes you made to the log settings back to their original values when you are done troubleshooting. Otherwise, the amount of detail will overwhelm your logging.

However, if more detail is required you can change settings such as:

- Lower the Log Frequency number in GTP Profiles so fewer or no log messages are dropped. This will allow a more accurate picture of everything happening on the network, where you may have had only a partial picture before.
- Ensure all the GTP log events are enabled to provide you with a complete picture.
- Ensure that all relevant event types are enabled under *Log & Report > Log Config > Event Log*.

For more information on GTP related logging, see “[Logging events on the Carrier-enabled FortiGate unit](#)” on page 743. See the log and report chapters of [Logging and Reporting Guide](#) and [FortiOS CLI Reference](#).

General information to look for in the logs includes:

- Are all packets having problems or just certain types?
- Are all devices on the network having problem, or just certain devices?
- Is it just GTP traffic that is having problems or are all types of traffic having the same problem?

## Check the routing table

On any network, the routing table determines how packets reach their destination. This is also true on a carrier network.

If the Carrier-enabled FortiGate unit is running in NAT mode, verify that all desired routes are in the routing table — local subnets, default routes, specific static routes, and dynamic routing protocols. For complete information, it is best to check the routing table in the CLI. This method provides more complete information.



If VDOMs are enabled on your Carrier-enabled FortiGate unit, all routing related CLI commands must be performed within a VDOM and not in the global context.

### To check the routing table using the CLI

```
get router info routing-table all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
 area
* - candidate default
```

```

S* 0.0.0.0/0 [10/0] via 192.168.183.254, port2
S 1.0.0.0/8 [10/0] via 192.168.183.254, port2
S 2.0.0.0/8 [10/0] via 192.168.183.254, port2
C 10.142.0.0/23 is directly connected, port3
B 10.160.0.0/23 [20/0] via 10.142.0.74, port3, 2d18h02m
C 192.168.182.0/23 is directly connected, port2

```

Examining an entry from the routing table above:

```

B 10.160.0.0/23 [20/0] via 10.142.0.74, port3, 2d18h02m

```

<b>B</b>	BGP. The routing protocol used.
<b>10.160.0.0/23</b>	The destination of this route including netmask.
<b>[20/0]</b>	20 indicates and administrative distance of 20 out of a range of 0 to 255.  0 is an additional metric associated with this route, such as in OSPF
<b>10.142.0.74</b>	The gateway, or next hop.
<b>port3</b>	The interface used by this route.
<b>2d18h02m</b>	How old this route is, in this case almost three days old.

## Perform a sniffer trace

When troubleshooting network traffic, it helps to look inside the headers of packets to determine if they are traveling along the route you expect. Packet sniffing can also be called a network tap, packet capture, or logic analyzing.



If your Carrier-enabled FortiGate unit has NP interfaces that are offloading traffic, this will change the sniffer trace. Before performing a trace on any NP interfaces, disable offloading on those interfaces.

## What can sniffing packets tell you

If you are running a constant traffic application such as ping, packet sniffing can tell you if the traffic is reaching the destination, what the port of entry is on the Carrier-enabled FortiGate unit, if the ARP resolution is correct, and if the traffic is being sent back to the source as expected.

Sniffing packets can also tell you if the Carrier-enabled FortiGate unit is silently dropping packets for reasons such as RPF (Reverse Path Forwarding), also called Anti Spoofing. This prevents an IP packet from being forwarded if its source IP address either does not belong to a locally attached subnet (local interface), or be a hop on the routing between the FortiOS Carrier and another source (static route, RIP, OSPF, BGP). Note that RPF can be disabled by turning on asymmetric routing in the CLI (`config system setting, set asymmetric enable`),

however this will disable stateful inspection on the Carrier-enabled FortiGate unit and consequently cause many features to be turned off.



If you configure virtual IP addresses on your Carrier-enabled FortiGate unit, the unit will use those addresses in preference to the physical IP addresses. If not configured properly, secondary IP addresses can cause a broadcast storm. You will notice the secondary address being preferred when you are sniffing packets because all the traffic will be using the virtual IP addresses. This is due to the ARP update that is sent out when the VIP address is configured.

## How to sniff packets

The general form of the internal FortiOS packet sniffer command is:

```
diag sniffer packet <interface_name> <'filter'> <verbose> <count>
```

To stop the sniffer, type CTRL+C.

<b>&lt;interface_name&gt;</b>	The name of the interface to sniff, such as <code>port1</code> or <code>internal</code> . This can also be <code>any</code> to sniff all interfaces.
<b>&lt;'filter'&gt;</b>	What to look for in the information the sniffer reads. <code>none</code> indicates no filtering, and all packets will be displayed as the other arguments indicate. The filter must be inside single quotes (').
<b>&lt;verbose&gt;</b>	The level of verbosity as one of: <b>1</b> - print header of packets <b>2</b> - print header and data from IP of packets <b>3</b> - print header and data from Ethernet of packets
<b>&lt;count&gt;</b>	The number of packets the sniffer reads before stopping. If you don't put a number here, the sniffer will run forever until you stop it with <code>&lt;CTRL C&gt;</code> .

For a simple sniffing example, enter the CLI command `diag sniffer packet port1 none 1 3`. This will display the next 3 packets on the `port1` interface using no filtering, and using verbose level 1. At this verbosity level you can see the source IP and port, the destination IP and port, action (such as `ack`), and sequence numbers.

In the output below, port 443 indicates these are HTTPS packets, and 172.20.120.17 is both sending and receiving traffic.

```
Head_Office_620b # diag sniffer packet port1 none 1 3
interfaces=[port1]
filters=[none]
0.545306 172.20.120.17.52989 -> 172.20.120.141.443: psh 3177924955
ack 1854307757

0.545963 172.20.120.141.443 -> 172.20.120.17.52989: psh 1854307757
ack 3177925808

0.562409 172.20.120.17.52988 -> 172.20.120.141.443: psh 4225311614
ack 3314279933
```

## Generate specific packets to test the network

If some packets are being delivered as expected while others are not, or after you believe you have fixed the problem, it is a good idea to generate specific traffic to test your network.

For example if you discover through log messages and packet sniffing that Create PDP Context Request messages are not being delivered between two SGSNs, you can generate those specific messages on your network to confirm they are the problem, and later that you have solved the problem and they are now being delivered as expected.

This step requires a third party traffic generation tool, either hardware or software. This is not supported by Fortinet.

# Chapter 5 Compliance

This handbook chapter contains the following sections:

[Configuring FortiGate units for PCI DSS compliance](#) explains the Payment Card Industry Data Security Standard (PCI DSS). It provides information about configuring your network and FortiGate unit to help you comply with PCI DSS requirements.

# Configuring FortiGate units for PCI DSS compliance

This chapter provides information about configuring your network and FortiGate unit to help you comply with PCI DSS requirements. The following topics are included in this section:

- [Introduction to PCI DSS](#)
- [Network topology](#)
- [Security policies for the CDE network](#)
- [Wireless network security](#)
- [Protecting stored cardholder data](#)
- [Protecting communicated cardholder data](#)
- [Protecting the CDE network from viruses](#)
- [Monitoring the network for vulnerabilities](#)
- [Restricting access to cardholder data](#)
- [Controlling access to the CDE network](#)

## Introduction to PCI DSS

The primary source of information for your PCI DSS compliance program is the *Payment Card Industry (PCI) Data Security Standard* itself. Version 3.0 of the standard was released in November 2013 and is active from January 1, 2014 to December 31, 2017. The following is only a brief summary of PCI DSS.

### What is PCI DSS?

The Payment Card Industry Data Security Standard (PCI DSS) sets data handling requirements for organizations that hold, process, or exchange cardholder information.

### What is the Customer Data Environment

Throughout the PCI DSS requirements, there are references to the Customer Data Environment (CDE). The CDE is the computer environment wherein cardholder data is transferred, processed, or stored, and any networks or devices directly connected to that environment.

### PCI DSS objectives and requirements

PCI DSS consists of 6 control objectives and 12 requirements.



**Table 39:** PCI DSS Control Objectives and Requirements

Control Objective	Requirement	Fortinet Solution
Build and Maintain a Secure Network	1) Install and maintain a firewall configuration to protect cardholder data	FortiGate firewall functionality. See <a href="#">“Security policies for the CDE network”</a> on page 773.
	2) Do not use vendor-supplied defaults for system passwords and other security parameters	FortiDB vulnerability assessment and auditing FortiScan OS vulnerability management FortiWeb web application password checking See <a href="#">“Password complexity and change requirements”</a> on page 779.
Protect Cardholder Data	3) Protect stored cardholder data	FortiDB vulnerability assessment and monitoring FortiWeb web application firewall See <a href="#">“Protecting stored cardholder data”</a> on page 776.
	4) Encrypt transmission of cardholder data across open, public networks	FortiGate IPsec VPN. See <a href="#">“Protecting communicated cardholder data”</a> on page 776.

**Table 39:** PCI DSS Control Objectives and Requirements (Continued)

Control Objective	Requirement	Fortinet Solution
Maintain a Vulnerability Management Program	5) Protect all systems against malware and regularly update anti-virus software or programs	FortiGate integrated AV FortiClient integrated AV FortiMobile integrated AV FortiMail integrated AV FortiGuard automated AV updates See <a href="#">“Protecting the CDE network from viruses”</a> on page 777.
	6) Develop and maintain secure systems and applications	FortiDB vulnerability assessment, auditing and monitoring FortiWeb web application security FortiScan OS vulnerability management See <a href="#">“Monitoring the network for vulnerabilities”</a> on page 778.
Implement Strong Access Control Measures	7) Restrict access to cardholder data by business need to know	FortiDB vulnerability assessment, auditing and monitoring. See <a href="#">“Restricting access to cardholder data”</a> on page 779.
	8) Identify and authenticate access to system components	FortiGate integrated database or hooks to Active Directory. See <a href="#">“Controlling access to the CDE network”</a> on page 779.
	9) Restrict physical access to cardholder data	Fortinet professional services in partnership with partner solutions

**Table 39:** PCI DSS Control Objectives and Requirements (Continued)

Control Objective	Requirement	Fortinet Solution
Regularly Monitor and Test Networks	10) Track and monitor all access to network resources and cardholder data	FortiDB auditing and monitoring  FortiAnalyzer event reporting  See <a href="#">“Monitoring the network for vulnerabilities” on page 778.</a>
	11) Regularly test security systems and processes	FortiDB vulnerability assessment  FortiScan OS vulnerability management. See <a href="#">“Monitoring the network for vulnerabilities” on page 778.</a>
Maintain an Information Security Policy	12) Maintain a policy that addresses information security for all personnel	FortiManager security policy management appliance

This chapter describes how the FortiGate unit’s features can help your organization to be compliant with PCI DSS. Requirements that the FortiGate cannot enforce need to be met through organization policies with some means determined for auditing compliance.

Be sure to read the section, [“Wireless guidelines”](#), below. Even if your organization does not use wireless networking, PCI DSS requires you to verify periodically that wireless networking has not been introduced into the CDE.

### Wireless guidelines

While wired networks usually connect fixed known workstations, wireless networks are more dynamic, introducing a different set of security concerns.

Even if your organization does not use wireless networking, PCI DSS requires you to verify periodically that unauthorized wireless networking has not been introduced into the CDE. Wireless networking could be introduced quite casually by adding a wireless device to a PC on the CDE network.

For all PCI DSS networks, whether they use wireless technology or not, the following requirement applies:

- Test for the presence of wireless access points (802.11), and detect and identify all authorized and unauthorized wireless access points on a quarterly basis (11.1)

If your organization uses wireless networking outside the CDE network and the firewall prevents communication with the CDE network, the wireless network is outside the PCI DSS scope, but the firewall configuration must meet PCI DSS requirements.

If your organization uses wireless networking inside the CDE network, the wireless network is within the PCI DSS scope. For information about wireless network requirements, see [“Wireless network security” on page 774.](#)


## Network topology

The cardholder data environment must be protected against unauthorized access from the Internet and from other networks in your organization. FortiGate unit firewall functionality provides tight control over the traffic that can pass between the following network interfaces:

- Internet
- CDE wired LAN
- CDE wireless LAN
- Other internal networks

Figure 154 shows how the Customer Data Environment can be delineated in a typical network.

**Figure 154:**Enterprise network with a customer data environment



## Internet

The FortiGate unit has at least one network interface connected to the Internet. If your organization uses more than one Internet service provider, there could be additional network interfaces that function as a route to the Internet.

## The CDE wired LAN

The CDE network typically contains point-of-sale (POS) terminals, databases, and servers. The only security policies between the CDE network and the Internet should be for encrypted connections. For remote point-of-sale terminals or off-site databases, VPN connections are required and they should use strong encryption. For a web server that handles online purchases, only HTTPS (SSL or TLS) connections can be permitted. The security policies that enable these connections should have the narrowest possible definitions for source address, destination address and service.

PCI DSS does not require the CDE network to be isolated from the rest of your corporate LAN. But isolating the CDE network reduces the scope of required data protection measures and may reduce the scope of PCI DSS assessments that are periodically required.

## The CDE wireless LAN

Wireless networking is a special issue. Even if you do not use wireless technology you must monitor to ensure that unauthorized wireless access has not been added to the CDE network. For this purpose, [Figure 154](#) shows a FortiAP device in the CDE. The FortiAP device can provide dedicated wireless monitoring, an access point, or both.

A small retail outlet could reduce costs by using a FortiWiFi unit, a FortiGate unit with integrated wireless networking. The FortiWiFi unit would have to be located where it could provide sufficient wireless monitoring (or access point) coverage for the entire premises.

## Other internal networks

Other internal networks such as your office LAN, unless they provide connection to the CDE, are not subject to PCI DSS requirements.

## Security policies for the CDE network

The FortiGate unit's firewall functionality is ideally suited to PCI DSS requirement 3.0, "Restrict inbound and outbound traffic to that which is necessary for the cardholder data environment." Security policies control the source, destination, and type of traffic passing between networks.

The PCI DSS standard includes requirements to document your network topology and configuration. As part of that requirement, and to assist the auditing of your network, make use of the *Comment* field available in FortiGate security policies. Describe the purpose of each policy.

## Controlling the source and destination of traffic

The source and destination are the first parameters you specify in a security policy. (Go to *Policy > Policy > Policy* and select *Create New*.)

Incoming Interface	<input type="text" value="Click to add..."/>
Source Address	<input type="text" value="Click to add..."/>
Outgoing Interface	<input type="text" value="Click to add..."/>
Destination Address	<input type="text" value="Click to add..."/>

The *Interface* settings depend on network topology. The *Address* settings define the IP addresses to which the policy applies. These should be as narrow as possible, so that only the appropriate hosts are included. For example, if the destination is a server with a single IP address, the named Destination Address should be defined as that single address, not the entire subnet on which the server resides.

Addresses are defined in *Firewall Objects > Address > Addresses*. You can also define a new address by selecting *Create* from either the *Source Address* or *Destination Address* drop-down lists in a security policy. Some addresses will be used in several security policies, so it is best to plan ahead and define the addresses first.


## Controlling the types of traffic in the CDE

The *Service* setting in each security policy determines which types of traffic can pass based on protocol.

You can select a single protocol from the *Service* drop-down list. To add another protocol, select the green “+” button to access the *Service* drop-down list again. If several security policies will need the same list of services, consider creating a named service group. (Go to *Firewall Objects > Service > Groups*.) In the security policy, service groups are available at the bottom of the *Service* drop-down list.

## The default deny policy

All traffic not specifically allowed by a security policy that you create is blocked by the Implicit policy listed at the bottom of the *Policy > Policy > Policy* page.



You cannot delete this policy and you can edit the policy only to enable or disable logging of the traffic that it handles.

## Wireless network security

Scanning for rogue access points is the minimum requirement for wireless security. Even if your organization does not use wireless networking, PCI DSS requires you to verify periodically that wireless networking has not been introduced into the CDE.

If you use wireless networking, the wireless network is only within the PCI DSS scope if it can connect to the CDE.

## On-wire detection of rogue APs

FortiGate units include an “on-wire” detection technique that correlates the SSID MAC addresses of the unknown access points with MAC addresses detected on your wired networks. This helps to differentiate unrelated neighboring APs from security-compromising unauthorized APs connected to your network.

## Setting up rogue access point scanning

A FortiGate unit with a connected FortiAP unit can perform wireless scanning. Each of the FortiAP radios can act as a dedicated monitor or can perform scanning in the background while acting as a wireless access point.

Radio 1 operates in the 2.4GHz band and Radio 2 operates in the 5GHz band. Both bands should be monitored. The FortiAP unit(s) used for scanning must be located within the coverage area that would result if an access point were added to the CDE.

### To configure rogue AP scanning in a custom AP profile

1. Go to *WiFi Controller > WiFi Network > Custom AP Profiles*.  
On some models, the menu is *WiFi & Switch Controller*.
2. Select an existing AP profile and edit it, or select *Create New*.
3. For each radio, select either *Access Point* or *Dedicated Monitor*, as required.

▼ **Radio 1**  
Mode  Disable  Access Point  Dedicated Monitor  
Rogue AP On-Wire Scan   
WIDS Profile default

▼ **Radio 2**  
Mode  Disable  Access Point  Dedicated Monitor  
Background Scan  Disable  Enable  
Rogue AP On-Wire Scan   
WIDS Profile default

4. If you selected *Access Point*, enable *Background Scan*.
5. Select *Rogue AP On-Wire Scan*.
6. If needed, modify other settings.
7. Select *OK*.

### To enable rogue AP scanning for the automatic AP profile

1. Go to *WiFi Controller > WiFi Network > Rogue AP Settings*.
2. Select *Enable Rogue AP Detection*.
3. Select *Enable On-wire Rogue AP Detection Technique* if you want to use that method of distinguishing rogues from neighbors.
4. Select *Apply*.

### Viewing the results of rogue AP scanning

Go to *WiFi Controller > Monitor > Rogue AP* to view information about detected wireless access points.

### Logging the results of rogue AP scanning

To ensure that detection of rogue access points is logged, go to *Log&Report > Log Config > Log Setting* and enable logging for *WiFi activity event*.

In the logs, the *Type* is event and the *Sub Type* is wireless.


## Securing a CDE network WAP

If your wireless network is within PCI DSS scope, it must meet the following requirements:

- Default settings such as SSID and passphrases must be changed.
- Use WPA security, not WEP.
- Log wireless activity.

### Setting wireless security

On FortiGate units, go to *WiFi Controller > WiFi Network > SSID* to configure wireless security settings for either a new or existing virtual access point.

WiFi Settings	
SSID	<input type="text" value="exampleco"/>
Security Mode	<input type="text" value="WPA/WPA2-Enterprise"/>
Data Encryption	<input checked="" type="radio"/> AES <input type="radio"/> TKIP <input type="radio"/> TKIP-AES
Authentication	<input type="radio"/> RADIUS Server <input checked="" type="radio"/> Usergroup
	<input type="text" value="pos_term"/> 

The default SSID for the FortiAP is “fortinet”. You must change this.

The *Security Mode* **must** be set to one of the WPA/WPA2 modes. Both WPA or WPA2 clients can be served. In the CLI, you can optionally select exclusively WPA or WPA2 operation.

AES is stronger *Data Encryption* than TKIP.

WPA/WPA2-Enterprise *Authentication* uses separate logon credentials for each user. Either FortiGate user group security or an external RADIUS server performs the authentication. Optionally, certificate-based security can also be applied. WPA/WPA2-Personal authentication requires a single pre-shared key that is used by all clients and is thus less secure.

For detailed information about wireless access points, see the *Deploying Wireless Networks* chapter of this FortiOS Handbook.

### Logging wireless network activity

To ensure that wireless network activity is logged, go to *Log&Report > Log Config > Log Settings* and enable logging for *WiFi activity event*. In the logs, the *Type* is event and the *Sub Type* is wireless.

## Protecting stored cardholder data

The Fortinet FortiDB and FortiWeb products can provide security for your sensitive cardholder data.

The Fortinet Database Security (FortiDB) device provides vulnerability assessment, database activity monitoring, auditing and monitoring. For more information about this product, see the Fortinet web site, [www.fortinet.com](http://www.fortinet.com).

The Fortinet FortiWeb Web Application Firewall deployed in front of public-facing web applications protects Web applications, databases, and the information exchanged between them. In particular, it addresses the PCI DSS requirements 6.5 and 6.6 regarding web application vulnerabilities such as cross-site scripting, SQL injection, and information leakage. For more information about this product, see the Fortinet web site, [www.fortinet.com](http://www.fortinet.com).

## Protecting communicated cardholder data

If cardholder data must be communicated over an untrusted network, such as the Internet, use the FortiGate unit’s IPsec VPN capability to exchange the data securely. If you support customer on-line transactions, use HTTPS (SSL or TLS encryption) for security. The relevant PCI DSS requirement is:

- Use strong cryptography and security protocols (for example, SSL/TLS, IPSEC, SSH, etc.) to safeguard sensitive cardholder data during transmission over open, public networks. (4.1)

This does not prescribe particular cryptography, but it can be interpreted as a requirement to follow industry best practices.



## Configuring IPsec VPN security

The security considerations for IPsec VPNs are encryption and authentication.

### Encryption

Go to *VPN > IPsec > Auto Key (IKE)* to configure an IPsec VPN. In both Phase 1 and Phase 2 parts of the configuration, you select the encryption to use.

1 - Encryption  Authentication   
2 - Encryption  Authentication

These are advanced settings, overriding defaults that are not necessarily the strongest algorithms. VPNs negotiate over standards, so you can list multiple proposed algorithms. The VPN will use the strongest encryption that both ends support.

Choose strong encryption. The available encryption algorithms in descending order of strength are AES256, AES192, AES128, 3DES, DES. DES encryption is the weakest with only a 64-bit key and does not meet the 80-bit key length minimum that PCI DSS requires. NULL means no encryption and must not be used.

The message digest (authentication) algorithms in descending order of strength are SHA512, SHA384, SHA256, SHA1 and MD5. MD5 is particularly weak and should be avoided. NULL means no message digest and must not be used.

### Authentication

VPN peers authenticate each other before establishing a tunnel. FortiGate units support two different authentication methods: pre-shared key and RSA signature (certificate). Certificates provide the best security. PCI DSS does not prohibit pre-shared keys, but you should limit access to the keys to the personnel who are responsible for the FortiGate units or other equipment at either end of the VPN.

## Configuring SSL VPN security

The SSL VPN configuration includes a choice of encryption algorithm. Go to *VPN > SSL > Config*. The *Default* selection, *RC4 (128 bits)* is acceptable, but the *High* option, *AES (128/256 bits)* and *3DES* is more secure. The *Low* option, *RC4 (64 bits)*, *DES* and *higher* does not meet PCI DSS requirements.

## Protecting the CDE network from viruses

PCI DSS requires the use of regularly updated antivirus protection. The antivirus functionality of the FortiGate unit protects both the FortiGate unit and the networks it manages. Workstations on these networks can be protected using FortiClient Endpoint Security. Both FortiGate and FortiClient antivirus protection can receive updates from Fortinet's FortiGuard service. Workstations can also use third-party antivirus applications with update services.

The FortiGate unit can enforce the use of antivirus software, denying unprotected workstations access to the network.

### Enabling FortiGate antivirus protection

The antivirus profile must apply AV scanning to all protocols. You also need to enable SSL inspection to include secure protocols in antivirus scanning. The extended AV database contains the largest number of virus signatures.

### To create the antivirus profile

1. Go to *Security Profiles > Antivirus > Profiles*.
2. Edit the *default* predefined profile or select *Create New*.
3. Set *Inspection Mode* to *Proxy*.
4. Ensure that all *Virus Scan and Removal* check boxes are selected.
5. Select *Apply*.

### To enable SSL inspection

1. Go to *Policy > Policy > SSL Inspection*.
2. Check that the *Enable* check box is selected for every protocol and then select *Apply*.

### To select the extended antivirus database

The antivirus database is selectable using the CLI:

```
config antivirus settings
 set default-db extended
end
```

For detailed information about the Antivirus feature, see the *Security Profiles* chapter of this FortiOS Handbook.

## Configuring antivirus updates

On the system dashboard, check the *License Information* widget. The *Support Contract* section should show *Valid Contract* and the contract expiry date. If your FortiGate unit is not registered, you need to visit the Fortinet Support web page (<http://support.fortinet.com/>) to register. Go to *Product Registration* and follow the instructions.

In the *FortiGuard Services* section, check the *Antivirus* field. If the service is unreachable, see the online Help for information about troubleshooting your connectivity to FortiGuard Services.

## Enforcing firewall use on endpoint PCs

PCI DSS requires you to “Install personal firewall software on any mobile and/or employee-owned devices that connect to the Internet when outside the network (for example, laptops used by employees), and which are also used to access the network. (1.4)” Consider using the *Endpoint Control* feature of the FortiGate unit to enforce use of this software.

## Monitoring the network for vulnerabilities

There are several tools that can assist you in monitoring your network for vulnerabilities and provide evidence to the PCI DSS auditor of such monitoring.

### FortiGate logs

FortiGate units can be configured to send logs to FortiAnalyzer unit. In a larger network, this enables you to collect log information in a central location from several FortiGate units.

### Using the FortiOS Network Vulnerability Scan feature

As part of its security features, FortiGate units provide a *Network Vulnerability Scan*. You define assets to monitor, such as servers, workstations, or point-of-sale terminals. Then, the FortiGate unit scans those devices on a regular schedule. The scan checks TCP and UDP ports against a

list of known vulnerabilities provided by FortiGuard Services. Scan settings determine how many of the ports are checked. Optionally, all ports are scanned.

To set up vulnerability scanning, go to *User & Device > Vulnerability Scan > Scan Definition*. To view scan logs, go to *User & Device > Vulnerability Scan > Vulnerability Result*. For more information, see Vulnerability Scan in this FortiOS Handbook.

## Monitoring with other Fortinet products

In addition to your FortiGate unit and its FortiOS firmware, there are several other Fortinet products that can assist your organization to comply with PCI DSS requirements.

### Fortinet Database Security (FortiDB)

A FortiDB appliance or FortiDB software can provide vulnerability scanning and activity monitoring for your databases. For more information, see the *FortiDB User Guide*.

### FortiScan Vulnerability and Compliance Management platform

The FortiScan Vulnerability and Compliance Management (VCM) platform combines a FortiScan appliance with FortiScan agent software to monitor your network assets such as servers, workstations, or point-of-sale terminals. This system can perform vulnerability scans and apply software patches provided by the software vendors. The scan profiles include a predefined one for PCI DSS. The FortiScan appliance produces compliance reports detailing the results of the vulnerability scan. For more information, see the *FortiScan Administration Guide*.

### FortiWeb Web Application Security

If your organization engages in e-Commerce, you can use FortiWeb Application Security to protect your web servers against attack. The FortiWeb application protects against HTTP and XML-based attacks, guards against attempts to deface your web sites, and scans web servers for vulnerabilities. For more information, see the *FortiWeb Web Application Security Administration Guide*.

## Restricting access to cardholder data

In addition to security policies and authentication governing access to the CDE, you can deploy the Fortinet Database Security (FortiDB) device, which provides vulnerability assessment, database activity monitoring, auditing and monitoring. For more information about this product, see the Fortinet web site, [www.fortinet.com](http://www.fortinet.com).

## Controlling access to the CDE network

PCI DSS requires each user to be uniquely identified and authenticated. On the FortiGate unit, this applies to administrators and to users of SSL VPN and IPsec VPNs.

### Password complexity and change requirements

By default, the FortiGate unit admin account has no password. Be sure to define a password.

PCI DSS password requirements are

- Require a minimum password length of at least seven characters. (8.2.3)
- Use passwords containing both numeric and alphabetic characters. (8.2.3)
- Change user passwords at least every 90 days. (8.2.4)

To facilitate creation of compliant administrator passwords, you can set a password policy. Go to *System > Admin > Settings*. In the *Password Policy* section, enter the following and then select OK at the bottom of the page.

**Enable Password Policy**

Minimum Length  (8-64 characters)

Must Contain

Upper Case Letters  Lower Case Letters

Numerical Digits  Non-alphanumeric Letters

Apply Password Policy to  Admin Password  IPsec Preshared Key

Enable Password Expiration   (days)

<b>Enable</b>	Select the check box.
<b>Minimum Length</b>	8 or more. (Field does not accept a value less than 8.)
<b>Must Contain</b>	At minimum, set a required number of <i>Numerical Digits</i> and either <i>Upper Case Letters</i> or <i>Lower Case Letters</i> . Also setting a required number of <i>Non-alphanumeric Letters</i> is acceptable.
<b>Apply Password Policy to</b>	Select <i>Admin Password</i> .
<b>Enable Password Expiration</b>	Set to 90 days or less. The default is 90 days.

Note that the FortiGate password policy does not apply to user passwords. Both password complexity and password expiry for users would need to be addressed by making them a policy in your organization.

## Password non-reuse requirement

PCI DSS requires that passwords are not re-used to satisfy the change requirement:

“Do not allow an individual to submit a new password that is the same as any of the last four passwords he or she has used.” (8.2.5)

FortiGate users don't set their own passwords. The super\_admin administrators can and so can admins with appropriate access. There is, however, no FortiGate-based mechanism to enforce non re-use of passwords.

## Administrator lockout requirement

PCI DSS requires a user account lockout for administrators to guard against unauthorized access attempts:

- Limit repeated access attempts by locking out the user ID after not more than six attempts. (8.1.6),
- Set the lockout duration to a minimum of 30 minutes or until administrator enables the user ID. (8.1.7)

You can meet these requirements with the following CLI commands:

```
config system global
 set admin-lockout-threshold 6
 set admin-lockout-duration 1800
end
```

The threshold can be less than 6 and the lockout duration can be more than 1800.

## Administrator timeout requirement

PCI DSS requires:

- If a session has been idle for more than 15 minutes, require the user to re-enter the password to reactivate the terminal. (8.1.8)

By default, the idle timeout is five minutes. You can go to *System > Admin > Settings* and change the *Idle Timeout* timeout to any value up to the permitted value of 15 minutes.

## Administrator access security

To accommodate the requirement for unique identification of each user, the generic admin account should either be assigned to only one administrator or not used at all. You can create an administrator account for each administrator in *System > Admin > Administrators*.

If an administrator always works from the same workstation, consider using the Trusted Host feature. The administrator will be able to log in only from a trusted IP address. You can define up to three trusted IP addresses per administrator.

Administrative access must also be enabled per network interface. Go to *System > Network > Interface* to edit the interface settings. Enable administrative access only on interfaces where you would expect the administrator to connect. Allow only secure connection protocols, HTTPS for web-based access, SSH for CLI access.

## Remote access security

For remote access, PCI DSS requires two-factor authentication: a password and some other authentication, such as a smart token or certificate. This applies to employees, administrators, and third parties.

### SSL VPN users

For SSL VPN users, implement two-factor authentication by requiring users to have a certificate in addition to the correct password. Go to *VPN > SSL > Config*, enable *Require Client Certificate*, and then select *Apply*. For more information, see the SSL VPN chapter of this FortiOS Handbook.

## IPsec VPN users

If users access your network using an IPsec VPN, you can implement two-factor authentication by enabling extended authentication (XAUTH). This requires the user to enter a password in addition to the VPN authentication provided by the certificate or pre-shared key. As PCI DSS requires each user to have a unique identifier, you should already have user accounts and user groups defined.

### To configure XAUTH on your VPN

1. Go to *VPN > IPsec > Auto Key (IKE)* and edit your Phase 1 configuration.
2. Select *Advanced*.
3. In *XAUTH*, select *Enable as Server*.  
*Enable as Server* is available only if *Remote Gateway* is *Dialup User*.
4. Set *Server Type* to *PAP*, *CHAP*, or *AUTO* as appropriate.
5. Select the *User Group* to which the VPN users belong.
6. Select *OK*.

# Chapter 6 Deploying Wireless Networks

## for FortiOS 5.0

This FortiOS Handbook chapter contains the following sections:

[Introduction to wireless networking](#) explains the basic concepts of wireless networking and how to plan your wireless network.

[Configuring a WiFi LAN](#) explains how to set up a basic wireless network, prior to deploying access point hardware.

[Access point deployment](#) explains how to deploy access point hardware and add it to your wireless network configuration.

[Wireless Mesh](#) explains how to configure a WiFi network where access points are connected to the WiFi controller wirelessly instead of by Ethernet.

[WiFi-Ethernet Bridge Operation](#) shows how to use the FortiAP WiFi-Ethernet bridge feature.

[Protecting the WiFi Network](#) explains the Wireless Intrusion Detection System (WIDS).

[Wireless network monitoring](#) explains how to monitor your wireless clients and how to monitor other wireless access points, potentially rogues, in your coverage area.

[Configuring wireless network clients](#) explains how to configure typical wireless clients to work with a WPA-Enterprise protected network.

[Wireless network examples](#) provides two examples. The first is a simple WiFi network using automatic configuration. The second is a more complex example of a business with two WiFi networks, one for employees and another for guests or customers.

[Using a FortiWiFi unit as a client](#) explains how to use a FortiWiFi unit as a wireless client to connect to other WiFi networks. This connection can take the place of an Ethernet connection where wired access to a network or to the Internet is not available.

[Reference](#) provides information about WiFi radio channels.

# Introduction to wireless networking

This chapter introduces some concepts you should understand before working with wireless networks, describes Fortinet's wireless equipment, and then describes the factors you need to consider in planning deployment of a wireless network.

The following topics are included in this section:

- [Wireless concepts](#)
- [Security](#)
- [Authentication](#)
- [Wireless networking equipment](#)
- [Deployment considerations](#)
- [Automatic Radio Resource Provisioning](#)

## Wireless concepts

Wireless networking is radio technology, subject to the same characteristics and limitations as the familiar audio and video radio communications. Various techniques are used to modulate the radio signal with a data stream.

### Bands and channels

Depending on the wireless protocol selected, you have specific channels available to you, depending on what region of the world you are in.

- IEEE 802.11b and g protocols provide up to 14 channels in the 2.400-2.500 GHz Industrial, Scientific and Medical (ISM) band.
- IEEE 802.11a,n (5.150-5.250, 5.250-5.350, 5.725–5.875 GHz, up to 16 channels) in portions of Unlicensed National Information Infrastructure (U-NII) band

Note that the width of these channels exceeds the spacing between the channels. This means that there is some overlap, creating the possibility of interference from adjacent channels, although less severe than interference on the same channel. Truly non-overlapping operation requires the use of every fourth or fifth channel, for example ISM channels 1, 6 and 11.

The capabilities of your wireless clients is the deciding factor in your choice of wireless protocol. If your clients support it, 5GHz protocols have some advantages. The 5GHz band is less used than 2.4GHz and its shorter wavelengths have a shorter range and penetrate obstacles less. All of these factors mean less interference from other access points, including your own.

When configuring your WAP, be sure to correctly select the Geography setting to ensure that you have access only to the channels permitted for WiFi use in your part of the world.

For detailed information about the channel assignments for wireless networks for each supported wireless protocol, see [“Wireless radio channels” on page 883](#).

### Power

Wireless LANs operate on frequencies that require no license but are limited by regulations to low power. As with other unlicensed radio operations, the regulations provide no protection against interference from other users who are in compliance with the regulations.



Power is often quoted in dBm. This is the power level in decibels compared to one milliwatt. 0dBm is one milliwatt, 10dBm is 10 milliwatts, 27dBm, the maximum power on Fortinet FortiAP equipment, is 500 milliwatts. The FortiGate unit limits the actual power available to the maximum permitted in your region as selected by the WiFi controller country setting.

Received signal strength is almost always quoted in dBm because the received power is very small. The numbers are negative because they are less than the one milliwatt reference. A received signal strength of -60dBm is one millionth of a milliwatt or one nanowatt.

## Antennas

Transmitted signal strength is a function of transmitter power and antenna gain. Directional antennas concentrate the signal in one direction, providing a stronger signal in that direction than would an omnidirectional antenna.

FortiWiFi units have detachable antennas. However, these units receive regulatory approvals based on the supplied antenna. Changing the antenna might cause your unit to violate radio regulations.

## Security

There are several security issues to consider when setting up a wireless network.

### Whether to broadcast SSID

Users who want to use a wireless network must configure their computers with the wireless service set identifier (SSID) or network name. Broadcasting the SSID makes connection to a wireless network easier because most wireless client applications present the user with a list of network SSIDs currently being received. This is desirable for a public network.

To obscure the presence of a wireless network, do not broadcast the SSID. This does not prevent attempts at unauthorized access, however, because the network is still detectable with wireless network “sniffer” software.

### Encryption

Wireless networking supports the following security modes for protecting wireless communication, listed in order of increasing security.

**None** — Open system. Any wireless user can connect to the wireless network.

**WEP64** — 64-bit Web Equivalent Privacy (WEP). This encryption requires a key containing 10 hexadecimal digits.

**WEP128** — 128-bit WEP. This encryption requires a key containing 26 hexadecimal digits.

**WPA** — 256-bit Wi-Fi Protected Access (WPA) security. This encryption can use either the TKIP or AES encryption algorithm and requires a key of either 64 hexadecimal digits or a text phrase of 8 to 63 characters. It is also possible to use a RADIUS server to store a separate key for each user.

**WPA2** — WPA with security improvements fully meeting the requirements of the IEEE 802.11i standard. Configuration requirements are the same as for WPA.

For best security use the WPA2 with AES encryption and a RADIUS server to verify individual credentials for each user. WEP, while better than no security at all, is an older algorithm that is easily compromised. With either WEP or WAP, changing encryption passphrases on a regular basis further enhances security.

## Separate access for employees and guests

Wireless access for guests or customers should be separate from wireless access for your employees. This does not require additional hardware. Both FortiWiFi units and FortiAP units support multiple wireless LANs on the same access point. Each of the two networks can have its own SSID, security settings, firewall policies, and user authentication.

A good practice is to broadcast the SSID for the guest network to make it easily visible to users, but not to broadcast the SSID for the employee network.

Two separate wireless networks are possible because multiple virtual APs can be associated with an AP profile. The same physical APs can provide two or more virtual WLANs.

## Captive portal

As part of authenticating your users, you might want them to view a web page containing your acceptable use policy or other information. This is called a captive portal. No matter what URL the user initially requested, the portal page is returned. Only after authenticating and agreeing to usage terms can the user access other web resources.

For information about setting up a captive portal, see [“Captive Portal security” on page 803](#).

## Power

Reducing power reduces unwanted coverage and potential interference to other WLANs. Areas of unwanted coverage are a potential security risk. There are people who look for wireless networks and attempt to access them. If your office WLAN is receivable out on the public street, you have created an opportunity for this sort of activity.

## Monitoring for rogue APs

It is likely that there are APs available in your location that are not part of your network. Most of these APs belong to neighboring businesses or homes. They may cause some interference, but they are not a security threat. There is a risk that people in your organization could connect unsecured WiFi-equipped devices to your wired network, inadvertently providing access to unauthorized parties. The optional On-Wire Rogue AP Detection Technique compares MAC addresses in the traffic of suspected rogues with the MAC addresses on your network. If wireless traffic to non-Fortinet APs is also seen on the wired network, the AP is a rogue, not an unrelated AP.

Decisions about which APs are rogues are made manually on the Rogue AP monitor page. For detailed information about monitoring rogue APs, see [“Monitoring rogue APs” on page 847](#).

## Suppressing rogue APs

When you have declared an AP to be a rogue, you have the option of suppressing it. To suppress an AP, the FortiGate WiFi controller sends reset packets to the rogue AP. Also, the MAC address of the rogue AP is blocked in the firewall policy. You select the suppression action on the Rogue AP monitor page. For more information, see [“Suppressing rogue APs” on page 851](#).



Rogue suppression is available only when there is a radio dedicated to scanning. It will not function during background scanning.

## Wireless Intrusion Detection (WIDS)

You can create a WIDS profile to enable several types of intrusion detection:

- Unauthorized Device Detection
- Rogue/Interfering AP Detection
- Ad-hoc Network Detection and Containment
- Wireless Bridge Detection
- Misconfigured AP Detection
- Weak WEP Detection
- Multi Tenancy Protection
- MAC OUI Checking

## Authentication

Wireless networks usually require authenticated access. FortiOS authentication methods apply to wireless networks the same as they do to wired networks because authentication is applied in the firewall policy.

The types of authentication that you might consider include:

- user accounts stored on the FortiGate unit
- user accounts managed and verified on an external RADIUS, LDAP or TACACS+ server
- Windows Active Directory authentication, in which users logged on to a Windows network are transparently authenticated to use the wireless network.

This Wireless chapter of the FortiOS Handbook will provide some information about each type of authentication, but more detailed information is available in the Authentication chapter.

What all of these types of authentication have in common is the use of user groups to specify who is authorized. For each wireless LAN, you will create a user group and add to it the users who can use the WLAN. In the identity-based firewall policies that you create for your wireless LAN, you will specify this user group.

Some access points, including FortiWiFi units, support MAC address filtering. You should not rely on this alone for authentication. MAC addresses can be “sniffed” from wireless traffic and used to impersonate legitimate clients.

## Wireless networking equipment

Fortinet produces two types of wireless networking equipment:

- [FortiWiFi units](#), which are FortiGate units with a built-in wireless access point/client
- [FortiAP units](#), which are wireless access points that you can control from any FortiGate unit that supports the WiFi Controller feature.

### FortiWiFi units

A FortiWiFi unit can:

- Provide an access point for clients with wireless network cards. This is called Access Point mode, which is the default mode.

or

- Connect the FortiWiFi unit to another wireless network. This is called Client mode. A FortiWiFi unit operating in client mode can only have one wireless interface.

or

- Monitor access points within radio range. This is called Monitoring mode. You can designate the detected access points as Accepted or Rogue for tracking purposes. No access point or client operation is possible in this mode. But, you can enable monitoring as a background activity while the unit is in Access Point mode.

FortiWiFi unit capabilities differ by model as follows:

**Table 40:** FortiWiFi model capabilities

<b>Model</b>	<b>Radio</b>	<b>Simultaneous SSIDs</b>
20C	802.11 b/g/n 2.4GHz 802.11 a/n 5GHz	7 for AP, 1 for monitoring
30B	802.11 b/g 2.4GHz	7 for AP, 1 for monitoring
40C	802.11 b/g/n 2.4GHz 802.11 a/n 5GHz	7 for AP, 1 for monitoring
50B	802.11 b/g 2.4GHz	7 for AP, 1 for monitoring
60B	802.11 b/g 2.4GHz 802.11 a 5GHz	7 for AP, 1 for monitoring
60C	802.11 b/g/n 2.4GHz 802.11 a/n 5GHz	7 for AP, 1 for monitoring
80/81CM	802.11 b/g/n 2.4GHz 802.11 a/n 5GHz	7 for AP, 1 for monitoring

## FortiAP units

FortiAP series wireless access points are controlled by a FortiGate unit over Ethernet. Capabilities differ by model as follows:

**Table 41:** FortiAP model capabilities

<b>Model</b>	<b>Radio 1</b>	<b>Radio 2</b>	<b>Simultaneous SSIDs</b>
210B (indoor)	802.11 b/g/n 2.4GHz 802.11 a/n 5GHz	N/A	7 for AP, 1 for monitoring
220A (indoor)	802.11 b/g/n 2.4GHz	802.11 a/n 5GHz	14 for AP, 2 for monitoring

**Table 41:** FortiAP model capabilities

220B (indoor)	802.11 b/g/n 2.4GHz 802.11 a/n 5GHz	802.11 b/g/n 2.4GHz	14 for AP, 2 for monitoring
222B (outdoor)	802.11 b/g/n 2.4GHz	802.11 a/n 5GHz	14 for AP, 2 for monitoring

Dual-band radios can function as an AP on either band or as a dual-band monitor. The monitoring function is also available during AP operation if Background Scan is enabled in the custom AP profile for the device.

## Deployment considerations

Several factors need to be considered when planning a wireless deployment.

### Types of wireless deployment

This Handbook chapter describes two main types of wireless deployment: single WAP and multiple WAP. You will know which type of deployment you need after you have evaluated the coverage area environment.

### Deployment methodology

1. Evaluate the coverage area environment.
2. Position access point(s).
3. Select access point hardware.
4. Install and configure the equipment.
5. Test and tune the network.

### Evaluating the coverage area environment

Consider the following factors:

- **Size of coverage area** — Even under ideal conditions, reliable wireless service is unlikely beyond 100 metres outdoors or 30 metres indoors. Indoor range can be further diminished by the presence of large metal objects that absorb or reflect radio frequency energy. If wireless users are located on more than one floor of a building, a minimum of one WAP for each floor will be needed.
- **Bandwidth required** — Wireless interface data rates are between 11 and 150 Mb/s, depending on the 802.11 protocol that is used. This bandwidth is shared amongst all users of the wireless data stream. If wireless clients run network-intensive applications, fewer of them can be served satisfactorily by a single WAP.
- Note that on some FortiWiFi units you can define up to four wireless interfaces, increasing the available total bandwidth.
- **Client wireless capabilities** — The 802.11n protocol provides the highest data rates and has channels in the less interference-prone 5GHz band, but it is supported only on the latest consumer devices. The 802.11g protocol is more common but offers lower bandwidth. Some older wireless client equipment supports only 802.11b with a maximum data rate of 11Mb/s. WAP radios support the protocol that you select with backward compatibility to older modes. For example, if you select 802.11n, clients can also connect using 802.11g or 802.11b.

The most important conclusion from these considerations is whether more than one WAP is required.

### Positioning access points


When placing the access point, your main concern is providing a strong signal to all users. A strong signal ensures a fast connection and efficient data transfer. A weaker signal means a greater chance of data transmission errors and the need to re-send information, slowing down data transfer.

Consider the following guidelines when placing access points:

- Physical barriers can impede the radio signals. Solid objects such as walls, furniture and people absorb radio waves, weakening the signal. Be aware of the physical barriers in your office space that may reduce a signal. If there is enough physical interference, you may encounter dead spots that receive no signal.
- Ensure the access point is located in a prominent location within a room for maximum coverage, rather than in a corner.
- Construction materials used in a building can also weaken radio signals. Rooms with walls of concrete or metal can affect the signal strength.

If you cannot avoid some of these impediments due to the shape of the office or building materials used, you may need to use multiple access points to help distribute the radio signal around the room. [Figure 155](#) shows how positioning two FortiAP-220A units within a uniquely shaped office space helps to distribute signals around the area.

**Figure 155:**Using multiple APs to provide a constant strong signal.



This sample office has washrooms, a stairwell and an elevator shaft in the center of the building, making it impossible to use a single access point effectively. The elevator shaft and multiple metal stalls in the washrooms can cause signal degradation. However, placing access points in diagonally opposite areas of the office provides maximum coverage.

When using multiple access points, set each access point to a different channel to avoid interference in areas where signals from both access points can be received.

### Selecting access point hardware

For a single WAP installation, you could deploy a single FortiWiFi unit. If the site already has a FortiGate unit that supports the WiFi controller feature, adding a FortiAP unit is the most economical solution.

For a multiple WAP deployment you need a FortiGate unit as a WiFi controller and multiple FortiAP units. A FortiWiFi unit can be used as a managed WAP, but it is more expensive.

The FortiAP unit offers more flexible placement. FortiWiFi units either sit on a shelf or are rack mounted. FortiAP units can be attached to any wall or ceiling, enabling you to locate them where they will provide the best coverage.

## Single access point networks

A single access point is appropriate for a limited number of users in a small area. For example, you might want to provide wireless access for a group of employees in one area on one floor of an office building.

A good rule of thumb is that one access point can serve 3000 to 4000 square feet of space, with no user more than 60 feet from the access point. Walls and floors reduce the coverage further, depending on the materials from which they are made.

## Multiple access point networks

To cover a larger area, such as multiple floors of a building, or multiple buildings, multiple access points are required.

In the WiFi controller, you configure a single virtual access point, but the controller manages multiple physical access points that share the same configuration. A feature known as “fast roaming” enables users to move from one physical access point coverage area to another while retaining their authentication.

### Fast Roaming

Users in a multi-AP network, especially with mobile devices, can move from one AP coverage area to another. But, the process of re-authentication can often take seconds to complete and this can impair wireless voice traffic and time sensitive applications. The FortiAP fast roaming feature solves this problem and is available only when moving between FortiAP units managed by the same FortiGate unit.

Fast roaming uses two standards-based techniques:

- Pairwise Master Key (PMK) Caching enables a RADIUS-authenticated user to roam away from an AP and then roam back without having to re-authenticate. To accomplish this, the FortiGate unit stores in a cache a master key negotiated with the first AP. This enables the 802.11i-specified method of “fast roam-back.”
- Pre-authentication or “fast-associate in advance” enables an 802.11 AP associated to a client to bridge to other APs over the wired network and pre-authenticate the client to the “next” AP to which the client might roam. This enables the PMK to be derived in advance of a roam and cached. When the client does roam, it will already have negotiated authentication in advance and will use its cached PMK to quickly associate to the next AP. This capability will ensure that wireless clients that support Pre-authentication to continue the data transfer without noticeable connection issues.

### WiFi Mesh Network

FortiAP units can be connected to the WiFi controller by Ethernet or by WiFi. In the latter case, you configure a special backhaul network, the mesh, that carries traffic and control signals between FortiAP units and the WiFi controller. Regular WiFi clients cannot connect to the mesh network, they can connect only to non-mesh SSIDs. The mesh network is useful when running Ethernet cables is not practical. For best results, the mesh network should use a dedicated radio at both the WiFi controller and FortiAP unit. Otherwise, the client SSIDs compete for bandwidth with the mesh backhaul.

## Automatic Radio Resource Provisioning

To prevent interference between APs, the FortiOS WiFi Controller includes the Automatic Radio Resource Provisioning (ARRP) feature. When enabled in an access point profile, *Radio Resource Provision* measures utilization and interference on the available channels and selects the clearest channel at each access point. The measurement can be repeated periodically to respond to changing conditions.



# Configuring a WiFi LAN

When working with a FortiGate WiFi controller, you can configure your wireless network before you install any access points. If you are working with a standalone FortiWiFi unit, the access point hardware is already present but the configuration is quite similar. Both are covered in this section.

The following topics are included in this section:

- [Overview of WiFi controller configuration](#)
- [Setting your geographic location](#)
- [Creating a custom AP Profile](#)
- [Defining a wireless network interface \(SSID\)](#)
- [Dynamic VLAN assignment](#)
- [Configuring user authentication](#)
- [Configuring firewall policies for the SSID](#)
- [Customizing captive portal pages](#)
- [Configuring the built-in access point on a FortiWiFi unit](#)



On FortiGate model 30D, web-based manager configuration of the WiFi controller is disabled by default. To enable it, enter the following CLI commands:

```
config system global
 set gui-wireless-controller enable
end
```

## Overview of WiFi controller configuration

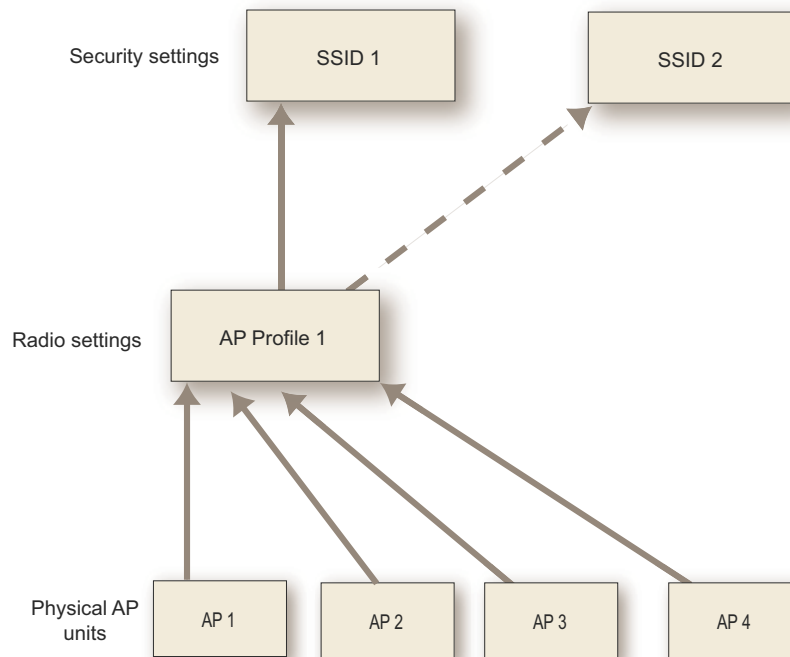
The FortiGate WiFi controller configuration is composed of three types of object, the SSID, the AP Profile and the physical Access Point.

- An **SSID** defines a virtual wireless network interface, including security settings. One SSID is sufficient for a wireless network, regardless how many physical access points are provided. You might, however, want to create multiple SSIDs to provide different services or privileges to different groups of users. Each SSID has separate firewall policies and authentication. Each radio in an access point can support up to 8 SSIDs.

A more common use of the term SSID is for the identifier that clients must use to connect to the wireless network. Each SSID (wireless interface) that you configure will have an SSID field for this identifier. In Managed Access Point configurations you choose wireless networks by SSID values. In firewall policies you choose wireless interfaces by their SSID name.

- An **AP Profile** defines the radio settings, such as band (802.11g for example) and channel selection. The AP Profile names the SSIDs to which it applies. Managed APs can use automatic profile settings or you can create custom AP profiles.
- **Managed Access Points** represent local wireless APs on FortiWiFi units and FortiAP units that the FortiGate unit has discovered. There is one managed access point definition for each AP device. An access point definition can use automatic AP profile settings or select a custom AP Profile. When automatic profile settings are used, the managed AP definition also selects the SSIDs to be carried on the AP.

**Figure 156:**Conceptual view of FortiGate WiFi controller configuration



### About SSIDs on FortiWiFi units

FortiWiFi units have a default SSID (wireless interface) named *wlan*. You can modify or delete this SSID as needed. As with external APs, the built-in wireless AP can be configured to carry any SSID.

The AP settings for the built-in wireless access point are located at *WiFi Controller > Managed Access Points > Local WiFi Radio*. The available operational settings are the same as those for external access points which are configured at *WiFi Controller > Managed Access Points > Managed FortiAP*.

### About automatic AP profile settings

FortiOS simplifies wireless network configuration by providing an automatic setting for the access point profile. You can enable wireless AP operation and Rogue AP scanning with the radios in the AP automatically allocated as follows:

**Table 42:** Radio functions in automatic profile

No. of Radios	Wireless Access enabled	Rogue AP Scan enabled	Wireless Access and Rogue AP Scan enabled
1	Radio 1 - AP	Radio 1 - scan	Radio 1 - AP + background scan
2	Radio 1 - AP Radio 2 - disabled	Radio 1 - disabled Radio 2 - scan	Radio 1 - AP Radio 2 - scan

You can select which SSIDs (wireless networks) will be available through the access point and adjust the wireless power level of the AP. Also, you can configure the unit's LAN ports, if you are using a FortiAP model that has them (see "[LAN port options](#)" on page 823).

## Process to create a wireless network

To set up your wireless network, you will need to perform the following steps.

- Make sure the FortiGate wireless controller is configured for your geographic location. This ensures that the available radio channels and radio power are in compliance with the regulations in your region.
- Optionally, if you don't want to use automatic AP profile settings, configure a custom Access Point (AP) profile, specifying the radio settings and the SSIDs to which they apply.
- Configure one or more SSIDs for your wireless network. The SSID configuration includes DHCP and DNS settings.
- Configure the user group and users for authentication on the WLAN.
- Configure the firewall policy for the WLAN.
- Optionally, customize the captive portal.
- Configure access points.

Configuration of the built-in AP on FortiWiFi units is described in this chapter. Connection and configuration of FortiAP units is described in the next chapter, "[Access point deployment](#)".

## Setting your geographic location

The maximum allowed transmitter power and permitted radio channels for Wi-Fi networks depend on the region in which the network is located. By default, the WiFi controller is configured for the United States. If you are located in any other region, you need to set your location before you begin configuring wireless networks.

### To change the location setting - CLI

To change the country to France, for example, enter

```
config wireless-controller setting
 set country FR
end
```

To see the list of country codes, enter a question mark ('?') instead of a country code.



Before changing the country setting, you must remove all Custom AP profiles. To do this, go to *WiFi Controller > WiFi Network > Custom AP Profile*.

---

## Creating a custom AP Profile

If the automatic AP profile settings don't meet your needs, you can define a custom AP Profile. For information about the automatic profile settings, see [“About automatic AP profile settings” on page 794](#).



On FortiGate model 30D web-based manager configuration of custom AP profiles is disabled by default. To enable AP profiles, enter the following CLI commands:

```
config system global
 set gui-ap-profile enable
end
```

An AP Profile configures radio settings and selects the Virtual APs to which the settings apply. FortiAP units contain two radio transceivers, making it possible, for example, to provide both 2.4GHz 802.11b/g/n and 5GHz 802.11a/n service from the same access point.

FortiAP units also provide a monitoring function for the Rogue AP feature.

### To configure an AP Profile - web-based manager

1. Go to *WiFi Controller > WiFi Network > Custom AP Profile* and select *Create New*.
2. Enter a *Name* for the AP Profile.
3. In *Platform*, select the FortiWiFi or FortiAP model to which this profile applies.
4. For each radio, enter:

<b>Mode</b>	Select the type of mode. <ul style="list-style-type: none"><li>• <i>Disable</i> – radio disabled</li><li>• <i>Access Point</i> – allows for the platform to be an access point</li><li>• <i>Dedicated Monitor</i> – allows for the platform to be a dedicated monitor. See <a href="#">“Wireless network monitoring” on page 846</a>.</li></ul>
<b>Background Scan</b>	Select to enable a background scan, which monitors other APs. This is needed for the Rogue AP feature. By default, background scan is disabled. See <a href="#">“Wireless network monitoring” on page 846</a> .
<b>Rogue AP On-Wire Scan</b>	If Mode is either <i>Dedicated Monitor</i> or <i>Access Point</i> with <i>Background Scan</i> enabled, you can enable on-wire scan to distinguish rogue APs from neighbors. For more information, see <a href="#">“On-wire rogue AP detection technique” on page 847</a> .
<b>WIDS Profile</b>	Optionally, select a Wireless Intrusion Detection (WIDS) profile. See <a href="#">“Wireless IDS” on page 843</a> .
<b>Radio Resource Provision</b>	Select to enable the radio resource provision feature. This feature measures utilization and interference on the available channels and selects the clearest channel at each access point. The measurement can be repeated periodically to respond to changing conditions.
<b>Client Load Balancing</b>	Select Frequency Handoff or AP Handoff as needed. See <a href="#">“Wireless client load balancing for high-density deployments” on page 822</a> .

<b>Band</b>	Select the wireless protocols that you want to support. The available choices depend on the radio's capabilities. Where multiple protocols are supported, the letter suffixes are combined: "802.11bg" means 802.11b <b>and</b> 802.11g. For 802.11n, 802.11n_2.4G indicates 2.4GHz, 802.11n_5G indicates 5GHz.  Note that on two-radio units such as the FortiAP-220B it is not possible to put both radios on the same band.
<b>20/40 Mhz Channel Width</b>	Select to enable 20/40 MHz channel width for 802.11n-5G.
<b>Short Guard Interval</b>	Select to enable the short guard interval for 802.11n_5G with 20/40 MHz channel width.
<b>Channel</b>	Select the channel or channels to include. The available channels depend on which IEEE wireless protocol you selected in <i>Band</i> . By default, all available channels are enabled.
<b>Auto Tx Power Control</b>	Optionally, enable automatic adjustment of transmit power, specifying minimum and maximum power levels.
<b>TX Power</b>	By default, the TX power is set to 100% of the maximum power permitted in your region. To change the level, drag the slider.
<b>SSID</b>	Choose the SSIDs (WiFi networks) that APs using this profile will carry.  Select the required SSIDs in the <i>Available</i> list and use the -> arrow to move them to the <i>Selected</i> list. To remove an SSID from the <i>Selected</i> list, select the SSID and then use the <- arrow to move it back to the <i>Available</i> list.

Radio 1 settings are the same as Radio 2 settings except for the options for *Channel*. Radio 2 settings are available only for FortiAP models with dual radios.

##### 5. Select OK.

#### To configure an AP Profile - CLI

This example configures a FortiAP-220B to use only Radio 2 for 802.11g operation applied to SSID example\_wlan.

```
config wireless-controller wtp-profile
 edit guest_prof
 config platform
 set type 220B
 end
 config radio-2
 set mode ap
 set band 802.11g
 set vaps example_wlan
 end
 end
end
```

## Defining a wireless network interface (SSID)

You begin configuring your wireless network by defining one or more SSIDs to which your users will connect. When you create an SSID, a virtual network interface is also created with the *Name* you specified in the SSID configuration. You can configure the settings of an existing SSID in either *WiFi Controller > WiFi Network > SSID* or *System > Network > Interface*.

### To create a new SSID

1. Go to *WiFi Controller > WiFi Network > SSID* and select *Create New*.
2. Fill in the SSID fields as described below.

### To configure the settings of an existing SSID

1. Either
  - Go to *WiFi Controller > WiFi Network > SSID*.
  - or
  - Go to *System > Network > Interfaces*.  
WiFi interfaces list the SSID beside the interface *Name*.
2. Edit a WiFi interface, modifying the SSID fields as needed.

### SSID fields

<b>Name</b>	Enter a name for the SSID interface.
<b>Type</b>	WiFi SSID.
<b>Traffic Mode</b>	<b>Tunnel to Wireless Controller</b> — Data for WLAN passes through WiFi Controller. This is the default. <b>Local bridge with FortiAP's Interface</b> — FortiAP unit Ethernet and WiFi interfaces are bridged. <b>Mesh Downlink</b> — Radio receives data for WLAN from mesh backhaul SSID.
<b>IP/Netmask</b>	Enter the IP address and netmask for the SSID.
<b>IPv6 Address</b>	Enter the IPv6 address. This is available only when IPv6 has been enabled on the unit.
<b>Administrative Access</b>	Select which types of administrative access are permitted on this SSID.
<b>IPv6 Administrative Access</b>	If you have IPv6 addresses, select the permitted IPv6 administrative access types for this SSID.
<b>DHCP Server</b>	Select to enable a DHCP server and define IP address ranges to assign to clients or to relay DHCP requests to another server.  If the unit is in transparent mode, the DHCP server settings will be unavailable.  For more information, see <a href="#">“Configuring DHCP for WiFi clients” on page 801</a> .

<b>WiFi Settings</b>	
<b>SSID</b>	Enter the SSID. By default, this field contains <code>fortinet</code> .
<b>Security Mode</b>	<p>Select the security mode for the wireless interface. Wireless users must use the same security mode to be able to connect to this wireless interface. Additional security mode options are available in the CLI. For more information, see <a href="#">“Configuring security” on page 801</a>.</p> <p><i>WPA/WPA2-Personal</i> – WPA or WPA2 security. WPA is WiFi protected access. WPA2 is WPA with additional security features. There is one shared key (password) that all users use.</p> <p><i>WPA/WPA2-Enterprise</i> – similar to WPA/WPA2-Personal, but is best used for enterprise networks. Each user is separately authenticated by user name and password.</p> <p><i>Captive Portal</i> – authenticates users through a customizable web page.</p>
<b>Pre-shared Key</b>	<p>Available only when <i>Security Mode</i> is <i>WPA/WPA2-Personal</i>.</p> <p>Enter the encryption key that the clients must use.</p>
<b>Data Encryption</b>	<p>Available only when <i>Security Mode</i> is <i>WPA/WPA2-Personal</i> or <i>WPA/WPA2-Enterprise</i>.</p> <p>Select <i>TKIP</i> or <i>AES</i> encryption as appropriate for the capabilities of your wireless clients. This is available for WPA/WPA2 security modes.</p>
<b>Authentication</b>	<p>Available only when <i>Security Mode</i> is <i>WPA/WPA2-Enterprise</i>.</p> <p>Select one of the following:</p> <p><i>RADIUS Server</i> – Select the RADIUS server that will authenticate the clients.</p> <p><i>Usergroup</i> – Select the user group(s) that can authenticate.</p>
<b>Customize Portal Messages</b>	<p>Available only when <i>Security Mode</i> is <i>Captive Portal</i>. Select to customize the endpoint replacement messages. When you select <i>Edit</i>, the Edit Message window appears. Within the window, you can modify each one of the endpoint replacement messages.</p>
<b>User Groups</b>	<p>Available only when <i>Security Mode</i> is <i>Captive Portal</i>. Select the user groups that can authenticate.</p> <p>To select a user group, select the group in <i>Available</i> and then use the <code>-&gt;</code> arrow to move that group to <i>Selected</i>. To remove a user group from <i>Selected</i>, select the group and then use the <code>&lt;-</code> arrow to move the group back to <i>Available</i>.</p>
<b>Block Intra-SSID Traffic</b>	Select to enable the unit to block intra-SSID traffic.
<b>Allow New WiFi Client Connections When Controller Is Down</b>	This option is available for local bridge SSIDs with WPA-Personal security. See <a href="#">“Continued FortiAP operation when WiFi controller connection is down” on page 841</a> .

<b>Maximum Clients</b>	Select to limit the number of clients permitted to connect simultaneously. Enter the limit value.
<b>Device Management</b>	Select <i>Detect and Identify Devices</i> if you want to monitor the device types using this interface or create device identity policies involving this interface. See <a href="#">“Managing “bring your own device”” on page 1996</a> .  Optionally, enable <i>Add New Devices to Vulnerability Scan List</i> .
<b>Enable Explicit Web Proxy</b>	Select to enable explicit web proxy for the SSID.
<b>Listen for RADIUS Accounting Messages</b>	This is required to permit RADIUS SSO authentication on this WiFi interface. See <a href="#">“SSO using RADIUS accounting records” on page 602</a> .
<b>Secondary IP Address</b>	Optionally, enable and define secondary IP addresses. Administrative access can be enabled on secondary interfaces.
<b>Comments</b>	Enter a description or comment for the SSID.

By default, the AP will broadcast its SSID. Optionally, you can disable SSID Broadcast in the CLI:

```
config wireless controller vap
 edit vap_name
 set broadcast-ssid disable
 end
```

For more information, see [“Whether to broadcast SSID” on page 785](#).

Each Virtual AP that you create is a wireless interface that establishes a wireless LAN. Go to *System > Network > Interfaces* to configure its IP address.

### To configure a virtual access point (SSID) - CLI

This example creates an access point with SSID “example” and WPA2-Personal security. The wireless interface is named example\_wlan.

```
config wireless-controller vap
 edit example_wlan
 set ssid "example"
 set broadcast-ssid enable
 set security wpa2-only-personal
 set passphrase "hardtoguess"
 set vdom root
 end
config system interface
 edit example_wlan
 set ip 10.10.120.1 255.255.255.0
 end
```



## Configuring DHCP for WiFi clients

Wireless clients need to have IP addresses. If you use RADIUS authentication, each user's IP address can be stored in the Framed-IP-Address attribute. Otherwise, you need to configure a DHCP server on the WLAN interface to assign IP addresses to wireless clients.

### To configure a DHCP server for WiFi clients - web-based manager

1. Go to *WiFi Controller > WiFi Network > SSID* and edit your SSID entry.
2. In the *WiFi Settings* section, enable *DHCP Server*.
3. In *Address Range*, select *Create New*.
4. In the *Starting IP* and *End IP* fields, enter the IP address range to assign.  
The address range needs to be in the same subnet as the wireless interface IP address, but not include that address.
5. Set the *Netmask* to an appropriate value, such as 255.255.255.0.
6. Set the *Default Gateway* to *Same as Interface IP*.
7. Set the *DNS Server* to *Same as System DNS*.
8. If you want to restrict access to the wireless network by MAC address, see [“Adding a MAC filter” on page 804](#).
9. Select *OK*.

### To configure a DHCP server for WiFi clients - CLI

In this example, WiFi clients on the `example_wlan` interface are assigned addresses in the 10.10.120.2-9 range to connect with the WiFi access point on 10.10.120.1.

```
config system dhcp server
 edit 0
 set default-gateway 10.10.120.1
 set dns-service default
 set interface example_wlan
 set netmask 255.255.255.0
 config ip-range
 edit 1
 set end-ip 10.10.120.9
 set start-ip 10.10.120.2
 end
 end
end
```



You cannot delete an SSID (wireless interface) that has DHCP enabled on it.

---

## Configuring security

Using the web-based manager, you can configure Open Portal security or Wi-Fi Protected Access (WPA) security modes WPA-Personal and WPA-Enterprise. The WPA options support both WPA and WPA2, which has additional security improvements. Using the CLI, you can also choose WPA-only and WPA2-only modes.

Using the CLI, you can also choose Wireless Equivalent Privacy (WEP) modes. WEP modes are much less secure and are provided for legacy support only. Wherever possible, use WPA security.

WPA security with a preshared key for authentication is called WPA-Personal. This can work well for one person a small group of trusted people. But, as the number of users increases, it is difficult to distribute new keys securely and there is increased risk that the key could fall into the wrong hands.

A more secure form of WPA security is WPA-Enterprise. Users each have their own authentication credentials, verified through an authentication server, usually RADIUS. FortiOS can also authenticate WPA-Enterprise users through its built-in user group functionality. FortiGate user groups can include RADIUS servers and can select users by RADIUS user group. This makes possible Role-Based Access Control (RBAC).

WPA security can encrypt communication with either Temporal Key Integrity Protocol (TKIP) or Advanced Encryption Standard (AES). AES is the preferred encryption, but some older wireless clients do not support it. You can select the encryption during setup.

Captive Portal security connects users to an open web portal defined in replacement messages. To navigate to any location beyond the web portal, the user must pass FortiGate user authentication.

## WPA-Personal security

WPA-Personal security setup requires only the preshared key that you will provide to your clients.

### To configure WPA-Personal security - web-based manager

1. Go to *WiFi Controller > WiFi Network > SSID* and edit your SSID entry.
2. In *Security Mode*, select *WPA/WPA2-Personal*.
3. In *Data Encryption*, select *AES*.  
If some of your wireless clients do not support AES, select *TKIP*.
4. In *Pre-shared Key*, enter a key between 8 and 63 characters long.
5. Select *OK*.

### To configure WPA-Personal security - CLI

```
config wireless-controller vap
 edit example_wlan
 set security wpa-personal
 set passphrase "hardtoguess"
 set encrypt AES
 end
```

## WPA-Enterprise security

If you will use FortiOS user groups for authentication, go to *User & Device > User > User Group* and create those groups first. The groups should be Firewall groups.

If you will use a RADIUS server to authenticate wireless clients, you must first configure the FortiGate unit to access the RADIUS server.

### To configure FortiGate unit access to the RADIUS server - web-based manager

1. Go to *User & Device > Authentication > RADIUS Server* and select *Create New*.
2. Enter a *Name* for the server.
3. In *Primary Server Name/IP*, enter the network name or IP address for the server.

4. In *Primary Server Secret*, enter the shared secret used to access the server.
5. Optionally, enter the information for a secondary or backup RADIUS server.
6. Select *OK*.

#### To configure the FortiGate unit to access the RADIUS server - CLI

```
config user radius
 edit exampleRADIUS
 set auth-type auto
 set server 10.11.102.100
 set secret aoewmntiasf
 end
```

#### To configure WPA-Enterprise security - web-based manager

1. Go to *WiFi Controller > WiFi Network > SSID* and edit your SSID entry.
2. In *Security Mode*, select *WPA/WPA2-Enterprise*.
3. In *Data Encryption*, select *AES*.  
If some of your wireless clients do not support AES, select *TKIP*.
4. In *Authentication*, do one of the following:
  - If you will use a RADIUS server for authentication, select *RADIUS Server* and then select the RADIUS server.
  - If you will use a local user group for authentication, select *Usergroup* and then select the user group that is permitted to use the wireless network.
5. Select *OK*.

#### To configure WPA-Enterprise security - CLI

```
config wireless-controller vap
 edit example_wlan
 set security wpa-enterprise
 set encrypt AES
 set auth radius
 set radius-server exampleRADIUS
 end
```

### Captive Portal security

Captive Portal security provides an access point that initially appears open. The wireless client can connect to the AP with no security credentials. The AP responds to the client's first HTTP request with a web page requesting user name and password. Until the user enters valid credentials, no communication beyond the AP is permitted.

The wireless controller authenticates users through the FortiGate user accounts. In the SSID configuration, you select the user groups that are permitted access through the captive portal.

The captive portal contains the following web pages:

- **Login page**—requests user credentials
- **Login failed page**—reports that the entered credentials were incorrect and enables the user to try again.
- **Disclaimer page**—is a statement of the legal responsibilities of the user and the host organization to which the user must agree before proceeding.
- **Declined disclaimer page**—is displayed if the user does not agree to the statement on the Disclaimer page. Access is denied until the user agrees to the disclaimer.

These pages are defined in replacement messages. Defaults are provided. In the web-based manager, you can modify the default messages in the SSID configuration by selecting *Customize Portal Messages*. Each SSID can have its own unique portal content.

#### **To configure Captive Portal security - web-based manager**

1. Configure user groups as needed in *User & Device > User > User Groups*.
2. Go to *WiFi Controller > WiFi Network > SSID* and edit your SSID entry.
3. In *Security Mode*, select *Captive Portal*.
4. Optionally, select *Customize Portal Messages* and modify the portal pages that users of this SSID will see.
5. In *User Groups*, select the group(s) that are allowed to use the wireless network and move them to the *Selected* list.
6. Select *OK*.

## **Adding a MAC filter**

On each SSID, you can create a MAC address filter list to either permit or exclude a list of clients identified by their MAC addresses.

This is actually not as secure as it appears. Someone seeking unauthorized access to your network can obtain MAC addresses from wireless traffic and use them to impersonate legitimate users. A MAC filter list should only be used in conjunction with other security measures such as encryption.

#### **To configure a MAC filter - web-based manager**

1. Go to *WiFi Controller > WiFi Network > SSID* and edit your SSID entry.
2. In the *DHCP Server* section, expand *Advanced*.
3. In *MAC Address Access Control List*, select *Create New*.
4. Enter a MAC address in the *MAC* field.
5. Double-click in *IP or Action*, and do one of:
  - Select *Reserve IP* and enter the IP address to assign to this MAC address.
  - Select *Assign IP*. This MAC address will be assigned an IP address automatically.
  - Select *Block*. This MAC address will not be assigned an IP address.
6. Double-click in the *Unknown MAC Addresses* line and select *Assign IP* or *Block*, as needed. By default, unlisted MAC addresses are assigned an IP address automatically.
7. Repeat steps 3 through 6 for each additional MAC address that you want to add.
8. Select *OK*.

#### **To configure a MAC filter - CLI**

1. Enter

```
config system dhcp server
show
```

2. Find the entry where `interface` is your WiFi interface. Edit that entry and configure the MAC filter. In this example, the MAC address `11:11:11:11:11:11` will be excluded. Unlisted MAC addresses will be assigned an IP address automatically.

```
edit 3
 config reserved-address
 edit 1
 set action block
 set mac 11:11:11:11:11:11
 end
 set mac-acl-default-action assign
end
```

## Multicast enhancement

FortiOS can translate multicast traffic into unicast traffic to send to clients, maintaining its own multicast client through IGMP snooping. You can configure this in the CLI:

```
config wireless-controller vap
 edit example_wlan
 set multicast-enhance enable
 set me-disable-thresh 32
 end
```

If the number of clients on the SSID is larger than `me-disable-thresh`, multicast enhancement is disabled.

## Dynamic VLAN assignment

You can assign each individual user to a VLAN based on information stored in the RADIUS authentication server. If the user's RADIUS record does not specify a VLAN ID, the user is assigned to the default VLAN for the SSID.

The RADIUS user attributes used for the VLAN ID assignment are:

IETF 64 (Tunnel Type)—Set this to VLAN.

IETF 65 (Tunnel Medium Type)—Set this to 802

IETF 81 (Tunnel Private Group ID)—Set this to the VLAN ID.

To configure dynamic VLAN assignment, you need to:

1. Configure access to the RADIUS server.
2. Create the SSID and enable dynamic VLAN assignment.
3. Create a custom AP profile and add the local bridge mode SSID to it.
4. Create the VLAN interfaces and their DHCP servers.
5. Create security policies to allow communication from the VLAN interfaces to the Internet.
6. Authorize the FortiAP unit and assign the custom profile to it.

### To configure access to the RADIUS server

1. Go to *User & Device > Authentication > RADIUS Server* and select *Create New*.
2. Enter a Name, the name or IP address in Primary Server Name/IP, and the server secret in Primary Server Secret.

### To create the dynamic VLAN SSID

1. Go to WiFi Controller > WiFi Network > SSID, select Create New and enter:

<b>Name</b>	An identifier, such as dynamic_vlan_ssid.
<b>Traffic Mode</b>	Local bridge or Tunnel, as needed.
<b>SSID</b>	An identifier, such as DYNSSID.
<b>Security Mode</b>	WPA/WPA2-Enterprise
<b>Authentication</b>	RADIUS Server. Select the RADIUS server that you configured.

2. Select OK.
3. Enable dynamic VLAN in the CLI. Optionally, you can also assign a VLAN ID to set the default VLAN for users without a VLAN assignment.

```
config wireless-controller vap
 edit dynamic_vlan_ssid
 set dynamic-vlan enable
 set vlanid 10
 end
```

### To create the custom AP profile for the dynamic VLAN SSID

1. Go to WiFi Controller > WiFi Network > Custom AP Profile, select Create New and enter:

<b>Name</b>	A name for the profile, such as dyn_vlan_profile.
<b>Platform</b>	The FortiAP model you are using. If you use more than one model of FortiAP, you will need a custom AP profile for each model.
<b>Radio 1 and Radio 2</b>	
<b>SSID</b>	In the <i>Available</i> column, select the SSID you created (example dynamic_vlan_ssid) and move it to the <i>Selected</i> column. There should be no other SSIDs in the <i>Selected</i> column.

2. Adjust other radio settings as needed.
3. Select OK.

### To create the VLAN interfaces

1. Go to System > Network > Interfaces and select Create New.
2. Enter:

<b>Name</b>	A name for the VLAN interface, such as VLAN100.
<b>Interface</b>	The physical interface associated with the VLAN interface.
<b>VLAN ID</b>	The numeric VLAN ID, for example 100.
<b>Addressing mode</b>	Select Manual and enter the IP address / Network Mask for the virtual interface.
<b>DHCP Server</b>	Enable and then select Create New to create an address range.

3. Select OK.
4. Repeat the preceding steps to create other VLANs as needed.

Security policies determine which VLANs can communicate with which other interfaces. These are the simple Firewall Address policy without authentication. Users are assigned to the appropriate VLAN when they authenticate.

#### **To connect and authorize the FortiAP unit**

1. Connect the FortiAP unit to the FortiGate unit.
2. Go to *WiFi Controller > Managed Access Points > Managed AP*.
3. When the FortiAP unit is listed, double-click the entry to edit it.
4. In *AP Profile*, select *Change*, then select the custom profile that you created. Select *Apply*.
5. Select *Authorize*.
6. Select OK.

## **Configuring user authentication**

You can perform user authentication when the wireless client joins the wireless network and when the wireless user communicates with another network through a firewall policy. WEP and WPA-Personal security rely on legitimate users knowing the correct key or passphrase for the wireless network. The more users you have, the more likely it is that the key or passphrase will become known to unauthorized people. WPA-Enterprise and captive portal security provide separate credentials for each user. User accounts can be managed through FortiGate user groups or an external RADIUS authentication server.

### **WPA-Enterprise authentication**

If your WiFi network uses WPA-Enterprise authentication verified by a RADIUS server, you need to configure the FortiGate unit to connect to that RADIUS server.

#### **Configuring connection to a RADIUS server - web-based manager**

1. Go to *User & Device > Authentication > RADIUS Server* and select *Create New*.
2. Enter a *Name* for the server.  
This name is used in FortiGate configurations. It is not the actual name of the server.
3. In *Primary Server Name/IP*, enter the network name or IP address for the server.
4. In *Primary Server Secret*, enter the shared secret used to access the server.
5. Optionally, enter the information for a secondary or backup RADIUS server.
6. Select *OK*.

#### **To configure the FortiGate unit to access the RADIUS server - CLI**

```
config user radius
 edit exampleRADIUS
 set auth-type auto
 set server 10.11.102.100
 set secret aoewmntiasf
 end
```

To implement WPA-Enterprise security, you select this server in the SSID security settings. See [“Configuring security” on page 801](#).

To use the RADIUS server for authentication, you can create individual FortiGate user accounts that specify the authentication server instead of a password, and you then add those accounts to a user group. Or, you can add the authentication server to a FortiGate user group, making all accounts on that server members of the user group.

## Creating a wireless user group

Most wireless networks require authenticated access. To enable creation of identity-based firewall policies, you should create at least one user group for your wireless users. You can add or remove users later. There are two types of user group to consider:

- A Firewall user group can contain user accounts stored on the FortiGate unit or external authentication servers such as RADIUS that contain and verify user credentials.
- A Directory Services user group is used for integration with Windows Active Directory or Novell eDirectory. The group can contain Windows or Novell user groups who will be permitted access to the wireless LAN. Fortinet Single Sign On (FSSO) agent must be installed on the network.

## MAC-based authentication

Wireless clients can also be supplementally authenticated by MAC address. A RADIUS server stores the allowed MAC address for each client and the wireless controller checks the MAC address independently of other authentication methods.

MAC-based authentication must be configured in the CLI. In the following example, MAC-based authentication is added to an existing access point “vap1” to use a RADIUS server at 192.168.1.95:

```
config wireless-controller vap
 edit vap1
 set radius-mac-auth enable
 set radius-mac-auth-server 192.168.1.95
 end
```

## Authenticating guest WiFi users

The FortiOS Guest Management feature enables you to easily add guest accounts to your FortiGate unit. These accounts are authenticate guest WiFi users for temporary access to a WiFi network managed by a FortiGate unit.

To implement guest access, you need to

1. Go to *User & Device > User > User Group* and create one or more guest user groups.
2. Go to *User & Device > User > Guest Management* to create guest accounts. You can print the guest account credentials or send them to the user as an email or SMS message.
3. Go to *WiFi Controller > WiFi Network > SSID* and configure your WiFi SSID to use captive portal authentication. Select the guest user group(s) that you created.

Guest users can log into the WiFi captive portal with their guest account credentials until the account expires. For more detailed information about creating guest accounts, see “Managing Guest Access” in the Authentication chapter of the this FortiOS Handbook.



## Configuring firewall policies for the SSID

For users on the WiFi LAN to communicate with other networks, firewall policies are required. Before you create firewall policies, you need to define any firewall addresses you will need. This section describes creating a WiFi network to Internet policy.

### To create a firewall address for WiFi users - web-based manager

1. Go to *Firewall Objects > Address > Addresses*.
2. Select *Create New*, enter the following information and select *OK*.

<b>Name</b>	Enter a name for the address, <i>wifi_net</i> for example.
<b>Type</b>	Select <i>Subnet</i> .
<b>Subnet / IP Range</b>	Enter the subnet address, <i>10.10.110.0/24</i> for example.
<b>Interface</b>	Select the interface where this address is used, e.g., <i>example_wifi</i>

### To create a firewall address for WiFi users - CLI

```
config firewall address
 edit "wifi_net"
 set associated-interface "example_wifi"
 set subnet 10.10.110.0 255.255.255.0
 end
```

### To create a firewall policy - web-based manager

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. In *Incoming Interface*, select the wireless interface.
3. In *Source Address*, select the address of your WiFi network, *wifi\_net* for example.
4. In *Outgoing Interface*, select the Internet interface, for example, *port1*.
5. In *Destination Address*, select *All*.
6. In *Service*, select *ALL*, or select the particular services that you want to allow, and then select the right arrow button to move the service to the *Selected Services* list.
7. In *Schedule*, select *always*, unless you want to define a schedule for limited hours.
8. In *Action*, select *ACCEPT*.
9. Select *Enable NAT*.
10. Optionally, set up UTM features for wireless users.
11. Select *OK*.

### To create a firewall policy - CLI

```
config firewall policy
 edit 0
 set srcintf "example_wifi"
 set dstintf "port1"
 set srcaddr "wifi_net"
 set dstaddr "all"
 set action accept
 set schedule "always"
 set service "ANY"
 set nat enable
 end
```

## Customizing captive portal pages

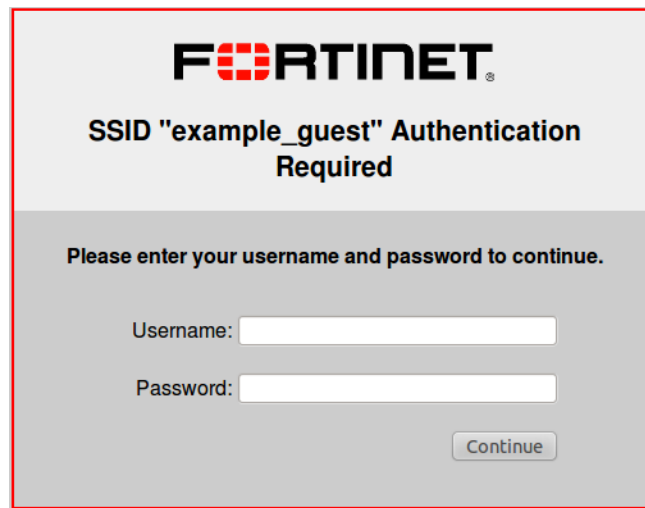
If you select Captive Portal authentication in the SSID, the wireless controller presents the user pages defined in Captive Portal Default replacement pages.

The captive portal contains the following web pages:

- **Captive Portal Login page**—requests user credentials
- **Captive Portal Login Failed page**—reports that the entered credentials were incorrect and enables the user to try again.
- **Captive Portal Disclaimer page**—is statement of the legal responsibilities of the user and the host organization to which the user must agree before proceeding.
- **Captive Portal Rejected page**—is displayed if the user does not agree to the statement on the Disclaimer page. Access is denied until the user agrees to the disclaimer.

These pages are defined in replacement messages. Defaults are provided. In the web-based manager, you can modify the default messages in the SSID configuration by selecting *Customize Portal Messages*. Each SSID can have its own unique portal content.

**Figure 157:**Default captive portal login page



### Modifying the login page

The login page requests the user's credentials. Typical modifications for this page would be to change the logo and modify some of the text.

#### Changing the logo

You can replace the default Fortinet logo with your organization's logo. First, import the logo file into the FortiGate unit and then modify the Login page code to reference your file.

##### To import a logo file

1. Go to *System > Config > Replacement Messages* and select *Manage Images*.
2. Select *Create New*.
3. Enter a *Name* for the logo and select the appropriate *Content Type*.  
The file must not exceed 6000 bytes.
4. Select *Browse*, find your logo file and then select *Open*.
5. Select *OK*.

### To specify the new logo in the replacement message

1. Go to *WiFi Controller > WiFi Network > SSID* and edit your SSID.  
The SSID *Security Mode* must be *Captive Portal*.
2. Make sure that *Customize Portal Messages* is selected and then select the adjacent Edit icon.
3. In the *Edit Message* window, select the *Login page* message.
4. In the Message HTML, find the %%IMAGE tag.  
By default it specifies the Fortinet logo:%%IMAGE:logo\_fw\_auth%%
5. Change the image name to the one you provided for your logo.  
The tag should now read, for example, %%IMAGE:mylogo%%
6. Select *OK*.

### Modifying text

You can change any text that is not part of the HTML code nor a special tag enclosed in double percent (%) characters. There are two exceptions to this rule:

- The line “Please enter your username and password to continue” is provided by the %%QUESTION%% tag. You can replace this tag with text of your choice.
- The line “SSID ... Authentication Required” includes the name of the SSID, provided by the %%CPAUTH\_SSID%% tag. You can remove or change the position of this tag.

Except for these items, you should not remove any tags because they may carry information that the FortiGate unit needs.

### To modify login page text

1. Go to *WiFi Controller > WiFi Network > SSID* and edit your SSID.  
The SSID *Security Mode* must be *Captive Portal*.
2. Make sure that *Customize Portal Messages* is selected and then select the adjacent Edit icon.
3. In the *Edit Message* window, select the *Login page* message.
4. In the *Message HTML* box, edit the text, then select *OK*.
5. Select *OK*.

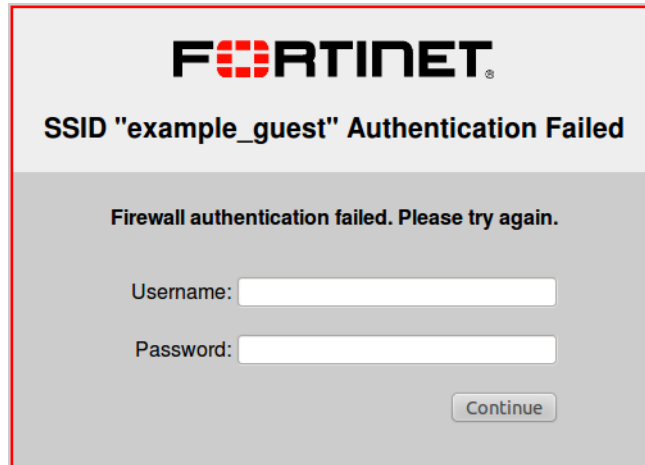
### Modifying the login failed page

The Login failed page is similar to the Login page. It even contains the same login form. You can change any text that is not part of the HTML code nor a special tag enclosed in double percent (%) characters. There are two exceptions to this rule:

- The line “Firewall authentication failed. Please try again.” is provided by the %%FAILED\_MESSAGE%% tag. You can replace this tag with text of your choice.
- The line “SSID ... Authentication Required” includes the name of the SSID, provided by the %%CPAUTH\_SSID%% tag. You can remove or change the position of this tag.

Except for these items, you should not remove any tags because they may carry information that the FortiGate unit needs.

**Figure 158:**Default login failed page



**FORTINET**

**SSID "example\_guest" Authentication Failed**

Firewall authentication failed. Please try again.

Username:

Password:

## Configuring the built-in access point on a FortiWiFi unit

Both FortiGate and FortiWiFi units have the WiFi controller feature. If you configure a WiFi network on a FortiWiFi unit, you can also use the built-in wireless capabilities in your WiFi network as one of the access points.

If Virtual Domains are enabled, you must select the VDOM to which the built-in access point belongs. You do this in the CLI. For example:

```
config wireless-controller global
 set local-radio-vdom vdom1
end
```

### To configure the FortiWiFi unit's built-in WiFi access point

1. Go to *WiFi Controller > Managed Access Points > Local WiFi Radio*.
2. Make sure that *AP Profile* is *Automatic*.
3. Make sure that *Enable WiFi Radio* is selected.
4. In *SSID*, if you do not want this AP to carry all SSIDs, select *Select SSIDs* and then select the required SSIDs.
5. Optionally, adjust the *TX Power* slider.  
If you have selected your location correctly (see [“Setting your geographic location” on page 795](#)), the 100% setting corresponds to the maximum power allowed in your region.
6. If you do not want the built-in WiFi radio to be used for rogue scanning, select *Do not participate in Rogue AP scanning*.
7. Select *OK*.

If you want to connect external APs, such as FortiAP units, see the next chapter, [“Access point deployment”](#).

# Access point deployment

This chapter describes how to configure access points for your wireless network. The following topics are included in this section:

- [Overview](#)
- [Network topology for managed APs](#)
- [Discovering and authorizing APs](#)
- [Advanced WiFi controller discovery](#)
- [Wireless client load balancing for high-density deployments](#)
- [LAN port options](#)
- [Preventing IP fragmentation of packets in CAPWAP tunnels](#)

## Overview

FortiAP units discover WiFi controllers. The administrator of the WiFi controller authorizes the FortiAP units that the controller will manage.

In most cases, FortiAP units can find WiFi controllers through the wired Ethernet without any special configuration. Review the following section, [“Network topology for managed APs”](#), to make sure that your method of connecting the FortiAP unit to the WiFi controller is valid. Then, you are ready to follow the procedures in [“Discovering and authorizing APs” on page 815](#).

If your FortiAP units are unable to find the WiFi controller, refer to [“Advanced WiFi controller discovery” on page 820](#) for detailed information about the FortiAP unit’s controller discovery methods and how you can configure them.

## Network topology for managed APs

The FortiAP unit can be connected to the FortiGate unit in any of the following ways:

**Direct connection:** The FortiAP unit is directly connected to the FortiGate unit with no switches between them. This configuration is common for locations where the number of FortiAP’s matches up with the number of ‘internal’ ports available on the FortiGate. In this configuration the FortiAP unit requests an IP address from the FortiGate unit, enters discovery mode and should quickly find the FortiGate WiFi controller. This is also known as a wirecloset deployment. See [Figure 159](#), below.

**Switched Connection:** The FortiAP unit is connected to the FortiGate WiFi controller by an Ethernet switch operating in L2 switching mode or L3 routing mode. There must be a routable path between the FortiAP unit and the FortiGate unit and ports 5246 and 5247 must be open. This is also known as a gateway deployment. See [Figure 159](#), below

**Connection over WAN:** The FortiGate WiFi controller is off-premises and connected by a VPN tunnel to a local FortiGate. In this method of connectivity its best to configure each FortiAP with the static IP address of the WiFi controller. Each FortiAP can be configured with three WiFi controller IP addresses for redundant failover. This is also known as a datacenter remote management deployment. See [Figure 160](#), below.

Figure 159: Wirecloset and Gateway deployments

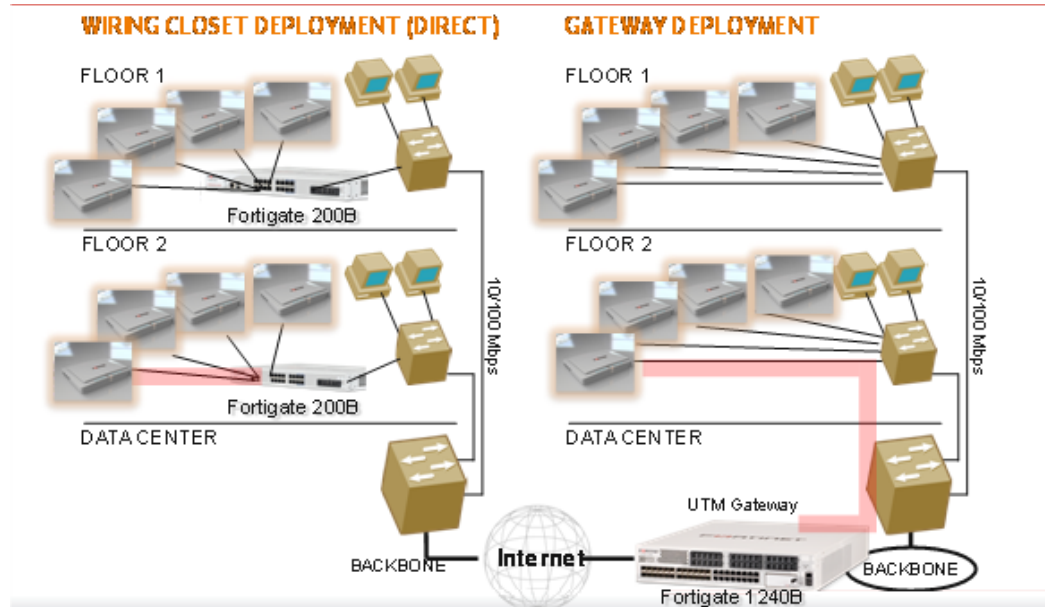
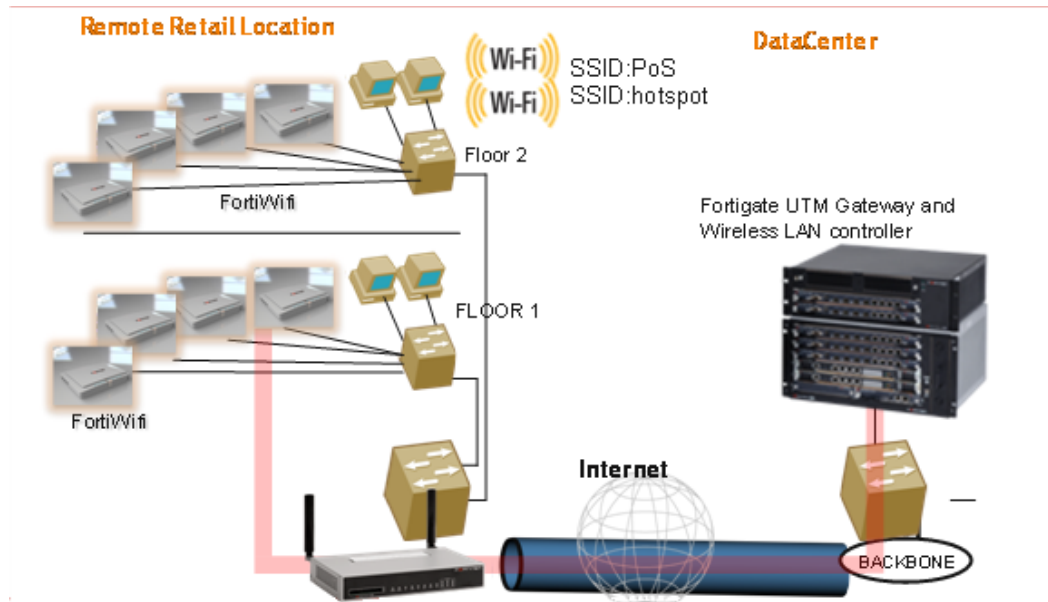


Figure 160: Remote deployment



## Discovering and authorizing APs

After you prepare your FortiGate unit, you can connect your APs to discover them using the discovery methods described earlier. To prepare the FortiGate unit, you need to

- Configure the network interface to which the AP will connect.
- Configure DHCP service on the interface to which the AP will connect.
- Optionally, preauthorize FortiAP units. They will begin to function when connected.
- Connect the AP units and let the FortiGate unit discover them.
- Enable each discovered AP and configure it or assign it to an AP profile.

## Configuring the network interface for the AP unit

The interface to which you connect your wireless access point needs an IP address. No administrative access, DNS Query service or authentication should be enabled.

### To configure the interface for the AP unit - web-based manager

1. Go to *System > Network > Interfaces* and edit the interface to which the AP unit connects.
2. Set *Addressing Mode* to *Dedicate to FortiAP*.
3. Enter the IP address and netmask to use.

This FortiGate unit automatically configures a DHCP server on the interface that will assign the remaining higher addresses up to .254 to FortiAP units. For example, if the IP address is 10.10.1.100, the FortiAP units will be assigned 10.10.1.101 to 10.10.1.254. To maximize the available addresses, use the .1 address for the interface: 10.10.1.1, for example.

4. Select OK.

### To configure the interface for the AP unit - CLI

In the CLI, you must configure the interface IP address and DHCP server separately.

```
config system interface
 edit port3
 set mode static
 set ip 10.10.70.1 255.255.255.0
 end
config system dhcp server
 edit 0
 set interface "dmz"
 config ip-range
 edit 1
 set end-ip 10.10.70.254
 set start-ip 10.10.70.2
 end
 set netmask 255.255.255.0
 set vci-match enable
 set vci-string "FortiAP"
 end
```

The optional `vci-match` and `vci-string` fields ensure that the DHCP server will provide IP addresses only to FortiAP units.

## Pre-authorizing a FortiAP unit

If you enter the FortiAP unit information in advance, it is authorized and will begin to function when it is connected.

### To pre-authorize a FortiAP unit

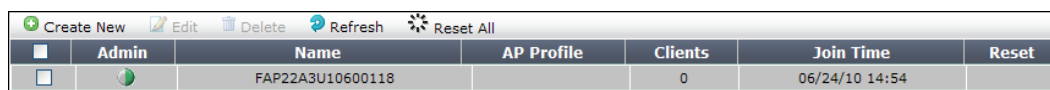
1. Go to *WiFi Controller > Managed Access Points > Managed FortiAPs* and select *Create New*.  
On some models the *WiFi Controller* menu is called *WiFi & Switch Controller*.
2. Enter the *Serial Number* of the FortiAP unit.
3. Configure the *Wireless Settings* as required.
4. Select OK.



## Enabling and configuring a discovered AP

Within two minutes of connecting the AP unit to the FortiGate unit, the discovered unit should be listed on *WiFi Controller > Managed Access Points > Managed FortiAP* page.

**Figure 161:**Discovered access point unit



Admin	Name	AP Profile	Clients	Join Time	Reset
<input type="checkbox"/>	FAP22A3U10600118		0	06/24/10 14:54	

When you authorize (enable) a FortiAP unit, it is configured by default to

- use the Automatic profile
- operate at the maximum radio power permitted in your region
- carry all SSIDs

You can change the radio power and selection of SSIDs or assign the unit to a custom AP profile which defines the entire configuration for the AP.

### To add and configure the discovered AP unit - web-based manager

1. Go to *WiFi Controller > Managed Access Points > Managed FortiAP*.  
This configuration also applies to local WiFi radio on FortiWiFi models.
2. Select the FortiAP unit from the list and edit it.
3. Optionally, enter a *Name*. Otherwise, the unit will be identified by serial number.
4. Select *Authorize*.
5. If you want to use the Automatic profile, adjust its settings if needed:

<b>Enable WiFi Radio</b>	This must be selected to enable operation of this AP.
<b>SSID</b>	<b>Automatically Inherit all SSIDs</b> — AP will carry all WiFi networks. <b>Select SSIDs</b> — select individual SSIDs for this AP to carry.
<b>Auto TX Power Control</b>	If you enable automatic transmitter power control, adjust TX Power Low and TX Power High to set the power range.
<b>Tx Power</b>	If you are not using automatic power control, adjust AP transmitter power. The 100% setting is the maximum permitted in your country. See <a href="#">“Setting your geographic location” on page 795</a>

6. If you want to use a custom AP profile, in *AP Profile* select *Change*, choose the AP Profile. Select *Apply*.
7. Optionally, select *Do not participate in Rogue AP scanning* if scanning adversely affects traffic.
8. Select *OK*.

The physical access point is now added to the system. If the rest of the configuration is complete, it should be possible to connect to the wireless network through the AP.

### To add the discovered AP unit - CLI

First get a list of the discovered access point unit serial numbers:

```
get wireless-controller wtp
```

Add a discovered unit and associate it with AP-profile1, for example:

```
config wireless-controller wtp
 edit FAP22A3U10600118
 set admin enable
 set wtp-profile AP-profile1
 end
```

To use the automatic profile, leave the wtp-profile field unset.

### To modify settings within the Automatic profile - CLI

When wtp-profile is unset (null value), the Automatic profile is in use and some of its settings can be adjusted. This example sets the AP to carry only the employee and guest SSIDs and operate at 80% of maximum power.

```
config wireless-controller wtp
 edit FAP22A3U10600118
 set radio-enable enable
 set vap-all disable
 set vaps employee guest
 set power-level 80
 end
```

### To select a custom AP profile - CLI

```
config wireless-controller wtp
 edit FAP22A3U10600118
 set wtp-profile AP-profile1
 end
```

### To select automatic AP profile - CLI

```
config wireless-controller wtp
 edit FAP22A3U10600118
 unset wtp-profile
 end
```

### To view the status of the added AP unit

```
config wireless-controller wtp
 edit FAP22A3U10600118
 get
```

The join-time field should show a time, not "N/A". See the preceding web-based manager procedure for more information.

## Assigning the same profile to multiple FortiAP units

The same profile can now be applied to multiple managed FortiAP units at the same time. To do this, do the following:

1. Go to *WiFi Controller > Managed Access Points > Managed FortiAPs* to view the AP list.
2. Select all FortiAP units you wish to apply the profile to.
3. Right click on one of the selected FortiAPs and select *Assign Profile*.
4. Choose the profile you wish to apply.

## Checking and updating FortiAP unit firmware

You can view and update the FortiAP unit's firmware from the FortiGate unit that acts as its WiFi controller.

### Checking the FortiAP unit firmware version

Go to *WiFi Controller > Managed Access Points > Managed FortiAP* to view the list of FortiAP units that the FortiGate unit can manage. The *OS Version* column shows the current firmware version running on each AP.

### Updating FortiAP firmware from the FortiGate unit

You can update the FortiAP firmware using either the web-based manager or the CLI. Only the CLI method can update all FortiAP units at once.

#### To update FortiAP unit firmware - web-based manager

1. Go to *WiFi Controller > Managed Access Points > Managed FortiAP*.
2. Select the FortiAP unit from the list and edit it.
3. In *FortiAP OS Version*, select *[Upgrade]*.
4. Select *Browse* and locate the firmware upgrade file.
5. Select *OK*.
6. When the upgrade process completes, select *OK*.

The FortiAP unit restarts.

#### To update FortiAP unit firmware - CLI

1. Upload the FortiAP image to the FortiGate unit.

For example, the Firmware file is `FAP_22A_v4.3.0_b0212_fortinet.out` and the server IP address is `192.168.0.100`.

```
execute wireless-controller upload-wtp-image tftp
 FAP_22A_v4.3.0_b0212_fortinet.out 192.168.0.100
```

If your server is FTP, change `tftp` to `ftp`, and if necessary add your user name and password at the end of the command.

2. Verify that the image is uploaded:

```
execute wireless-controller list-wtp-image
```

3. Upgrade the FortiAP units:

```
exec wireless-controller reset-wtp all
```

If you want to upgrade only one FortiAP unit, enter its serial number instead of `all`.

### Updating FortiAP firmware from the FortiAP unit

You can connect to a FortiAP unit's internal CLI to update its firmware from a TFTP server on the same network. This method does not require access to the wireless controller.

1. Place the FortiAP firmware image on a TFTP server on your computer.
2. Connect the FortiAP unit to a separate private switch or hub or directly connect to your computer via a cross-over cable.
3. Change your computer's IP address to `192.168.1.3`.
4. Telnet to IP address `192.168.1.2`.

This IP address is overwritten if the FortiAP is connected to a DHCP environment. Ensure that the FortiAP unit is in a private network with no DHCP server.

5. Login with the username "admin" and no password.

6. Enter the following command.

For example, the FortiAP image file name is FAP\_22A\_v4.3.0\_b0212\_fortinet.out.

```
restore FAP_22A_v4.3.0_b0212_fortinet.out 192.168.1.3
```

## Advanced WiFi controller discovery

A FortiAP unit can use any of four methods to locate a controller. By default, FortiAP units cycle through all four of the discovery methods. In most cases there is no need to make configuration changes on the FortiAP unit.

There are exceptions. The following section describes the WiFi controller discovery methods in more detail and provides information about configuration changes you might need to make so that discovery will work.

### Controller discovery methods

There are four methods that a FortiAP unit can use to discover a WiFi controller.

#### Static IP configuration

If FortiAP and the controller are not in the same subnet, broadcast and multicast packets cannot reach the controller. The admin can specify the controller's static IP on the AP unit. The AP unit sends a discovery request message in unicast to the controller. Routing must be properly configured in both directions.

##### To specify the controller's IP address on a FortiAP unit

```
cfg -a AC_IPADDR_1="192.168.0.1"
```

By default, the FortiAP unit receives its IP address by DHCP. If you prefer, you can assign the AP unit a static IP address.

##### To assign a static IP address to the FortiAP unit

```
cfg -a ADDR_MODE=STATIC
cfg -a AP_IPADDR="192.168.0.100"
cfg -a AP_NETMASK="255.255.255.0"
```

For information about connecting to the FortiAP CLI, see [“Connecting to the FortiAP CLI” on page 821](#).

#### Broadcast request

The AP unit broadcasts a discovery request message to the network and the controller replies. The AP and the controller must be in the same broadcast domain. No configuration adjustments are required.

#### Multicast request

The AP unit sends a multicast discovery request and the controller replies with a unicast discovery response message. The AP and the controller do not need to be in the same broadcast domain if multicast routing is properly configured.

The default multicast destination address is 224.0.1.140. It can be changed through the CLI. The address must be same on the controller and AP. For information about connecting to the FortiAP CLI, see [“Connecting to the FortiAP CLI” on page 821](#).

### To change the multicast address on the controller

```
config wireless-controller global
 set discovery-mc-addr 224.0.1.250
end
```

### To change the multicast address on a FortiAP unit

```
cfg -a AC_DISCOVERY_MC_ADDR="224.0.1.250"
```

For information about connecting to the FortiAP CLI, see [“Connecting to the FortiAP CLI”](#) on page 821.

## DHCP

If you use DHCP to assign an IP address to your FortiAP unit, you can also provide the WiFi controller IP address at the same time. This is useful if the AP is located remotely from the WiFi controller and other discovery techniques will not work.

When you configure the DHCP server, configure Option 138 to specify the WiFi controller IP address. You need to convert the address into hexadecimal. Convert each octet value separately from left to right and concatenate them. For example, 192.168.0.1 converts to C0A80001.

If Option 138 is used for some other purpose on your network, you can use a different option number if you configure the AP units to match.

### To change the FortiAP DHCP option code

To use option code 139 for example, enter

```
cfg -a AC_DISCOVERY_DHCP_OPTION_CODE=139
```

For information about connecting to the FortiAP CLI, see [“Connecting to the FortiAP CLI”](#) below.

## Connecting to the FortiAP CLI

The FortiAP unit has a CLI through which some configuration options can be set.

### To access the FortiAP unit CLI

1. Connect your computer to the FortiAP directly with a cross-over cable or through a separate switch or hub.
2. Change your computer's IP address to 192.168.1.3
3. Telnet to IP address 192.168.1.2.  
Ensure that FortiAP is in a private network with no DHCP server for the static IP address to be accessible.
4. Login with user name admin and no password.
5. Enter commands as needed.
6. Optionally, use the `passwd` command to assign an administrative password for better security.
7. Save the configuration by entering the following command:

```
cfg -c .
```

8. Unplug the FortiAP and then plug it back in, in order for the configuration to take effect.



When a WiFi controller has taken control of the FortiAP unit, Telnet access to the FortiAP unit's CLI is no longer available.

## Wireless client load balancing for high-density deployments

Wireless load balancing allows your wireless network to distribute wireless traffic more efficiently among wireless access points and available frequency bands. FortiGate wireless controllers support the following types of client load balancing:

- Access Point Hand-off - the wireless controller signals a client to switch to another access point.
- Frequency Hand-off - the wireless controller monitors the usage of 2.4GHz and 5GHz bands, and signals clients to switch to the lesser-used frequency.

Load balancing is not applied to roaming clients.

### Access point hand-off

Access point handoff wireless load balancing involves the following:

- If the load on an access point (ap1) exceeds a threshold (of for example, 30 clients) then the client with the weakest signal will be signaled by wireless controller to drop off and join another nearby access point (ap2).
- When one or more access points are overloaded (for example, more than 30 clients) and a new client attempts to join a wireless network, the wireless controller selects the least busy access point that is closest to the new client and this access point is the one that responds to the client and the one that the client joins.

### Frequency hand-off or band-steering

Encouraging clients to use the 5GHz WiFi band if possible enables those clients to benefit from faster interference-free 5GHz communication. The remaining 2.4GHz clients benefit from reduced interference.

The WiFi controller probes clients to determine their WiFi band capability. It also records the RSSI (signal strength) for each client on each band.

If a new client attempts to join the network, the controller looks up that client's MAC address in its wireless device table and determines if it's a dual band device. If it is not a dual band device, then its allowed to join. If it is a dual band device, then its RSSI on 5GHz is used to determine whether the device is close enough to an access point to benefit from movement to 5GHz frequency.

If both conditions of 1) dual band device and 2) RSSI value is strong, then the wireless controller does not reply to the join request of the client. This forces the client to retry a few more times and then timeout and attempt to join the same SSID on 5GHz. Once the Controller see this new request on 5GHz, the RSSI is again measured and the client is allowed to join. If the RSSI is below threshold, then the device table is updated and the controller forces the client to timeout again. A client's second attempt to connect on 2.4GHz will be accepted.

## Configuration

From the web-based manager edit a custom AP profile and select *Frequency Handoff* and *AP Handoff* as required for each radio on the AP.

From the CLI, you configure wireless client load balancing thresholds for each custom AP profile. Enable access point hand-off and frequency hand-off separately for each radio in the custom AP profile.

```
config wireless-controller wtp-profile
 edit new-ap-profile
 set handoff-rssi <rssi_int>
 set handoff-sta-thresh <clients_int>
 config radio-1
 set frequency-handoff {disable | enable}
 set ap-handoff {disable | enable}
 end
 config radio-2
 set frequency-handoff {disable | enable}
 set ap-handoff {disable | enable}
 end
 end
end
```

Where:

- `handoff-rssi` is the RSSI threshold. Clients with a 5 GHz RSSI threshold over this value are load balanced to the 5GHz frequency band. Default is 25. Range is 20 to 30.
- `handoff-sta-thresh` is the access point handoff threshold. If the access point has more clients than this threshold it is considered busy and clients are changed to another access point. Default is 30, range is 5 to 25.
- `frequency-handoff` enable or disable frequency handoff load balancing for this radio. Disabled by default.
- `ap-handoff` enable or disable access point handoff load balancing for this radio. Disabled by default.

Frequency handoff must be enabled on the 5GHz radio to learn client capability.

## LAN port options

Some FortiAP models have one or more ethernet interfaces marked LAN. These ports can provide wired network access. LAN ports are bridged to either the wired WAN interface or to one of the WiFi SSIDs that the FortiAP unit carries.

### Bridging a LAN port with a FortiAP SSID

Bridging a LAN port with a FortiAP SSID combines traffic from both sources to provide a single broadcast domain for wired and wireless users.

In this configuration

- The IP addresses for LAN clients come from the DHCP server that serves the wireless clients.
- Traffic from LAN clients is bridged to the SSID's VLAN. Dynamic VLAN assignment for hosts on the LAN port is not supported.
- Wireless and LAN clients are on the same network and can communicate locally, via the FortiAP.
- Any host connected to the LAN port will be taken as authenticated. RADIUS MAC authentication for hosts on the LAN port is not supported.

For configuration instructions, see [“Configuring FortiAP LAN ports”](#), below.

## Bridging a LAN port with the WAN port

Bridging a LAN port with the WAN port enables the FortiAP unit to be used as a hub which is also an access point.

In this configuration

- The IP addresses for LAN clients come from the WAN directly and will typically be in the same range as the AP itself.
- All LAN client traffic is bridged directly to the WAN interface.
- Communication between wireless and LAN clients can only occur if a policy on the FortiGate unit allows it.

For configuration instructions, see [“Configuring FortiAP LAN ports”](#), below.

## Configuring FortiAP LAN ports

How you configure FortiAP LAN ports depends on whether you use the automatic AP profile or create custom AP profiles. The Automatic AP profile enables APs to be individually configured. Custom AP profiles are useful if you have multiple APs that are the same model and share the same configuration.

### Configuring LAN ports for an FortiAP unit - web-based manager

If your FortiAP units use the automatic AP profile, you configure each individual unit's LAN ports after you connect it to the FortiGate unit and the FortiAP discovers the WiFi controller. The AP is then listed in the Managed FortiAPs list.

#### To configure a FortiAP unit's LAN ports - web-based manager

1. Go to *WiFi Controller > Managed Access Points > Managed FortiAPs*.  
On FortiGate models 100D, 600C, 800C, and 1000C, go to *WiFi & Switch Controller > Managed Devices > Managed FortiAPs*.
2. Select the FortiAP unit from the list and select *Edit*.
3. In the *LAN Port* section, set *Mode* to *Bridge to* and select an SSID or *WAN Port* as needed.  
On some models with multiple LAN ports, you can set *Mode* to *Custom* and configure the LAN ports individually. Enable each port that you want to use and select an SSID or *WAN Port* as needed.
4. Select *OK*.



## To configure a FortiAP unit's LAN ports - CLI

In this example, a FortiAP unit's configuration is modified to bridge the LAN port to the office SSID.

```
config wireless-controller wtp
 edit FAP11C3X13000412
 config lan
 set port-mode bridge-to-ssid
 set port-ssid office
 end
 end
end
```

In this example, a FortiAP-28C unit's configuration is modified to bridge LAN port1 to the office SSID and to bridge the other LAN ports to the WAN port.

```
config wireless-controller wtp
 edit FAP28C3X13000412
 config lan
 set port1-mode bridge-to-ssid
 set port1-ssid office
 set port2-mode bridge-to-wan
 set port3-mode bridge-to-wan
 set port4-mode bridge-to-wan
 set port5-mode bridge-to-wan
 set port6-mode bridge-to-wan
 set port7-mode bridge-to-wan
 set port8-mode bridge-to-wan
 end
 end
end
```

## Configuring LAN ports in a custom AP profile - web-based manager

When multiple FortiAP units will have the same LAN port configuration, you can create a custom AP profile or edit the default profile for that model.

### To configure FortiAP LAN ports in a custom AP profile - web-based manager

1. Go to *WiFi Controller > WiFi Network > Custom AP Profiles*.
2. Either
  - Edit the default profile for your model,
  - or
  - Select *Create New* and then in *Platform*, select your FortiAP model.
3. In the *LAN Port* section, set mode to *Bridge to* and select an SSID or *WAN Port* as needed.  
On some models with multiple LAN ports, you can select *Custom* to configure the ports individually. Enable each port that you want to use and select an SSID or *WAN Port* as needed.
4. Optionally, adjust other settings. See [“Creating a custom AP Profile” on page 796](#).
5. Select *OK*.

## To configure FortiAP LAN ports in a custom AP profile - CLI

In this example, the default FortiAP-11C profile is configured to bridge the LAN port to the office SSID.

```
config wireless-controller wtp-profile
 edit FAP11C-default
 config lan
 set port-mode bridge-to-ssid
 set port-ssid office
 end
 end
end
```

In this example, the default FortiAP-28C profile is configured to bridge LAN port1 to the office SSID and to bridge the other LAN ports to the WAN port.

```
config wireless-controller wtp-profile
 edit FAP28C-default
 config lan
 set port1-mode bridge-to-ssid
 set port1-ssid office
 set port2-mode bridge-to-wan
 set port3-mode bridge-to-wan
 set port4-mode bridge-to-wan
 set port5-mode bridge-to-wan
 set port6-mode bridge-to-wan
 set port7-mode bridge-to-wan
 set port8-mode bridge-to-wan
 end
 end
end
```

## Preventing IP fragmentation of packets in CAPWAP tunnels

A common problem with controller-based WiFi networks is reduced performance due to IP fragmentation of the packets in the CAPWAP tunnel.

Fragmentation can occur because of CAPWAP tunnel overhead increasing packet size. If the original wireless client packets are close to the maximum transmission unit (MTU) size for the network (usually 1500 bytes for Ethernet networks unless jumbo frames are used) the resulting CAPWAP packets may be larger than the MTU, causing the packets to be fragmented. Fragmented packets can result in data loss, jitter, and decreased throughput.

The FortiOS/FortiAP solution to this problem is to cause wireless clients to send smaller packets to FortiAP devices, resulting in 1500-byte CAPWAP packets and no fragmentation. The following options configure CAPWAP IP fragmentation control:

```
config wireless-controller wtp
 edit new-wtp
 set ip-fragment-preventing {tcp-mss-adjust | icmp-unreachable}
 set tun-mtu-uplink {0 | 576 | 1500}
 set tun-mtu-downlink {0 | 576 | 1500}
 end
end
```

By default, `tcp-mss-adjust` is enabled, `icmp-unreachable` is disabled, and `tun-mtu-uplink` and `tun-mtu-downlink` are set to 0.

To set `tun-mtu-uplink` and `tun-mtu-downlink`, use the default TCP MTU value of 1500. This default configuration prevents packet fragmentation because the FortiAP unit limits the size of TCP packets received from wireless clients so the packets don't have to be ed before CAPWAP encapsulation.

The `tcp-mss-adjust` option causes the FortiAP unit to limit the maximum segment size (MSS) of TCP packets sent by wireless clients. The FortiAP does this by adding a reduced MSS value to the SYN packets sent by the FortiAP unit when negotiating with a wireless client to establish a session. This results in the wireless client sending packets that are smaller than the `tun-mtu-uplink` setting, so that when the CAPWAP headers are added, the CAPWAP packets have an MTU that matches the `tun-mtu-uplink` size.

The `icmp-unreachable` option affects all traffic (UDP and TCP) between wireless clients and the FortiAP unit. This option causes the FortiAP unit to drop packets that have the "Don't Fragment" bit set in their IP header and that are large enough to cause fragmentation and then send an ICMP packet -- type 3 "ICMP Destination unreachable" with code 4 "Fragmentation Needed and Don't Fragment was Set" back to the wireless controller. This should cause the wireless client to send smaller TCP and UDP packets.

# Wireless Mesh

The access points of a WiFi network are usually connected to the WiFi controller through Ethernet wiring. A wireless mesh eliminates the need for Ethernet wiring by connecting WiFi access points to the controller by radio. This is useful where installation of Ethernet wiring is impractical.


The following topics are included in this section:

- [Overview of Wireless Mesh](#)
- [Configuring a meshed WiFi network](#)
- [Configuring a point-to-point bridge](#)

## Overview of Wireless Mesh

Figure 162 shows a wireless mesh topology.

**Figure 162:**Wireless mesh topology



The AP that is connected to the network by Ethernet is called the Mesh Root node. It is configured with an SSID (also called a virtual access point or VAP) dedicated to backhaul communication with the remote FortiAP units. The backhaul SSID carries CAPWAP discovery, configuration, and other communications that would usually be carried on an Ethernet connection. Regular WiFi clients cannot connect to the backhaul SSID. They connect to the regular SSIDs carried on the access points.

The root node can be a FortiAP unit or the built-in AP of a FortiWiFi unit. APs that serve only regular WiFi clients are called Leaf nodes. Leaf APs that also carry the mesh SSID for more distant Leaf nodes are called Leaf/branch nodes.

All access points in a wireless mesh configuration must have at least one of their radios configured to provide mesh backhaul communication. As with wired APs, when mesh APs start up they can be discovered by a FortiGate or FortiWiFi unit WiFi controller and authorized to join the network.

The backhaul SSID delivers the best performance when it is carried on a dedicated radio. On a two-radio FortiAP unit, for example, the 5GHz radio could carry only the backhaul SSID while the 2.4GHz radio carries one or more SSIDs that serve users. Background WiFi scanning is possible in this mode.

The backhaul SSID can also share the same radio with SSIDs that serve users. Performance is reduced because the backhaul and user traffic compete for the available bandwidth. Background WiFi scanning is not available in this mode. One advantage of this mode is that a two-radio AP can offer WiFi coverage on both bands.

The root mesh AP is the AP unit that has a wired Ethernet connection to the WiFi controller. The AP units that are wirelessly linked to the controller over the backhaul SSID are called branch or leaf APs.

## Wireless mesh deployment modes

There are two common wireless mesh deployment modes:

<b>Wireless Mesh</b>	Access points are wirelessly connected to a FortiGate or FortiWiFi unit WiFi controller. WiFi users connect to wireless SSIDs in the same way as on non-mesh WiFi networks.
<b>Wireless bridging</b>	Two LAN segments are connected together over a wireless link (the backhaul SSID). On the leaf AP, the Ethernet connection can be used to provide a wired network. Both WiFi and wired users on the leaf AP are connected to the LAN segment to which the root AP is connected.


## Firmware requirements

All FortiAP units that will be part of the wireless mesh network must be upgraded to FAP firmware version 5.0 build 003. FortiAP-222B units must have their BIOS upgraded to version 400012. The FortiWiFi or FortiGate unit used as the WiFi controller must be running FortiOS 5.0.

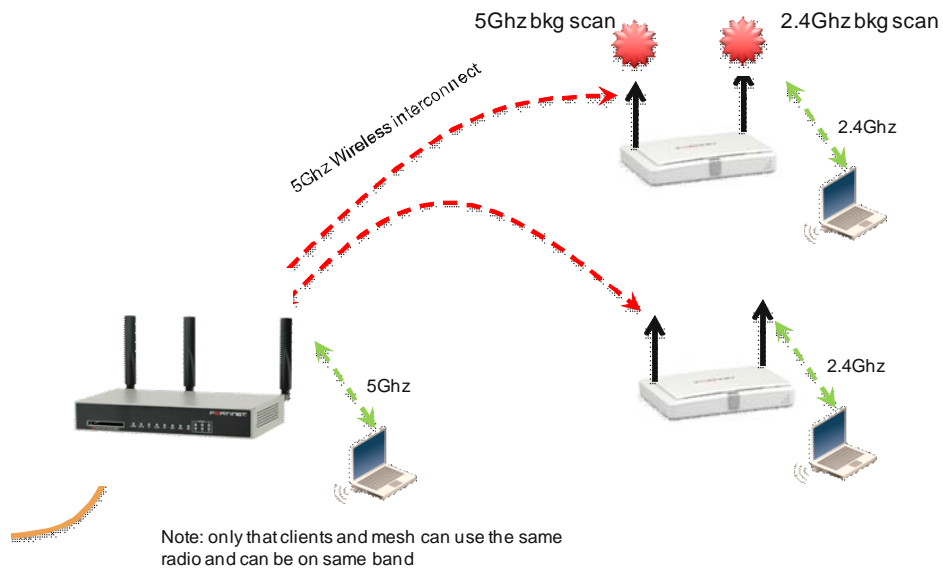
## Types of wireless mesh

A WiFi mesh can provide access to widely-distributed clients. The root mesh AP which is directly connected to the WiFi controller can be either a FortiAP unit or the built-in AP of a FortiWiFi unit that is also the WiFi controller.

**Figure 163:** FortiAP units used as both mesh root AP and leaf AP




**Figure 164:** FortiWiFi unit as root mesh AP with FortiAP units as leaf APs



An alternate use of the wireless mesh functionality is as a point-to-point relay. Both wired and WiFi users on the leaf AP side are connected to the LAN segment on the root mesh side.

**Figure 165:**Point-to-point wireless mesh



## Configuring a meshed WiFi network

Each VDOM on the FortiGate unit contains a predefined WiFi mesh interface named `wl.mesh` and a predefined SSID (which cannot be deleted) named `fortinet.mesh.<vdom-name>`. You can create additional mesh SSIDs. Create the SSID with *Traffic Mode* set to *Mesh Downlink*.

You need to:

- Create custom AP profiles, if you are not using the automatic AP profile.
- Configure the mesh root AP, either a FortiWiFi unit's Local Radio or a FortiAP unit.
- Configure mesh branch/leaf AP units.
- Authorize the mesh branch/leaf units when they connect to the WiFi Controller.

### Creating custom AP profiles

You can apply the automatic AP profile or create one or more custom AP profiles for the mesh root and branch/leaf APs. A custom profile provides more control over which radio channels are used, intrusion protection, load balancing, background rogue AP scanning, and so on. Typically, the custom profiles are configured so that Radio 1 (5GHz) carries the mesh backhaul SSID while Radio 2 (2.4GHz) carries the SSIDs to which users connect.

For more information, see [“Creating a custom AP Profile” on page 796](#).

### Configuring the mesh root AP

The mesh root AP can be either a FortiWiFi unit's built-in AP or a FortiAP unit.

#### To enable a FortiWiFi unit's Local Radio as mesh root - web-based manager

1. Go to *WiFi Controller > Managed Access Points > Local WiFi Radio*.
2. Select *Enable WiFi Radio*.
3. In *SSID*, select *Select SSIDs*, then select *fortinet.mesh.root*.
4. Optionally, adjust *TX Power* or select *Auto Tx Power Control*.

5. Select *Apply*.



In a network with multiple wireless controllers, you need to change the mesh SSID so that each mesh root has a unique SSID. Other controllers using the same mesh root SSID might be detected as fake or rogue APs. Go to *WiFi Controller > WiFi Network > SSID* to change the SSID.

Fortinet also recommends that you create a new preshared key instead of using the default.

### To configure a network interface for the FortiAP unit

1. On the FortiGate unit, go to *System > Network > Interfaces*.
2. Select the interface where you will connect the FortiAP unit and edit it.
3. In *Addressing mode*, select *Manual*.
4. In *IP/Network Mask*, enter an IP address and netmask for the interface.  
To maximize the number of addresses available for clients, the interface address should end with 1, for example 192.168.10.1.
5. In *DHCP Server* select *Enable*.  
An *Address Range* is entered automatically. It consists of the subnet address space above the interface address. For example, if the interface IP/mask is 192.168.10.100/24, the DHCP address range is 192.168.10.101 through 192.168.10.254.
6. Select *OK*.

### To enable a FortiAP unit as mesh root - web-based manager

1. Connect the root FortiAP unit's Ethernet port to the FortiGate network interface that you configured for it. Connect the FortiAP unit to its power source.
2. Go to *WiFi Controller > Managed Access Points > Managed FortiAP*.  
If the root FortiAP unit is not listed, wait 15 seconds and select *Refresh*. Repeat if necessary. If the unit is still missing after a minute or two, power cycle the root FortiAP unit and try again.
3. Select the discovered FortiAP unit and edit its settings.
  - To use the Automatic profile enter:

<b>AP Profile</b>	Automatic
<b>Enable WiFi Radio</b>	Selected
<b>SSID</b>	Select SSID, then enable fortinet.root.mesh.
<b>Tx Power</b>	Optionally, adjust <i>TX Power</i> or select <i>Auto Tx Power Control</i> .

or

- In *AP Profile*, select *Change* and then select the custom AP profile you created for the mesh root AP.
4. In *State*, select *Authorize*.
  5. Select *OK*.

You need to create firewall policies to permit traffic to flow from the network interface where the FortiAP unit is connected to the network interfaces for the Internet and other networks. Enable NAT.



## Configuring the mesh branches or leaves

The FortiAP units that will serve as branch/leaf nodes must be preconfigured.

1. Connect to the FortiAP unit web-based manager on its default Ethernet interface IP address, 192.168.1.2.
2. In the *Connectivity* section enter:

<b>Uplink</b>	Mesh
<b>Mesh AP SSID</b>	fortinet.mesh.<vdom-name> For example, for the root domain, fortinet.mesh.root.
<b>Mesh AP Password</b>	Same as Mesh AP SSID.
<b>Ethernet Bridge</b>	Select

3. Select *Apply* and then select *Logout*.

## Authorizing mesh branch/leaf APs

The pre-configured branch/leaf FortiAP units will connect themselves wirelessly to the WiFi Controller through the mesh network. You must authorize each unit

1. Go to *WiFi Controller > Managed Access Points > Managed FortiAP*. Periodically select *Refresh* until the FortiAP unit is listed.

The *State* of the FortiAP unit should be *Waiting for Authorization*.

2. Open the FortiAP entry for editing.
  - To use the Automatic profile enter:

<b>AP Profile</b>	Automatic
<b>Enable WiFi Radio</b>	Selected
<b>SSID</b>	Optionally, select the SSIDs to make available to users.
<b>Tx Power</b>	Optionally, adjust <i>TX Power</i> or select <i>Auto Tx Power Control</i> .

or

- In *AP Profile*, select *Change* and then select the custom AP profile that you created for the branch/leaf APs and then select *[Apply]*.
3. Select *Authorize*.
  4. Select *OK*.

Initially, the *State* of the FortiAP unit is *Offline*. Periodically select *Refresh* to update the status. Within about two minutes, the state changes to *Online*.

**Figure 166:**FortiWiFi unit as root mesh with FortiAP unit as branch/leaf node

Access Point	State	Connected Via	SSIDs	Channel	Clients
Local WiFi Radio	Online	Ethernet (127.0.0.1)	All	Radio 1: 1	Radio 1: 1
FAP22B3U11005354	Online	Mesh (192.168.3.110)	All	Radio 2: 1	Radio 2: 0

## Viewing the status of the mesh network

Go to *WiFi Controller > Managed Access Points > Managed FortiAP* to view the list of APs. The *Connected Via* field shows *Mesh* for mesh-connected units and lists the IP address to which they connect.

In the FortiAP CLI, you can check the `main ip` field in the output from the command

```
cw_diag -c mesh
```

## Configuring a point-to-point bridge

You can create a point-to-point bridge to connect two wired network segments using a WiFi link. The effect is the same as connecting the two network segments to the same wired switch.

You need to:

- Configure a backhaul link and root mesh AP as described in [“Configuring a meshed WiFi network” on page 831](#). **Note:** The root mesh AP for a point-to-point bridge must be a FortiAP unit, not the internal AP of a FortiWiFi unit.
- Configure bridging on the leaf AP unit.

### To configure the leaf AP unit for bridged operation - FortiAP web-based manager

1. With your browser, connect to the FortiAP unit web-based manager.

You can temporarily connect to the unit's Ethernet port and use its default address: 192.168.1.2.

2. Enter:

<b>Operation Mode</b>	Mesh
<b>Mesh AP SSID</b>	fortinet-ap
<b>Mesh AP Password</b>	fortinet
<b>Ethernet Bridge</b>	Select

3. Select *Apply*.

4. Connect the local wired network to the Ethernet port on the FortiAP unit.

Users are assigned IP addresses from the DHCP server on the wired network connected to the root mesh AP unit.

### To configure a FortiAP unit as a leaf AP - FortiAP CLI

```
cfg -a MESH_AP_SSID=fortinet-ap
cfg -a MESH_AP_PASSWD=fortinet
cfg -a MESH_ETH_BRIDGE=1
cfg -a MESH_AP_TYPE=1
cfg -c
```

# WiFi-Ethernet Bridge Operation

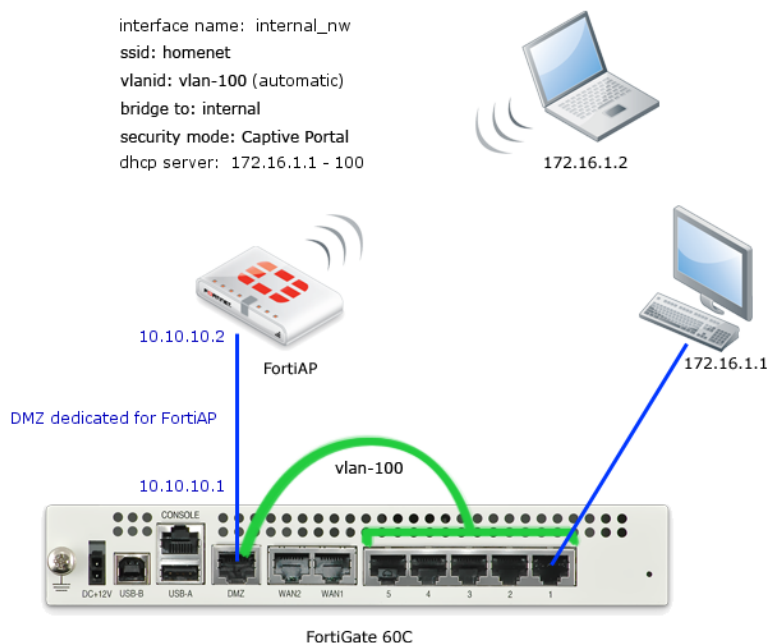
The following topics are included in this section:

- Bridge SSID to FortiGate wired network
- FortiAP local bridging (Private Cloud-Managed AP)
- Using bridged FortiAPs to increase scalability

## Bridge SSID to FortiGate wired network

A WiFi network can be combined with a wired LAN so that WiFi and wired clients are on the same subnet. This is a convenient configuration for users.

**Figure 167:**A FortiAP unit bridged with the internal network



This configuration cannot be used in conjunction with Wireless Mesh features because it enables the FortiAP Local Bridge option.

To create the bridged WiFi and wired LAN configuration, you need to

- Configure the SSID with the Local Bridge option so that traffic is sent directly over the FortiAP unit's Ethernet interface to the FortiGate unit, instead of being tunneled to the WiFi controller.
- Configure a software switch interface on the FortiGate unit with the WiFi and Internal network interfaces as members.
- Configure Captive Portal security for the software switch interface.

### To configure the SSID - web-based manager

1. Go to *WiFi Controller > WiFi Network > SSID* and select *Create New*.
2. Enter:

<b>Interface name</b>	A name for the new WiFi interface, <i>homenet_if</i> for example.
<b>Traffic Mode</b>	Local bridge with FortiAP's Interface
<b>SSID</b>	The SSID visible to users, <i>homenet</i> for example.
<b>Security Mode Data Encryption Preshared Key</b>	Configure security as you would for a regular WiFi network.

3. Select OK.
4. Go to *WiFi Controller > Managed Access Points > Managed FortiAP*, select the FortiAP unit for editing.
5. Authorize the FortiAP unit.
6. The FortiAP unit can carry regular SSIDs in addition to the Bridge SSID.

**Figure 168:** SSID configured with Local Bridge option

The screenshot shows the 'Edit SSID' configuration window. The 'Interface Name' is 'homenet\_if'. The 'Status' is 'Enabled'. The 'Traffic Mode' is 'Local bridge with FortiAP's Interface', which is circled in red. Under 'WiFi Settings', the 'SSID' is 'homenet', 'Security Mode' is 'WPA/WPA2-Personal', 'Data Encryption' is 'AES', and 'Pre-shared Key' is masked with dots. There are checkboxes for 'Block Intra-SSID Traffic' and 'Limit Concurrent WiFi Clients', both of which are unchecked. A 'Comments' field is at the bottom with a character count of 0/255. 'OK' and 'Cancel' buttons are at the bottom right.

### To configure the SSID - CLI

This example creates a WiFi interface "homenet\_if" with SSID "homenet" using WPA-Personal security, passphrase "Fortinet1".

```
config wireless-controller vap
 edit "homenet_if"
 set vdom "root"
 set ssid "homenet"
 set local-bridging enable
 set security wpa-personal
 set passphrase "Fortinet1"
 end
config wireless-controller wtp
 edit FAP22B3U11005354
 set admin enable
 set vaps "homenet_if"
 end
```

### To configure the FortiGate unit - web-based manager

1. Go to *System > Network > Interfaces* and select *Create New*.
2. Enter:

<b>Name</b>	A name for the new interface, homenet_nw for example.
<b>Type</b>	Software Switch
<b>Interface Members</b>	Move internal and homenet_if into the <i>Selected Interfaces</i> list.
<b>Addressing Mode</b>	Select Manual and enter an address, for example 172.16.96.32/255.255.255.0
<b>Enable DHCP Server</b>	Enable.
<b>Security Mode</b>	Select <i>Captive Portal</i> . Add the permitted <i>User Groups</i> .

3. Select OK.

### To configure the FortiGate unit - CLI

```
config system interface
 edit homenet_nw
 set ip 172.16.96.32 255.255.255.0
 set type switch
 set security-mode captive-portal
 set security-groups "Guest-group"
 end
config system interface
 edit homenet_nw
 set member "homenet_if" "internal"
 end
```

## VLAN configuration

If your environment uses VLAN tagging, you assign the SSID to a specific VLAN in the CLI. For example, to assign the `homenet_if` interface to VLAN 100, enter:

```
config wireless-controller vap
 edit "homenet_if"
 set vlanid 100
 end
```

## Additional configuration

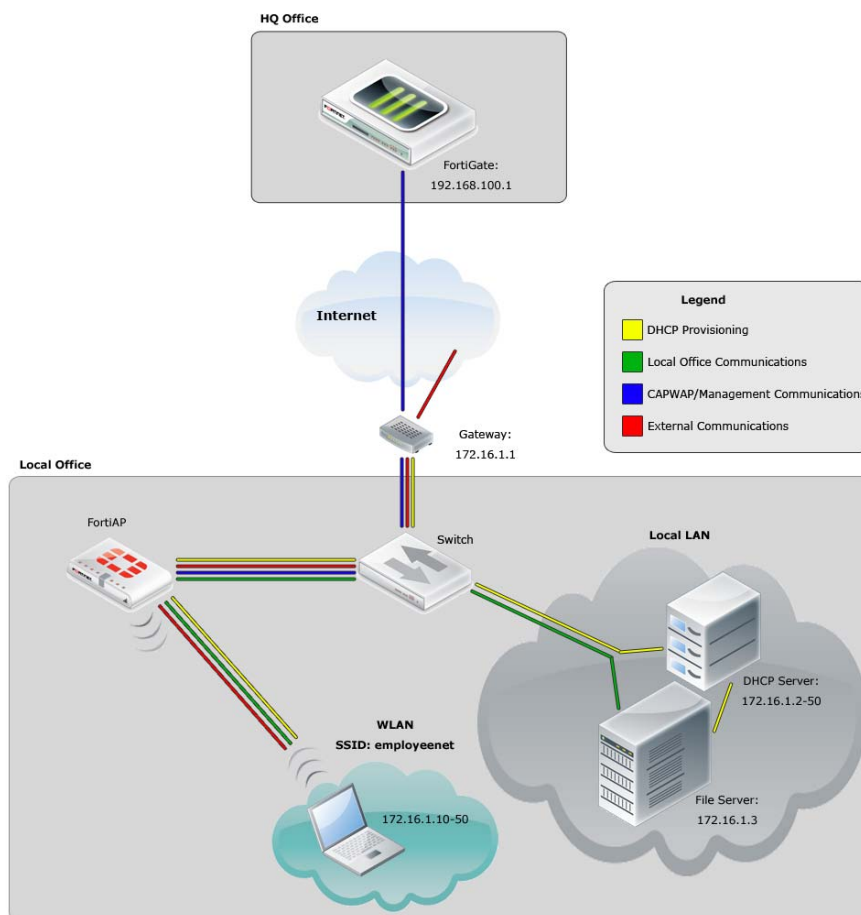
The configuration described above provides communication between WiFi and wired LAN users only. To provide access to other networks, create appropriate firewall policies between the software switch and other interfaces.

## FortiAP local bridging (Private Cloud-Managed AP)

A FortiAP unit can provide WiFi access to a LAN, even when the wireless controller is located remotely. This configuration is useful for the following situations:

- Installations where the WiFi controller is remote and most of the traffic is local or uses the local Internet gateway
- Wireless-PCI compliance with remote WiFi controller
- Telecommuting, where the FortiAP unit has the WiFi controller IP address pre-configured and broadcasts the office SSID in the user's home or hotel room. In this case, data is sent in the wireless tunnel across the Internet to the office and you should enable encryption using DTLS.

**Figure 169:**Remotely-managed FortiAP providing WiFi access to local network



On the remote FortiGate wireless controller, the WiFi SSID is created with the *Bridge with FortiAP Interface* option selected. In this mode, no IP addresses are configured. The FortiAP unit's WiFi and Ethernet interfaces behave as a switch. WiFi client devices obtain IP addresses from the same DHCP server as wired devices on the LAN.

There can be only one Bridge mode SSID per FortiAP unit.



The Local Bridge feature cannot be used in conjunction with Wireless Mesh features. *Block-Intra-SSID Traffic* is not available in Bridge mode.

### To configure a FortiAP local bridge - web-based manager

1. Go to *WiFi Controller > WiFi Network > SSID* and select *Create New*.
2. Enter:

<b>Interface name</b>	A name for the new WiFi interface.
<b>Traffic Mode</b>	Local bridge with FortiAP's Interface
<b>SSID</b>	The SSID visible to users.
<b>Security Mode Data Encryption Preshared Key</b>	Configure security as you would for a regular WiFi network.

3. Select *OK*.
4. Go to *WiFi Controller > Managed Access Points > Managed FortiAP*, select the FortiAP unit for editing.
5. Authorize the FortiAP unit.
6. The FortiAP unit can carry regular SSIDs in addition to the Bridge SSID.

**Figure 170:**SSID configured for Local Bridge operation

The screenshot shows the 'New Interface' configuration dialog. The 'Traffic Mode' dropdown menu is highlighted with a red circle and contains the option 'Local bridge with FortiAP's Inter...'. Other visible settings include: Name: branchbridge, Type: WiFi SSID, SSID: LANbridge, Security Mode: WPA/WPA2-Personal, Data Encryption: AES (selected), Pre-shared Key: 8-63 characters, and a Comments field at the bottom.



## To configure a FortiAP local bridge - CLI

This example creates a WiFi interface “branchbridge” with SSID “LANbridge” using WPA-Personal security, passphrase “Fortinet1”.

```
config wireless-controller vap
 edit "branchbridge"
 set vdom "root"
 set ssid "LANbridge"
 set local-bridging enable
 set security wpa-personal
 set passphrase "Fortinet1"
 end
config wireless-controller wtp
 edit FAP22B3U11005354
 set admin enable
 set vaps "branchbridge"
 end
```

## Continued FortiAP operation when WiFi controller connection is down

The wireless controller, or the connection to it, might occasionally become unavailable. During such an outage, clients already associated with a bridge mode FortiAP unit continue to have access to the WiFi and wired networks. Optionally, the FortiAP unit can also continue to authenticate users if the SSID meets these conditions:

- *Traffic Mode is Local bridge with FortiAP's Interface.*  
In this mode, the FortiAP unit does not send traffic back to the wireless controller.
- *Security Mode is either WPA/WPA2-Personal or Open.*  
These modes do not require the user database. In WPA/WPA2-Personal authentication, all clients use the same pre-shared key which is known to the FortiAP unit.
- *Allow new client association when controller connection is down is enabled.*  
This field is available only if the other conditions have been met.

The “LANbridge” SSID example would be configured like this in the CLI:

```
config wireless-controller vap
 edit "branchbridge"
 set vdom "root"
 set ssid "LANbridge"
 set local-bridging enable
 set security wpa-personal
 set passphrase "Fortinet1"
 set local-authentication enable
 end
```

## Using bridged FortiAPs to increase scalability

The FortiGate wireless controller can support more FortiAP units in local bridge mode than in the normal mode. But this is only true if you configure some of your FortiAP units to operate in remote mode, which supports only local bridge mode SSIDs.

The Managed FortiAP page (*WiFi Controller > Managed Devices > Managed FortiAP*) shows at the top right the current number of Managed FortiAPs and the maximum number that can be managed, “5/64” for example. The maximum number, however, is true only if all FortiAP units operate in remote mode. For more detailed information, consult the [Maximum Values Table](#). For each FortiGate model, there are two maximum values for managed FortiAP units: the total number of FortiAPs and the number of FortiAPs that can operate in normal mode.

### To configure FortiAP units for remote mode operation

1. Create at least one SSID with *Traffic Mode* set to *Local Bridge*.
2. Create a custom AP profile that includes only local bridge SSIDs.
3. Configure each managed FortiAP unit to use the custom AP profile. You also need to set the FortiAP unit’s `wtp-mode` to `remote`, which is possible only in the CLI. The following example uses the CLI both to set `wtp-mode` and select the custom AP profile:

```
config wireless-controller wtp
 edit FAP22B3U11005354
 set wtp-mode remote
 set wtp-profile 220B_bridge
 end
```

# Protecting the WiFi Network

The FortiGate unit provides WiFi-specific network protection. The following topics are included in this section:

- [Wireless IDS](#)
- [WiFi data channel encryption](#)

## Wireless IDS

The FortiGate Wireless Intrusion Detection System (WIDS) monitors wireless traffic for a wide range of security threats by detecting and reporting on possible intrusion attempts. When an attack is detected the FortiGate unit records a log message.

You can create a WIDS profile to enable the types of intrusion detection:

- **Asleep Attack—ASLEAP** is a tool used to perform attacks against LEAP authentication.
- **Association Frame Flooding**—A Denial of Service attack using association requests. The default detection threshold is 30 requests in 10 seconds.
- **Authentication Frame Flooding**—A Denial of Service attack using association requests. The default detection threshold is 30 requests in 10 seconds.
- **Broadcasting De-authentication**—This is a type of Denial of Service attack. A flood of spoofed de-authentication frames forces wireless clients to de-authenticate, then re-authenticate with their AP.
- **EAPOL Packet Flooding**—Extensible Authentication Protocol over LAN (EAPOL) packets are used in WPA and WPA2 authentication. Flooding the AP with these packets can be a denial of service attack. Several types of EAPOL packets are detected: EAPOL-FAIL, EAPOL-LOGOFF, EAPOL-Start, EAPOL-SUCC.
- **Invalid MAC OUI**—Some attackers use randomly-generated MAC addresses. The first three bytes of the MAC address are the Organizationally Unique Identifier (OUI), administered by IEEE. Invalid OUIs are logged.
- **Long Duration Attack**—To share radio bandwidth, WiFi devices reserve channels for brief periods of time. Excessively long reservation periods can be used as a denial of service attack. You can set a threshold between 1000 and 32 767 microseconds. The default is 8200.
- **Null SSID Probe Response**—When a wireless client sends out a probe request, the attacker sends a response with a null SSID. This causes many wireless cards and devices to stop responding.
- **Spoofed De-authentication**—Spoofed de-authentication frames form the basis for most denial of service attacks.
- **Weak WEP IV Detection**—A primary means of cracking WEP keys is by capturing 802.11 frames over an extended period of time and searching for patterns of WEP initialization vectors (IVs) that are known to be weak. WIDS detects known weak WEP IVs in on-air traffic.
- **Wireless Bridge**—WiFi frames with both the fromDS and ToDS fields set indicate a wireless bridge. This will also detect a wireless bridge that you intentionally configured in your network.

You can enable wireless IDS by going to *WiFi Controller > WiFi Network > Custom AP Profile* and editing an access point profile. Inside the profile set *WIDS Profile* to the name of a wireless IDS profile to apply wireless IDS protection to the access points that uses the profile.

**To create a WIDS Profile**

1. Go to *WiFi Controller > WiFi Network > WIDS Profile*.
2. Select a profile to edit or select *Create New*.
3. Select the types of intrusion to protect against.  
By default, all types are selected.
4. Select *Apply*.

You can also configure a WIDS profile in the CLI using the `config wireless-controller wids-profile` command.

## WiFi data channel encryption

Optionally, you can apply DTLS encryption to the data channel between the wireless controller and FortiAP units. This enhances security.

There are data channel encryption settings on both the FortiGate unit and the FortiAP units. At both ends, you can enable Clear Text, DTLS encryption, or both. The settings must agree or the FortiAP unit will not be able to join the WiFi network. By default, both Clear Text and DTLS-encrypted communication are enabled on the FortiAP unit, allowing the FortiGate setting to determine whether data channel encryption is used. If the FortiGate unit also enables both Clear Text and DTLS, Clear Text is used.

Data channel encryption settings are located in the Custom AP profile. If you use Automatic profile, only Clear Text is supported.



Data channel encryption is software-based and can affect performance. Verify that the system meets your performance requirements with encryption enabled.

### Configuring encryption on the FortiGate unit

You can use the CLI to configure data channel encryption.

#### Enabling encryption

In the CLI, the `wireless wtp-profile` command contains a new field, `dtls-policy`, with options `clear-text` and `dtls-enabled`. To enable encryption in `profile1` for example, enter:

```
config wireless-controller wtp-profile
 edit profile1
 set dtls-policy dtls-enabled
 end
```

### Configuring encryption on the FortiAP unit

The FortiAP unit has its own settings for data channel encryption.

#### Enabling CAPWAP encryption - FortiAP web-based manager

- 1 On the *System Information* page, in *WTP Configuration > AC Data Channel Security*, select one of:
  - Clear Text
  - DTLS Enabled
  - Clear Text or DTLS Enabled (default)
- 2 Select *Apply*.

#### Enabling encryption - FortiAP CLI

You can set the data channel encryption using the `AC_DATA_CHAN_SEC` variable: 0 is Clear Text, 1 is DTLS Enabled, 2 (the default) is Clear Text or DTLS Enabled.

For example, to set security to DTLS and then save the setting, enter

```
cfg -a AC_DATA_CHAN_SEC=1
cfg -c
```

# Wireless network monitoring

You can monitor both your wireless clients and other wireless networks that are available in your coverage area.

The following topics are included in this section:

- [Monitoring wireless clients](#)
- [Monitoring rogue APs](#)
- [Suppressing rogue APs](#)
- [Monitoring wireless network health](#)

## Monitoring wireless clients

**To view connected clients on a FortiWiFi unit**

- Go to *WiFi Controller > Monitor > Client Monitor*.

The following information is displayed:

<b>SSID</b>	The SSID that the client connected to.
<b>FortiAP</b>	The serial number of the FortiAP unit to which the client connected.
<b>User</b>	User name
<b>IP</b>	The IP address assigned to the wireless client.
<b>Device</b>	
<b>Auth</b>	The type of authentication used.
<b>Channel</b>	WiFi radio channel in use.
<b>Bandwidth Tx/Rx</b>	Client received and transmitted bandwidth, in Kbps.
<b>Signal Strength / Noise</b>	The signal-to-noise ratio in deciBels calculated from signal strength and noise level.
<b>Signal Strength</b>	
<b>Association Time</b>	How long the client has been connected to this access point.

Results can be filtered. Select the filter icon on the column you want to filter. Enter the values to include or select NOT if you want to exclude the specified values.

## Monitoring rogue APs

The access point radio equipment can scan for other available access points, either as a dedicated monitor or as a background scan performed while the access point is idle.

Discovered access points are listed in the *Rogue AP Monitor* list. You can then mark them as either Accepted or Rogue access points. This designation helps you to track access points. It does not affect anyone's ability to use these access points.

Rogue AP Settings (*WiFi Controller > WiFi Network > Rogue AP Settings*) control rogue AP detection for APs that use the Automatic profile. For APs that use a custom AP profile, rogue AP detection settings are contained in the custom AP profile (*WiFi Controller > WiFi Network > Custom AP Profile*).

It is also possible to suppress rogue APs. See [“Suppressing rogue APs” on page 851](#).

### On-wire rogue AP detection technique

Other APs that are available in the same area as your own APs are not necessarily rogues. A neighboring AP that has no connection to your network might cause interference, but it is not a security threat. A rogue AP is an unauthorized AP connected to your wired network. This can enable unauthorized access. When rogue AP detection is enabled, the *On-wire* column in the *Rogue AP Monitor* list shows a green up-arrow on detected rogues.

Rogue AP monitoring of WiFi client traffic builds a table of WiFi clients and the Access Points that they are communicating through. The FortiGate unit also builds a table of MAC addresses that it sees on the LAN. The FortiGate unit's on-wire correlation engine constantly compares the MAC addresses seen on the LAN to the MAC addresses seen on the WiFi network.

There are two methods of Rogue AP on-wire detection operating simultaneously: Exact MAC address match and MAC adjacency.

#### Exact MAC address match

If the same MAC address is seen on the LAN and on the WiFi network, this means that the wireless client is connected to the LAN. If the AP that the client is using is not authorized in the FortiGate unit configuration, that AP is deemed an 'on-wire' rogue. This scheme works for non-NAT rogue APs.

#### MAC adjacency

If an access point is also a router, it applies NAT to WiFi packets. This can make rogue detection more difficult. However, an AP's WiFi interface MAC address is usually in the same range as its wired MAC address. So, the MAC adjacency rogue detection method matches LAN and WiFi network MAC addresses that are within a defined numerical distance of each other. By default, the MAC adjacency value is 7. If the AP for these matching MAC addresses is not authorized in the FortiGate unit configuration, that AP is deemed an 'on-wire' rogue.

#### Limitations

On-wire rogue detection has some limitations. There must be at least one WiFi client connected to the suspect AP and continuously sending traffic. If the suspect AP is a router, its WiFi MAC address must be very similar to its Ethernet port MAC address.

#### Logging

Information about detected rogue APs is logged and uploaded to your FortiAnalyzer unit, if you have one. By default, rogue APs generate an alert level log, unknown APs generate a warning level log. This log information can help you with PCI-DSS compliance requirements.

## Rogue AP scanning as a background activity

Each WiFi radio can perform monitoring of radio channels in its operating band while acting as an AP. It does this by briefly switching from AP to monitoring mode. By default, a scan period starts every 300 seconds. Each second a different channel is monitored for 20ms until all channels have been checked.

During heavy AP traffic, it is possible for background scanning to cause lost packets when the radio switches to monitoring. To reduce the probability of lost packets, you can set the CLI `ap-bgscan-idle` field to delay the switch to monitoring until the AP has been idle for a specified period. This means that heavy AP traffic may slow background scanning.

The following CLI example configures default background rogue scanning operation except that it sets `ap-bgscan-idle` to require 100ms of AP inactivity before scanning the next channel.

```
config wireless-controller wtp-profile
 edit ourprofile
 config radio-1
 set ap-bgscan enable
 set rogue-scan enable
 set ap-bgscan-period 300
 set ap-bgscan-intv 1
 set ap-bgscan-duration 20
 set ap-bgscan-idle 100
 end
 end
```

## Configuring rogue scanning

All APs using the Automatic profile share the same rogue scanning settings. APs using a custom AP profile follow the settings in the custom profile.

### To enable rogue AP scanning for the automatic AP profile

1. Go to *WiFi Controller > WiFi Network > Rogue AP Settings*.
2. Select *Enable Rogue AP Detection*.
3. Select *Enable On-wire Rogue AP Detection Technique* if you want to use that method of distinguishing rogues from neighbors.
4. Select *Apply*.

### To enable the rogue AP scanning feature for the automatic AP profile - CLI

```
config wireless-controller setting
 set ap-scan enable
 set on-wire-scan enable
end
```

### To configure rogue AP scanning in a custom AP profile

1. Go to *WiFi Controller > WiFi Network > Custom AP Profiles*.  
On some models, the menu is *WiFi & Switch Controller*.
2. Select an existing AP profile and edit it, or select *Create New*.
3. For each radio, select either *Access Point* or *Dedicated Monitor*, as required.
4. If you selected *Access Point*, enable *Background Scan*.
5. Select *Rogue AP On-Wire Scan*.



6. If needed, modify other settings.
7. Select *OK*.

#### To enable the rogue AP scanning feature in a custom AP profile - CLI

```
config wireless-controller wtp-profile
 edit FAP220B-default
 config radio-1
 set mode ap
 set ap-bgscan enable
 set rogue-scan enable
```

### Exempting an AP from rogue scanning

By default, if Rogue AP Detection is enabled, it is enabled on all managed FortiAP units. Optionally, you can exempt an AP from scanning. You should be careful about doing this if your organization must perform scanning to meet PCI-DSS requirements.

#### To exempt an AP from rogue scanning - web-based manager

1. Go to *WiFi Controller > Managed Access Points > Managed FortiAP*.
2. Select which AP to edit.
3. Select *Do not participate in Rogue AP Scanning* and then select *OK*.

#### To exempt an AP from rogue scanning - CLI

This example shows how to exempt access point AP1 from rogue scanning.

```
config wireless-controller wtp
 edit AP1
 set ap-scan disable
 end
```

### MAC adjacency

You can adjust the maximum WiFi to Ethernet MAC difference used when determining whether an suspect AP is a rogue.

#### To adjust MAC adjacency

For example, to change the adjacency to 8, enter








```
config wireless-controller global
 set rogue-scan-mac-adjacency 8
end
```

## Using the Rogue AP Monitor

Go to *WiFi Controller > Monitor > Rogue AP Monitor* to view the list of other wireless access points that are receivable at your location.

### Information Columns

Actual columns displayed depends on *Column Settings*.

<b>State</b>	 Rogue AP — Use this status for unauthorized APs that <i>On-wire</i> status indicates are attached to your wired networks.  Accepted AP — Use this status for APs that are an authorized part of your network or are neighboring APs that are not a security threat. To see accepted APs in the list, select <i>Show Accepted</i> .  Unclassified — This is the initial status of a discovered AP. You can change an AP back to unclassified if you have mistakenly marked it as Rogue or Accepted.
<b>Online Status</b>	 Active AP  Inactive AP  Active ad-hoc WiFi device  Inactive ad-hoc WiFi device
<b>SSID</b>	The wireless service set identifier (SSID) or network name for the wireless interface.
<b>Security Type</b>	The type of security currently being used.
<b>Channel</b>	The wireless radio channel that the access point uses.
<b>MAC Address</b>	The MAC address of the Wireless interface.
<b>Vendor Info</b>	The name of the vendor.
<b>Signal Strength</b>	The relative signal strength of the AP. Mouse over the symbol to view the signal-to-noise ratio.
<b>Detected By</b>	The name or serial number of the AP unit that detected the signal.
<b>On-wire</b>	A green up-arrow indicates a suspected rogue, based on the on-wire detection technique. A red down-arrow indicates AP is not a suspected rogue.
<b>First Seen</b>	How long ago this AP was first detected.
<b>Last Seen</b>	How long ago this AP was last detected.
<b>Rate</b>	Data rate in bps.

To change the Online Status of an AP, right-click it and select *Mark Accepted* or *Mark Rogue*.

## Suppressing rogue APs

In addition to monitoring rogue APs, you can actively prevent your users from connecting to them. When suppression is activated against an AP, the FortiGate WiFi controller sends deauthentication messages to the rogue AP's clients, posing as the rogue AP, and also sends deauthentication messages to the rogue AP, posing as its clients. This is done using the monitoring radio.

To enable rogue AP suppression, you must enable monitoring of rogue APs with the on-wire detection technique. See [“Monitoring rogue APs” on page 847](#). The monitoring radio must be in the Dedicated Monitor mode.

### To activate AP suppression against a rogue AP

1. Go to *WiFi Controller > Monitor > Rogue AP Monitor*.
2. When you see an AP listed that is a rogue detected “on-wire”, select it and then select *Mark > Mark Rogue*.
3. To suppress an AP that is marked as a rogue, select it and then select *Suppress AP*.

### To deactivate AP suppression

1. Go to *WiFi Controller > Monitor > Rogue AP Monitor*.
2. Select the suppressed rogue AP and then select *Suppress AP > Unsuppress AP*.

## Monitoring wireless network health

The Wireless Health Dashboard provides a comprehensive view of the health of your network's wireless infrastructure. The dashboard includes widgets to display

- AP Status - Active, Down or missing, up for over 24 hours, rebooted in past 24 hours
- Client Count Over Time - viewable for past hour, day, or 30 days
- Top Client Count Per-AP - separate widgets for 2.4GHz and 5GHz bands
- Top Wireless Interference - separate widgets for 2.4GHz and 5GHz bands
- Login Failures Information

To view the Wireless Health dashboard, go to *WiFi Controller > Monitor > Wireless Health*.

# Configuring wireless network clients

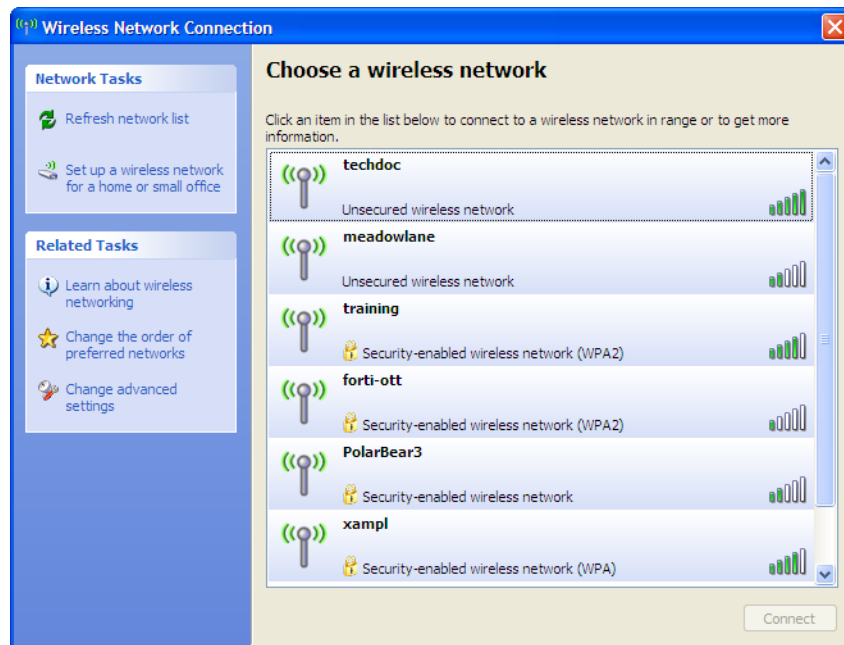
This chapter shows how to configure typical wireless network clients to connect to a wireless network with WPA-Enterprise security. The following topics are included in this section:

- [Windows XP client](#)
- [Windows 7 client](#)
- [Mac OS client](#)
- [Linux client](#)
- [Troubleshooting](#)

## Windows XP client

### To configure the WPA-Enterprise network connection

1. In the Windows Start menu, go to *Control Panel > Network Connections > Wireless Network Connection* or select the wireless network icon in the Notification area of the Taskbar. A list of available networks is displayed.

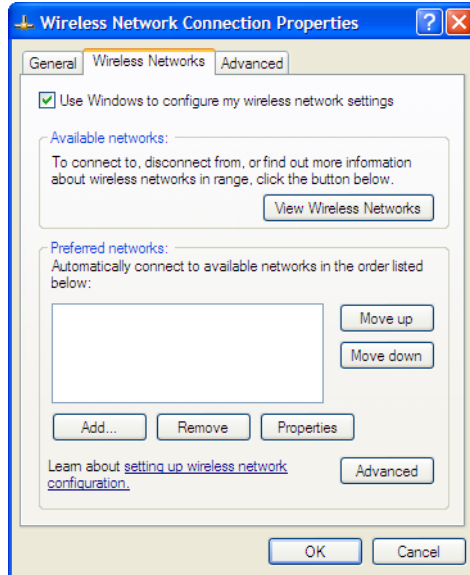


If you are already connected to another wireless network, the Connection Status window displays. Select *View Wireless Networks* on the *General* tab to view the list.

If the network broadcasts its SSID, it is listed. But do not try to connect until you have completed the configuration step below. Because the network doesn't use the Windows XP default security configuration, configure the client's network settings manually before trying to connect.

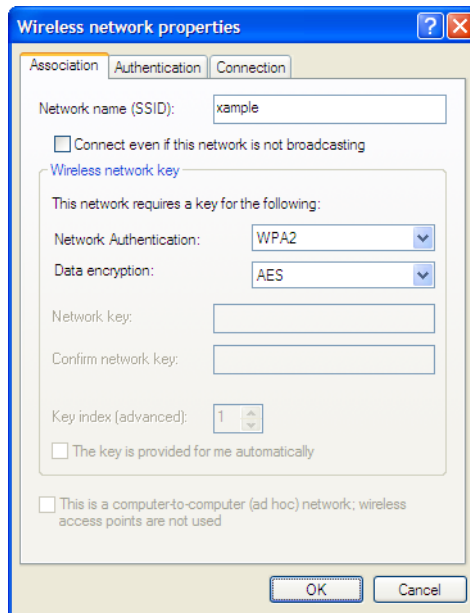
2. You can configure the WPA-Enterprise network to be accessible from the *View Wireless Networks* window even if it does not broadcast its SSID.

3. Select *Change Advanced Settings* and then select the *Wireless Networks* tab.



Any existing networks that you have already configured are listed in the *Preferred Networks* list.

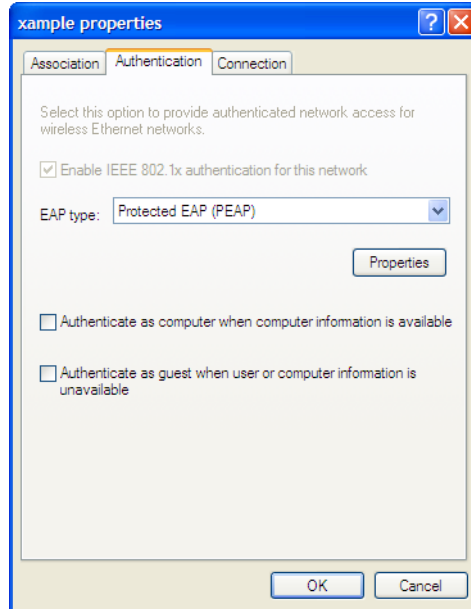
4. Select *Add* and enter the following information:



<b>Network Name (SSID)</b>	The SSID for your wireless network
<b>Network Authentication</b>	WPA2
<b>Data Encryption</b>	AES

5. If this wireless network does not broadcast its SSID, select *Connect even if this network is not broadcasting* so that the network will appear in the *View Wireless Networks* list.

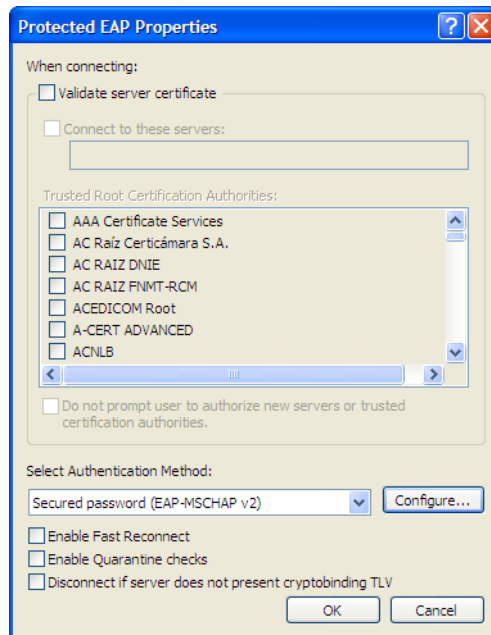
6. Select the *Authentication* tab.



7. In *EAP Type*, select *Protected EAP (PEAP)*.

8. Make sure that the other two authentication options are not selected.

9. Select *Properties*.




10. Make sure that *Validate server\_certificate* is selected.

11. Select the server certificate *UTN-USERFirst-Hardware*.

12. In *Select Authentication Method*, select *Secured Password (EAP-MSCHAPv2)*.

13. Ensure that the remaining options are not selected.

14. Select *Configure*.



15. If your wireless network credentials are the same as your Windows logon credentials, select *Automatically use my Windows logon name and password*. Otherwise, make sure that this option is not selected.


16. Select OK. Repeat until you have closed all of the *Wireless Network Connection Properties* windows.

**To connect to the WPA-Enterprise wireless network**

1. Select the wireless network icon in the Notification area of the Taskbar.
2. In the *View Wireless Networks* list, select the network you just added and then select *Connect*.

You might need to log off of your current wireless network and refresh the list.

3. When the following popup displays, click on it.




4. In the *Enter Credentials* window, enter your wireless network *User name*, *Password*, and *Logon domain* (if applicable). Then, select OK.



In future, Windows will automatically send your credentials when you log on to this network.

# Windows 7 client

1. In the Windows Start menu, go to *Control Panel > Network and Internet > Network and Sharing Center > Manage Wireless Networks* or select the wireless network icon in the Notification area of the Taskbar. A list of available networks is displayed.



2. Do one of the following:
  - If the wireless network is listed (it broadcasts its SSID), select it from the list.
  - Select *Add > Manually create a network profile*.
3. Enter the following information and select *Next*.




<b>Network name</b>	Enter the SSID of the wireless network. (Required only if you selected <i>Add</i> .)
<b>Security type</b>	WPA2-Enterprise
<b>Encryption type</b>	AES
<b>Start this connection automatically</b>	Select
<b>Connect even if the network is not broadcasting.</b>	Select

The Wireless Network icon will display a popup requesting that you click to enter credentials for the network. Click on the popup notification.

4. In the *Enter Credentials* window, enter your wireless network *User name*, *Password*, and *Logon domain* (if applicable). Then, select *OK*.



5. Select *Change connection settings*.
6. On the *Connection* tab, select *Connect automatically when this network is in range*.
7. On the *Security* tab, select the Microsoft PEAP authentication method and then select *Settings*.
8. Make sure that *Validate server\_certificate* is selected.
9. Select the server certificate *UTN-USERFirst-Hardware*.
10. In *Select Authentication Method*, select *Secured Password (EAP-MSCHAPv2)*.
11. Select *Configure*.

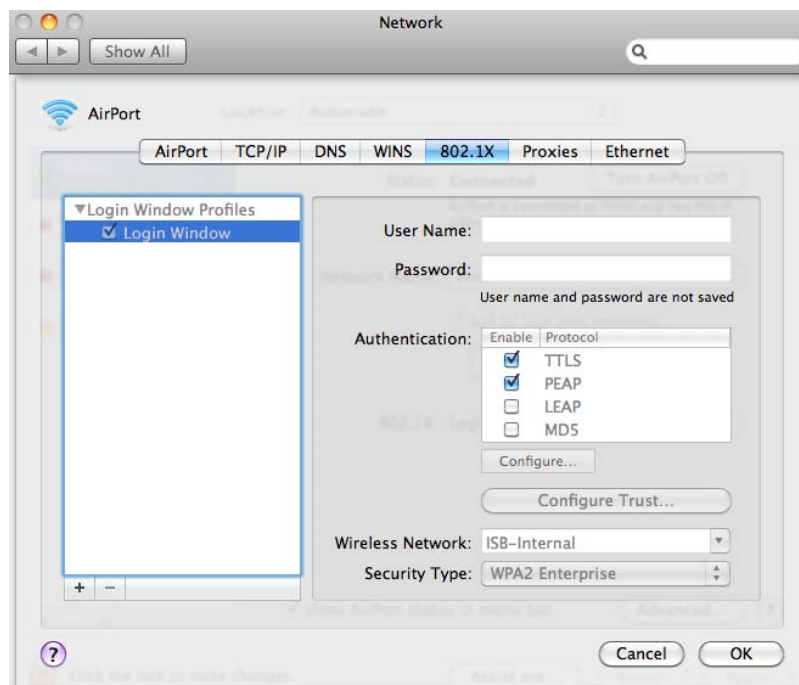


12. If your wireless network credentials are the same as your Windows logon credentials, select *Automatically use my Windows logon name and password*. Otherwise, make sure that this option is not selected.
13. Ensure that the remaining options are not selected.
14. Select OK. Repeat until you have closed all of the *Wireless Network Properties* windows.

## Mac OS client

### To configure network preferences

1. Right-click the *AirPort* icon in the toolbar and select *Open Network Preferences*.
2. Select *Advanced* and then select the *802.1X* tab.



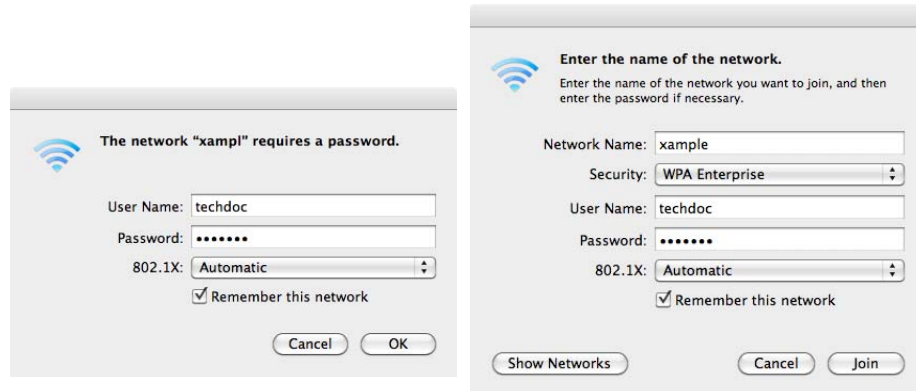
3. If there are no Login Window Profiles in the left column, select the + button and then select *Add Login Window Profile*.

4. Select the Login Window Profile and then make sure that both TTLS and PEAP are selected in *Authentication*.

**To configure the WPA-Enterprise network connection**

1. Select the *AirPort* icon in the toolbar.
2. Do one of the following:
  - If the network is listed, select the network from the list.
  - Select *Connect to Other Network*.

One of the following windows opens, depending on your selection.



3. Enter the following information and select *OK* or *Join*:

<b>Network name</b>	Enter the SSID of your wireless network. (Other network only)
<b>Wireless Security</b>	WPA Enterprise
<b>802.1X</b>	Automatic
<b>Username Password</b>	Enter your logon credentials for the wireless network.
<b>Remember this network</b>	Select.

You are connected to the wireless network.



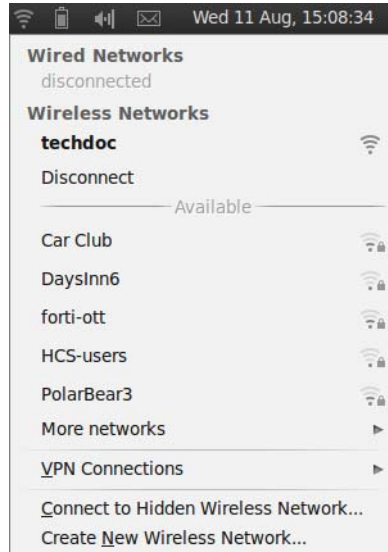
Mac OS supports only PEAP with MSCHAPv2 authentication and therefore can authenticate only to a RADIUS server, not an LDAP or TACACS+ server.

# Linux client

This example is based on the Ubuntu 10.04 Linux wireless client.

## To connect to a WPA-Enterprise network

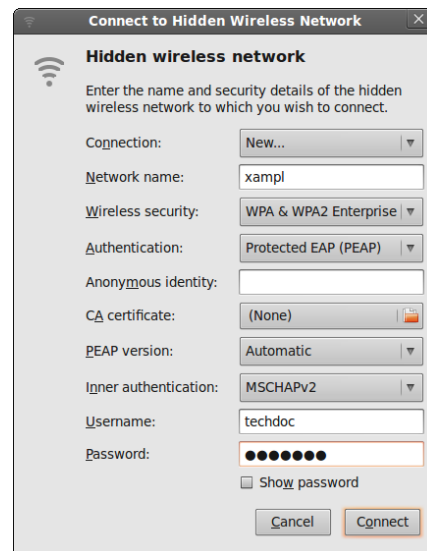
1. Select the Network Manager icon to view the Wireless Networks menu.



Wireless networks that broadcast their SSID are listed in the *Available* section of the menu. If the list is long, it is continued in the *More Networks* submenu.

2. Do one of the following:
  - Select the network from the list (also check *More Networks*).
  - Select *Connect to Hidden Wireless Network*.

One of the following windows opens, depending on your selection.



3. Enter the following information:

<b>Connection</b>	Leave as <i>New</i> . (Hidden network only)
<b>Network name</b>	Enter the SSID of your wireless network. (Hidden network only)
<b>Wireless Security</b>	WPA & WPA2 Enterprise
<b>Authentication</b>	Protected EAP (PEAP) for RADIUS-based authentication Tunneled TLS for TACACS+ or LDAP-based authentication
<b>Anonymous identity</b>	This is not required.
<b>CA Certificate</b>	If you want to validate the AP's certificate, select the UTN-USERFirst-Hardware root certificate. The default location for the certificate is /usr/share/ca-certificates/mozilla/.
<b>PEAP version</b>	Automatic (applies only to PEAP)
<b>Inner authentication</b>	MSCHAPv2 for RADIUS-based authentication PAP or CHAP for TACACS+ or LDAP-based authentication
<b>Username Password</b>	Enter your logon credentials for the wireless network.

4. If you did not select a CA Certificate above, you are asked to do so. Select Ignore.



5. Select *Connect*. You are connected to the wireless network.

#### To connect to a WPA-Enterprise network

1. Select the Network Manager icon to view the Wireless Networks menu.
2. Select the network from the list (also check *More Networks*).

If your network is not listed (but was configured), select *Connect to Hidden Wireless Network*, select your network from the Connection drop-down list, and then select *Connect*.

# Troubleshooting

Using tools provided in your operating system, you can find the source of common wireless networking problems.

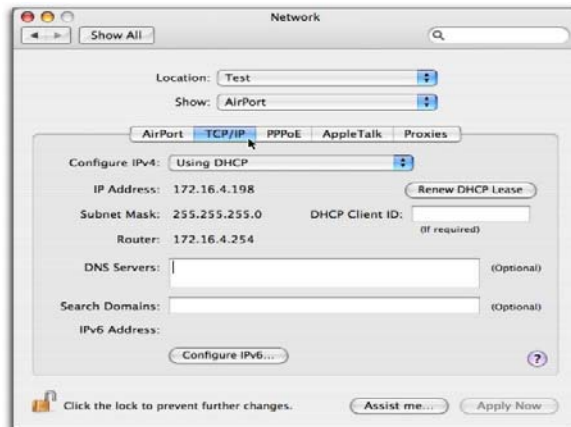
## Checking that the client has received IP address and DNS server information

### Windows XP

1. Double-click the network icon in the taskbar to display the *Wireless Network Connection Status* window. Check that the correct network is listed in the *Connection* section.
2. Select the *Support* tab.  
Check that the *Address Type* is *Assigned by DHCP*. Check that the *IP Address*, *Subnet Mask*, and *Default Gateway* values are valid.
3. Select *Details* to view the DNS server addresses.  
The listed address should be the DNS servers that were assigned to the WAP. Usually a wireless network that provides access to the private LAN is assigned the same DNS servers as the wired private LAN. A wireless network that provides guest or customer users access to the Internet is usually assigned public DNS servers.
4. If any of the addresses are missing, select *Repair*.  
If the repair procedure doesn't correct the problem, check your network settings.

### Mac OS

1. From the Apple menu, open *System Preferences > Network*.
2. Select *AirPort* and then select *Configure*.
3. On the *Network* page, select the *TCP/IP* tab.



4. If there is no IP address or the IP address starts with 169, select *Renew DHCP Lease*.
5. To check DNS server addresses, open a terminal window and enter the following command:  

```
cat /etc/resolv.conf
```

  
Check the listed nameserver addresses. A network for employees should use the wired private LAN DNS server. A network for guests should specify a public DNS server.

## Linux

This example is based on the Ubuntu 10.04 Linux wireless client.

1. Right-click the Network Manager icon and select *Connection Information*.



2. Check the IP address, and DNS settings. If they are incorrect, check your network settings.

# Wireless network examples

This chapter provides an example wireless network configuration. The following topics are included in this section:

- [Basic wireless network](#)
- [A more complex example](#)

## Basic wireless network

This example uses automatic configuration to set up a basic wireless network.

To configure this wireless network, you must:

- Configure authentication for wireless users
- Configure the SSID (WiFi network interface)
- Configure the firewall policy
- Configure and connect FortiAP units

### Configuring authentication for wireless users

You need to configure user accounts and add the users to a user group. This example shows only one account, but multiple accounts can be added as user group members.

#### To configure a WiFi user - web-based manager

1. Go to *User & Device > User > User Definition* and select *Create New*.
2. Enter a *User Name* and *Password* and then select *OK*.

#### To configure the WiFi user group - web-based manager

1. Go to *User & Device > User > User Group* and select *Create New*.
2. Enter the following information and then select *OK*:

<b>Name</b>	wlan_users
<b>Type</b>	Firewall
<b>Available Users/ Members</b>	Move users to the <i>Members</i> list.

#### To configure a WiFi user and the WiFi user group - CLI

```
config user user
 edit "user01"
 set type password
 set passwd "asdf12ghjk"
 end
config user group
 edit "wlan_users"
 set member "user01"
 end
```

## Configuring the SSID

First, establish the SSID (network interface) for the network. This is independent of the number of physical access points that will be deployed. The network assigns IP addresses using DHCP.

### To configure the SSID - web-based manager

1. Go to *WiFi Controller > WiFi Network > SSID* and select *Create New*.
2. Enter the following information and select OK:

<b>Name</b>	example_wifi
<b>IP/Netmask</b>	10.10.110.1/24
<b>Administrative Access</b>	Ping (to assist with testing)
<b>SSID</b>	example_wifi
<b>Enable DHCP Server</b>	Enable
<b>Address Range</b>	10.10.110.2 - 10.10.110.199
<b>Netmask</b>	255.255.255.0
<b>Default Gateway</b>	Same As Interface IP
<b>DNS Server</b>	Same as System DNS
<b>Security Mode</b>	WPA/WPA2-Enterprise
<b>Data Encryption</b>	AES
<b>Authentication</b>	Usergroup, select <i>wlan_users</i> .
Leave other settings at their default values.	



### To configure the SSID - CLI

```
config wireless-controller vap
 edit example_wifi
 set ssid "example_wifi"
 set broadcast-ssid enable
 set security wpa-enterprise
 set auth usergroup
 set usergroup wlan_users
 end
config system interface
 edit example_wifi
 set ip 10.10.110.1 255.255.255.0
 end
config system dhcp server
 edit 0
 set default-gateway 10.10.110.1
 set dns-service default
 set interface "example_wifi"
 config ip-range
 edit 1
 set end-ip 10.10.110.199
 set start-ip 10.10.110.2
 end
 set netmask 255.255.255.0
 end
end
```

## Configuring firewall policies

A firewall policy is needed to enable WiFi users to access the Internet on port1. First you create firewall address for the WiFi network, then you create the example\_wifi to port1 policy.

### To create a firewall address for WiFi users - web-based manager

1. Go to *Firewall Objects > Address > Addresses*.
2. Select *Create New*, enter the following information and select *OK*.

<b>Address Name</b>	wlan_user_net
<b>Type</b>	Subnet / IP Range
<b>Subnet / IP Range</b>	10.10.110.0/24
<b>Interface</b>	example_wifi

### To create a firewall address for WiFi users - CLI

```
config firewall address
 edit "wlan_user_net"
 set associated-interface "example_wifi"
 set subnet 10.10.110.0 255.255.255.0
 end
```

### To create a firewall policy for WiFi users - web-based manager

1. Go to *Firewall Objects > Policy* and select *Create New*.
2. Enter the following information and select *OK*:

<b>Incoming Interface</b>	example_wifi
<b>Source Address</b>	wlan_user_net
<b>Outgoing Interface</b>	port1
<b>Destination Address</b>	All
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Selected. Select <i>Use Destination Interface Address</i> (default).

Leave other settings at their default values.

### To create a firewall policy for WiFi users - CLI

```
config firewall policy
 edit 0
 set srcintf "example_wifi"
 set dstintf "port1"
 set srcaddr "wlan_user_net"
 set dstaddr "all"
 set schedule always
 set service ALL
 set action accept
 set nat enable
 end
```

## Connecting the FortiAP units

You need to connect each FortiAP unit to the FortiGate unit, wait for it to be recognized, and then assign it to the AP Profile. But first, you must configure the interface to which the FortiAP units connect and the DHCP server that assigns their IP addresses.

In this example, the FortiAP units connect to port 3 and are controlled through IP addresses on the 192.168.8.0/24 network.

### To configure the interface for the AP unit - web-based manager

1. Go to *System > Network > Interfaces* and edit the port3 interface.
2. Set the *Addressing mode* to *Dedicate to FortiAP* and set the *IP/Network Mask* to 192.168.8.1/255.255.255.0.
3. Select *OK*.

This procedure automatically configures a DHCP server for the AP units. You can see this configuration in *System > Network > DHCP Server*.

### To configure the interface for the AP unit - CLI

```
config system interface
 edit port3
 set mode static
 set ip 192.168.8.1 255.255.255.0
 end
```

### To configure the DHCP server for AP units - CLI

```
config system dhcp server
 edit 0
 set interface port3
 config exclude-range
 edit 1
 set end-ip 192.168.8.1
 set start-ip 192.168.8.1
 end
 config ip-range
 edit 1
 set end-ip 192.168.8.254
 set start-ip 192.168.8.2
 end
 set netmask 255.255.255.0
 set vci-match enable
 set vci-string "FortiAP"
 end
```

### To connect a FortiAP unit - web-based manager

1. Go to *WiFi Controller > Managed Access Points > Managed FortiAP*.
2. Connect the FortiAP unit to port 3.
3. Periodically select *Refresh* while waiting for the FortiAP unit to be listed.  
Recognition of the FortiAP unit can take up to two minutes.  
If FortiAP units are connected but cannot be recognized, try disabling VCI-Match in the DHCP server settings.
4. When the FortiAP unit is listed, select the entry to edit it.  
The *Edit Managed Access Point* window opens.
5. In *State*, select *Authorize*.
6. Make sure that AP Profile is set to *Automatic*.
7. In *SSID*, select *Automatically Inherit all SSIDs*.
8. Select *OK*.
9. Repeat Steps 2 through 8 for each FortiAP unit.

### To connect a FortiAP unit - CLI

1. Connect the FortiAP unit to port 3.
2. Enter

```
config wireless-controller wtp
```

3. Wait 30 seconds, then enter `get`.  
Retry the `get` command every 15 seconds or so until the unit is listed, like this:  

```
== [FAP22A3U10600118]
wtp-id: FAP22A3U10600118
```
4. Edit the discovered FortiAP unit like this:  

```
edit FAP22A3U10600118
 set admin enable
end
```
5. Repeat Steps 2 through 4 for each FortiAP unit.

## A more complex example

This example creates multiple networks and uses custom AP profiles.

### Scenario

In this example, Example Co. provides two wireless networks, one for its employees and the other for customers or other guests of its business. Guest users have access only to the Internet, not to the company's private network. The equipment for these WiFi networks consists of FortiAP-220A units controlled by a FortiGate unit.

The employee network operates in 802.11n mode on both the 2.4GHz and 5GHz bands. Client IP addresses are in the 10.10.120.0/24 subnet, with 10.10.120.1 the IP address of the WAP. The guest network also operates in 802.11n mode, but only on the 2.4GHz band. Client IP addresses are on the 10.10.115.0/24 subnet, with 10.10.115.1 the IP address of the WAP.

On FortiAP-220A units, the 802.11n mode also supports 802.11g and 802.11b clients on the 2.4GHz band and 802.11a clients on the 5GHz band.

The guest network WAP broadcasts its SSID, the employee network WAP does not.

The employees network uses WPA-Enterprise authentication through a FortiGate user group. The guest network features a captive portal. When a guest first tries to connect to the Internet, a login page requests logon credentials. Guests use numbered guest accounts authenticated by RADIUS. The captive portal for the guests includes a disclaimer page.

In this example, the FortiAP units connect to port 3 and are assigned addresses on the 192.168.8.0/24 subnet.

### Configuration

To configure these wireless networks, you must:

- Configure authentication for wireless users
- Configure the SSIDs (network interfaces)
- Configure the AP profile
- Configure the WiFi LAN interface and a DHCP server
- Configure firewall policies

## Configuring authentication for employee wireless users

Employees have user accounts on the FortiGate unit. This example shows creation of one user account, but you can create multiple accounts and add them as members to the user group.

### To configure the user group for employee access - web-based manager

1. Go to *User & Device > User > User Group* and select *Create New*.
2. Enter the following information and then select OK:

<b>Name</b>	employee-group
<b>Type</b>	Firewall
<b>Available Users Members</b>	Move appropriate user accounts to the <i>Members</i> list.

### To configure the user group for employee access - CLI

```
config user group
 edit "employee-group"
 set member "user01"
 end
```

The user authentication setup will be complete when you select the employee-group in the SSID configuration.

## Configuring authentication for guest wireless users

Guests are assigned temporary user accounts created on a RADIUS server. The RADIUS server stores each user's group name in the Fortinet-Group-Name attribute. Wireless users are in the group named "wireless".

The FortiGate unit must be configured to access the RADIUS server.

### To configure the FortiGate unit to access the guest RADIUS server - web-based manager

1. Go to *User & Device > Authentication > RADIUS Server* and select *Create New*.
2. Enter the following information and select OK:

<b>Name</b>	guestRADIUS
<b>Primary Server Name / IP</b>	10.11.102.100
<b>Primary Server Secret</b>	grikfwpfdg
<b>Secondary Server Name / IP</b>	Optional
<b>Secondary Server Secret</b>	Optional
<b>Authentication Scheme</b>	Use default, unless server requires otherwise.

Leave other settings at their default values.

### To configure the FortiGate unit to access the guest RADIUS server - CLI

```
config user radius
 edit guestRADIUS
 set auth-type auto
 set server 10.11.102.100
 set secret grikfwpfdfg
 end
```

### To configure the user group for guest access - web-based manager

1. Go to *User & Device > User > User Group* and select *Create New*.
2. Enter the following information and then select OK:

<b>Name</b>	guest-group
<b>Type</b>	Firewall
<b>Available Users / Members</b>	Move <i>guestRADIUS</i> to the <i>Members</i> list.
<b>Match one of these group names</b>	Select Add and fill in the following fields:
<b>Remote Server</b>	Select <i>guestRADIUS</i> .
<b>Group Name</b>	Enter wireless

3. Select *Add*.
4. Enter

<b>Remote Server</b>	Select <i>guestRADIUS</i> .
<b>Group Name</b>	Select <i>Specify</i> and then enter wireless

### To configure the user group for guest access - CLI

```
config user group
 edit "guest-group"
 set member "guestRADIUS"
 config match
 edit 0
 set server-name "guestRADIUS"
 set group-name "wireless"
 end
 end
end
```

The user authentication setup will be complete when you select the guest-group user group in the SSID configuration.

## Configuring the SSIDs

First, establish the SSIDs (network interfaces) for the employee and guest networks. This is independent of the number of physical access points that will be deployed. Both networks assign IP addresses using DHCP.

### To configure the employee SSID - web-based manager

1. Go to *WiFi Controller > WiFi Network > SSID* and select *Create New*.
2. Enter the following information and select OK:

<b>Interface Name</b>	example_inc
<b>IP/Netmask</b>	10.10.120.1/24
<b>Administrative Access</b>	Ping (to assist with testing)
<b>SSID</b>	example_inc
<b>Enable DHCP</b>	Enable
<b>Address Range</b>	10.10.120.2 - 10.10.120.199
<b>Netmask</b>	255.255.255.0
<b>Default Gateway</b>	Same As Interface IP
<b>DNS Server</b>	Same as System DNS
<b>Security Mode</b>	WPA/WPA2-Enterprise
<b>Data Encryption</b>	AES
<b>Authentication</b>	Select <i>Usergroup</i> , then select <i>employee-group</i> .

Leave other settings at their default values.

### To configure the employee SSID - CLI

```
config wireless-controller vap
 edit example_inc
 set ssid "example_inc"
 set security wpa-enterprise
 set auth usergroup
 set usergroup employee-group
 end
config system interface
 edit example_inc
 set ip 10.10.120.1 255.255.255.0
 end
config system dhcp server
 edit 0
 set default-gateway 10.10.120.1
 set dns-service default
 set interface example_inc
```

```

config ip-range
 edit 1
 set end-ip 10.10.120.199
 set start-ip 10.10.120.2
 end
set lease-time 7200
set netmask 255.255.255.0
end

```

### To configure the example\_guest SSID - web-based manager

1. Go to *WiFi Controller > WiFi Network > SSID* and select *Create New*.
2. Enter the following information and select OK:

<b>Name</b>	example_guest
<b>IP/Netmask</b>	10.10.115.1/24
<b>Administrative Access</b>	Ping (to assist with testing)
<b>SSID</b>	example_guest
<b>Enable DHCP</b>	Enable
<b>Address Range</b>	10.10.115.2 - 10.10.115.50
<b>Netmask</b>	255.255.255.0
<b>Default Gateway</b>	Same as Interface IP
<b>DNS Server</b>	Same as System DNS
<b>Security Mode</b>	Captive Portal
<b>Customize Portal Messages</b>	Select
<b>User Groups</b>	Select <i>guest-group</i>
Leave other settings at their default values.	

### To configure the example\_guest SSID - CLI

```

config wireless-controller vap
 edit example_guest
 set ssid "example_guest"
 set security captive-portal
 set selected-usergroups guest-group
 end
config system interface
 edit example_guest
 set ip 10.10.115.1 255.255.255.0
 end
config system dhcp server
 edit 0
 set default-gateway 10.10.115.1
 set dns-service default

```



```

set interface "example_guest"
config ip-range
 edit 1
 set end-ip 10.10.115.50
 set start-ip 10.10.115.2
 end
set lease-time 7200
set netmask 255.255.255.0
end

```

## Configuring the custom AP profile

The custom AP Profile defines the radio settings for the networks. The profile provides access to both Radio 1 (2.4GHz) and Radio 2 (5GHz) for the employee virtual AP, but provides access only to Radio 1 for the guest virtual AP.

### To configure the AP Profile - web-based manager

1. Go to *WiFi Controller > Managed Access Points > Custom AP Profile* and select *Create New*.
2. Enter the following information and select OK:

<b>Name</b>	example_AP
<b>Platform</b>	FAP220A
<b>Radio 1</b>	
<b>Mode</b>	Access Point
<b>Background Scan</b>	Enable
<b>Rogue AP On-wire Scan</b>	Enabled
<b>Radio Resource Provision</b>	Not enabled
<b>Band</b>	802.11n
<b>Short Guard Interval</b>	Not enabled
<b>Channel</b>	Select 1, 6, and 11.
<b>Tx Power</b>	100%
<b>SSID</b>	Select <i>example_inc</i> and <i>example_guest</i> .
<b>Radio 2</b>	
<b>Mode</b>	Access Point
<b>Background Scan</b>	Enable
<b>Rogue AP On-wire Scan</b>	Enabled
<b>Radio Resource Provision</b>	Enabled
<b>Band</b>	802.11n_5G

<b>Short Guard Interval</b>	Not enabled
<b>20/40 MHz Channel Width</b>	Not enabled
<b>Channel</b>	Select all.
<b>Tx Power</b>	100%
<b>SSID</b>	Select <i>example_inc</i> .

### To configure the AP Profile - CLI

```

config wireless-controller wtp-profile
 edit "example_AP"
 config platform
 set type 220A
 end
 config radio-1
 set ap-bgscan enable
 set band 802.11n
 set channel "1" "6" "11"
 set rogue-scan enable
 set vaps "example_inc" "example_guest"
 end
 config radio-2
 set ap-bgscan enable
 set band 802.11n-5G
 set channel "36" "40" "44" "48" "149" "153" "157" "161" "165"
 set rogue-scan enable
 set vaps "example_inc"
 end
 end

```

## Configuring firewall policies

Identity-based firewall policies are needed to enable the WLAN users to access the Internet on Port1. First you create firewall addresses for employee and guest users, then you create the firewall policies.

### To create firewall addresses for employee and guest WiFi users

1. Go to *Firewall Objects > Address > Addresses*.
2. Select *Create New*, enter the following information and select *OK*.

<b>Address Name</b>	employee-wifi-net
<b>Type</b>	Subnet / IP Range
<b>Subnet / IP Range</b>	10.10.120.0/24
<b>Interface</b>	example_inc

3. Select *Create New*, enter the following information and select *OK*.

<b>Address Name</b>	guest-wifi-net
<b>Type</b>	Subnet / IP Range
<b>Subnet / IP Range</b>	10.10.115.0/24
<b>Interface</b>	example_guest

### To create firewall policies for employee WiFi users - web-based manager

1. Go to *Policy > Policy* and select *Create New*.
2. Enter the following information and select *OK*:

<b>Source Interface/Zone</b>	example_inc
<b>Source Address</b>	employee-wifi-net
<b>Destination Interface/Zone</b>	port1
<b>Destination Address</b>	all
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT
<b>NAT</b>	Enable NAT

3. Optionally, select *UTM* and set up UTM features for wireless users.
4. Select *OK*.
5. Repeat steps 1 through 4 but select *Internal* as the *Destination Interface/Zone* to provides access to the *ExampleCo* private network.

### To create firewall policies for employee WiFi users - CLI

```
config firewall policy
 edit 0
 set srcintf "employee_inc"
 set dstintf "port1"
 set srcaddr "employee-wifi-net"
 set dstaddr "all"
 set action accept
 set schedule "always"
 set service "ANY"
 set nat enable
 set schedule "always"
 set service "ANY"
 next
 edit 0
 set srcintf "employee_inc"
 set dstintf "internal"
```

```

set srcaddr "employee-wifi-net"
set dstaddr "all"
set action accept
set schedule "always"
set service "ANY"
set nat enable
set schedule "always"
set service "ANY"
end

```

### To create a firewall policy for guest WiFi users - web-based manager

1. Go to *Policy > Policy* and select *Create New*.
2. Enter the following information and select OK:

<b>Source Interface/Zone</b>	example_guest
<b>Source Address</b>	guest-wifi-net
<b>Destination Interface/Zone</b>	port1
<b>Destination Address</b>	all
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT
<b>NAT</b>	Enable NAT

3. Optionally, select *UTM* and set up UTM features for wireless users.
4. Select *OK*.

### To create a firewall policy for guest WiFi users - CLI

```

config firewall policy
edit 0
set srcintf "example_guest"
set dstintf "port1"
set srcaddr "guest-wifi-net"
set dstaddr "all"
set action accept
set schedule "always"
set service "ANY"
set nat enable
end

```

## Connecting the FortiAP units

You need to connect each FortiAP-220A unit to the FortiGate unit, wait for it to be recognized, and then assign it to the AP Profile. But first, you must configure the interface to which the FortiAP units connect and the DHCP server that assigns their IP addresses.

In this example, the FortiAP units connect to port 3 and are controlled through IP addresses on the 192.168.8.0/24 network.

### To configure the interface for the AP unit - web-based manager

1. Go to *System > Network > Interfaces* and edit the port3 interface.
2. Set the *Addressing mode* to *Manual* and set the *IP/Netmask* to 192.168.8.1.
3. Enable *Connect FortiAP to this interface* and set *Reserve IP addresses for FortiAP* to 192.168.8.2 - 192.168.8.9.

This step automatically configures a DHCP server for the AP units.

4. Select *OK*.

### To configure the interface for the AP unit - CLI

```
config system interface
 edit port3
 set mode static
 set ip 192.168.8.1 255.255.255.0
 end
```

### To configure the DHCP server for AP units - CLI

```
config system dhcp server
 edit 0
 set interface port3
 config ip-range
 edit 1
 set end-ip 192.168.8.9
 set start-ip 192.168.8.2
 end
 set netmask 255.255.255.0
 set vci-match enable
 set vci-string "FortiAP"
 end
```

### To connect a FortiAP-220A unit - web-based manager

1. Go to *WiFi Controller > Managed Access Points > Managed FortiAP*.
2. Connect the FortiAP unit to port 3.
3. Periodically select *Refresh* while waiting for the FortiAP unit to be listed.

Recognition of the FortiAP unit can take up to two minutes.

If there is persistent difficulty recognizing FortiAP units, try disabling VCI-Match in the DHCP server settings.

4. When the FortiAP unit is listed, select the entry to edit it.  
The *Edit Managed Access Point* window opens.
5. In *State*, select *Authorize*.
6. In the *AP Profile*, select *[Change]* and then select the *example\_AP* profile.
7. Select *OK*.
8. Repeat Steps 2 through 8 for each FortiAP unit.

### To connect a FortiAP-220A unit - CLI

1. Connect the FortiAP unit to port 3.
2. Enter  

```
config wireless-controller wtp
```

3. Wait 30 seconds, then enter `get`.

Retry the `get` command every 15 seconds or so until the unit is listed, like this:

```
== [FAP22A3U10600118]
wtp-id: FAP22A3U10600118
```

4. Edit the discovered FortiAP unit like this:

```
edit FAP22A3U10600118
 set admin enable
 set wtp-profile example_AP
end
```

5. Repeat Steps 2 through 4 for each FortiAP unit.

# Using a FortiWiFi unit as a client

A FortiWiFi unit by default operates as a wireless access point. But a FortiWiFi unit can also operate as a wireless client, connecting the FortiGate unit to another wireless network.

This section includes the following topics:

- [Use of client mode](#)
- [Configuring client mode](#)

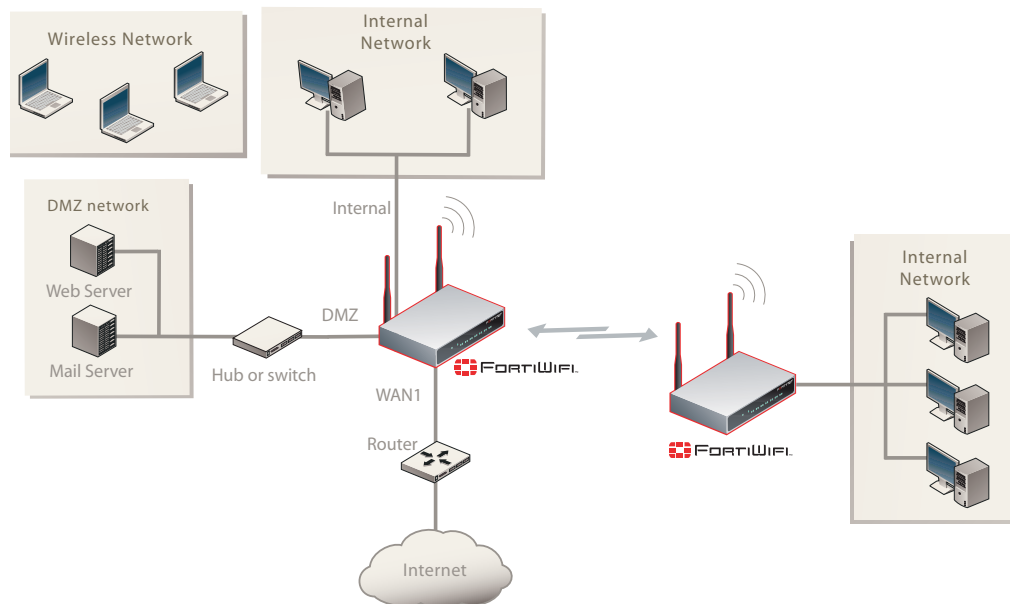
## Use of client mode

In client mode, the FortiWiFi unit connects to a remote WiFi access point to access other networks or the Internet. This is most useful when the FortiWiFi unit is in a location that does not have a wired infrastructure.

For example, in a warehouse where shipping and receiving are on opposite sides of the building, running cables might not be an option due to the warehouse environment. The FortiWiFi unit can support wired users using its Ethernet ports and can connect to another access point wirelessly as a client. This connects the wired users to the network using the 802.11 WiFi standard as a backbone.

Note that in client mode the FortiWiFi unit cannot operate as an AP. WiFi clients cannot see or connect to the FortiWiFi unit in Client mode.

**Figure 171:**Fortinet unit in Client mode



## Configuring client mode

To set up the FortiAP unit as a WiFi client, you must use the CLI. Before you do this, be sure to remove any AP WiFi configurations such as SSIDs, DHCP servers, policies, and so on.

### To configure wireless client mode

1. Change the WiFi mode to client.

In the CLI, enter the following commands:

```
config system global
 set wireless-mode client
end
```

Respond “y” when asked if you want to continue. The FortiWiFi unit will reboot.

2. Configure the WiFi interface settings.

For example, to configure the client for WPA-Personal authentication on the *our\_wifi* SSID with passphrase *justforus*, enter the following in the CLI:

```
config system interface
 edit wifi
 set mode dhcp
 config wifi-networks
 edit 0
 set wifi-ssid our_wifi
 set wifi-security wpa-personal
 set wifi-passphrase "justforus"
 end
 end
 end
```

The WiFi interface *client\_wifi* will receive an IP address using DHCP.

3. Configure a wifi to port1 policy.

You can use either CLI or web-based manager to do this. The important settings are:

<b>Incoming Interface (srcintf)</b>	wifi
<b>Source Address (srcaddr)</b>	all
<b>Outgoing Interface (dstintf)</b>	port1
<b>Destination Address (dstaddr)</b>	all
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Selected



# Support for location-based services

FortiOS supports location-based services by collecting information about WiFi devices near FortiGate-managed access points, even if the devices don't associate with the network.

The following topics are included in this section:

- [Overview](#)
- [Configuring location tracking](#)
- [Viewing device location data on the FortiGate unit](#)

## Overview

WiFi devices broadcast packets as they search for available networks. The FortiGate WiFi controller can collect information about the interval, duration, and signal strength of these packets. The Euclid Analytics service uses this information to track the movements of the device owner. A typical application of this technology is to analyze shopper behavior in a shopping center. Which stores do people walk past? Which window displays do they stop to look at? Which stores do they enter and how long do they spend there? The shoppers are not personally identified, each is known only by the MAC address of their WiFi device.

After enabling location tracking on the FortiGate unit, you can confirm that the feature is working by using a specialized diagnostic command to view the raw tracking data. The Euclid Analytics service obtains the same data in its proprietary format using a JSON inquiry through the FortiGate unit's web-based manager interface.

## Configuring location tracking

You can enable location tracking in any custom AP profile, using the CLI. For each radio, set the `station-locate` field to `enable`. For example:

```
config wireless-controller wtp-profile
 edit "FAP220B-locate"
 set ap-country US
 config platform
 set type 220B
 end
 config radio-1
 set station-locate enable
 end
 config radio-2
 set station-locate enable
 end
 end
end
```

## Viewing device location data on the FortiGate unit

You can use the FortiGate CLI to list located devices. This is mainly useful to confirm that the location data feature is working, You can also reset device location data.

### To list located devices

```
diag wireless-controller wlac -c sta-locate
```

### To reset device location data

```
diag wireless-controller wlac -c sta-locate-reset
```

### Example output

The following output shows data for three WiFi devices.

```
FWF60C3G11004319 # diagnose wireless-controller wlac -c sta-locate
sta_mac vfid rid base_mac freq_lst frm_cnt
 frm_fst frm_last intv_sum intv2_sum intv3_sum intv_min
 intv_max signal_sum signal2_sum signal3_sum sig_min sig_max
 sig_fst sig_last ap
00:0b:6b:22:82:61 0
FAP22B3U11005354 0 0 00:09:0f:f1:bb:e4 5745 257
 708 56 651 1836 6441 0 12
 -21832 1855438 -157758796 -88 -81 -84 -88
 0
00:db:df:24:1a:67 0
FAP22B3U11005354 0 0 00:09:0f:f1:bb:e4 5745 42
 1666 41 1625 97210 5831613 0 60
 -3608 310072 -26658680 -90 -83 -85 -89
 0
10:68:3f:50:22:29 0
FAP22B3U11005354 0 0 00:09:0f:f1:bb:e4 5745 102
 1623 58 1565 94136 5664566 0 60
 -8025 631703 -49751433 -84 -75 -78 -79
 0
```

The output for each device appears on two lines. The first line contains only the device MAC address and the VLAN ID. The second line begins with the ID (serial number) of the FortiWiFi or FortiAP unit that detected the device, the AP's MAC address, and then the fields that the Euclid service uses. Because of its length, this line wraps around and displays as multiple lines.

# Reference

This chapter provides some reference information pertaining to wireless networks. The following topics are included in this section:

- [Wireless radio channels](#)
- [FortiAP CLI](#)

## Wireless radio channels

### IEEE 802.11a/n channels

Table 43 lists the channels supported on FortiWiFi products that support the IEEE 802.11a and 802.11n wireless standards. 802.11a is available on FortiWiFi models 60B and higher. 802.11n is available on FortiWiFi models 80CM and higher.

All channels are restricted to indoor usage except in the Americas, where both indoor and outdoor use is permitted on channels 52 through 64 in the United States.

**Table 43:** IEEE 802.11 a/n (5-GHz Band) channel numbers

Channel number	Frequency (MHz)	Regulatory Areas				
		Americas	Europe	Taiwan	Singapore	Japan
34	5170					•
36	5180	•	•		•	
38	5190					
40	5200	•	•		•	•
42	5210					
44	5220	•	•		•	•
46	5230					
48	5240	•	•		•	•
149	5745	•		•	•	
153	5765	•		•	•	
157	5785	•		•	•	
161	5805	•		•	•	
165	5825	•			•	

## IEEE 802.11b/g/n channel numbers

Table 44 lists IEEE 802.11b/g/n channels. All FortiWiFi units support 802.11b and 802.11g. Newer models also support 802.11n.

Mexico is included in the Americas regulatory domain. Channels 1 through 8 are for indoor use only. Channels 9 through 11 can be used indoors and outdoors. You must make sure that the channel number complies with the regulatory standards of Mexico.

**Table 44:** IEEE 802.11b/g/n (2.4-GHz Band) channel numbers

Channel number	Frequency (MHz)	Regulatory Areas			
		Americas	EMEA	Israel	Japan
1	2412	•	•	•?	•
2	2417	•	•	•?	•
3	2422	•	•	•?	•
4	2427	•	•	•	•
5	2432	•	•	•	•
6	2437	•	•	•	•
7	2442	•	•	•	•
8	2447	•	•	•	•
9	2452	•	•	•	•
10	2457	•	•	•	•
11	2462	•	•	•?	•
12	2467		•	•?	•
13	2472		•	•?	•
14	2484				b only

## FortiAP CLI

The FortiAP CLI now includes more configuration commands and a complete set of diagnose commands.

Configuration commands include the following

<code>cfg -h</code>	Display help for all commands.
<code>cfg -r var</code>	Remove variables.
<code>cfg -e</code>	Export variables.
<code>cfg -s</code>	List variables.
<code>cfg -x</code>	Reset to factory defaults.
<code>cfg -c</code>	Commit the change to flash.
<code>cfg -a var=value</code>	Add or change variables.

Diagnose commands include:

<code>cw_diag help</code>	Display help for all diagnose commands.
<code>cw_diag uptime</code>	Show daemon uptime.
<code>cw_diag --tlog &lt;on off&gt;</code>	Turn on/off telnet log message.
<code>cw_diag --clog &lt;on off&gt;</code>	Turn on/off console log message.
<code>cw_diag baudrate [9600   19200   38400   57600   115200]</code>	Set the console baud rate.
<code>cw_diag plain-ctl [0 1]</code>	Show or change current plain control setting.
<code>cw_diag sniff-cfg ip port</code>	Set sniff server ip and port.
<code>cw_diag sniff [0 1 2]</code>	Enable/disable sniff packet.
<code>cw_diag stats wl_intf</code>	Show wl_intf status.
<code>cw_diag admin-timeout [30]</code>	Set shell idle timeout in minutes.
<code>cw_diag -c wtp-cfg</code>	Show current wtp config parameters in control plane.
<code>cw_diag -c radio-cfg</code>	Show current radio config parameters in control plane.
<code>cw_diag -c vap-cfg</code>	Show current vaps in control plane.
<code>cw_diag -c ap-rogue</code>	Show rogue APs pushed by AC for on-wire scan.
<code>cw_diag -c sta-rogue</code>	Show rogue STAs pushed by AC for on-wire scan.
<code>cw_diag -c arp-req</code>	Show scanned arp requests.
<code>cw_diag -c ap-scan</code>	Show scanned APs.

<code>cw_diag -c sta-scan</code>	Show scanned STAs.
<code>cw_diag -c sta-cap</code>	Show scanned STA capabilities.
<code>cw_diag -c wids</code>	Show scanned WIDS detections.
<code>cw_diag -c darrp</code>	Show darrp radio channel.
<code>cw_diag -c mesh</code>	Show mesh status.
<code>cw_diag -c mesh-veth-acinfo</code>	Show mesh veth ac info, and mesh ether type.
<code>cw_diag -c mesh-veth-vap</code>	Show mesh veth vap.
<code>cw_diag -c mesh-veth-host</code>	Show mesh veth host.
<code>cw_diag -c mesh-ap</code>	Show mesh ap candidates.
<code>cw_diag -c scan-clr-all</code>	Flush all scanned AP/STA/ARPs.
<code>cw_diag -c ap-suppress</code>	Show suppressed APs.
<code>cw_diag -c sta-deauth</code>	De-authenticate an STA.

## FortiAP web-based manager

You can access the FortiAP unit's built-in web-based manager. This is useful to adjust settings that are not available through the FortiGate unit's WiFi Controller. Logging into the FortiAP web-based manager is similar to logging into the FortiGate web-based manager.

The advanced settings are in the WTP Configuration section on the System Information tab.

<b>Uplink</b>	The FortiAP unit can be connected to the controller by Ethernet, WiFi mesh, or Ethernet with mesh fallback.
<b>AC Discovery Type</b>	Static, DHCP, DNS, Broadcast, Multicast, Auto
<b>AC Control Port</b>	Default port is 5246.
<b>AC IP Address 1</b> <b>AC IP Address 2</b> <b>AC IP Address 3</b>	You enter up to three WiFi controller IP addresses for static discovery. Routing must be properly configured in both directions.
<b>AC Host Name 1</b> <b>AC Host Name 2</b> <b>AC Host Name 3</b>	As an alternative to AC IP addresses, you can enter their fully qualified domain names (FQDNs).
<b>AC Discovery Multicast Address</b>	224.0.1.140
<b>AC Discovery DHCP Option Code</b>	When using DHCP discovery, you can configure the DHCP server to provide the controller address. By default the FortiAP unit expects this in option 138.
<b>AC Data Channel Security</b>	Select whether the FortiAP unit requires Clear Text data, DTLS encrypted data, or will accept either.



# Chapter 7 Firewall for FortiOS 5.0

[Firewall concepts](#) explains the ideas behind the components, techniques and processes that are involved in setting up and running a firewall in general and the FortiGate firewall in particular. The premise here is that regardless of how experienced someone is with firewalls as they go through the process of configuring a firewall that is new to them they are likely to come across a term or setting that they may not be familiar with even if it is only in the context of the setting they are working in at the moment. FortiGate firewalls are quite comprehensive and can be very granular in the functions that they perform, so it makes sense to have a consistent frame of reference for the ideas that we will be working with.

Some examples of the concepts that will be addressed here are:

- [What is a Firewall?](#)
- [NAT](#)
- [IPv6](#)

[Firewall objects](#) describes the following firewall objects:

- Addressing
- Services
- Firewall Policies

[Network defense](#) describes various methods of defending your Network using the abilities of the FortiGate Firewall.

[GUI & CLI - What You May Not Know](#) helps you navigate and find the components in the Web-based Manager that you will need to build the functions. This section does not include any in-depth explanations of what each object does as that is covered in the concepts section. This section is for showing you where you need to input your information and let you know what format the interface expects to get that information.

[Building firewall objects and policies](#) is similar to a cookbook in that it will refer to a number of common tasks that you will likely perform to get the full functionality out of your FortiGate firewall. Because of the way that firewalls are designed, performing many of the tasks requires that firewall components be set up in a number of different sections of the interface and be configured to work together to achieve the desired result. This section will bring those components all together as a straight forward series of instructions.

[Multicast forwarding](#) is a reference guide including the concepts and examples that are involved in the use of multicast addressing and policy forwarding as it is used in the FortiGate firewall.

## FortiGate Firewall Components

The FortiGate firewall is made up of a number of different components that are used to build an impressive list of features that have flexibility of scope and granularity of control that provide protection that is beyond that provided by the basic firewalls of the past.



Some of the components that FortiOS uses to build features are:

- Interfaces
- VLANs
- Soft Switches
- Zones
- Predefined Addresses
- IP address based
- FQDN based
- Geography based
- Access Schedules
- Authentication
- Local User based
- Authentication Server based (Active Directory, Radius, LDAP)
- Device Based
- Configurable Services
- IPv4 and IPv6 protocol support

The features of FortiOS include but are not limited to:

- Security profiles, sometimes referred to as Unified Threat Management (UTM) or Next Generation Firewall (NGFW)
- Predefined firewall addresses (this includes IPv4 and IPv6, IP pools, . wildcard addresses and netmasks, and geography-based addresses)
- Monitoring traffic
- Traffic shaping and per-IP traffic shaping (advanced)
- Firewall schedules
- Services (such as AOL, DHCP and FTP)
- Logging traffic
- Quality of Service (QoS)
- Identity-based policies
- Endpoint security

The [Firewall concepts](#) expands on what each of the features does and how they relate to the administration of the FortiGate firewall. The section will also try to explain some of the common firewall concepts that will be touched on in the implementing of these features.

[Building firewall objects and policies](#) shows how to perform specific tasks with the FortiGate firewall.

## How does a FortiGate Protect Your Network

The FortiGate firewall protects your network by taking the various components and using them together to build a kind of wall or access control point so that anyone that is not supposed to be on your network is prevented from accessing your network in anyway other than those approved by you. It also protects your network from itself by keeping things that shouldn't happen from happening and optimizing the flow of traffic so that the network is protected from traffic congestion that would otherwise impede traffic flow.

Most people have at one time or another played with the children's toy system that is made up of interlocking blocks. The blocks come in different shapes and sizes so that you can build

structures to suit your needs and in your way. The components of the FortiGate firewall are similar. You are not forced to use all of the blocks all of the time. You mix and match them to get the results that you are looking for. You can build a very basic structure that's only function is to direct traffic in and out to the correct subnets or you can build a fortress that only allows specific traffic to specific hosts from specific hosts at specific times of day and that is only if they provide the credentials that have been pre-approved and all of the traffic is encrypted so that even when the traffic is out on the Internet it is private from the world. Just like the interlocking blocks, what you build is up to you, but chances are if you put them together the right way there isn't much that can't be built.

Here is one example of how the components could be put together to support the requirements of a network infrastructure design.

- Off the Internal interface you could have separate VLANs. One for each for the departments of Sales, Marketing and Engineering so that the traffic from the users on one VLAN does not intrude upon the hosts of the other VLANs and the department are isolated from one another for security reasons.
- To ease in the administration each of the VLAN sub-interfaces is made a member of a zone so that security policies that apply to all of the hosts on all of the VLANs can be applied to all of them at once.
- Using the addresses component each of the IP address ranges could be assigned a user friendly name so that they could be referred to individually and then for policies that would refer to them all as a whole the individual ranges to be made members of an address group.
- Firewall schedules could be created to address the differing needs of each of the groups so that Sales and Marketing could be allowed access to the Internet during regular business hours and the Engineering department could be allowed access during the lunch break.
- By setting up the outgoing policies to use FortiGuard Web-filtering the employees could be prevented from visiting inappropriate sites and thus enforcing the policies of the HR department.
- A couple of virtual IP addresses with port forwarding could be configured to allow users on the Internet to access a web server on the DMZ subnet using the company's only Public IP address without affecting the traffic that goes to the company's mail server that is hosted on a complete different computer.
- Even though the Web server on the same DMZ has an FTP service to allow for the uploading of web pages to the web server from the Marketing and Engineer teams, by placing a DENY policy on any FTP traffic from the Internet malicious users are prevented from abusing the FTP service.
- By monitoring the traffic as it goes through the policies you can verify that the policies are in working order.
- By using a combination of ALLOW and DENY policies and placing them in the correct order you could arrange for an outside contractor to be allowed to update the web site as well

These set of configurations is not extensive but it does give an idea of how different components can be mixed and matched to build a configuration that meets an organization's needs but at the same time protect it from security risks.

# Firewall concepts

There are a number of concepts that are consistent throughout the firewall industry and having a solid grasp of these ideas and terms can give you a better idea of what your FortiGate firewall is capable of and how it will be able to fit within your networks architecture.

This chapter describes the following firewall concepts:

- [What is a Firewall?](#)
- [IPv6](#)
- [NAT](#)
- [How Packets are handled by FortiOS](#)
- [FortiGate Modes](#)
- [Quality of Service](#)
- [Interfaces and Zones](#)

## What is a Firewall?

The term firewall originally referred to a wall intended to confine a fire or potential fire within a building. Later uses refer to similar structures, such as the metal sheet separating the engine compartment of a vehicle or aircraft from the passenger compartment.

A firewall can either be software-based or hardware-based and is used to help keep a network secure. Its primary objective is to control the incoming and outgoing network traffic by analyzing the data packets and determining whether it should be allowed through or not, based on a predetermined rule set. A network's firewall builds a bridge between an internal network that is assumed to be secure and trusted, and another network, usually an external (inter)network, such as the Internet, that is not assumed to be secure and trusted.

## Network Layer or Packet Filter Firewalls

### Stateless Firewalls

Stateless firewalls are the oldest form of these firewalls. They are faster and simple in design requiring less memory because they process each packet individually and don't require the resources necessary to hold onto packets like stateful firewalls. Stateless firewalls inspect each packet individually and check to see if it matches a predetermined set of rules. According to the matching rule the packet is either be allowed, dropped or rejected. In the case of a rejection an error message is sent to the source of the traffic. Each packet is inspected in isolation and information is only gathered from the packet itself. Simply put, if the packets were not specifically allowed according to the list of rules held by the firewall they were not getting through.

### Stateful Firewalls

Stateful firewalls retain packets in memory so that they can maintain context about active sessions and make judgements about the state of an incoming packet's connection. This enables Stateful firewalls to determine if a packet is the start of a new connection, a part of an existing connection, or not part of any connection. If a packet is part of an existing connection based on comparison with the firewall's state table, it will be allowed to pass without further processing. If a packet does not match an existing connection, it will be evaluated according to

the rules set for new connections. Predetermined rules are used in the same way as a stateless firewall but they can now work with the additional criteria of the state of the connection to the firewall.



Best Practices Tip for improving performance:

Blocking the packets in a denied session can take more CPU processing resources than passing the traffic through. By putting denied sessions in the session table, they can be kept track of in the same way that allowed sessions are so that the FortiGate unit does not have to redetermine whether or not to deny all of the packets of a session individually. If the session is denied all packets of that session are also denied.

In order to configure this you will need to use 2 CLI commands

```
config system setting
 set ses-denied-traffic enable
 set block-session-timer <integer 1 - 300> (this determines in seconds
 how long, in seconds, the session is kept in the table)
end
```

---

## Application Layer Firewalls

Application layer filtering is yet another approach and as the name implies it works primarily on the Application Layer of the OSI Model.

Application Layer Firewalls actually, for lack of a better term, understand certain applications and protocols. Examples would be FTP, DNS and HTTP. This form of filtration is able to check to see if the packets are actually behaving incorrectly or if the packets have been incorrectly formatted for the protocol that is indicated. This process also allows for the use of deep packet inspection and the sharing of functionality with Intrusion Prevention Systems (IPS).

Application-layer firewalls work on the application level of the TCP/IP stack (i.e., all browser traffic, or all telnet or ftp traffic), and may intercept all packets traveling to or from an application. They block other packets (usually dropping them without acknowledgment to the sender). Application firewalls work much like a packet filter but application filters apply filtering rules (allow/block) on a per process basis instead of filtering connections on a per port basis.

On inspecting all packets for improper content, firewalls can restrict or prevent outright the spread of networked computer worms and trojans. The additional inspection criteria can add extra latency to the forwarding of packets to their destination.

## Proxy Servers

A proxy server is an appliance or application that acts as an intermediary for communicating between computers. A computer has a request for information. The packets are sent to the designated resource but before they can get there they are blocked by the proxy server saying that it will take the request and pass it on. The Proxy Server processes the request and if it is valid it passes onto the designated computer. The designated computer gets the packet and processes the request, sending the answer back to the proxy server. The proxy server sends the information back to the originating computer. It's all a little like a situation with two people who refuse to talk directly with each other using someone else to take messages back and forth.

From a security stand point a Proxy Server can serve a few purposes:

- Protects the anonymity of the originating computer
- The two computers never deal directly with each other
- Packets that are not configured to be forwarded are dropped before reaching the destination computer.
- If malicious code is sent it will affect the Proxy server with out affecting the originating or sending computer.

Proxies can perform a number of roles including:

- Content Filtering
- Caching
- DNS proxy
- Bypassing Filters and Censorship
- Logging and eavesdropping
- Gateways to private networks
- Accessing service anonymously

## Security Profiles

Unified Threat Management and Next Generation Firewall are terms originally coined by market research firms and refer to the concept of a comprehensive security solution provided in a single package. It is basically combining of what used to be accomplished by a number of different security technologies all under a single umbrella or in this case, a single device. On the FortiGate firewall this is achieved by the use of Security Profiles and optimized hardware.

In effect it is going from a previous style of firewall that included among its features:

- Gateway Network Firewall
- Routing
- VPN

To a more complete system that includes:

- Gateway Network Firewall
- Routing
- VPN
- Traffic Optimization
- Proxy Services
- Content Filtering
- Application Control
- Intrusion Protection
- Denial of Service Attack Protection
- Anti-virus
- Anti-spam
- Data Leak Prevention
- Endpoint Control of Security Applications
- Load Balancing
- WiFi Access Management
- Authentication Integration into Gateway Security
- Logging
- Reporting

### Advantages of using Security Profiles

- Avoidance of multiple installations.
- Hardware requirements are fewer.
- Fewer hardware maintenance requirements.
- Less space required.
- Compatibility - multiple installations of products increase the probability of incompatibility between systems.
- Easier support and management.
- There is only one product to learn therefore a reduced requirement of technical knowledge.
- Only a single vendor so there are fewer support contracts and Service Level Agreements.
- Easier to incorporated into existing security architecture.
- Plug and play architecture.
- Web based GUI for administration.

## IPv6

### What is IPv6?

Internet Protocol version 6 (IPv6) will succeed IPv4 as the standard networking protocol of the Internet. IPv6 provides a number of advances over IPv4 but the primary reason for its replacing IPv4 is its limitation in addresses. IPv4 uses 32 bit addresses which means there is a theoretical limit of 2 to the power of 32. The IPv6 address scheme is based on a 128 bit address or a theoretical limit of 2 to the power of 128.

Possible Addresses:

- IPv4 = 4,294,967,296 (over 4 billion)
- IPv6 = 340,282,366,920,938,463,374,607,431,768,211,456 (over 340 undecillion - We had to look that term up. We didn't know what a number followed by 36 digits was either)

Assuming a world population of approximately 8 billion people, IPv6 would allow for each individual to have approximately 42,535,295,865,117,200,000,000,000,000 devices with an IP address. That's 42 quintillion devices.

There is little likelihood that you will ever need to worry about these numbers as any kind of serious limitation in addressing but they do give an idea of the scope of the difference in the available addressing.

Aside from the difference of possible addresses there is also the different formatting of the addresses that will need to be addressed.

A computer would view an IPv4 address as a 32 bit string of binary digits made up of 1s and 0s, broken up into 4 octets of 8 digits separated by a period “.”

Example:

```
10101100.00010000.11111110.00000001
```

To make number more user friendly for humans we translate this into decimal, again 4 octets separated by a period “.” which works out to:

```
172.16.254.1
```

A computer would view an IPv6 address as a 128 bit string of binary digits made up of 1s and 0s, broken up into 8 octets of 16 digits separated by a colon “:”

```
1000000000000001:0000110110111000:101011000001000:1111111000000001:00000000
00000000:0000000000000000:0000000000000000:0000000000000000
```

To make number a little more user friendly for humans we translate this into hexadecimal, again 8 octets separated by a colon “:” which works out to:

```
8001:0DB8:AC10:FE01:0000:0000:0000:0000:
```

Because any four-digit group of zeros within an IPv6 address may be reduced to a single zero or altogether omitted, this address can be shortened further to:

```
8001:0DB8:AC10:FE01:0:0:0:0
```

or

```
8001:0DB8:AC10:FE01::
```

Some of the other benefits of IPv6 include:

- More efficient routing
- Reduced management requirement
- Stateless auto-reconfiguration of hosts
- Improved methods to change Internet Service Providers
- Better mobility support
- Multi-homing
- Security
- Scoped address: link-local, site-local and global address space

## IPv6 in FortiOS

From an administrative point of view IPv6 works almost the same as IPv4 in FortiOS. The primary difference is the use IPv6 format for addresses. There is also no need for NAT if the FortiGate firewall is the interface between IPv6 networks. If the subnets attached to the FortiGate firewall are IPv6 and IPv4 NAT can be configured between the 2 different formats. This will involve either configuring a dual stack routing or IPv4 tunnelling configuration. The reason for this is simple. NAT was developed primarily for the purpose of extending the number of usable IPv4 addresses. IPv6's addressing allows for enough available addresses so the NAT is no longer necessary.

When configuring IPv6 in FortiOS, you can create a dual stack route or IPv4-IPv6 tunnel. A dual stack routing configuration implements dual IP layers, supporting both IPv4 and IPv6, in both hosts and routers. An IPv4-IPv6 tunnel is essentially similar, creating a tunnel that encapsulates IPv6 packets within IPv4 headers that carry these IPv6 packets over IPv4 tunnels. The FortiGate unit can also be easily integrated into an IPv6 network. Connecting the FortiGate unit to an IPv6 network is exactly the same as connecting it to an IPv4 network, the only difference is that you are using IPv6 addresses.

By default the IPv6 settings are not displayed in the Web-based Manager. It is just a matter of enabling the display of these feature to use them through the web interface. To enable them just go to System > Admin > Settings and select IPv6 Support on GUI. Once enabled, you will be able to use IPv6 addresses as well as the IPv4 addressing for the following FortiGate firewall features:

- Static routing
- Policy Routing
- Packet and network sniffing
- Dynamic routing (RIPv6, BGP4+, and OSPFv3)
- IPsec VPN
- DNS
- DHCP
- SSL VPN
- Network interface addressing
- Security Profiles protection
- Routing access lists and prefix lists
- NAT/Route and Transparent mode
- NAT 64 and NAT 66
- IPv6 tunnel over IPv4 and IPv4 tunnel over IPv6
- Logging and reporting
- Security policies
- SNMP
- Authentication
- Virtual IPs and groups
- IPv6 over SCTP
- IPv6-specific troubleshooting, such as ping6

## Dual Stack routing configuration

Dual stack routing implements dual IP layers in hosts and routers, supporting both IPv6 and IPv4. A dual stack architecture supports both IPv4 and IPv6 traffic and routes the appropriate



traffic as required to any device on the network. Administrators can update network components and applications to IPv6 on their own schedule, and even maintain some IPv4 support indefinitely if that is necessary. Devices that are on this type of network, and connect to the Internet, can query Internet DNS servers for both IPv4 and IPv6 addresses. If the Internet site supports IPv6, the device can easily connect using the IPv6 address. If the Internet site does not support IPv6, then the device can connect using the IPv4 addresses. In the FortiOS dual stack architecture it is not just the basic addressing functions that operate in both versions of IP. The other features of the appliance such as Security Profiles and routing can also use both IP stacks.

If an organization with a mixed network uses an Internet service provider that does not support IPv6, they can use an IPv6 tunnel broker to connect to IPv6 addresses that are on the Internet. FortiOS supports IPv6 tunnelling over IPv4 networks to tunnel brokers. The tunnel broker extracts the IPv6 packets from the tunnel and routes them to their destinations.

## IPv6 Tunnelling

IPv6 Tunnelling is the act of tunnelling IPv6 packets from an IPv6 network through an IPv4 network to another IPv6 network. This is different than Network Address Translation (NAT) because once the packet reaches its final destination the true originating address of the sender will still be readable. The IPv6 packets are encapsulated within packets with IPv4 headers, which carry their IPv6 payload through the IPv4 network. This type of configuration is more appropriate for those who have completely transitional over to IPv6, but need an Internet connection, which is still mostly IPv4 addresses.

The key to IPv6 tunnelling is the ability of the 2 devices, whether they are a host or a network device, to be dual stack compatible. They have to be able to work with both IPv4 and IPv6 at the same time. In the process the entry node of the tunnel portion of the path will create an encapsulating IPv4 header and transmit the encapsulated packet. The exit node at the end of the tunnel receives the encapsulated packet. The IPv4 header is removed. The IPv6 header is updated and the IPv6 packet is processed.

There are two types of tunnels in IPv6:

---

<b>Automatic tunnels</b>	Automatic tunnels are configured by using IPv4 address information embedded in an IPv6 address – the IPv6 address of the destination host includes information about which IPv4 address the packet should be tunnelled to.
--------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

---

<b>Configured tunnels</b>	Configured tunnels must be configured manually. These tunnels are used when using IPv6 addresses that do not have any embedded IPv4 information. The IPv6 and IPv4 addresses of the endpoints of the tunnel must be specified.
---------------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

---

## Tunnel Configurations

There are a few ways in which the tunnelling can be performed depending on which segment of the path between the end points of the session the encapsulation takes place.

---

<b>Network Device to Network Device</b>	Dual Stack capable devices connected by an IPv4 infrastructure can tunnel IPv6 packets between themselves. In this case, the tunnel spans one segment of the path taken by the IPv6 packets.
-----------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

---

<b>Host to Network Device</b>	Dual Stack capable hosts can tunnel IPv6 packets to an intermediary IPv6 or IPv4 network device that is reachable through an IPv4 infrastructure. This type of tunnel spans the first segment of the path taken by the IPv6 packets.
-------------------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

---

<b>Host to Host</b>	Dual Stack capable hosts that are interconnected by an IPv4 infrastructure can tunnel IPv6 packets between themselves. In this case, the tunnel spans the entire path taken by the IPv6 packets.
<b>Network Device to Host</b>	Dual Stack capable network devices can tunnel IPv6 packets to their final destination IPv6 or IPv4 host. This tunnel spans only the last segment of the path taken by the IPv6 packets.

Regardless of whether the tunnel starts at a host or a network device, the node that does the encapsulation needs to maintain soft state information, such as the maximum transmission unit (MTU), about each tunnel in order to process the IPv6 packets.

## Tunnelling IPv6 through IPsec VPN

A variation on the tunnelling IPv6 through IPv4 is using an IPsec VPN tunnel between two FortiGate devices. FortiOS supports IPv6 over IPsec. In this sort of scenario, two networks using IPv6 behind FortiGate units are separated by the Internet, which uses IPv4. An IPsec VPN tunnel is created between the two FortiGate units and a tunnel is created over the IPv4 based Internet but the traffic in the tunnel is IPv6. This has the additional advantage of making the traffic secure as well.

## NAT

### What is NAT?

NAT or Network Address Translation is the process that enables a single device such as a router or firewall to act as an agent between the Internet or Public Network and a local or private network. This “agent”, in real time, translates the source IP address of a device on one network interface, usually the Internal, to a different IP address as it leaves another interface, usually the interface connected to the ISP and the Internet. This enables a single public address to represent a significantly larger number of private addresses.

### The Origins of NAT

In order to understand NAT it helps to know why it was created. At one time, every computer that was part of a network had to have its own addresses so that the other computers could talk to it. There were a few protocols in use at the time, some of which were only for use on a single network, but of those that were routable, the one that had become the standard for the Internet was IP (Internet Protocol) version 4.

When IP version 4 addressing was created nobody had any idea how many addresses would be needed. The total address range was based on the concept of 2 to the 32nd power, which works out to be 4 294 967 296 potential addresses. Once you eliminate some of those for reserved addresses, broadcast addresses, network addresses, multicasting, etc., you end up with a workable scope of about 3.2 million addressees. This was thought to be more than enough at the time. The designers were not expecting the explosion of personal computing, the World Wide Web or smart phones. As of the beginning of 2012, some estimate the number of computers in the world in the neighborhood of 1 billion, and most of those computer users are going to want to be on the Internet or Search the World Wide Web. In short, we ran out of addresses.

This problem of an address shortage was realized before we actually ran out, and in the mid 1990s two technical papers called RFCs numbered 1631 (<http://www.ietf.org/rfc/rfc1631.txt>) and 1918 (<http://tools.ietf.org/html/rfc1918>), proposed components of a method that would be used

as a solution until a new addressing methodology could be implemented across the Internet infrastructure. For more information on this you can look up IP version 6.

RFC 1631 described a process that would allow networking devices to translate a single public address to multiple private IP addresses and RFC 1918 laid out the use of the private addresses. The addresses that were on the Internet (Public IP addresses) could not be duplicated for them to work as unique addresses, but behind a firewall, which most large institutions had, they could use their own Private IP addresses for internal use and the internal computers could share the external or Public IP address.

To give an idea on a small scale how this works, image that a company has a need for 200 computer addresses. Before Private IP addresses and NAT the company would have purchased a full Class C address range which would have been 254 usable IP addresses; wasting about 50 addresses. Now with NAT, that company only needs 1 IP address for its 200 computers and this leaves the rest of the IP addresses in that range available for other companies to do the same thing.

NAT gives better value than it would first appear because it is not 253 companies that can use 254 addresses but each of those 254 companies could set up their networking infrastructures to use up to thousands of Private IP addresses, more if they don't all have to talk to the Internet at the same time. This process enabled the Internet to keep growing even though we technically have many more computers networked than we have addresses.

## Static NAT

In Static NAT one internal IP address is always mapped to the same public IP address.

In FortiGate firewall configurations this is most commonly done with the use of Virtual IP addressing.

An example would be if you had a small range of IP addresses assigned to you by your ISP and you wished to use one of those IP address exclusively for a particular server such as an email server.

Say the internal address of the Email server was 192.168.12.25 and the Public IP address from your assigned addresses range from 256.16.32.65 to 256.16.32.127. Many readers will notice that because one of the numbers is above 255 that this is not a real Public IP address. The Address that you have assigned to the interface connected to your ISP is 256.16.32.66, with 256.16.32.65 being the remote gateway. You wish to use the address of 256.16.32.70 exclusively for your email server.

When using a Virtual IP address you set the external IP address of 256.16.32.70 to map to 192.168.12.25. This means that any traffic being sent to the public address of 256.16.32.70 will be directed to the internal computer at the address of 192.168.12.25

When using a Virtual IP address, this will have the added function that when ever traffic goes from 192.168.12.25 to the Internet it will appear to the recipient of that traffic at the other end as coming from 256.16.32.70.

You should note that if you use Virtual IP addressing with the Port Forwarding enabled you do not get this reciprocal effect and must use IP pools to make sure that the outbound traffic uses the specified IP address.

## Dynamic NAT

Dynamic NAT maps the private IP addresses to the first available Public Address from a pool of possible Addresses. In the FortiGate firewall this can be done by using IP Pools.

## Overloading

This is a form of Dynamic NAT that maps multiple private IP address to a single Public IP address but differentiates them by using a different port assignment. This is probably the most widely used version of NAT. This is also referred to as PAT (Port Address Translation) or Masquerading.

An example would be if you had a single IP address assigned to you by your ISP but had 50 or 60 computers on your local network.

Say the internal address of the interface connected to the ISP was 256.16.32.65 (again an impossible address) with 256.16.32.64 being the remote gateway. If you are using this form of NAT any time one of your computers accesses the Internet it will be seen from the Internet as 256.16.32.65. If you wish to test this go to 2 different computers and verify that they each have a different private IP address then go to a site that tells you your IP address such as [www.ipchicken.com](http://www.ipchicken.com). You will see that the site gives the same result of 256.16.32.65, if it existed, as the public address for both computers.

As mentioned before this is sometimes called Port Address Translation because network device uses TCP ports to determine which internal IP address is associated with each session through the network device. For example, if you have a network with internal addresses ranging from 192.168.1.1 to 192.168.1.255 and you have 5 computers all trying to connect to a web site which is normally listening on port 80 all of them will appear to the remote web site to have the IP address of 256.16.32.65 but they will each have a different sending TCP port, with the port numbers being somewhere between 1 and 65 535, although the port numbers between 1 to 1024 are usually reserved or already in use. So it could be something like the following:

192.168.1.10	256.16.32.65:	port 486
192.168.1.23	256.16.32.65:	port 2409
192.168.1.56	256.16.32.65:	port 53763
192.168.1.109	256.16.32.65:	port 5548
192.168.1.201	256.16.32.65:	port 4396

And the remote web server would send the responding traffic back based on those port numbers so the network device would be able to sort through the incoming traffic and pass it on to the correct computer.

## Overlapping

Because everybody is using the relative same small selection of Private IP addresses it is inevitable that there will be two networks that share the same network range that will need to talk with each other. This happens most often over Virtual Private Networks or when one organization ends up merging with another. This is a case where a private IP address may be translated into a different private IP address so there are no issues with conflict of addresses or confusion in terms of routing.

An example of this would be when you have a Main office that is using an IP range of 172.16.0.1 to 172.20.255.255 connecting through a VPN to a recently acquired branch office that is already running with an IP range of 172.17.1.1 to 172.17.255.255. Both of these ranges are perfectly valid but because the Branch office range is included in the Main Office range any time the system from the Main office try to connect to an address in the Branch Office the routing the system will not send the packet to the default gateway because according to the routing table the address is in its own subnet.

The plan here would be to NAT in both directions so that traffic from neither side of the firewall would be in conflict and they would be able to route the traffic. Everything coming from the Branch Office could be assigned an address in the 192.168.1.1 to 192.168.1.255 range and everything from the Main office going to the Branch Office could be assigned to an address in the 192.168.10.1 to 192.168.10.255 range.

## Benefits of NAT

### More IP addresses Available while Conserving Public IP Addresses

As explained earlier, this was the original intent of the technology and does not need to be gone into further.

### Financial Savings

Because an organization does not have to purchase IP addresses for every computer in use there is a significant cost savings due to using the process of Network Address Translation.

### Security Enhancements

One of the side benefits of the process of NAT is an improvement in security. Individual computers are harder to target from the outside and if port forwarding is being used computers on the inside of a firewall are less likely to have unmonitored open ports accessible from the Internet.

### Ease of Compartmentalization of Your Network

With a large available pool of IP addresses to use internally a network administrator can arrange things to be compartmentalized in a rational and easily remembered fashion and networks can be broken apart easily to isolate for reasons of network performance and security.

#### Example:

You have a large organization that for security reasons has certain departments that do not share network resources.

You can have the main section of the organization set up as follows;

<b>Network Devices</b>	192.168.1.1 to 192.168.1.25
<b>Internal Servers</b>	192.168.1.26 to 192.168.1.50
<b>Printers</b>	192.168.1.51 to 192.168.1.75
<b>Administration Personnel</b>	192.168.1.76 to 192.168.1.100
<b>Sales People</b>	192.168.1.101 to 192.168.1.200
<b>Marketing</b>	192.168.1.201 to 192.168.1.250

You could then have the following groups broken off into separate subnets:

<b>Accounting</b>	192.168.100.1 to 192.168.100.255
<b>Research and Development</b>	172.16.1.1 to 172.16.255.255
<b>Executive Management</b>	192.168.50.1 to 192.168.50.255
<b>Web sites and Email Servers</b>	10.0.50.1 to 10.0.50.255

These addresses do not have to be assigned right away but can be used as planned ranges.

## NAT in Transparent Mode

Similar to operating in NAT mode, when operating a FortiGate unit in Transparent mode you can add security policies and:

- Enable NAT to translate the source addresses of packets as they pass through the FortiGate unit.
- Add virtual IPs to translate destination addresses of packets as they pass through the FortiGate unit.
- Add IP pools as required for source address translation

A FortiGate unit operating in Transparent mode normally has only one IP address - the management IP. To support NAT in Transparent mode, you can add a second management IP. These two management IPs must be on different subnets. When you add two management IP addresses, all FortiGate unit network interfaces will respond to connections to both of these IP addresses.

Use the following steps to configure NAT in Transparent mode

1. Add two management IPs
2. Add an IP pool to the WAN1 interface
3. Add an Internal to WAN1 security policy

You can add the security policy from the web-based manager and then use the CLI to enable NAT and add the IP pool.

The usual practice of NATing in transparent mode makes use of two management IP addresses that are on different subnets, but this is not an essential requirement in every case.

If there is a router between the client systems and the FortiGate unit you can use the router's capabilities of tracking sessions to assign NATed addresses from an IP pool to the clients even if the assigned address don't belong to a subnet on your network.

### Example:

Client computer has an IP address of 1.1.1.33 on the subnet 1.1.1.0/24

Router "A" sits between the client computer and the FortiGate (in Transparent mode) with the IP address of 1.1.1.1 on the client's side of the router and the IP address of 192.168.1.211 on the FortiGate's side of the router.

Use NAT to assign addresses from an address pool of 9.9.9.1 to 9.9.9.99 to traffic coming from gateway of 192.168.1.211.

To enable the return traffic to get to the original computer, set up a static route than assigns any traffic with a destination of 9.9.9.0/24 to go through the 192.168.1.211 gateway. As long as the session for the outgoing traffic has been maintained, communication between the client computer and the external system on the other side of the FortiGate will work.

## Central NAT Table

The central NAT table enables you to define, and control with more granularity, the address translation performed by the FortiGate unit. With the NAT table, you can define the rules which dictate the source address or address group and which IP pool the destination address uses.

While similar in functionality to IP pools, where a single address is translated to an alternate address from a range of IP addresses, with IP pools there is no control over the translated port. When using the IP pool for source NAT, you can define a fixed port to guarantee the source port number is unchanged. If no fix port is defined, the port translation is randomly chosen by the FortiGate unit. With the central NAT table, you have full control over both the IP address and port translation.

The FortiGate unit reads the NAT rules in a top-down methodology, until it hits a matching rule for the incoming address. This enables you to create multiple NAT policies that dictate which IP pool is used based on the source address. The NAT policies can be rearranged within the policy list as well. NAT policies are applied to network traffic after a security policy.

## NAT 64 and NAT46

NAT64 and NAT46 are the terms used to refer to the mechanism that allows IPv6 addressed hosts to communicate with IPv4 addressed hosts and vice-versa. Without such a mechanism an IPv6 node on a network such as a corporate LAN would not be able to communicate with a web site that was still in a IPv4 only environment and IPv4 environments would not be able to connect to IPv6 networks.

One of these setups involves having at least 2 interfaces, 1 on an IPv4 network and 1 on an IPv6 network. The NAT64 server synthesizes AAAA records, used by IPv6 from A records used by IPv4. This way client-server and peer to peer communications will be able to work between an IPv6 only client and an IPv4 server without making changes to either of the end nodes in the communication transaction. The IPv6 network attached to the FortiGate unit should be a 32 bit segment, (for instance 64:ff9b::/96, see RFC 6052[Set up a link to <http://tools.ietf.org/html/rfc6052>], RFC 6146[set up a link <http://tools.ietf.org/html/rfc6146>]). IPv4 address will be embedded into the communications from the IPv6 client.

Because the IPv6 range of addresses is so much larger than the IPv4 range, a one to one mapping is not feasible. Therefore the NAT64 function is required to maintain any IPv6 to IPv4 mappings that it synthesizes. This can be done either statically by the administrator or automatically by the service as the packets from the IPv6 network go through the device. The first method would be a stateless translation and the second would be a stateful translation. NAT64 is designed for communication initiated from IPv6 hosts to IPv4 addresses. It is address mapping like this that allows the reverse to occur between established connections. The stateless or manual method is an appropriate solution when the NAT64 translation is taking place in front of legacy IPv4 servers to allow those specific servers to be accessed by remote IPv6-only clients. The stateful or automatic solution is best used closer to the client side when you have to allow some specific IPv6 clients to talk to any of the IPv4-only servers on the Internet.

There are currently issues with NAT64 not being able to make everything accessible. Examples would be SIP, Skype, MSN, Goggle talk, and sites with IPv4 literals. IPv4 literals being IPv4 addresses that are imbedded into content rather than a FQDN.

Policies that employ NAT64 or NAT46 can be configured from the web-based manager as long as the feature is enabled using the Features setting found at *System > Config > Features*.

- To create a NAT64 policy go to *Policy > Policy > NAT64 Policy* and select *Create New*.
- To create a NAT46 policy go to *Policy > Policy > NAT46 Policy* and select *Create New*.

The difference between these NAT policies and regular polices is that there is no option to use the security profiles and sensors.

## NAT 66

NAT 66 is Network Address Translation between 2 IPv6 network. The basic idea behind NAT 66 is no different than the regular NAT between IPv4 networks that we are all used to. The difference are in the mechanics of how it is performed, mainly because of the complexity and size of the addresses that are being dealt with.

In an IPv4 world, the reason for the use of NAT was usually one or a combination of the following 3 reasons:

- Improved security - actual addresses behind NAT are virtually hidden
- Amplification of addresses - hundreds of computers can use as little as a single public IP address
- Internal address stability - there is control of internal addressing. The addresses can stay the same even if Internet Service Providers change.

In these days of security awareness the protective properties of NAT are not something that are not normally depended on by themselves to defend a network and with the vastly enlarged IPv6 address scope there is no longer a need to amplify the available addresses. However, the desire to have internal address control still exists. The most common reason for using NAT66 is likely to be the maintaining of the existing address scheme of the internal network despite changes outside of it. Imagine that you have an internal network of 2000 IP addresses and one day the company changes its ISP and thus the addresses assigned to it. Even if most of the addressing is handled by DHCP, changing the address scheme is going to have an impact on operations.

Addressing stability can be achieved by:

- Keeping the same provider - this would depend on the reason for the change. If the cost of this provider has become too expensive this is unlikely. If the ISP is out of business it becomes impossible.
- Transfer the addresses from the old provider to the new one - There is little motivation for an ISP to do you a favor for not doing business with them.
- Get your own autonomous system number - this can be too expensive for smaller organizations.
- NAT - this is the only one on the list that is in the control of IT.

There are differences between NAT66 and IPv4 NAT. Because there is no shortage of addresses most organizations will be given a /48 network that can be translated into another /48 network. This allows for a one to one translation, no need for port forwarding. This is a good thing because port forwarding is more complicated in IPv6. In fact, NAT66 will actually just be the rewriting of the prefix on the address.

Example:

If your current IPv6 address is

```
2001:db8:cafe::/48
```

you could change it to

```
2001:db8:fea7::/48
```

There is an exception to the one to one translation. NAT66 cannot translate internal networks that contain 0xffff in bits 49 through 63 - this is due to the way checksums are calculated in TCP/IP: they use the one's-complement representation of numbers which assigns the value zero to both 0x0000 and 0xffff.

## How Packets are handled by FortiOS

To give you idea of what happens to a packet as it makes its way through the FortiGate unit here is a brief overview. This particular trip of the packet is starting on the Internet side of the FortiGate firewall and ends with the packet exiting to the Internal network. An outbound trip would be similar. At any point in the path if the packet is going through what would be considered a filtering process and if fails the filter check the packet is dropped and does not continue any further down the path.



This information is covered in more detail in other in the Troubleshooting chapter of the FortiOS Handbook in the Life of a Packet section.

The incoming packet arrives at the external interface. This process of entering the device is referred to as **ingress**.

#### **Step #1 - Ingress**

1. Denial of Service Sensor
2. IP integrity header checking
3. IPSec connection check
4. Destination NAT
5. Routing

#### **Step #2 - Stateful Inspection Engine**

1. Session Helpers
2. Management Traffic
3. SSL VPN
4. User Authentication
5. Traffic Shaping
6. Session Tracking
7. Policy lookup

#### **Step #3 - Security Profiles scanning process**

1. Flow-based Inspection Engine
2. IPS
3. Application Control
4. Data Leak Prevention
5. Email Filter
6. Web Filter
7. Anti-virus
8. Proxy-based Inspection Engine
9. VoIP Inspection
10. Data Leak Prevention
11. Email Filter
12. Web Filter
13. Anti-virus
14. ICAP

#### **Step #4 - Egress**

1. IPSec
2. Source NAT
3. Routing

## **FortiGate Modes**

The FortiGate unit has a choice of modes that it can be used in, either NAT/Route mode or Transparent mode. The FortiGate unit is able to operate as a firewall in both modes, but some of

its features are limited in Transparent mode. It is always best to choose which mode you are going to be using at the beginning of the set up. Once you start configuring the device, if you want to change the mode you are going to lose all configuration settings in the change process.

## NAT/Route Mode

NAT/Route mode is the most commonly used mode by a significant margin and is thus the default setting on the device. As the name implies the function of NAT is commonly used in this mode and is easily configured but there is no requirement to use NAT. The FortiGate unit performs network address translation before IP packets are sent to the destination network.

These are some of the characteristics of NAT/Route mode:

- Typically used when the FortiGate unit is a gateway between private and public networks.
- Can act as a router between multiple networks within a network infrastructure.
- When used, the FortiGate unit is visible to the networks that is connected to.
- Each logical interface is on a distinct subnet.
- Each Interface needs to be assigned a valid IP address for the subnet that it is connected to it.

## Transparent Mode

Transparent mode is so named because the device is effectively transparent in that it does not appear on the network in the way that other network devices show as a nodes in the path of network traffic. Transparent mode is typically used to apply the FortiOS features such as Security Profiles etc. on a private network where the FortiGate unit will be behind an existing firewall or router.

These are some of the characteristics of Transparent mode:

- The FortiGate unit is invisible to the network.
- All of its interfaces are on the same subnet and share the same IP address.
- The FortiGate unit uses a Management IP address for the purposes of Administration.
- Still able to use NAT to a degree, but the configuration is less straightforward

In Transparent mode, you can also perform NAT by creating a security policy or policies that translates the source addresses of packets passing through the FortiGate unit as well as virtual IP addresses and/or IP pools.

## Quality of Service

The Quality of Service (QoS) feature allows the management of the level of service and preference given to the various types and sources of traffic going through the firewall so that the traffic that is important to the services and functions connecting through the firewall gets the treatment required to ensure the level of quality that is required. QoS can be helpful for organizations that are trying to manage their voice and streaming multi-media traffic, which can rapidly consume bandwidth. Both voice and streaming multi-media are sensitive to latency. FortiGate units support QoS using traffic policing, traffic shaping, and queuing.

### Traffic policing

Packets are dropped that do not conform to bandwidth limitations

## Traffic Shaping

Assigning minimum levels of bandwidth to be allocated to specific traffic flows to guarantee levels of service or assigning maximum levels of bandwidth to be allocated to specific traffic flows so that they do not impede other flows of traffic.

This helps to ensure that the traffic may consume bandwidth at least at the guaranteed rate by assigning a greater priority queue if the guarantee is not being met. Traffic shaping also ensures that the traffic cannot consume bandwidth greater than the maximum at any given instant in time. Flows that are greater than the maximum rate are subject to traffic policing.

## Queuing

Assigning differing levels of priority to different traffic flows so that traffic flows that are adversely effected by latency are prevented from being effected by traffic flows that are not subject to the effects of latency. All traffic in a higher priority traffic queue must be completely transmitted before traffic in lower priority queues will be transmitted.

An example of where you would want to use something like this is if you had competing traffic flows of Voice over IP traffic and email traffic. The VoIP traffic is highly susceptible to latency issues. If you have a delay of a few seconds it is quickly noticeable when it is occurring. Email on the other hand can have a time delay of much longer and it is highly unlikely that it will be noticed at all.



By default, the priority given to any traffic is high, so if you want to give one type of traffic priority over all other traffic you will need to lower the priority of all of the other traffic.

---

For additional information about QoS, see the Traffic Shaping chapter of the FortiOS Handbook.

## Interfaces and Zones

A Firewall is a gateway device that may be the nexus point for more than 2 networks. The interface that the traffic is coming in on and should be going out on is a fundamental concern for the purposes of routing as well as security. Routing, policies and addresses are all associated with interfaces. The interface is essentially the connection point of a subnet to the FortiGate unit and once connected can be connected to other subnets.

Physical interfaces or not the only ones that need to be considered. There are also virtual interfaces that can be applied to security policies. VLANs are one such virtual interface. Interfaces if certain VPN tunnels are another.

Policies are the foundation of the traffic control in a firewall and the Interfaces and addressing is the foundation that policies are based upon. Using the identity of the interface that the traffic connects to the FortiGate unit tells the firewall the initial direction of the traffic. The direction of the traffic is one of the determining factors in deciding how the traffic should be dealt with. You can tell that interfaces are a fundamental part of the policies because, by default, this is the criteria that the policies are sorted by.

Zones are a mechanism that was created to help in the administration of the firewalls. If you have a FortiGate unit with a large number of ports and a large number of nodes in you network the chances are high that there is going to be some duplication of policies. Zones provide the option of logically grouping multiple virtual and physical FortiGate firewall interfaces. The zones can then be used to apply security policies to control the incoming and outgoing traffic on those interfaces. This helps to keep the administration of the firewall simple and maintain consistency.

For example you may have several floors of people and each of the port interfaces could go to a separate floor where it connects to a switch controlling a different subnet. The people may be on different subnets but in terms of security they have the same requirements. If there were 4 floors and 4 interfaces a separate policy would have to be written for each floor to be allowed out on to the Internet off the WAN1 interface. This is not too bad if that is all that is being done, but now start adding the use of more complicated policy scenarios with Security Profiles, then throw in a number of Identity based issues and then add the complication that people in that organization tend to move around in that building between floors with their notebook computers. Each time a policy is created for each of those floors there is a chance of an inconsistency cropping up. Rather than make up an additional duplicate set of policies for each floor, a zone can be created that combines multiple interfaces. And then a single policy can be created that uses that zone as one side of the traffic connection.

# Firewall objects

As was mentioned earlier, the components of the FortiGate firewall go together like interlocking building blocks. The Firewall objects are a prime example of those building blocks. They are something that can be configured once and then used over and over again to build what you need. They can assist in making the administration of the FortiGate unit easier and more intuitive as well as easier to change. By configuring these objects with their future use in mind as well as building in accurate descriptions the firewall will become almost self documenting. That way, months later when a situation changes, you can take a look at a policy that needs to change and use a different firewall object to adapt to the new situation rather than build everything new from the ground up to accommodate the change.

This chapter includes information about the following Firewall objects:

- [Addresses](#)
- [Services and TCP ports](#)
- [Firewall schedules](#)
- [Security profiles](#)

## Addresses

Firewall addresses define sources and destinations of network traffic and are used when creating policies. When properly set up these firewall objects can be used with great flexibility to make the configuration of firewall policies simpler and more intuitive. The FortiGate unit compares the IP addresses contained in packet headers with a security policy's source and destination addresses to determine if the security policy matches the traffic.

The addresses in the FortiGate unit can include:

- IPv4 addresses
- IPv6 addresses
- IPv4 Address Groups
- IPv6 Address Groups
- IP Pools
- Virtual IP Addresses
- Geography based addresses
- Wildcard addresses and netmasks
- Fully Qualified Domain Name addresses

When setting up an address one of the parameters that is asked for is the interface. This means that the system will expect to see that address only on the interface that you select. You can only select one interface. If you expect that the address may be seen at more than one interface you can choose the "any" interface option. Whenever, possible it is best to choose a more specific interface than the "any" option because in the GUI configuration of firewall policies there is a drop down field that will show the possible addresses that can be used. The drop down will only show those addresses that can be on the interface assigned for that interface in the policy.

Example:

- You have an address called “XYZ”
- “XYZ” is set to the WAN1 interface because that is the only interface that will be able to access that address.
- When you are selecting a Source Address in the Web-based Manager for a policy that is using the DMZ the address “XYZ” will not be in the drop-down menu.

When there are only 10 or 20 addresses this is not a concern, but if there are a few hundred addresses configured it can make your life easier.

Addresses, address groups, and virtual IPs must have unique names to avoid confusion in firewall policies. If an address is selected in a policy, the address cannot be deleted until it is deselected from the policy.



#### Addressing Best Practices Tip

The other reason to assign a specific interface to addresses is that it will prevent you from accidentally assigning an address where it will not work properly. Using the example from earlier, if the “XYZ” address was assigned to the “Any” interface instead of WAN1 and you configure the “XYZ” address.

---

## IPv4 Address and Net Mask

When representing hosts by an IP address with a netmask, the IP address can represent one or more hosts. For example, a firewall address can be:

- A single host such as a single computer with the address 192.45.46.45
- A range of hosts such as all of the hosts on the subnet 192.45.46.1 to 192.45.46.255
- All hosts, represented by 0.0.0.0 which matches any IP address

The netmask corresponds to the subnet class of the address being added, and can be represented in either dotted decimal or CIDR format. The FortiGate unit automatically converts CIDR formatted netmasks to dotted decimal format. Example formats:

- Netmask for a class A subnet of 16,777,214 usable addresses: 255.0.0.0, or /8
- Netmask for a class B subnet of 65,534 usable addresses: 255.255.0.0, or /16
- Netmask for a class C subnet of 254 usable addresses: 255.255.255.0, or /24
- Netmask for subnetted class C of 126 usable addresses: 255.255.255.128, or /25
- Netmask for subnetted class C of 62 usable addresses: 255.255.255.128, or /26
- Netmask for subnetted class C of 30 usable addresses: 255.255.255.128, or /27
- Netmask for subnetted class C of 14 usable addresses: 255.255.255.128, or /28
- Netmask for subnetted class C of 6 usable addresses: 255.255.255.128, or /29
- Netmask for subnetted class C of 2 usable addresses: 255.255.255.128, or /30
- Netmask for a single computer: 255.255.255.255, or /32
- Netmask used with 0.0.0.0 to include all IP addresses: 0.0.0.0, or /0

So for a single host or subnet the valid format of IP address and netmask could be either:

x.x.x.x/x.x.x.x, such as 192.168.1.0/255.255.255.0

or

x.x.x.x/x, such as 192.168.1.0/24

It is also possible to describe a range of IP addresses with a subnet. This would be a continuous set of IP addresses within a subnet. It does not have to include the whole subnet. The formats would be (the \* being used as a wildcard character):

x.x.x.x-x.x.x.x, such as 192.168.110.100-192.168.110.120

x.x.x.[x-x], such as 192.168.110.[100-120]

x.x.x.\*, such as 192.168.110.\*

## FQDN Addressing

By using Fully Qualified Domain Name (FQDN) addressing you can take advantage of the dynamic ability of the service to keep up with address changes without having to manually change the addresses on the FortiGate. FQDN addresses are most often used with external web sites but they can be used for internal web sites as well if there is a trusted DNS server that can be accessed. FQDN addressing also comes in handy for large web sites that may use multiple addresses and load balancers for their web sites. The FortiGate firewall automatically maintains a cached record of all the addresses resolved by the DNS for the FQDN addresses used.

For example, if you were doing this manually and you wanted to have a security policy that involved Google you could track down all of the IP addresses that they use across multiple countries. Using the FQDN address is simpler and more convenient.

When representing hosts by an FQDN, the domain name can also be a subdomain, such as mail.example.com.

Valid FQDN formats include:

- <host\_name>.<top\_level\_domain\_name> such as example.com
- <host\_name>.<second\_level\_domain\_name>.<top\_level\_domain\_name>, such as mail.example.com

When creating FQDN entries it is important to remember that:

- Wildcards are not supported in FQDN address objects
- While there is a level of convention that would imply it, “www.example.com” is not necessarily the same address of “example.com”. they will each have their own records on the DNS server.

The FortiGate firewall keeps track of the DNS TTLs so as the entries change on the DNS servers the IP address will effectively be updated for the FortiGate. As long as the FQDN address is used in a security policy, it stores the address in the DNS cache.



There is a possible security downside to using FQDN addresses. Using a fully qualified domain name in a security policy means that your policies are relying on the DNS server to be accurate and correct. DNS servers in the past were not seen as potential targets because the thinking was that there was little of value on them and therefore are often not as well protected as some other network resources. People are becoming more aware that the value of the DNS server is that in many ways it controls where users and computers go on the Internet. Should the DNS server be compromised, security policies requiring domain name resolution may no longer function properly.

---

## Geography Based Addressing

Geography based addressing enables you to design policies based on addresses that are associated with a country.

This feature can be used to make either inclusive or exclusive policies. For instance, if you have a SSL VPN where the users will only be connecting from a single country, but you don't know

from where in that country, you can filter out any connections coming in that are from outside that country.

On the other side of the equation, if you find that you are constantly being attacked by malicious intruders from a few countries that you have no dealings with you can block access to them before any traffic comes through.

The matching of geographical country designations to an IP address is achieved by collecting data from any IP addresses that connect to any of the FortiGuard Servers throughout the world. As a secondary task, when a FortiGuard server connects to an IP it also does a search on the Country of origin for the address and updates the database.

There is no single comprehensive list of IP addresses and their locations available because IP addresses can be transferred between ISPs or countries and some organization may not keep complete or up-to-date records regarding locations. FortiGuard Services are constantly updating their database of addresses matched to locations, but the database is dynamic and there may be addresses that have not been resolved to a location. While this means that there can be gaps in the completeness of the database it is possible to fill them in manually by means local to your FortiGate unit.

IPv6 does not support geography-based addressing. This feature is for IPv4 addresses only.



#### Best Practices Tip:

Based on the limitation of the IP address matched to country database, it is best to use this type of address in a group with other addresses to fill in the gaps. For instance, if you are a company in Country “A” and all of your employees that will be using the SSL-VPN connection are in that country, the best practice would be to create an address group that includes the geographical address of Country “A”. As valid addresses appear that are not allowed, you can add these other IPs to that group using IP addresses or IP range addresses without having to change the policy itself.

---

If you are trying to block addresses the principle works just the same. Your logs show that someone from IP address x.x.x.x has been trying to connect inappropriately to your network. You use a IP locator web site to determine that they have been attempting to connect from Country “X”. Up until now they have not been successful, but you don’t deal with the country they are connecting from so don’t mind blocking the whole country. Create an address group that is designed for Blocking Access to any addresses in it, then add the geographical address for Country “X”. Even if the policy does not block every single IP address from that country you have greatly increased your odds of blocking potential intrusion attempts. As your logs show other attempts you can look them up in an IP locator web site and if they are from the same country you can add the IP address for the subnet that they are connecting from.

## Address Groups

Address groups are designed for ease of use in the administration of the device. If you have a number of addresses or address ranges that will commonly be treated the same or require the same security policies, you can put them into address groups, rather than entering multiple individual addresses in each policy refers to them.

The use of groups is not required. If you have a number of different addresses you could add them individually to a policy and the FortiGate firewall will process them just as quickly and efficiently as if they were in a group, but the chances are that if you have used a group once you could need to use it again and depending on the number of addresses involved entering them individually for each policy can become tedious and the likelihood of an address being missed becomes greater. If you have a number of policies using that combination of addresses it is much easier to add or subtract addresses from the group than to try and remember all of the firewall policies that combination of addresses was used in. With the group, you only have to make the one edit and it is used by any firewall policy using that address group.



Because security policies require addresses with homogenous network interfaces, address groups should contain only addresses bound to the same network interface, or to Any.

For example, if address 1.1.1.1 is associated with port1, and address 2.2.2.2 is associated with port2, they cannot be in the same group. However, if 1.1.1.1 and 2.2.2.2 are configured with an interface of Any, they can be grouped, even if the addresses involve different networks.

IPv4 address groups and IPv6 address groups are created and treated separately. You cannot mix IPv4 firewall addresses and IPv6 firewall addresses in the same address group. Because the Internet is currently based on IPv4 addresses IPv6 address groups cannot include FQDN or Geography based addresses.

## Wildcard Addressing

Wildcard addresses are addresses that identify ranges of IP addresses, reducing the amount of firewall addresses and security policies required to match some of the traffic on your network. Wildcard addresses are an advanced feature, usually required only for complex networks with complex firewall filtering requirements. By using these wildcard addresses in the firewall configuration, administrators can eliminate creating multiple, separate IP based address objects and then grouping them to then apply to multiple security policies.

A wildcard address consists of an IP address and a wildcard netmask, for example, 192.168.0.56 255.255.0.255. In this example, the IP address is 192.168.0.56 and the wildcard netmask is 255.255.0.255. The IP address defines the networks to match and the wildcard netmask defines the specific addresses to match on these networks.

In a wildcard netmask, zero means ignore the value of the octet in the IP address, which means the wildcard firewall address matches any number in this address octet. This also means that the number included in this octet of IP address is ignored and can be any number. Usually, if the octet in the wildcard netmask is zero, the corresponding octet in the IP address is also zero.

In a wildcard netmask, a number means match addresses according to how the numbers translate into binary addresses. For example, the wildcard netmask is 255; the wildcard address will only match addresses with the value for this octet that is in the IP address part of the wildcard address. So, if the first octet of the IP address is 192 and the first octet of the wildcard netmask is 255, the wildcard address will only match addresses with 192 in the first octet.

In the above example, the wildcard address 192.168.0.56 255.255.0.255 would match the following IP addresses:

192.168.0.56, 192.168.1.56, 192.168.2.56, ..., 192.168.255.56

The wildcard addresses 192.168.0.56 255.255.0.255 and 192.168.1.56 255.255.0.255 define the same thing since the 0 in the wildcard mask means to match any address in the third octet.

If we use the wildcard address 172.0.20.10 255.0.255.255, it would match the following IP addresses:

172.1.20.10, 172.2.20.10, 172.3.20.10, ..., 172.255.20.10

In a wildcard netmask, a number other than 255 matches multiple addresses for this octet. You can perform a binary conversion to calculate the addresses that would be matched by a given

value. For example, to create the IP address and wildcard netmask to match the following network addresses:

```
192.168.32.0/24
192.168.33.0/24
192.168.34.0/24
192.168.35.0/24
192.168.36.0/24
192.168.37.0/24
192.168.38.0/24
192.168.39.0/24
```

Table 45 shows how to write the third octet for these networks according to the octet bit position and address value for each bit.

**Table 45:** Octet bit position and address value for each bit

Decimal	128	64	32	16	8	4	2	1
32	0	0	1	0	0	0	0	0
33	0	0	1	0	0	0	0	1
34	0	0	1	0	0	0	1	0
35	0	0	1	0	0	0	1	1
36	0	0	1	0	0	1	0	0
37	0	0	1	0	0	1	0	1
38	0	0	1	0	0	1	1	0
39	0	0	1	0	0	1	1	1
	M	M	M	M	M	D	D	D

Since the first five bits match, the networks can be summarized into one network (192.168.32.0/21 or 192.168.32.0 255.255.248.0). All eight possible combinations of the three low-order bits are relevant for the network ranges. The wildcard address that would match all of these subnet addresses can be written as 192.168.32.0 255.255.248.0.

Wildcard addresses are similar to routing access list wildcard masks. You add routing access lists containing wildcard masks using the `config router access list` command. However, router access list wildcard masks use the inverse of the masking system used for firewall wildcard addresses. For the router access list wildcard masks, zero (0) means match all IP addresses and one (1) means ignore all IP addresses. So to match IP addresses 192.168.0.56, 192.268.1.56, 192.168.2.56,... 192.168.255.56 you would use the following router access IP address prefix and wildcard mask: 192.168.0.56 0.0.255.0.

Wildcard firewall addresses are configured only in the CLI. The following is an example of how to configure a wildcard firewall address.

```
config firewall address
 edit example_wildcard_address
 set type wildcard
 set wildcard 192.168.0.56 255.255.0.255
 end
```

## Virtual IP Addresses (VIPs)

The mapping of a specific IP address to another specific IP address is usually referred to as Destination NAT. FortiOS has a component that is a bit more specialized along this line called a Virtual IP Address, sometimes referred to as a VIP. FortiOS uses a Virtual IP address to map an External IP address to an IP address. This address does not have to be an individual host, it can also be an address range. This mapping can include all TCP/UDP ports or if Port Forwarding is enabled it will only refer to the specific ports configured.

Virtual IP addresses are typically used to NAT external or Public IP addresses to internal or Private IP addresses. Using a Virtual IP address between 2 internal Interfaces made up of Private IP addresses is possible but there is rarely a reason to do so as the 2 networks can just use the IP addresses of the networks without the need for any address translation. Using a Virtual IP address for traffic going from the inside to the Internet is even less likely to be a requirement, but it is supported.

Something that needs to be considered when there are multiple Public IP addresses on the external interface(s) is that when a Virtual IP address is used without Port Forwarding enabled there is a reciprocal effect as far as traffic flow is concerned. Normally, on a firewall policy where NAT is enabled, for outgoing traffic the internal address is translated to the Public address that is assigned to the FortiGate, but if there is a Virtual IP address with no port forwarding enabled, then the Internal IP address in the Mapped field would be translated to the IP address configured as the External Address in the VIP settings.



**Best practice:** Put any policies with a VIP right at the beginning of the policy list, with nothing before them. VIP traffic is processed first, before the regular rules of order in policies are applied. The sequence of the policies not containing VIPS will not effect those that do contain VIPS, regardless of the order they are in. Put VIP policies before any others to remind yourself of where they really are in the sequence.

For more on this topic, read the [“Exception to policy order \(VIPs\)”](#) on page 956.

### Example:

- The assigned External address (WAN1) of the FortiGate unit is 172.12.96.3 with a subnet mask of 255.255.255.128
- There is a Virtual IP address set up to map the external address 172.12.96.127 on WAN1 to the internal IP address of 192.168.1.127
- Port Forwarding is not enabled because you want all allowed traffic going to the external IP address to go to this server.

In this case any outbound traffic from 192.168.1.127 will go out on WAN1 with the IP address of 172.12.96.127 as the source IP address.

In terms of actually using the Virtual IP address, they would be using in the security policies in the same places that other addresses would be used, usually as a Destination Address.



There is another type of address that the term “virtual IP address” commonly refers to which is used in load balancing and other similar configurations. In those cases, a number of devices share a separately created virtual IP address that can be sent to multiple possible devices. In FortiOS these are referred to as Virtual Servers and are configured in the “Load Balance” section.

## Virtual IP Groups

Just like other address, Virtual IP addresses can be organized into groups for ease of administration. If you have multiple virtual IPs that are likely to be associated to common firewall policies rather than add them individually to each of the policies you can add the instead. That

way, if the members of the group change then any changes made to the group will propagate to all of the policies using that group.

When using a Virtual IP address group the firewall policy will take into account all of the configured parameters of the Virtual IPs: IP addresses, Ports and port types.

## IP Pools

IP Pools are a mechanism that allow sessions leaving the FortiGate Firewall to use NAT. An IP pool defines a single IP address or a range of IP addresses to be used as the source address for the duration of the session. These assigned addresses will be used instead of the IP address assigned to that FortiGate interface.



When using IP pools for NATing, there is a limitation that must be taken into account when configuring the pool. If the IP address(es) within the pool are different from the IP address(es) that are assigned to the interface communications based on those IP addresses will fail. For example if the IP addresses assigned to an interface are 172.16.100.1 - 172.16.100.14, you cannot choose 10.11.12.50 - 10.11.12.59 for the IP pool.

There are 4 types of IP Pools that can be configured on the FortiGate firewall:

- One-to-One - in this case the only internal address used by the external address is the internal address that it is mapped to.
- Overload - this is the default setting. Internal addresses other than the one designated in the policy can use this address for the purposes of NAT.
- Fixed Port Range - rather than a single address to be used, there is a range of addresses that can be used as the NAT address. These addresses are randomly assigned as the connections are made.
- Port Block Allocation - this setting is used to allocate a block of port numbers for IP pool users. Two variables will also have to be set. The block size can be set from 64 to 4096 and as the name implies describes the number of ports in one block of port numbers. The number of blocks per user determines how many of these blocks will be assigned. This number can range from 1 to 128.



Be careful when calculating the values of the variables. The maximum number of ports that are available on an address is 65,536. If you chose the maximum value for both variables you will get a number far in excess of the available port numbers.

$$4096 \times 128 = 524,288$$

One of the more common examples is when you have an email server behind your FortiGate firewall and the range of IP addresses assigned to you by your ISP is more than one. If an organization is assigned multiple IP addresses it is normally considered a best practice to assign a specific address other than the one used for the Firewall to the mail server. However, when normal NAT is used the address assigned to the firewall is also assigned to any outbound sessions. Anti-spam services match the source IP address of mail traffic that they receive to the MX record on DNS servers as an indicator for spam. If there is a mismatch the mail may not get through so there is a need to make sure that the NATed address assigned matches the MX record.

You can also use the Central NAT table as a way to configure IP pools.

### Source IP address and IP pool address matching when using a range

When the source addresses are translated to an IP pool that is a range of addresses, one of the following three cases may occur:

### Scenario 1:

The number of source addresses equals that of IP pool addresses

In this case, the FortiGate unit always matches the IP addressed one to one.

If you enable fixed port in such a case, the FortiGate unit preserves the original source port. This may cause conflicts if more than one security policy uses the same IP pool, or the same IP addresses are used in more than one IP pool.

### Scenario 2:

The number of source addresses is more than that of IP pool addresses

In this case, the FortiGate unit translates IP addresses using a wrap-around mechanism. If you enable fixed port in such a case, the FortiGate unit preserves the original source port. But conflicts may occur since users may have different sessions using the same TCP 5 tuples.

### Scenario 3:

The number of source addresses is fewer than that of IP pool addresses

In this case, some of the IP pool addresses are used and the rest of them are not be used.

## ARP Replies

If a FortiGate firewall interface IP address overlaps with one or more IP pool address ranges, the interface responds to ARP requests for all of the IP addresses in the overlapping IP pools. For example, consider a FortiGate unit with the following IP addresses for the port1 and port2 interfaces:

- port1 IP address: 1.1.1.1/255.255.255.0 (range is 1.1.1.0-1.1.1.255)
- port2 IP address: 2.2.2.2/255.255.255.0 (range is 2.2.2.0-2.2.2.255)

And the following IP pools:

- IP\_pool\_1: 1.1.1.10-1.1.1.20
- IP\_pool\_2: 2.2.2.10-2.2.2.20
- IP\_pool\_3: 2.2.2.30-2.2.2.40

The port1 interface overlap IP range with IP\_pool\_1 is:

$(1.1.1.0-1.1.1.255) \text{ and } (1.1.1.10-1.1.1.20) = 1.1.1.10-1.1.1.20$

The port2 interface overlap IP range with IP\_pool\_2 is:

$(2.2.2.0-2.2.2.255) \text{ \& } (2.2.2.10-2.2.2.20) = 2.2.2.10-2.2.2.20$

The port2 interface overlap IP range with IP\_pool\_3 is:

$(2.2.2.0-2.2.2.255) \text{ \& } (2.2.2.30-2.2.2.40) = 2.2.2.30-2.2.2.40$

And the result is:

- The port1 interface answers ARP requests for 1.1.1.10-1.1.1.20
- The port2 interface answers ARP requests for 2.2.2.10-2.2.2.20 and for 2.2.2.30-2.2.2.40

Select Enable NAT in a security policy and then select Dynamic IP Pool. Select an IP pool to translate the source address of packets leaving the FortiGate unit to an address randomly selected from the IP pool. Whether or not the external address of an IP Pool will respond to an ARP request can be disabled. You might want to disable the ability to responded to ARP requests so that these address cannot be used as a way into your network or show up on a port scan.

## IP pools and zones

Because IP pools are associated with individual interfaces

IP pools cannot be set up for a zone. IP pools are connected to individual interfaces.

## Fixed Port

Some network configurations do not operate correctly if a NAT policy translates the source port of packets used by the connection. NAT translates source ports to keep track of connections for a particular service.

However, enabling the use of a fixed port means that only one connection can be supported through the firewall for this service. To be able to support multiple connections, add an IP pool, and then select Dynamic IP pool in the policy. The firewall randomly selects an IP address from the IP pool and assigns it to each connection. In this case, the number of connections that the firewall can support is limited by the number of IP addresses in the IP pool.

## Match-VIP

The match-vip feature allows the FortiGate unit to log virtual IP traffic that gets implicitly dropped. This feature eliminates the need to create two policies for virtual IPs; one that allows the virtual IP, and the other to get proper log entry for DROP rules.

For example, you have a virtual IP security policy and enabled the match-vip feature; the virtual IP traffic that is not matched by the policy is now caught.

The match-vip feature is available only in the CLI. By default, the feature is disabled.

## Services and TCP ports

There are a number of different services and protocols in use on the Internet. The most commonly known is HTTP which is used by web servers to transmit requests and responses for unencrypted web pages. These services are set up to listen for requests on a numbered port. These services and protocols can use any port from 1 to 65,535. To keep things simple for everyone a large number of the more commonly used services started using a standardized list of ports. For instance, though it is not required, by default, most web servers listen for HTTP requests on port 80 and by default, web browsers will send HTTP traffic to port 80. If you wish to use another port such as 8080 you would put “:8080” at the end of the URL to indicate that you want the browser to use 8080 instead of the default port.

### Example:

Default URL for HTTP traffic when the web server is listening on the standard HTTP port:

```
http://fortinet.com
```

URL to the same address when the web server is listening for HTTP traffic on port 8080

```
http://fortinet.com:8080
```

Services represent typical traffic types and application packets that pass through the FortiGate unit. Firewall services define one or more protocols and port numbers associated with each service. Security policies use service definitions to match session types. You can organize related services into service groups to simplify your security policy list.

Many well-known traffic types have been predefined on the FortiGate unit. If there is a service that does not appear on the list you can create a service or edit an existing one. You need to

know the ports, IP addresses or protocols of that particular service or application uses, to create a service.



### Best Practices

While you can edit a predefined service it is best to leave those ones alone and create a new service and name it something similar such as the same service name with a descriptive identifier appended.

Based on the previous example, instead of the name “HTTP” you could name the service “HTTP8080” or use the application that is using that port, “HTTP-Application”.

---

## Categories

In order to make sorting through the services easier there is a field to categorize the services. The services can be sorted into the following groups:

- Uncategorized
- General
- Web Access
- File Access
- Email
- Network Services
- Authentication
- Remote Access
- Tunnelling
- VoIP, Messaging and Other Applications

## Protocol Types

One of the fundamental aspects of a service is the type of protocol that use used to define it. When a service is defined one of the following categories of protocol needs to be determined:

- TCP/UDP/SCTP
- ICMP
- ICMP6
- IP

Depending on which of these protocol categories is choose another set of specifications will can also be defined.

### TCP/UDP/SCTP

This is the most commonly used service protocol category. Once this category has been selected the other available options to choose are an address, either IP or FQDN, and the protocol and port number.

The protocol will be TCP, UDP or SCTP.

### ICMP or ICMP6

When ICMP or ICMP6 is chosen the available options are the ICMP Type and its code.

## IP

When IP is the chosen protocol type the addition option is the Protocol Number.

## TCP

Transmission Control Protocol (TCP) is one of the core or fundamental protocols of the Internet. It is part of the Transport Layer of the OSI Model. It is designed to provide reliable delivery of data from a program on one device on the network or Internet to another program on another device on the network or Internet. TCP achieves its reliability because it is a connection based protocol. TCP is stream-oriented. It transports streams of data reliably and in order.

TCP establishes a prior connection link between the hosts before sending data. This is often referred to as the handshake. Once the link is established the protocol uses checks to verify that the data transmitted. If an error check fails the data is retransmitted. This makes sure that the data is getting to the destination error free and in the correct order so that it can be put back together into a form that is identical to the way they were sent.

TCP is configured more for reliability than for speed and because of this TCP will likely be slower than a connectionless protocol such as UDP. This is why TCP is generally not used for real time applications such as voice communication or online gaming.

Some of the applications that use TCP are:

- World Wide Web (HTTP and HTTPS)
- Email (SMTP, POP3, IMAP4)
- Remote administration (RDP)
- File transfer (FTP)

## UDP

User Datagram Protocol (UDP) like TCP is one of the core protocols of the Internet and part of the Transport Layer of the OSI Model. UDP is designed more for speed than reliability and is generally used for different applications than TCP. UDP sends messages, referred to as datagrams across the network or Internet to other hosts without establishing a prior communication link. In other words, there is no handshake.

UDP is an unreliable service as the datagrams can arrive out of order, duplicated or go missing without any mechanism to verify them. UDP works on the assumption that any error checking is done by the application or is not necessary for the function of the application. This way it avoids the overhead that is required to verify the integrity of the data.

This lack of overhead improves the speed of the data transfer and is why UDP is often used by applications that are time sensitive in nature. UDP's stateless nature is also great for applications that answer a large number of small queries from a large number of clients.

Common uses for UDP are:

- Domain Name Resolution (DNS)
- Time (NTP)
- Streaming media (RTSP, RTP and RTCP)
- Telephone of the Internet (VoIP)
- File Transfer (TFTP)
- Logging (SNMP)
- Online games (GTP and OGP)



## SCTP

Stream Control Transmission Protocol (SCTP) is part of the Transport Layer of the OSI Model just like TCP and UDP and provides some of the features of both of those protocols. It is message or datagram orientated like UDP but it also ensures reliable sequential transport of data with congestion control like TCP.

SCTP provides the following services:

- Acknowledged error-free non-duplicated transfer of user data
- Data fragmentation to conform to discovered path MTU size
- Sequenced delivery of user messages within multiple streams, with an option for order-of-arrival delivery of individual user messages
- Optional bundling of multiple user messages into a single SCTP packet
- Network-level fault tolerance through supporting of multi-homing at either or both ends of an association
- Congestion avoidance behavior and resistance to flooding and masquerade attacks

SCTP uses multi-streaming to transport its messages which means that there can be several independent streams of messages traveling in parallel between the points of the transmission. The data is sent out in larger chunks of data than is used by TCP just like UDP but the messages include a sequence number within each message in the same way that TCP does so that the data can be reassembled at the other end of the transmission in the correct sequence without the data having to arrive in the correct sequence.

SCTP is effective as the transport protocol for applications that require monitoring and session-loss detection. For such applications, the SCTP path and session failure detection mechanisms actively monitor the connectivity of the session. SCTP differs from TCP in having multi-homing capabilities at either or both ends and several streams within a connection, typically referred to as an association. A TCP stream represents a sequence of bytes; an SCTP stream represents a sequence of messages.

Some common applications of SCTP include supporting transmission of the following protocols over IP networks:

- SCTP is important in 3G and 4G/LTE networks (for example, HomeNodeB = FemtoCells)
- SS7 over IP (for example, for 3G mobile networks)
- SCTP is also defined and used for SIP over SCTP and H.248 over SCTP
- Transport of Public Switched Telephone Network (PSTN) signaling messages over IP networks.

SCTP is a much newer protocol. It was defined by the IETF Signaling Transport (SIGTRAN) working group in 2000. It was introduced by [RFC 3286](#) and more fully define by [RFC 4960](#).

The FortiGate firewall can apply security policies to SCTP sessions in the same way as TCP and UDP sessions. You can create security policies that accept or deny SCTP traffic by setting the service to "ALL". FortiOS does not include pre-defined SCTP services. To configure security policies for traffic with specific SCTP source or destination ports you must create custom firewall services for SCTP.

FortiGate units route SCTP traffic in the same way as TCP and UDP traffic. You can configure policy routes specifically for routing SCTP traffic by setting the protocol number to 132. SCTP policy routes can route SCTP traffic according to the destination port of the traffic if you add a port range to the policy route.

You can configure a FortiGate unit to perform stateful inspection of different types of SCTP traffic by creating custom SCTP services and defining the port numbers or port ranges used by

those services. FortiGate units support SCTP over IPv4. The FortiGate unit performs the following checks on SCTP packets:

- Source and Destination Port and Verification Tag.
- Chunk Type, Chunk Flags and Chunk Length
- Verify that association exists
- Sequence of Chunk Types (INIT, INIT ACK, etc)
- Timer checking
- Four way handshake checking
- Heartbeat mechanism
- Protection against INIT/ACK flood DoS attacks, and long-INIT flooding
- Protection against association hijacking

FortiOS also supports SCTP sessions over IPSec VPN tunnels, as well as full traffic and event logging for SCTP sessions.

### Specific Addresses in TCP/UDP/SCTP

In the TCP/UDP/SCTP services it is also possible to set the parameter for a specific IP or Fully Qualified Domain Name address. The IP/FQDN field refers to the destination address of the traffic, not the source. This means for example, that you can set up a custom service that will describe in a policy the TCP traffic over port 80 going to the web site example.com, but you cannot set up a service that describes the TCP traffic over port 80 that is coming from the computer with the address 192.168.29.59.

### Protocol Port Values

The source and destination ports for TCP/UDP/SCTP services are important to get correct. If they are reversed the service will not work. The destination port(s) are the ones that refer to the ports that the computer will be listening on. These are the port numbers that most people are familiar with when they associate a port number to a protocol. In most cases the source port will be one that is randomly assigned by the computer that is not being already used by another service.

Most people associate HTTP with port 80. This means that a web-server will be listening on port 80 for any http requests being sent to the computer. The computer that is sending the request can use any port that is not already assigned to another service or communication session. There are 65,535 ports that it can randomly assign, but because the ports from 1 to 1024 are normally used for listening for incoming communications it is usually not in that range. It is unless there is a specific instance when you know that a communication will be coming from a predefined source port it is best practice to set the source port range from 1 to 65,535.

### ICMP

The Internet Control Message Protocol (ICMP) is a protocol layered onto the Internet Protocol Suite to provide error reporting flow control and first-hop gateway redirection. It is normally used by the operating systems of networked computers to send connectivity status query, response and error messages. It is assigned protocol number 1. There is a version of the protocol for both IPv4 and for IPv6. It is not designed to be absolutely reliable like TCP.

ICMP is not typically used for transporting data or for end-user network applications with the exception of some diagnostic utilities such as ping and traceroute.

ICMP messages are sent in several situations, for example:

- when a datagram cannot reach its destination,
- time exceeded messages
- redirect messages
- when the gateway does not have the buffering capacity to forward a datagram
- when the gateway can direct the host to send traffic on a shorter route.

Some of the specific ICMP message types are:

- ICMP\_ECHO
- ICMP\_TIMESTAMP
- ICMP\_INFO\_REQUEST
- ICMP\_ADDRESS

For ICMP error messages, only those reporting an error for an existing session can pass through the firewall. The security policy will allow traffic to be routed, forwarded or denied. If allowed, the ICMP packets will start a new session. Only ICMP error messages of a corresponding security policy is available will be sent back to the source. Otherwise, the packet is dropped. That is, only ICMP packets for a corresponding security policy can traverse the FortiGate unit.

### ICMP Types and Codes

ICMP has a number of messages that are identified by the “Type” field. Some of these types have assigned “Code” fields as well. The table below shows the different types of ICMP Types with their associated codes if there are any.

**Table 46:** ICMP Types and Codes

Type Number	Type Name	Code
0	Echo Reply	
1	Unassigned	
2	Unassigned	

**Table 46:** ICMP Types and Codes (continued)

Type Number	Type Name	Code
3	Destination Unreachable	0 Net Unreachable 1 Host Unreachable 2 Protocol Unreachable 3 Port Unreachable 4 Fragmentation Needed and Don't Fragment was Set 5 Source Route Failed 6 Destination Network Unknown 7 Destination Host Unknown 8 Source Host Isolated 9 Communication with Destination Network is Administratively Prohibited 10 Communication with Destination Host is Administratively Prohibited 11 Destination Network Unreachable for Type of Service 12 Destination Host Unreachable for Type of Service 13 Communication Administratively Prohibited 14 Host Precedence Violation 15 Precedence cutoff in effect
4	Source Quench	
5	Redirect	0 Redirect Datagram for the Network (or subnet) 1 Redirect Datagram for the Host 2 Redirect Datagram for the Type of Service and Network 3 Redirect Datagram for the Type of Service and Host
6	Alternate Host Address	
7	Unassigned	
8	Echo	
9	Router Advertisement	
10	Router Selection	
11	Time Exceeded	0 Time to Live exceeded in Transit 1 Fragment Reassembly Time Exceeded
12	Parameter Problem	0 Pointer indicates the error 1 Missing a Required Option 2 Bad Length

**Table 46:** ICMP Types and Codes (continued)

Type Number	Type Name	Code
13	Timestamp	
14	Timestamp Reply	
15	Information Request	
16	Information Reply	
17	Address Mask Request	
18	Address Mask Reply	
19	REserved (for Security)	
20 - 29	Reserved (for Robustness Experiment)	
30	Traceroute	
31	Datagram Conversion Error	
32	Mobile Host Redirect	
33	IPv6 Where-Are-You	
34	IPv6 I-Am-Here	
35	Mobile Registration	
36	Mobile Registration Reply	
37	Domain Name Request	
38	Domain Name Reply	
39	SKIP	
40	Photuris	
41 - 255	Reserved	

## ICMPv6

Internet Control Message Protocol version 6 (ICMPv6) is the new implementation of the Internet Control Message Protocol (ICMP) that is part of Internet Protocol version 6 (IPv6). The ICMPv6 protocol is defined in [RFC 4443](#).

ICMPv6 is a multipurpose protocol. It performs such things as:

- error reporting in packet processing
- diagnostic functions
- Neighbor Discovery process
- IPv6 multicast membership reporting

It also designed as a framework to use extensions for use with future implementations and changes.

Examples of extensions that have already been written for ICMPv6:

- Neighbor Discovery Protocol (NDP) - a node discovery protocol in IPv6 which replaces and enhances functions of ARP.
- Secure Neighbor Discovery Protocol (SEND) - an extension of NDP with extra security.
- Multicast Router Discovery (MRD) - allows discovery of multicast routers.

ICMPv6 messages use IPv6 packets for transportation and can include IPv6 extension headers. ICMPv6 includes some of the functionality that in IPv4 was distributed among protocols such as ICMPv4, ARP (Address Resolution Protocol), and IGMP (Internet Group Membership Protocol version 3).

ICMPv6 has simplified the communication process by eliminating obsolete messages.

ICMPv6 messages are subdivided into two classes: error messages and information messages.

Error Messages are divided into four categories:

1. Destination Unreachable
2. Time Exceeded
3. Packet Too Big
4. Parameter Problems

Information messages are divided into three groups:

1. Diagnostic messages
2. Neighbor Discovery messages
3. Messages for the management of multicast groups.

### ICMPv6 Types and Codes

ICMPv6 has a number of messages that are identified by the “Type” field. Some of these types have assigned “Code” fields as well. The table below shows the different types of ICMP Types with their associated codes if there are any.

Type codes 0 – 127 are error messages and type codes 128 – 255 are for information messages.

**Table 47:** ICMPv6 Types and Codes

Type Number	Type Name	Code
0	Reserved	0 - no route to destination 1 - communication with destination administratively prohibited 2 - beyond scope of source address 3 - address unreachable 4 - port unreachable 5 - source address failed ingress/egress policy 6 - reject route to destination 7 - Error in Source Routing Header
1	Destination Unreachable	
2	Packet Too Big	
3	Time Exceeded	0 - hop limit exceeded in transit 1 - fragment reassembly time exceeded
4	Parameter Problem	0 - erroneous header field encountered 1 - unrecognized Next Header type encountered 2 - unrecognized IPv6 option encountered
100	Private Experimentation	
101	Private Experimentation	
102 - 126	Unassigned	
127	Reserved for expansion if ICMPv6 error messages	
128	Echo Request	
129	Echo Replay	
130	Multicast Listener Query	
131	Multicast Listener Report	
132	Multicast Listener Done	
133	Router Solicitation	
134	Router Advertisement	

**Table 47:** ICMPv6 Types and Codes (continued)

Type Number	Type Name	Code
135	Neighbor Solicitation	
136	Neighbor Advertisement	
137	Redirect Message	
138	Router Renumbering	0 - Router Renumbering Command 1 - Router Renumbering Result 255 - Sequence Number Reset
139	ICMP Node Information Query	0 - The Data field contains an IPv6 address which is the Subject of this Query. 1 - The Data field contains a name which is the Subject of this Query, or is empty, as in the case of a NOOP. 2 - The Data field contains an IPv4 address which is the Subject of this Query. 140 ICMP Node Information Response 0 - A successful reply. The Reply Data field may or may not be empty. 1 - The Responder refuses to supply the answer. The Reply Data field will be empty. 2 - The Qtype of the Query is unknown to the Responder. The Reply Data field will be empty.
140	ICMP Node Information Response	0 - A successful reply. The Reply Data field may or may not be empty. 1 - The Responder refuses to supply the answer. The Reply Data field will be empty. 2 - The Qtype of the Query is unknown to the Responder. The Reply Data field will be empty.
141	Inverse Neighbor Discovery Solicitation Message	
142	Inverse Neighbor Discovery Advertisement Message	
143	Version 2 Multicast Listener Report	
144	Home Agent Address Discovery Request Message	



**Table 47:** ICMPv6 Types and Codes (continued)

Type Number	Type Name	Code
145	Home Agent Address Discovery Reply Message	
146	Mobile Prefix Solicitation	
147	Mobile Prefix Advertisement	
148	Certification Path Solicitation Message	
149	Certification Path Advertisement Message	
150	ICMP messages utilized by experimental mobility protocols such as Seamoby	
151	Multicast Router Advertisement	
152	Multicast Router Solicitation	
153	Multicast Router Termination	
154	FMIPv6 Messages	
155	RPL Control Message	
156	ILNIPv6 Locator Update Message	
157	Duplicate Address Request	
158	Duplicate Address Confirmation	
159 – 199	Unassigned	
200	Private experimentation	
201	Private experimentation	
255	Reserved for expansion of ICMPv6 informational messages	

## IP

Internet Protocol (IP) is the primary part of the Network Layer of the OSI Model that is responsible for routing traffic across network boundaries. It is the protocol that is responsible for addressing. IPv4 is probably the version that most people are familiar with and it has been around since 1974. IPv6 is its current successor and due to a shortage of available IPv4 addresses compared to the explosive increase in the number of devices that use IP addresses, IPv6 is rapidly increasing in use.

When IP is chosen as the protocol type the available option to further specify the protocol is the protocol number. This is used to narrow down which protocol within the Internet Protocol Suite and provide a more granular control.

### Protocol Number

IP is responsible for more than the address that it is most commonly associated with and there are a number of associated protocols that make up the Network Layer. While there are not 256 of them, the field that identifies them is a numeric value between 0 and 256.

In the Internet Protocol version 4 (IPv4) [RFC791] there is a field called "Protocol" to identify the next level protocol. This is an 8 bit field. In Internet Protocol version 6 (IPv6) [RFC2460], this field is called the "Next Header" field.

**Table 48:** Protocol Numbers

#	Protocol	Protocol's Full Name
0	HOPOPT	IPv6 Hop-by-Hop Option
1	ICMP	Internet Control Message Protocol
2	IGMP	Internet Group Management
3	GGP	Gateway-to-Gateway
4	IPv4	IPv4 encapsulation Protocol
5	ST	Stream
6	TCP	Transmission Control Protocol
7	CBT	CBT
8	EGP	Exterior Gateway Protocol
9	IGP	Any private interior gateway (used by Cisco for their IGRP)
10	BBN-RCC-MON	BBN RCC Monitoring
11	NVP-II	Network Voice Protocol
12	PUP	PUP
13	ARGUS	ARGUS
14	EMCON	EMCON
15	XNET	Cross Net Debugger
16	CHAOS	Chaos

**Table 48:** Protocol Numbers

#	Protocol	Protocol's Full Name
17	UDP	User Datagram Protocol
18	MUX	Multiplexing
19	DCN-MEAS	DCN Measurement Subsystems
20	HMP	Host Monitoring
21	PRM	Packet Radio Measurement
22	XNS-IDP	XEROX NS IDP
23	TRUNK-1	Trunk-1
24	TRUNK-2	Trunk-2
25	LEAF-1	Leaf-1
26	LEAF-2	Leaf-2
27	RDP	Reliable Data Protocol
28	IRTP	Internet Reliable Transaction
29	ISO-TP4	ISO Transport Protocol Class 4
30	NETBLT	Bulk Data Transfer Protocol
31	MFE-NSP	MFE Network Services Protocol
32	MERIT-INP	MERIT Internodal Protocol
33	DCCP	Datagram Congestion Control Protocol
34	3PC	Third Party Connect Protocol
35	IDPR	Inter-Domain Policy Routing Protocol
36	XTP	XTP
37	DDP	Datagram Delivery Protocol
38	IDPR-CMTP	IDPR Control Message Transport Proto
39	TP++	TP++ Transport Protocol
40	IL	IL Transport Protocol
41	IPv6	IPv6 encapsulation
42	IPv6	SDRPSource Demand Routing Protocol
43	IPv6-Route	Routing Header for IPv6
44	IPv6-Frag	Fragment Header for IPv6
45	IDRP	Inter-Domain Routing Protocol

**Table 48:** Protocol Numbers

#	Protocol	Protocol's Full Name
46	RSVP	Reservation Protocol
47	GRE	General Routing Encapsulation
48	DSR	Dynamic Source Routing Protocol
49	BNA	BNA
50	ESP	Encap Security Payload
51	AH	Authentication Header
52	I-NLSP	Integrated Net Layer Security TUBA
53	SWIPE	IP with Encryption
54	NARP	NBMA Address Resolution Protocol
55	MOBILE	IP Mobility
56	TLSP	Transport Layer Security Protocol using Kryptonet key management
57	SKIP	SKIP
58	IPv6-ICMP	ICMP for IPv6
59	IPv6-NoNxt	No Next Header for IPv6
60	IPv6-Opts	Destination Options for IPv6
61		any host internal protocol
62	CFTP	CFTP
63		any local network
64	SAT-EXPAK	SATNET and Backroom EXPAK
65	KRYPTOLAN	Kryptolan
66	RVD	MIT Remote Virtual Disk Protocol
67	IPPC	Internet Pluribus Packet Core
68		any distributed file system
69	SAT-MON	SATNET Monitoring
70	VISA	VISA Protocol
71	IPCV	Internet Packet Core Utility
72	CPNX	Computer Protocol Network Executive
73	CPHB	Computer Protocol Heart Beat

**Table 48:** Protocol Numbers

#	Protocol	Protocol's Full Name
74	WSN	Wang Span Network
75	PVP	Packet Video Protocol
76	BR-SAT-MON	Backroom SATNET Monitoring
77	SUN-ND	SUN ND PROTOCOL-Temporary
78	WB-MON	WIDEBAND Monitoring
79	WB-EXPAK	WIDEBAND EXPAK
80	ISO-IP	ISO Internet Protocol
81	VMTP	VMTP
82	SECURE-VMTP	SECURE-VMTP
83	VINES	VINES
84	TTP	TTP
84	IPTM	Protocol Internet Protocol Traffic
85	NSFNET-IGP	NSFNET-IGP
86	DGP	Dissimilar Gateway Protocol
87	TCF	TCF
88	EIGRP	EIGRP
89	OSPFIGP	OSPFIGP
90	Sprite-RPC	Sprite RPC Protocol
91	LARP	Locus Address Resolution Protocol
92	MTP	Multicast Transport Protocol
93	AX.25	AX.25 Frames
94	IPIP	IP-within-IP Encapsulation Protocol
95	MICP	Mobile Internetworking Control Pro.
96	SCC-SP	Semaphore Communications Sec. Pro.
97	ETHERIP	Ethernet-within-IP Encapsulation
98	ENCAP	Encapsulation Header
99		any private encryption scheme
100	GMTP	GMTP
101	IFMP	Ipsilon Flow Management Protocol

**Table 48:** Protocol Numbers

#	Protocol	Protocol's Full Name
102	PNNI	PNNI over IP
103	PIM	Protocol Independent Multicast
104	ARIS	ARIS
105	SCPS	SCPS
106	QNX	QNX
107	A/N	Active Networks
108	IPComp	IP Payload Compression Protocol
109	SNP	Sitara Networks Protocol
110	Compaq-Peer	Compaq Peer Protocol
111	IPX-in-IP	IPX in IP
112	VRRP	Virtual Router Redundancy Protocol
113	PGM	PGM Reliable Transport Protocol
114		any 0-hop protocol
115	L2TP	Layer Two Tunneling Protocol
116	DDX	D-II Data Exchange (DDX)
117	IATP	Interactive Agent Transfer Protocol
118	STP	Schedule Transfer Protocol
119	SRP	SpectraLink Radio Protocol
120	UTI	UTI
121	SMP	Simple Message Protocol
122	SM	SM
123	PTP	Performance Transparency Protocol
124	ISIS over IPv4	
125	FIRE	
126	CRTP	Combat Radio Transport Protocol
127	CRUDP	Combat Radio User Datagram
128	SSCOPMCE	
129	IPLT	
130	SPS	Secure Packet Shield

**Table 48:** Protocol Numbers

#	Protocol	Protocol's Full Name
131	PIPE	Private IP Encapsulation within IP
132	SCTP	Stream Control Transmission Protocol
133	FC	Fibre Channel
134	RSVP-E2E-IGNORE	
135	Mobility Header	
136	UDPLite	
137	MPLS-in-IP	
138	manet	
139	HIP	
140	Shim6	
141	WESP	
142	ROHC	
143 – 252	Unassigned	Unassigned
253		Use for experimentation and testing
254		Use for experimentation and testing
255	Reserved	

Further information can be found by researching [RFC 5237](#).

## Service Groups

Just like some of the other firewall components, services can also be bundled into groups for ease of administration.

## Example Scenario: Using FortiGate services to support Audio/Visual Conferencing



The feature, and the transmitting of data for the purpose of, Tele-conferencing or Audio/Visual Conferencing is covered by a number of standards:

- The IETF standard known as the Binary Floor Control Protocol (BFCP).
- RFC 4582, for SIP-based video devices
- The ITU standard H.239 (for H.323-based video devices)

While these standards have been set up by various authoritative bodies and can take place on different layers of the OSI model, they share common requirements that are addressed by the FortiGate firewall's ability to manage the traffic and the protocols involved. This means that the same ability that make the device RFC 4582 compliant makes it compliant with H.239 as well.

To demonstrate how services and service groups are used we show the setup of a firewall that will need to support the connectivity of a video conferencing unit. The FortiGate does not manipulate or change the content of the traffic but it does allow for the traffic to pass through the device. In this case it allow for only the needed traffic to pass through the device so as to allow the functionality of Audio Visual Conference call but not to allow other traffic through.

The theoretical location for this scenario is a hospital that hosts conferences and lectures from doctors from all over the world, sometimes from multiple locations, using video conferencing technology such as a Polycom Video Conference system. There is a special room set up with dedicated Ethernet connectivity to the Internet. A hospital has a lot of sensitive information going over its network so the setup has to be secure to prevent any chance of penetration.

The approach is fairly simple. The conference room has a dedicated port on the FortiGate (port #7) and its own LAN. We will assume that the interface has already been configured properly. Video conference traffic can come from the Internet to the Polycom in that room and traffic can get out to the Internet, but traffic going to other areas of the hospital network have to go through the FortiGate and traffic going from the Video Conference LAN is thoroughly filtered.

To give an idea of how extensive this can be, we will use an extreme case and include just about all of the services that could be commonly used in one of these setups. The protocols listed here may differ from other setups. It will depend on which features are being used and which equipment is within the network. Always check the documentation that comes with the set up before opening ports into your network.

### VIP

In this particular case there is an IP address set aside for the conferencing system so a separate VIP is not needed for every port. One Virtual IP will be created for the system and then only the approved of protocols will be allows through the firewall.

<b>Name</b>	Vid-Conf_Room216
<b>External Interface</b>	wan1
<b>External IP Address/Range</b>	256.87.212.51 – 256.87.212.51
<b>Mapped IP Address/Range</b>	192.168.7.25 – 192.168.7.25
<b>Port Forwarding</b>	not selected



## Creating an address for the subnet

In the same way that the VIP was created to identify and direct incoming traffic an address should be created to identify the addresses of computer that will be in the Conference room. This included computers on the LAN as well as the Teleconferencing equipment.

Go to *Firewall Objects* -> *Address* -> *Addresses*

*Create New*

Fill out the fields with the following information:

<b>Category</b>	Address
<b>Name</b>	Port7_subnet
<b>Type</b>	Subnet
<b>Subnet/IP Range</b>	192.168.7.0/255.255.255.0
<b>Interface</b>	port7
<b>Show in address list</b>	checked

## Configuring the services

### Services already created:

The following are standard services that have already been created by default:

<b>HTTP</b>	TCP 80
<b>SNMP</b>	TCP 161-162/UDP 161-162
<b>LDAP</b>	TCP 389
<b>HTTPS</b>	TCP 443
<b>SYSLOG</b>	UDP 514

### Existing Services to be edited:

There are a few services that have already been created for you, but they need to be expanded to accommodate the list of protocols listed for this scenario.

#### The default h323 contains:

- TCP 1503
- UDP 1719
- TCP 1720

#### We need to add:

- TCP1719

#### The default SIP contains:

- UDP 5060

**We need to add:**

- TCP 5060

**To edit an existing service:**

- H323

Go to *Firewall Objects -> Service -> Services*

Scroll down to the section: *VoIP, Messaging & Other Applications*

Select *H323*

Select *Edit*

In the Protocol section add the additional protocol:

<b>Protocol Type</b>	TCP
<b>Destination port /Low</b>	1719

Select *OK* to save

- SIP

Select *SIP*

Select *Edit*

In the Protocol section add the additional protocol:

<b>Protocol Type</b>	TCP
<b>Destination port / Low:</b>	5060

**Custom Services that need to be created:**

There are a number of possible services that may need to be added from scratch rather than editing existing ones. While it is possible to create a single custom service that contains all of the open ports needed, it make more sense to make this modular in case only a small subset of the service needs to be added to another policy.

- Polycom API

Go to *Firewall Objects -> Service -> Services*

*Create New*

Fill in the fields of the new service with the following information:

<b>Name</b>	Polycom API
<b>Service Type</b>	Firewall
<b>Category</b>	VoIP, Messaging & Other
<b>Protocol Type</b>	TCP/UDP/SCTP
<b>Protocol</b>	TCP/UDP/SCTP
<b>Protocol</b>	TCP

<b>Destination Port - Low:</b>	24
<b>Destination Port - High:</b>	<leave blank>

Select *OK*

- Polycom Endpoints

Go to *Firewall Objects -> Service -> Services*

*Create New*

Fill in the fields of the new service with the following information:

<b>Name</b>	Polycom Endpoints
<b>Service Type</b>	Firewall
<b>Category</b>	VoIP, Messaging & Other
<b>Protocol Type</b>	TCP/UDP/SCTP
<b>Protocol</b>	TCP
<b>Destination - Low:</b>	3230
<b>Destination - High:</b>	3253

Select *OK*

Other Services to add in the same way:

**Table 49:**

Name of Service	Category	Protocol & Port #
LDAP secure communications	Authentication	TCP 636
Win 2000 ILS Registration	Network Services	TCP 1002
Gatekeeper discovery	VoIP, Messaging & Other Applications	TCP 1718
Audio Call Control	VoIP, Messaging & Other Applications	TCP 1731
Polycom proprietary Global directory data	VoIP, Messaging & Other Applications	TCP 3601
Polycom People+Content	VoIP, Messaging & Other Applications	TCP 5001
HTTP Server Push	Web Access	TCP 8080

## Creating the Service Group

Go to *Firewall Objects -> Service -> Groups*

*Create New*

Build the Service group by filling in the fields with the following information

<b>Group Name</b>	A-V_Conference
<b>Type</b>	Firewall
<b>Members</b> (click in the drop down menu to add the following services)	<ul style="list-style-type: none"><li>•HTTP</li><li>•SNMP</li><li>•LDAP</li><li>•HTTPS</li><li>•SYSLOG</li><li>•Polycom API</li><li>•Polycom Endpoints</li><li>•LDAP secure communications</li><li>•Win 2000 ILS Registration</li><li>•Gatekeeper discovery</li><li>•Audio Call Control</li><li>•Polycom proprietary Global directory data</li><li>•Polycom People+Content</li><li>•HTTP Server Push</li></ul>

## Creating the IPS Security Profile

This is by no means the only way to set up this IPS filter, but it is the way that the fictional System Administrator wants it set up. Yours may be different.

Go to *Security Profiles -> Intrusion Protection -> IPS Sensors*.

Create a new sensor

<b>Name</b>	A-V_Conference-incoming
-------------	-------------------------

Select OK

In the newly created sensor, create a new IPS filter.

<b>Sensor Type</b>	Filter Based
<b>Filter Options</b>	Advanced
<b>Severity</b>	<ul style="list-style-type: none"><li>• Critical</li><li>• High</li><li>• Medium</li><li>• Low</li></ul>
<b>Target</b>	<ul style="list-style-type: none"><li>• Server</li></ul>
<b>OS</b>	<ul style="list-style-type: none"><li>• Windows</li></ul>
<b>Application</b>	<ul style="list-style-type: none"><li>• IIS</li><li>• other</li></ul>

<b>Protocol</b> Use the [Show more...] option	<ul style="list-style-type: none"> <li>• HTTP</li> <li>• LDAP</li> <li>• SIP</li> <li>• SSL</li> <li>• H323</li> </ul>
--------------------------------------------------	------------------------------------------------------------------------------------------------------------------------

<b>Packet logging</b>	<ul style="list-style-type: none"> <li>• enabled</li> </ul>
-----------------------	-------------------------------------------------------------

Based on these filters there should be somewhere in the neighborhood of 750 signatures that the FortiGate will run traffic against in the IPS engine.

## Policies

### Incoming Policy

A policy has to be made to allow the traffic to come in from the Internet to connect to the Tele-conferencing server equipment.

Go to *Policy* -> *Policy* -> *Policy*.

Create New

Fill out the fields with the following information:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	wan1
<b>Source Address</b>	all
<b>Outgoing Interface</b>	port7
<b>Destination Address</b>	Vid-Conf_Room216
<b>Schedule</b>	always
<b>Service</b>	A-V_Conference
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	<not enabled>
<b>Logging Options</b>	Logging is a good idea but how much will depend on storage capabilities.
<b>Security Profiles</b>	Turn on IPS and choose "A-V_Conference-incoming"
<b>Traffic Shaping, Web cache, WAN Optimization, Disclaimer:</b>	The use of these features will depend on your network environment and should be decided by the network architect, as the decision will largely be based on network bandwidth, usage and importance of Video conferencing compared to other traffic.

Select OK.

The policy will then need to be put in the correct position in the sequence of the policies. Because it is a rather focused policy it should be acceptable to place it near the top of the policy order sequence.

## Outgoing Policy

A policy has to be made to allow the traffic to leave from the subnet in the conference room to the Internet, not only for the traffic for the Tele-conferencing equipment but for normal traffic of users on the Internet such as web research and email. The traffic is outgoing so there is less of a need for an Intrusion Protection System filter, but check with the network architect in case there is a need for using one of the other security profiles.

Go to *Policy* -> *Policy* -> *Policy*.

Create New

Fill out the fields with the following information:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	port7
<b>Source Address</b>	Port7_subnet
<b>Outgoing Interface</b>	wan1
<b>Destination Address</b>	all
<b>Schedule</b>	always
<b>Service</b>	any
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	enabled Use Destination Interface Address
<b>Logging Options</b>	Logging is a good idea but how much will depend on storage capabilities.
<b>Security Profiles</b>	<see above>
<b>Traffic Shaping, Web cache, WAN Optimization, Disclaimer:</b>	The use of these features will depend on your network environment and should be decided by the network architect, as the decision will largely be based on network bandwidth, usage and importance of Video conferencing compared to other traffic.

Select *OK*.

The policy will then need to be put in the correct position in the sequence of the policies.

## Firewall schedules

Firewall schedules control when policies are in effect. When you add a security policy on a FortiGate unit you need to set a schedule to determine the time frame in which that the policy

will be functioning. While it is not set by default, the normal schedule would be always. This would mean that the policy that has been created is always function and always policing the traffic going through the FortiGate.

The time component of the schedule is based on a 24 hour clock notation or military time as some people would say.

There are two types of schedules: One-time schedules and recurring schedules.

**One-Time** schedules are in effect only once for the period of time specified in the schedule. This can be useful for testing to limit how long a policy will be in effect in case it is not removed, or it can be used for isolated events such as a conference where you will only need a temporary infrastructure change for a few days.

The time frame for a One-time schedule is configured by using a start time which includes, Year | Month | Day | Hour | Minute and a Stop time which includes the same variables. So while the frequency of the schedule is only once it can last anywhere from 1 minute to multiple years.

**Recurring** schedules are in effect repeatedly at specified times of specified days of the week. The Recurring schedule is based on a repeating cycle of the days of the week as opposed to every x days or days of the month. This means that you can configure the schedule to be in effect on Tuesday, Thursday, and Saturday but not every 2 days or on odd numbered days of the month.

If a recurring schedule has a stop time that is earlier than the start time, the schedule will take effect at the start time but end at the stop time on the next day. You can use this technique to create recurring schedules that run from one day to the next.

### Example

You want to schedule the use of Skype to only between noon (12:00) and 1 p.m. (13:00).

You could create a schedule that allows Skype traffic:

- Starting at Hour:12 and Minute: 00
- Stopping at Hour:13 and Minute: 00
- Set for days of the week: Sunday | Monday |Tuesday |Wednesday | Thursday | Friday | Saturday

Or you could have a schedule that blocks Skype traffic:

- Starting at Hour:13 and Minute: 00 (and goes to the next day)
- Stopping at Hour:12 and Minute: 00
- Set for days of the week: Sunday | Monday |Tuesday |Wednesday | Thursday | Friday | Saturday

Either way is effective for the task but other factors may make one method work better than another in certain situations of it could be just a preference in approach.

## Schedule Groups

You can organize multiple firewall schedules into a schedule group to simplify your security policy list. The schedule parameter in the policy configuration does not allow for the entering of multiple schedules into a single policy so if you have a combination of time frames that you want to schedule the policy for then the best approach, rather than making multiple policies is to use a schedule group.

### Example

Your Internet policy allows employees to visit Social Media sites from company computers but not during what is considered working hours. The offices are open a few hours before working

hours and the doors are not locked until a few hours after official closing so work hours are from 9 to 5 with a lunch break from Noon to 1:00 p.m.

Your approach is to block the traffic between 9 and noon and between 1:00 p.m. and 5:00 p.m. This means you will need two schedules for a single policy and the schedule group handles this for you.

Schedule groups can contain both recurring and one-time schedules. Schedule groups cannot contain other schedule groups.

## Schedule Expiration

The schedule in a security policy enables certain aspects of network traffic to occur for a specific length of time. What it does not do however, is police that time. That is, the policy is active for a given time frame, and as long as the session is open, traffic can continue to flow.

For example, in an office environment, Skype use is allowed between noon and 1pm. During that hour, any Skype traffic continues. As long as that session is open, after the 1pm end time, the Skype conversations can continue, yet new sessions will be blocked. Ideally, the Skype session should close at 1pm.

Using a CLI command you can set the schedule to terminate all sessions when the end time of the schedule is reached. Within the config firewall command enter the command:

```
set schedule-timeout enable
```

By default, this is set to `disable`.

## Security profiles

Where security policies provide the instructions to the FortiGate unit for controlling what traffic is allowed through the device, the Security profiles provide the screening that filters the content coming and going on the network. Security profiles enable you to instruct the FortiGate unit about what to look for in the traffic that you don't want, or want to monitor, as it passes through the device.

A security profile is a group of options and filters that you can apply to one or more firewall policies. Security profiles can be used by more than one security policy. You can configure sets of security profiles for the traffic types handled by a set of security policies that require identical protection levels and types, rather than repeatedly configuring those same security profile settings for each individual security policy.

For example, while traffic between trusted and untrusted networks might need strict antivirus protection, traffic between trusted internal addresses might need moderate antivirus protection. To provide the different levels of protection, you might configure two separate profiles: one for traffic between trusted networks, and one for traffic between trusted and untrusted networks.

Security profiles are available for various unwanted traffic and network threats. Each are configured separately and can be used in different groupings as needed. You configure security profiles in the Security Profiles menu and applied when creating a security policy by selecting the security profile type.

There is a separate handbook for the topic of the Security Profiles, but because the Security Profiles are applied through the Firewall policies it makes sense to have at least a basic idea of what the security profile do and how they integrate into the FortiGate's firewall policies. The following is a listing and a brief description of what the security profiles offer by way of functionality and how they can be configured into the firewall policies.



## AntiVirus

Antivirus is used as a catch all term to describe the technology for protection against the transmission of malicious computer code sometimes referred to as malware. As anyone who has listened to the media has heard that the Internet can be a dangerous place filled with malware of various flavours. Currently, the malware that is most common in the Internet, in descending order, is Trojan horses, viruses, worms, adware, back door exploits, spyware and other variations. In recent years, not only has the volume of malicious software become greater than would have been believed when it first appeared but the level of sophistication has risen as well.

The Antivirus Filter works by inspecting the traffic that is about to be transmitted through the FortiGate. To increase the efficiency of effort it only inspects the traffic being transmitted via the protocols that it has been configured to check. Before the data moves across the FortiGate firewall from one interface to another it is checked for attributes or signatures that have been known to be associated with malware. If malware is detected, it is removed.

## Web Filtering

Malicious code is not the only thing to be wary of on the Internet. There is also the actual content. While the content will not damage or steal information from your computer there is still a number of reasons that would require protection from it.

In a setting where there are children or other sensitive people using the access provided by a connected computer there is a need to make sure that images or information that is not appropriate is not inadvertently displayed to them. Even if there is supervision, in the time it takes to recognize something that is inappropriate and then properly react can expose those we wish to protect. It is more efficient to make sure that the content cannot reach the screen in the first place.

In an organizational setting, there is still the expectation that organization will do what it can to prevent inappropriate content from getting onto the computer screens and thus provoking an Human Resources incident. There is also the potential loss of productivity that can take place if people have unfiltered access to the Internet. Some organizations prefer to limit the amount of distractions available to tempt their workers away from their duties.

The Web filter works primarily by looking at the destination location request for a HTTP(S) request made by the sending computer. If the URL is on a list that you have configured to list unwanted sites, the connection will be disallowed. If the site is part of a category of sites that you have configured to deny connections to the session will also be denied. You can also configure the content filter to check for specific key strings of data on the actual web site and if any of those strings of data appear the connection will not be allowed.

## Application Control

Application Control is designed to allow you to determine what applications are operating on your network and to also filter the use of these applications as required. Application control is also for outgoing traffic to prevent the use of applications that are against an organization's policy from crossing the network gateway to other networks. An example of this would be the use of proxy servers to circumvent the restrictions put in place using the Web Filtering.

## Intrusion Protection (IPS)

Intrusion Prevention System is almost self explanatory. In the same way that there is malware out on the Internet that the network needs to be protected from there are also people out there that take a more targeted approach to malicious cyber activity. No operating system is perfect and new vulnerabilities are being discovered all of the time. An intrusion prevention system is designed to look for activity or behavior that is consistent with attacks against your network.

When attack like behavior is detected it can either be dropped or just monitored depending on the approach that you would like to take.

As new vulnerabilities are discovered they can be added to the IPS database so that the protection is current.

## Email Filtering

Spam or unsolicited bulk email is said to account for approximately 90% of the email traffic on the Internet. Sorting through it is both time consuming and frustrating. By putting an email filter on policies that handle email traffic, the amount of spam that users have to deal with can be greatly reduced.

## Data Leak Prevention (DLP)

Data Leak Prevention is used to prevent sensitive information from leaving your network. When people think of security in the cyber-world one of the most common images is that of a hacker penetrating your network and making off with your sensitive information, but the other way that you can lose sensitive data is if someone already on the inside of your network sends it out. This does not have to be an act of industrial espionage. It can just be a case of not knowing the policies of the organization or a lack of knowledge of security or laws concerning privacy.

For instance, a company may have a policy that they will not reveal anyone's Social Security number, but an employee emails a number of documents to another company that included a lengthy document that has a Social Security number buried deep within it. There is not malicious intent but if the information got out there could be repercussions.

If an organization has any information in a digital format that it cannot afford for financial or legal reasons, to leave its network, it makes sense to have Data Leak Prevention in place as an additional layer of protection.

## VoIP

Voice over IP is essentially the protocols for transmitting voice or other multimedia communications over Internet Protocol networks such as the Internet. The Security Profiles VoIP options apply the SIP Application Level Gateway (ALG) to support SIP through the FortiGate unit. The SIP ALG can also be used to protect networks from SIP-based attacks.

## ICAP

Internet Content Adaptation Protocol (ICAP) off loads HTTP traffic to another location for specialized processing. The purpose of this module when triggered is to send the incoming HTTP traffic over to a remote server to be processed thus taking some of the strain off of the resources of the FortiGate unit. The reasons for the specialized process could be anything from more sophisticated Antivirus to manipulation of the HTTP headers and URLs.

## EndPoint Control

EndPoint Control makes sure that certain standards are kept. When a computer on the Internet becomes connected to the FortiGate unit by VPN that computer is now part of the same network and there for needs to be subject to the same levels of protection, not only to protect the computer but the network. In the EndPoint Control section you can set the minimum standards for thins like AntiVirus software and VPN software.

## Proxy Option Components

Any time a security profile that requires the use of a proxy is enabled the Proxy Options field will be displayed. Certain inspections defined in security profiles require that the traffic be held in proxy while the inspection is carried out and so the Proxy Options are there to define the parameters of how the traffic will be processed and to what level the traffic will be processed. In the same way that there can be multiple security profiles of a single type there can also be a number of unique Proxy Option profiles so that as the requirements for a policy differ from one policy to the next you can also configure a different Proxy Option profile for each individual policy or you can use one profile repeatedly.

The Proxy Options refer to the handling of the following protocols:

- HTTP
- HTTPS
- FTP
- FTPS
- IMAP
- IMAPS
- POP3
- POP3S
- SMTP
- SMTPS
- DNS
- IM
- NNTP

The configuration for each of these protocols is handled separately.

It should also be noted that these configuration apply to only the Security Profiles Proxy-based processes and not the Flow-based processes.

## The use of different proxy profiles and profile options

Just like other components of the FortiGate, there is the option for different Proxy Option profiles so that you can be very granular in your control of the workings of the FortiGate. In the case of the Proxy Option profiles the thing that you will want to focus on is the matching up of the correct profile to a firewall policy that is using the appropriate protocols. If you are creating a Proxy Option profile that is designed for policies that control SMTP traffic into your network you only want to configure the settings that apply to SMTP. You do not need or want to configure the HTTP components.

### Oversized File Log

This setting is for those that would like to log the occurrence of oversized files being processed. It does not change how they are processed it only enables the FortiGate unit to log that they were either blocked or allowed through. A common practice is to allow larger files through without antivirus processing. This allows you to get an idea of how often this happens and decide on whether or not to alter the settings relating to the treatment of oversized files.

The setting of the threshold for what is considered to be an oversized file is located in the Oversized File / Email Threshold that is found in some of the protocol options for the Proxy Options.

## Invalid Certificate Log

This setting enables the logging of occurrences that have occurred in policies governed by the proxy option profile that this setting is enabled in. As stated elsewhere in this document there can be a number of reasons that a web site may have a certificate that registers as being invalid. Because some administrators decide to allow users to access these sites it makes sense to keep track of when it happens. This does not prevent the accessing of web sites with an invalid certificate. It just allows for logging when it happens.

## Port

While each of the protocols listed has a default TCP port that is commonly used, the level of granularity of control on the FortiGate firewall allows that the port used by the protocols can be individually modified in each separate Profile.

## Comfort Clients

When proxy-based antivirus scanning is enabled, the FortiGate unit buffers files as they are downloaded. Once the entire file is captured, the FortiGate unit begins scanning the file. During the buffering and scanning procedure, the user must wait. After the scan is completed, if no infection is found, the file is sent to the next step in the process flow. If the file is a large one this part of the process can take some time. In some cases enough time that some users may get impatient and cancel the download.

The comfort client feature to mitigates this potential issue by feeding a trickle of data while waiting for the scan to complete so as to let the user know that processing is taking place and that there hasn't been a failure in the transmission. This slow transfer rate continues until the antivirus scan is complete. Once the file has been successfully scanned without any indication of viruses the transfer will proceed at full speed.

If there is evidence of an infection the FortiGate unit caches the URL and drops the connection. The client does not receive any notification of what happened because the download to the client had already started. Instead, the download stops and the user is left with a partially downloaded file. If the user tries to download the same file again within a short period of time, the cached URL is matched and the download is blocked. The client receives the Infection cache message replacement message as a notification that the download has been blocked. The number of URLs in the cache is limited by the size of the cache.

Client comforting is available for HTTP and FTP traffic. If your FortiGate unit supports SSL content scanning and inspection, you can also configure client comforting for HTTPS and FTPS traffic.



Buffering the entire file allows the FortiGate unit to eliminate the danger of missing an infection due to fragmentation because the file is reassembled before examination. Client comforting can send unscanned and therefore potentially infected content to the client. You should only enable client comforting if you are prepared to accept this risk. Keeping the client comforting interval high and the amount low will reduce the amount of potentially infected data that is downloaded.

---

## Oversized File/Email Threshold

This is another feature that is related to antivirus scanning. The FortiGate unit has a finite amount of resources that can be used to buffer and scan a file. If a large file such as an ISO image or video file was to be downloaded this could not only overwhelm the memory of the FortiGate, especially if there were other large files being downloaded at the same time, but could exceed it as well. For this reason, how to treat large files needs to be addressed.

A threshold is assigned to determine what should be considered an oversized file or email. This can be set at any size from 1 MB to 50 MB. Any file or email over this threshold will not be

processed by the Antivirus Security Profiles. Once a file is determined to be oversized it must be then determined whether to allow it or to block it.

These settings are not a technical decision but a policy one that will depend on your comfort level with letting files into your network. As there often is, there is a compromise between convenience or ease of use and security. If you want to go for a high peace of mind level you can configure the firewall to block oversized files and thus no files would be coming into the network that have not been scanned. If you are looking for optimizing the memory of the FortiGate unit and making sure that everybody is getting the files they want, you can lower the threshold and allow files that are over the threshold.



It should be noted that in terms of probability that malware is more likely to be found in smaller files than in larger files. A number of administrators take this into account when they lower the default threshold so as to lessen the impact on memory if they see the FortiGate unit going into conserve mode on a regular basis.

---

### Chunked Bypass

The HTTP section allows the enabling of “Chunked Bypass”. This refers to the mechanism in version 1.1 of HTTP that allows a web server to start sending chunks of dynamically generated output in response to a request before actually knowing the actual size of the content. Where dynamically generated content is concerned this means that there is a faster initial response to HTTP requests. From a security stand point it means that the content will not be held in the proxy as an entire file before proceeding.

### Allow Fragmented Messages

The specifications of RFC 2046 allow for the breaking up of emails and sending the fragments in parallel to be rebuilt and read at the other end by the mail server. It was originally designed to increase the performance over slower connections where larger email messages were involved. It will depend on your mail configuration if this is even possible for your network but outside of Microsoft Outlook and Outlook Express, not many email clients are set up to break up messages like this. The drawback of allowing this feature is that if malware is broken up between multiple fragments of the message the risk is run that it will not be detected by some antivirus configurations because the code may not all be present at the same time to identify.

### Append Email Signature

The Append Email Signature is used when an organization would like to ensure that over and above our in this case underneath the existing personal signatures of the sender, all of the emails going out of their network have the appropriate “boilerplate”, for lack of a better term. These appended emails do not replace existing signatures. They are as the feature states, appended to the email.

Examples could include things like:

- Without prior approval the email should not be forwarded.
- Please be environmentally friendly and don't print out emails
- For questions regarding the purchasing of our products please call...

It can be anything that the organization would like as long as it is in text format. The use of this feature usually works best in an environment where there is some standardization of what goes into the personal signatures of the senders so that there is no duplication or contradiction of information in the signatures.

## SSL/SSH Inspection

While the profile configuration for this is not found in the Security Profiles section but in the Policy Section, it is set in the policy along with the security profiles. This sort of analysis is some times referred to as deep scanning.

Deep Inspection works along the following lines. If your FortiGate unit has the correct chipset it will be able to scan HTTPS traffic in the same way that HTTP traffic can be scanned. The FortiGate firewall will essentially receive the traffic on behalf of the client and open up the encrypted traffic. Once it is finished it re-encrypts the traffic and sends it on to its intended recipient. It is very similar to a man-in-the-middle attack. By enabling this feature, it allows the FortiGate firewall to filter on traffic that is using the HTTPS protocol.

Sometimes the regular web filter can be circumvented by using https:// instead of http:// in the URL and this would prevent that circumvention. However, because when the encrypted traffic is decrypted it has to be re-encrypted with the FortiGate's certificate rather than the original certificate it can cause errors because the name on the certificate does not match the name on the web site.

At one point deep inspection was something that was either turned on or off. Now individual deep inspection profiles can be created depending on the requirements of the policy. Depending on the Inspection Profile, you can:

- Configure which CA certificate will be used to decrypt the SSL encrypted traffic.
- Configure which SSL protocols will be inspected.
- Configure which ports will be associated with which SSL protocols for the purpose of inspection.
- Configure whether or not to allow invalid SSL certificates.
- Configure whether or not SSH traffic will be inspected.

### Allow Invalid SSL Certificate

This setting was something that used to be part of the *Proxy Options*, but now that SSL inspection has its own configuration setting it is configured with those. It might seem like a straight forward decision that the allowing of invalid SSL certificates must be bad and therefore should not be allowed, but there can be some reasons that should be considered. The issues at hand are the reasons to use a SSL certificate and the reasons that a certificate will be considered invalid.

At a purely technical level, a properly formed certificate will encrypt the data so that it can only be read by the intended parties and not be read by anyone sniffing traffic on the network. For this reason, people will often use self-signed certificates. These self signed certificates are free and will encrypt the data just as well as those purchased from any of the big vendors of certificates, but if they are not listed as an approved Certificate Authority (CA) the certificates will be considered invalid.

On the other hand, one of the services the vendors provide is verification of identity of those that purchase their certificates. This means that if you see a valid certificate from a site that identified itself as being from "valid-company.com" that you can be reasonably sure that the site does belong to that company and not a false site masquerading as being part of that company.

### Creating a new SSL/SSH Inspection profile

1. Go to *Policy > Policy > SSL/SSH Inspection*.
2. In the Name field give the profile a name.
3. In the Comments field you can optionally include an brief description of the profile.

### SSL Inspection Options

4. Use the drop down menu for the CA Certificate field to choose the SSL Certificate to be used by Policies that are associated with this profile.
5. Choose between
  - a. Inspecting all SSL protocol ports -- enable the check box
  - b. Enabling only specific SSL protocol ports -- enable which of the following protocol you intend to inspect:
    - HTTPS
    - SMTPS
    - POP3S
    - IMAPS
    - FTPS

You can optionally edit the TCP/IP port numbers that you expect the traffic to be travelling over.

### SSH Inspection Options

6. Choose whether or not to enable SSH Deep Scan. If yes, enable the check box.  
Once the check box is enabled a window will appear to be used in the configuring of:
  - SSH - across any port or only the specified one.
  - Exec - *Block, Log* or neither. Select using check boxes.
  - Port-Forward - *Block, Log* or neither. Select using check boxes.
  - SSH-Shell - *Block, Log* or neither. Select using check boxes.
  - X11-Filter - *Block, Log* or neither. Select using check boxes.

### Common Options

7. Choose whether to *Allow Invalid SSL Certificates*. If yes, enable the check box.
8. Select *OK*.



The *Enable SSH Deep Scan* feature is enabled by default when creating a new SSL/SSH Inspection profile. There are situations where this feature can cause issues so be sure that you would like it enabled before applying it.

---

# Security policies

One of the foundations upon which a firewall works is the use of policies. These are what bring the other firewall objects and components together into an elegant mechanism for the governing of the traffic going through the network.

This Chapter includes information on the following topics:

- [Firewall policies](#)
- [Identity Based Policies](#)
- [Device Identity Policies](#)
- [VPN Policies](#)
- [Interface Policies](#)
- [One-Arm IDS](#)
- [Local-In Policies](#)
- [Security Policy 0](#)
- [Deny Policies](#)
- [Accept Policies](#)
- [IPv6 Policies](#)
- [Fixed Port](#)
- [Endpoint Security](#)
- [Traffic Logging](#)
- [Quality of Service](#)
- [Policy Monitor](#)

## Firewall policies

The firewall policy is the axis around which most of the other features of the FortiGate firewall revolve. A large portion of the settings in the firewall at some point will end up relating to or being associated with the firewall policies and the traffic that they govern. Any traffic going through a FortiGate unit has to be associated with a policy. These policies are essentially discrete compartmentalized sets of instructions that control the traffic flow going through the firewall. These instructions control where the traffic goes, how it's processed, if it's processed and even whether or not it's allowed to pass through the FortiGate.

When the firewall receives a connection packet, it analyzes the packet's source address, destination address, and service (by port number). It also registers the incoming interface, the outgoing interface it will need to use and the time of day. Using this information the FortiGate firewall attempts to locate a security policy that matches the packet. If it finds a policy that matches the parameters it then looks at the action for that policy. If it is ACCEPT the traffic is allowed to proceed to the next step. If the Action is DENY or a match cannot be found the traffic is not allowed to proceed.

The 2 basic actions at the initial connection are either ACCEPT or DENY:

- If the Action is ACCEPT, the policy action permits communication sessions. There may be other packet processing instructions, such as requiring authentication to use the policy.



While you may not see it in the configuration there is the implied subset of the ACCEPT Action that include VPN policies, whether they be an IPSec VPN or SSL.

- If the Action is DENY, the policy action blocks communication sessions, and you can optionally log the denied traffic. If no security policy matches the traffic, the packets are dropped. A DENY security policy is needed when it is required to log the denied traffic, also called “violation traffic”.

The policy may contain a number of instructions for the FortiGate firewall in addition to the ACCEPT or DENY actions, some of which are optional. Instructions on how to process the traffic can also include such things as:

- Logging Traffic
- Authentication
- Network Address Translation or Port Address Translation
- Use Virtual IPs or IP Pools
- Caching
- Whether to use address or Identity based rules
- Whether to treat as regular traffic or VPN traffic
- What certificates to use
- Security profiles to apply
- Proxy Options
- Traffic Shaping

As mentioned before, for traffic to flow through the FortiGate firewall there must be a policy that matches its parameters:

- Source Interface
- Destination Interface
- Source Address
- Destination Address
- Service or TCP/IP suite port number
- Schedule and time of the session’s initiation

Without all five of these things matching the traffic will be declined. Each traffic flow requires a policy and the direction is important as well. Just because packets can go from point A to point B on port X does not mean that the traffic can flow from point B to point A on port X. A policy must be configured for each direction.

When designing a policy there is often reference to the traffic flow, but most communication is a two way connection so trying to determine the direction of the flow can be somewhat confusing. If traffic is HTTP web traffic the user sends a request to the web site, but most of the traffic flow will be coming from the web site to the user. Is the traffic flow considered to be from the user to the web site, the web site to the user or in both directions? For the purposes of determining the direction for a policy the important factor is the direction of the initiating communication. The user is sending a request to the web site so this is the initial communication and the web site is just responding to it so the traffic will be from the users network to the Internet.

A case where either side can initiate the communication like between two internal interfaces on the FortiGate unit would be a more likely situation to require a policy for each direction.

## What is not expressly allowed is denied

One of the fundamental ideas that can be found in just about any firewall is the rule than anything that is not expressly allowed is by default denied. This is the foundation for any

strategy of protecting your network. Right out of the box, once you have your FortiGate device connected into your network and hooked up with your ISP your network is protected. Nothing is getting out or in so it is not very convenient, but you don't have to worry that between the time you hooked it up and the point that you got all of the policies in place that someone could have gotten in and done something to your resources. The reason that this needs to be kept in mind when designing policies is because you cannot assume that any traffic will be allowed just because it makes sense to do so. If you want any kind of traffic to make it past the FortiGate firewall you need to create a policy that will allow that traffic. To maintain the protection of the network should also make sure that the any policy you create allows only the traffic you intend to go only to where you specifically want it to go and when you want it to go there.

### Example

You have a web server on your network that is meant to provide a collaborative work environment web site for your employees and a partner company for a project over the course of the next 3 months.

It is theoretically possible to allow connections into your network to any device on that network for any service and at any time. The problem with this is that we might not want just anybody looking at those resources. Sadly, no matter how much it is wished otherwise, not everybody on the Internet can be trusted. Which means we now have to be very specific in our instructions as to what traffic to allow into the network. Each step that we take towards being more specific as to what we allow means that there is that much more that is not allowed and the level of protection of a resources is directly proportional to the amount of traffic that is not allowed. If somebody can't get at it they can't damage or steal it.

Limiting where the traffic is allowed to go to means that other computers on your network besides the web-server are protected.

- Limiting where the traffic is allowed to come from means that, if feasible, you can limit the systems that can access the web server to just employees or the partner company computers.
- Limiting the services to just web traffic means that a malicious person, even if they were connection from a computer at the partner organization could only use the features of web traffic to do anything malicious.
- Limiting the policy to the time span of the project would mean that even if the IT department forgot to remove the policy after the end of the project than no computer from the other company could be used to do anything malicious through the policy that allowed the traffic.

This is just a very basic example but it shows the underlying principles of how the idea that anything not expressly allowed is by default denied can be used to effectively protect your network.

## Policy order

Another important factor in how firewall policies work is the concept of precedence of order or if you prefer a more recognizable term, "first come, first served".

It is highly likely that even after only a relatively small number of policies have been created that there will be some that overlap or are subsets of the parameters that the policies use to determine which policy should be matched against the incoming traffic. When this happens there has to be a method to determine which policy should be applied to the packet. The method which is used by most firewalls it based on the order of the sequence of the policies.

If all of the policies were placed in a sequential list the process to match up the packet would start at the top of the list and work its way down. It would compare information about the packet, specifically these points of information:

1. The interface the packet connected to the FortiGate firewall
2. The source address of the packet

3. The destination address of the packet
4. The interface the packet would need to use to get to the destination address based on the routing table
5. The port the packet is destined for
6. The time that the packet connected to the FortiGate

As soon as the a policy is reached that matched all six (6) of these parameters, the instructions of that policy are applied and the search for any other matching policies is stopped. All subsequent policies are disregarded. Only 1 policy is applied to the packet.

If there is no matching policy among the policies that have been configured for traffic the packet finally drops down to what is always the last policy. It is an implicit policy. One of a few that are referred to by the term “policy0”. This policy denies everything.

The implicit policy is made up of the following settings:

- Incoming Interface: any
- Source Address: any
- Outgoing Interface: any
- Destination Address: any
- Action: DENY

The only setting that is editable in the implicit policy is the logging of violation traffic.

A logical best practice that comes from the knowledge of how this process works is to make sure that the more specific or specialized a policy is, the closer to the beginning of the sequence it should be. The more general a policy is the higher the likelihood that it could include in its range of parameters a more specifically targeted policy. The more specific a policy is, the higher the probability that there is a requirement for treating that traffic in a specific way.

### Example

For security reasons there is no FTP traffic allowed out of a specific subnet so there is a policy that states that any traffic coming from that subnet is denied if the service is FTP, so the following policy was created:

Policy #1

<b>Source Interface</b>	Internal1
<b>Source Address</b>	192.168.1.0/24
<b>Outgoing Interface</b>	WAN1
<b>Destination Address</b>	0.0.0.0/0.0.0.0
<b>Service</b>	FTP
<b>Schedule</b>	always
<b>Action</b>	deny

Now as these things usually go it turns out that there has to be an exception to the rule. There is one very secure computer on the subnet that is allowed to use FTP and once the content has been checked it can then be distributed to the other computer on the subnet. So a second firewall policy is created.

## Policy #2

<b>Source Interface</b>	Internal1
<b>Source Address</b>	192.168.1.38/32
<b>Outgoing Interface</b>	WAN1
<b>Destination Address</b>	0.0.0.0/0.0.0.0
<b>Service</b>	FTP
<b>Schedule</b>	always
<b>Action</b>	Allow

By default, a policy that has just been created will be placed last in the sequence so that it is less likely to interfere with existing policies before it can be moved to its intended position. If you look at Policy #2 you will notice that it is essentially the same as Policy #1 except for the Source Address and the Action. You will also notice that the Source Address of the Policy #2 is a subset of the Source address in policy #1. This means that if nothing further is done, Policy #2 will never see any traffic because the traffic will always be matched by Policy #1 and processed before it has a chance to reach the second policy in the sequence. For both policies to work as intended Policy #2 needs to be moved to before Policy #1 in the sequence.

### Exception to policy order (VIPs)

There is a relevant exception to the normal policy order. Policies with VIPs don't appear to behave the same way. Traffic that is handled by VIPs is processed through the associated policy before the traffic is checked against other policies in the usual top down order.

This only appears inconsistent the policy order rule because when handling the traffic, firewall policies are not the first thing checked. VIP translations are checked first, and if there is more than one VIP that the traffic fits, it is handled in the top down order that is followed by policies. If the traffic is not claimed by a policy in the VIP translation phase, it is checked against the routing rules. If it passes the routing checks, the traffic is allowed to be controlled by the policies.

This processing of traffic targeted a VIPs only applies if there is a policy that included the VIP and the traffic matches all of the criteria checks. There is no need to worry about creating VIPs that are not controlled by a policy.

There are security implications associated with this behavior. Administrators could assume that a policy will process traffic before it drops down to a policy with a VIP in it. This can allow traffic to pass through the firewall into a part of the network that it was not intended for, if it was to be allowed in at all. The way to prevent traffic being incorrectly allowed through a policy containing a VIP is to have that policy be more restrictive or to have a separate policy containing the same VIP deny the traffic earlier in the sequence.

As proof of the behavior, look at the following traffic analysis of a packet sent to a VIP. You will see that the packet is translated even before it is allowed to pass through the firewall by a policy.

```
2015-06-10 06:33:21 id=20085 trace_id=1 func=print_pkt_detail line=4373
 msg="vd-root received a packet(proto=6,
```

```
24.114.222.34:51434->24.212.230.77:3389) from wan1. flag [S], seq
1579917634, ack 0, win 8192"
2015-06-10 06:33:21 id=20085 trace_id=1 func=init_ip_session_common
line=4522 msg="allocate a new session-01480894"
2015-06-10 06:33:21 id=20085 trace_id=1 func=fw_pre_route_handler
line=174 msg="VIP-10.10.66.2:3389, outdev-wan1"
2015-06-10 06:33:21 id=20085 trace_id=1 func=__ip_session_run_tuple
line=2534 msg="DNAT 24.212.230.77:3389->10.10.66.2:3389"
2015-06-10 06:33:21 id=20085 trace_id=1 func=vf_ip4_route_input
line=1596 msg="find a route: flags=00000000 gw-10.10.66.2 via
internal5"
2015-06-10 06:33:21 id=20085 trace_id=1 func=fw_forward_handler
line=670 msg="Allowed by Policy-15:"
```

## Viewing Firewall Policies

When you first go into the Policy window, found by going to Policy > Policy > Policy, you will see a table with a menu bar across the top. The menu bar will have the following items:

At the top left

- Create New (with a “+” sign on the left and a downward pointing triangle on the right)
- Clone
- Delete
- Column Settings
- Filter Settings

At the top right

- Section View
- Global View

The items at the top right with their radio buttons represent the 2 potential views that the policies can be displayed in.

The Global View shows all of the policies in the order of their sequence. With the default settings you will be able to see the sequence number in a column close to the left side of the table.

The Section view is similar to the Global View except that as the name implies it is divided into sections. By default the sections are based on the paths between the interfaces. These can be referred to as “interface pairings”. For instance, all of the policies referencing traffic from WAN1 to DMZ will be in one section. The policies referencing traffic from DMZ to WAN1 will be in another section.

The sections are collapsible so that you only need to look at the sections with policies you are interested in. It is possible to add customized subsections within the default sections of interface pairings. This would be useful in a situation where you have a lot of policies and would like to further compartmentalize them by common attributes so that things are easier to find.

The default column headings are:

- [Check box icon]
- Seq.#
- Source
- Destination
- Authentication
- Schedule
- Service
- Action
- Log

The columns that are shown are configurable. All but the first 2 can be removed or their position changed. There are also a number of other columns that display information about the policies that can be added. One of the more useful ones that can be added is the ID column. The reason for adding this one is that policies are referenced by their ID number for simplicity and ease of administration. If you are looking in the CLI you will see that the only designation for a policy is its number and if you wish to change the order of a policy you will be asked to move it before or after another policy by referencing its number.

## How “Any” policy can remove the Section View

The FortiGate unit will automatically change the view on the policy list page to Global View whenever a policy containing “any” in the Source interface/zone or Destination interface/zone is created. If the Section View is greyed out it is likely that one or more of the policies has “any” as a Source or Destination interface.

With the use of the “any” the policy should go into multiple sections because it could effectively be any of a number of interface pairings. As mentioned, policies are sectioned by using the interface pairings (for example, port1 -> port2) and each section has its own specific policy order. The order in which a policy is checked for matching criteria to a packet’s information is based solely on the position of the policy within its section or within the entire list of policies as a whole but if the policy is in multiple sections at the same time there is no mechanism for placing the policy in a proper order within all of those sections at the same time because it is a manual process and there is no parameter to compare the precedence of one section or policy over the other. Thus a conflict is created. In order to resolve the conflict the FortiGate firewall removes that aspect of the sections so that there is no need to compare and find precedence between the sections and it therefore has only the Global View to work with.

## Security policy configuration extensions

When first creating the policy the configuration form will ask for a choice between the policy types of Firewall or VPN, Firewall being the default. Choosing whether or not to leave the selection as Firewall is straight forward. If the policy is not a policy based VPN policy then it is a Firewall policy type.

There are essentially 2 types of VPN connections, Interface Based and Policy Based. In an Interface Based VPN tunnel a logical interface is created that can be seen as an interface by the policies in the same way that any of the physical interfaces can be seen. Therefore to govern the traffic a regular policy will work. The policy based VPN tunnels work slightly different and therefore need a slightly different policy configuration. For a more detail explanation of the difference between the types of VPN tunnels refer to the VPN documentation found in the VPN handbooks or in the VPN section of the Complete Administration Guide.

Once either the Firewall or the VPN type has been chosen there is then a choice between one of subtypes for each of the Policy types. For the Firewall type of policy the subtypes are:

- Address
- User Identity
- Device Identity

The Address subtype refers to policies where access through the FortiGate firewall is dependant on the source location of the addresses of the devices involved in the traffic matched to the policy.

The User Identity subtype refers to polices where access through the FortiGate firewall is dependant on the users credentials or Identity.

The Device Identity subtype refers to policy where access through the FortiGate firewall is dependant on the specific device being used based on the MAC address of the device or belonging to a group of devices that are based on device types or belonging to custom made groups.

For the VPN type the subtypes are:

- IPSec
- SSL-VPN

As expected the two subtypes are the two different types of VPN tunnels that the FortiGate firewall supports in a policy based configuration.

## Identity Based Policies

Identity-based security policies, also known as authentication policies, match traffic that requires a supported authentication protocol to trigger the firewall authentication challenge and successfully authenticate network users. Network users authentication can occur using HTTP, HTTPS, FTP, and Telnet protocols as well as through automatic login using NTLM and FSSO, to bypass user intervention.

Identity-based security policies are usually configured for IPSec or SSL VPN traffic since this type of traffic usually requires authentication from network users.

When configuring identity-based policies, you can use schedules to limit network users authentication sessions. For example, example.com has a schedule policy to use P2P applications between noon and 1:00 pm, and a user authentication timeout of 30

minutes. When a user logs in at 12:15 pm, their authentication time logs them off at 12:45 (30 minutes later). You can configure this type of authentication by using the `scheduletimeout` field in the `config firewall policy` command in the CLI.

### Identity-based policy positioning

With identity-based security policies, positioning is extremely important. For a typical security policy, the FortiGate unit matches the source, destination and service of the policy. If matched, it acts on that policy. If not, the FortiGate unit moves to the next policy.

With identity-based policies, once the FortiGate unit matches the source and destination addresses, it processes the identity sub-rules for the user groups and services. That is, it acts on the authentication and completes the remainder of that policy and goes no further in the policy list.

The way identity based policies work is that once `src/dest` are matched, it will process the identity based sub-rules (for lack of a better term) around the user groups and services. It will

never process the rest of your rule base. For this reason, unique security policies should be placed before an identity-based policy.

For example, consider the following policies:

DNS traffic goes through successfully as does any HTTP traffic after being authenticated. However, if there was FTP traffic, it would not get through. As the FortiGate unit processes FTP traffic, it skips rule one since it's matching the source, destination and service. When it moves to rule two it matches the source and destination, it determines there is a match and, sees there are also processes the group/service rules, which requires authentication and acts on those rules. Once satisfied, the FortiGate unit will never go to rule three.

In this situation, where you would want FTP traffic to traverse the FortiGate unit, create a security policy specific to the services you require and place it above the authentication policy.

## Identity-based sub-policies

When adding authentication to a security policy, you can add multiple authentication rules, or sub-policies. Within these policies you can include additional security profiles, traffic shaping and so on, to take affect on the selected services.

These sub-policies work on the same principle as normal security policies, that is, top down until the criteria has been met. As such, if there is no matching policy within the list, the packet can still be dropped even after authentication is successful.

## Identity policies an unauthenticated users

One of the previous drawbacks with User Identity policies is that if traffic from an unauthenticated user enters the policy it will be denied by default because it doesn't match up with any of the user group and therefore falls to policy 0 which denies access to any traffic that reaches it. Allowances have been made so that if you are using User Identity based policies you are not forced to authenticate all users and create a subpolicy for all of the users.

When configuring User Identity policies you can select the option to *Skip this policy for unauthenticated user*. This policy will only apply to user traffic where the user has already authenticated with the FortiGate unit. As the name of the option implies, this policy will not apply to unauthenticated users and any traffic from unauthenticated uses that makes it through the sequence to this policy will continue on the next policy.

The command line syntax for using this feature is:

```
config firewall policy
 edit <id>
 set identity-based enable
 set fall-through-unauthenticated enable
 next
end
```

Because of this option, User Identity policies can be placed much higher in the sequence than they once were. Now that the policy will no longer interfere with unauthenticated traffic it can be placed so that non user specific policies will not act upon the traffic before it reaches its intended policy.

## Device Identity Policies

Device identity policies are designed to assist in accommodating the BYOD trends.



These policies will share many of the same characteristics and field values as the other firewall policies. The point at which a Device Policy defers from other policies on the FortiGate firewall is that when creating one of these policies the criteria that the authentication will be based on is the MAC address of the device making the connection. With the MAC addresses being unique to a specific networkable device there is a great deal of control that can be exercised with these policies.

In cooperation with the MAC scanning capabilities of the FortiGate it is relatively simple to create a profile that will allow access to the wireless network to personal devices such as smart phones, tablets and personal laptops that are brought into work by employees or even contractors and other guests.

The process of authentication is similar to the processes taking place in the Identity based policies. The FortiGate firewall checks the incoming traffic for parameters such as interfaces, addresses, times and services and then matches them with the values associated with the traffic that is associated with the MAC addresses listed in the policy. If everything matches up correctly traffic is allowed to proceed.

The device addresses can be listed by individual MAC addresses, predefined groups based on device type or custom made groupings of the MAC addresses.

## VPN Policies

At one point, if you wanted to have secure digital communications between 2 points a private network would be created. This network would only allow the people that were intended to get the communications on it. This is very straightforward if the 2 points are in the same room or even in the same building. It can all be done physically. If you are supposed to be on the secure network

VPNs are an answer to one of today's biggest concerns, how to make digital communications secure between to points that must communicate over the Internet which anybody can have access to

## IPSec Policies

IPSec policies allow IPSec VPN traffic access to the internal network from a remote location. These policies include authentication information that authenticates users and user group or groups. These policies specify the following:

- the FortiGate firewall interface that provides the physical connection to the remote VPN gateway, usually an interface connected to the Internet
- the FortiGate firewall interface that connects to the private network
- IP addresses associated with data that has to be encrypted and decrypted
- optional: a schedule that restricts when the VPN can operate, and services (or types of data) that can be sent.

For a route-based (interface mode) VPN, you do not configure an IPSec security policy. Instead, you configure two regular ACCEPT security policies, one for each direction of communication, with the IPSec virtual interface as the source or destination interface, as appropriate.

## SSL VPN Policies

SSL VPN security policies are created for permitting SSL VPN clients, web-mode or tunnel-mode, access to the protected network behind the FortiGate unit. These security policies also contain authentication information that will authenticate the users and user group or groups.

## Interface Policies

Interface policies are implemented before the “security” policies and are only flow based.

This feature allows you to attach a set of IPS policies with the interface instead of the forwarding path, so packets can be delivered to IPS before entering firewall. This feature is used for following IPS deployments:

- One-Arm: by defining interface policies with IPS and DoS anomaly checks and enabling sniff-mode on the interface, the interface can be used for one-arm IDS;
- IPv6 IPS: IPS inspection can be enabled through interface IPv6 policy. Only IPS signature scan is supported in FortiOS 4.0. IPv6 DoS protection is not supported;
- Scan traffics that destined to FortiGate;
- Scan and log traffics that are silently dropped or flooded by Firewall or Multicast traffic.

IPS sensors can be assigned to an interface policy. Both incoming and outgoing packets are inspected by IPS sensor (signature).

Here is an example of an interface policy,

```
config firewall interface-policy
 edit 1
 set status enable
 set interface "port14"
 set srcaddr "all"
 set dstaddr "all"
 set service "ALL"
 set application-list-status disable
 set ips-sensor-status enable
 set ips-sensor "default"
 set av-profile-status disable
 set webfilter-profile-status disable
 set spamfilter-profile-status disable
 set dlp-sensor-status disable
 set label "Port 14 Interface Policy"
 next
end
```

## DoS Protection

Denial of Service (DoS) policies are primarily used to apply DoS anomaly checks to network traffic based on the FortiGate interface it is entering as well as the source and destination addresses. DoS checks are a traffic anomaly detection feature to identify network traffic that does not fit known or common traffic patterns and behavior. A common example of anomalous traffic is the denial of service attack. A denial of service occurs when an attacking system starts an abnormally large number of sessions with a target system. The large number of sessions slows down or disables the target system, so that legitimate users can no longer use it.

DoS policies are similar to firewall policies except that instead of defining the way traffic is allowed to flow, they keep track of certain traffic patterns and attributes and will stop traffic displaying those attributes. Further, DoS policies affect only incoming traffic on a single interface. You can further limit a DoS policy by source address, destination address, and service.

DoS configurations have been changed a couple of times in the past. In FortiOS 4.0, DoS protection is moved to the interface policy, so when it is enabled, it is the first thing checked

when a packet enters FortiGate. Because of this early detection, DoS policies are a very efficient defence that uses few resources. Denial of service attacks, for example, are detected and its packets dropped before requiring security policy look-ups, antivirus scans, and other protective but resource-intensive operations.

A DoS policy examines network traffic arriving at an interface for anomalous patterns usually indicating an attack. This does not mean that all anomalies experience by the firewall are the result of an intentional attack.

Because an improperly configured DoS anomaly check can interfere with network traffic, no DoS checks are preconfigured on a factory default FortiGate unit. You must create your own before they will take effect. Thresholds for newly created sensors are preset with recommended values that you can adjust to meet the needs of your network.

To create a Denial of Service policy determine if it needs to be an IPv4 or IPv6 policy, then go to:

*Policy > Policy > DoS Policy for IPv4.*

*Policy > Policy > IPv6 DoS Policy for IPv6.*



It is important to know normal and expected network traffic before changing the default anomaly thresholds. Setting the thresholds too low could cause false positives, and setting the thresholds too high could allow otherwise avoidable attacks.

---

## Settings used in configuring DoS

### Incoming Interface

The interface to which this security policy applies. It will be the that the traffic is coming into the firewall on.

### Source Address

This will be the address that the traffic is coming from and must be a address listed in the Address section of the Firewall Objects. This can include the predefined “all” address which covers any address coming in on any interface. Multiple addresses or address groups can be chosen

### Destination Address

This will be the address that the traffic is addressed to. In this case it must be an address that is associated with the firewall itself. For instance it could be one of the interface address of the firewall, a secondary IP address or the interface address assigned to a Virtual IP address. Just like with the Source Address this address must be already configured before being used in the DoS policy. Multiple addresses, virtual IPs or virtual IP groups can be chosen.

### Service

While the Service field allows for the use of the ALL service some administrators prefer to optimize the resources of the firewall and only check on the services that will be answered on an interface. Multiple services or service groups can be chosen.

### Anomalies

The anomalies can not be configured by the user. They are predefined sensors set up for specific patterns of anomalous traffic

The anomalies that have been predefined for use in the DoS Policies are:

<b>Anomaly Name</b>	<b>Description</b>	<b>Recommended Threshold</b>
<b>tcp_syn_flood</b>	If the SYN packet rate of new TCP connections, including retransmission, to one destination IP address exceeds the configured threshold value, the action is executed.	2000 packets per second.
<b>tcp_port_scan</b>	If the SYN packet rate of new TCP connections, including retransmission, from one source IP address exceeds the configured threshold value, the action is executed.	1000 packets per second.
<b>tcp_src_session</b>	If the number of concurrent TCP connections from one source IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.
<b>tcp_dst_session</b>	If the number of concurrent TCP connections to one destination IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.
<b>udp_flood</b>	If the UDP traffic to one destination IP address exceeds the configured threshold value, the action is executed.	2000 packets per second.
<b>udp_scan</b>	If the number of UDP sessions originating from one source IP address exceeds the configured threshold value, the action is executed.	2000 packets per second.
<b>udp_src_session</b>	If the number of concurrent UDP connections from one source IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.
<b>udp_dst_session</b>	If the number of concurrent UDP connections to one destination IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.
<b>icmp_flood</b>	If the number of ICMP packets sent to one destination IP address exceeds the configured threshold value, the action is executed.	250 packets per second.
<b>icmp_sweep</b>	If the number of ICMP packets originating from one source IP address exceeds the configured threshold value, the action is executed.	100 packets per second.
<b>icmp_src_session</b>	If the number of concurrent ICMP connections from one source IP address exceeds the configured threshold value, the action is executed.	300 concurrent sessions
<b>icmp_dst_session</b>	If the number of concurrent ICMP connections to one destination IP address exceeds the configured threshold value, the action is executed.	3000 concurrent sessions
<b>ip_src_session</b>	If the number of concurrent IP connections from one source IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.

<b>Anomaly Name</b>	<b>Description</b>	<b>Recommended Threshold</b>
<b>ip_dst_session</b>	If the number of concurrent IP connections to one destination IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.
<b>sctp_flood</b>	If the number of Sctp packets sent to one destination IP address exceeds the configured threshold value, the action is executed.	2000 packets per second
<b>sctp_scan</b>	If the number of Sctp sessions originating from one source IP address exceeds the configured threshold value, the action is executed.	1000 packets per second
<b>sctp_src_session</b>	If the number of concurrent Sctp connections from one source IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions
<b>sctp_dst_session</b>	If the number of concurrent Sctp connections to one destination IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions

### Status

The status field is enabled to enable the sensor for the associated anomaly. In terms of actions performed there is no difference between disabling a sensor and having the action as “Pass” but by disabling sensors that are not being used for blocking or logging you can save some resources of the firewall that can be better used elsewhere.

### Logging

Regardless of whether the traffic is blocked or passed through the anomalous traffic will be logged.

### Pass

Allows the anomalous traffic to pass through unimpeded.

### Block

For Thresholds based on the number of concurrent sessions blocking the anomaly will not allow more than the number of concurrent sessions set as the threshold.

For rate based thresholds where the threshold is measured in packets per second, the Action setting “Block” prevents the overwhelming of the firewall by anomalous traffic in one of 2 ways. Setting which of those 2 ways will be issued is determined in the CLI.

- continuous - blocks any packets that match the anomaly criteria once the threshold has been reached
- periodical - allows matching anomalous traffic up to the rate set by the threshold.

To set the type of block action for the rate based anomaly sensors:

```
config ips global
 set anomaly-mode continuous
 set anomaly-mode periodical
end
```

### Threshold

The threshold can be either in terms of concurrent session or in packets per second depending on which sensor is being referred to.

## One-Arm IDS

Interface-based policy only defines what and how IPS functions are applied to the packets transmitted by the interface. It works no matter if the port is used in a forwarding path or used as an One-Arm device.

To enable One-Arm IDS, the user should first enable sniff-mode on the interface,

```
config system interface
 edit port2
 set ips-sniffer-mode enable
 next
end
```

Once sniff-mode is turned on, both incoming and outgoing packets will be dropped after IPS inspections. The port can be connected to a hub or a switch's SPAN port. Any packet picked up by the interface will still follow the interface policy so different IPS and DoS anomaly checks can be applied.

## IPv6 IPS

IPv6 IPS signature scan can be enabled by interface policy. The user can create a normal IPS sensor and assign it to the IPv6 interface policy.

```
config firewall interface-policy6
 edit 1
 set interface "port1"
 set srcaddr6 "all"
 set dstaddr6 "all"
 set service6 "ANY"
 set ips-sensor-status enable
 set ips-sensor "all_default"
 next
end
```

## Traffic Destined to the FortiGate unit

IPS enabled in firewall policies can only inspect the traffic pass through FortiGate unit, not the traffic destined to FortiGate unit. Enabling IPS in interface-policy allows IPS to pick up any packet on the interface so it is able to inspect attacks targeting FGT.

## Dropped, Flooded, Broadcast, Multicast and L2 packets

In many evaluation or certification tests, FortiGate firewall is often required to log any packets dropped by the firewall. In most of cases, these packets are of invalid headers so firewall just drops them silently. It is natural to forward all these packets to IPS first so FortiGate firewall is able to generate logs for invalid packets.

Flooded, broadcast and multicast traffics do not reach any of services in the forwarding path. They can be inspected by the interface policy as long as they match the addresses defined.

Potentially, L2 packets can also be sent to IPS for inspection through interface-policy, but it is not enabled.

## GUI and CLI

Now in FortiGate, there are two places that IPS can be enabled, in a firewall policy and in an interface policy. In the firewall policy implementation, IPS sensor can be configured in both CLI and GUI. When adding an IPS sensor to an interface policy it must be done through the CLI. There is no GUI input window for the “Interface Policy”. There is however, a DoS Policy section in the GUI.

## Local-In Policies

On the FortiGate unit, there are a number of protocols and traffic that is specific to the internal workings of FortiOS. For many of these traffic sources, you can identify a specific port/IP address for this self-originating traffic. The following traffic can be configured to a specific port/IP address:

- SNMP
- Syslog
- alert email
- FortiManager connection IP
- FortiGuard services
- FortiAnalyzer logging
- NTP
- DNS
- Authorization requests such as RADIUS
- FSSO

Security policies control the flow of traffic through the FortiGate unit. The FortiGate unit also includes the option of controlling internal traffic, that is, management traffic.

Each interface includes an allow access configuration to allow management access for specific protocols. Local policies are set up automatically to allow all users all access. Local-in policies takes this a step further, to enable or restrict the user with that access. This also extends beyond the allow access selection.

Local-in policies are configured in the CLI with the commands:

```
config firewall local-in-policy
 edit <policy_number>
 set intf <source_interface>
 set srcaddr <source_address>
 set dstaddr <destination_address>
 set action {accept | deny}
 set service <service name>
 set schedule <schedule_name>
 end
```

For example, you can configure a local-in policy so that only administrators can access the FortiGate unit on weekends from a specific management computer at 192.168.21.12, represented by the address object mgmt-comp1, using SSH on port 3 (192.168.21.77

represented by the address object FG-port3) using the Weekend schedule which defines the time the of access.

```
config firewall local-in-policy
 edit <1>
 set intf port3
 set srcaddr mgmt-comp1
 set dstaddr FG-port3
 set action accept
 set service SSH
 set schedule Weekend
 end
```

You can also disable a policy should there be a requirement to turn off a policy for troubleshooting or other purpose. To disable a policy enter the commands:

```
config firewall local-in-policy
 edit <policy_number>
 set status disable
 end
```

Use the same commands with a status of enable to use the policy again.

Local-in policies are also supported for IPv6 by entering the command

```
config firewall local-in-policy6.
```

## Security Policy 0

Any security policy that is automatically added by the FortiGate unit has a policy ID number of zero (0). The most common reasons the FortiGate unit creates this policy is:

- The IPSec policy for FortiAnalyzer (and FortiManager version 3.0) is automatically added when an IPSec connection to the FortiAnalyzer unit or FortiManager is enabled.
- The policy to allow FortiGuard servers to be automatically added has a policy ID number of zero.
- The (default) drop rule that is the last rule in the policy and that is automatically added has a policy ID number of zero.
- When a network zone is defined within a VDOM, the intra-zone traffic set to allow or block is managed by policy 0 if it is not processed by a configured security policy.

This policy can appear in logs but will never appear in the security policy list, and therefore, can never be repositioned in the list.

When viewing the FortiGate firewall logs, you may find a log field entry indicating policyid=0. The following log message example indicates the log field policyid=0 in bold.

```
2008-10-06 00:13:49 log_id=0022013001 type=traffic subtype=violation
pri=warning vd=root SN=179089 duration=0 user=N/A group=N/A rule=0
policyid=0 proto=17 service=137/udp app_type=N/A status=deny
src=10.181.77.73 srcname=10.181.77.73 dst=10.128.1.161
dstname=10.128.1.161 src_int=N/A dst_int="Internal" sent=0 rcvd=0
src_port=137 dst_port=137 vpn=N/A tran_ip=0.0.0.0 tran_port=0
```



## Deny Policies

Deny security policies deny traffic that is coming into the network. The FortiGate unit automatically blocks traffic that is associated with a deny security policy.

Deny security policies are usually configured when you need to restrict specific traffic, for example, SSH traffic. Deny security policies can also help when you want to block a service, such as DNS, but allow a specific DNS server.



There is a disparity in the effectiveness of deny policies. Only deny policies that contain VIPs will block traffic directed at those VIPs. Policies with VIPs are processed before other policies, so using a deny policy earlier in the list will not work. For more on this topic, see [“Exception to policy order \(VIPs\)”](#) on page 956.

---

## Accept Policies

Accept security policies accept traffic that is coming into the network. These policies allow traffic through the FortiGate unit, where the packets are scanned, translated if NAT is enabled, and then sent out to its destination.

Accept security policies are the most common security policies that are created in FortiOS. These security policies are basic policies, such as allowing Internet access, as well as complex policies, such as IPSec VPN.

## IPv6 Policies

IPv6 security policies are created both for an IPv6 network, and a transitional network. A transitional network is a network that is transitioning over to IPv6, but must still have access to the Internet or must connect over an IPv4 network.

These policies allow for this specific type of traffic to travel between the IPv6 and IPv4 networks. The IPv6 options for creating these policies is hidden by default. You must enable this feature in *System > Config > Settings*.

## Fixed Port

Some network configurations do not operate correctly if a NAT policy translates the source port of packets used by the connection. NAT translates source ports to keep track of connections for a particular service.

From the CLI you can enable `fixedport` when configuring a security policy for NAT policies to prevent source port translation.

```
config firewall policy
 edit <policy-id>
 ...
 set fixedport enable
 ...
 end
```

However, enabling `fixedport` means that only one connection can be supported through the firewall for this service. To be able to support multiple connections, add an IP pool, and then select Dynamic IP pool in the policy. The firewall randomly selects an IP address from the IP

pool and assigns it to each connection. In this case, the number of connections that the firewall can support is limited by the number of IP addresses in the IP pool.

## Endpoint Security

Endpoint security enforces the use of the FortiClient End Point Security (FortiClient and FortiClient Lite) application on your network. It can also allow or deny endpoints access to the network based on the application installed on them.

By applying endpoint security to a security policy, you can enforce this type of security on your network. FortiClient enforcement can check that the endpoint is running the most recent version of the FortiClient application, that the antivirus signatures are up-to-date, and that the firewall is enabled. An endpoint is usually often a single PC with a single IP address being used to access network services through a FortiGate unit.

With endpoint security enabled on a policy, traffic that attempts to pass through, the FortiGate unit runs compliance checks on the originating host on the source interface. Non-compliant endpoints are blocked. If someone is browsing the web, the endpoints are redirected to a web portal which explains the non-compliance and provides a link to download the FortiClient application installer. The web portal is already installed on the FortiGate unit, as a replacement message, which you can modify if required.

Endpoint Security requires that all hosts using the security policy have the FortiClient Endpoint Security agent installed.

For more information about endpoint security, see the Security Profiles chapter in the FortiOS Handbook.

## Traffic Logging

When you enable logging on a security policy, the FortiGate unit records the scanning process activity that occurs, as well as whether the FortiGate unit allowed or denied the traffic according to the rules stated in the security policy. This information can provide insight into whether a security policy is working properly, as well as if there needs to be any modifications to the security policy, such as adding traffic shaping for better traffic performance.

Depending on what the FortiGate unit has in the way of resources, there may be advantages in optimizing the amount of logging taking places. This is why in each policy you are given 3 options for the logging:

- *No Log* - Does not record any log messages about traffic accepted by this policy.
- *Log Security Events* - records only log messages relating to security events caused by traffic accepted by this policy.
- *Log all Sessions* - records all log messages relating to all of the traffic accepted by this policy.

Depending on the model, if the Log all Sessions option is selected there may be 2 additional options. These options are normally available in the GUI on the higher end models such as the FortiGate 600C or larger.

- *Generate Logs when Session Starts*
- *Capture Packets*

You can also use the CLI to enter the following command to write a log message when a session starts:

```
config firewall policy
 edit <policy-index>
 set logtraffic-start
 end
```

Traffic is logged in the traffic log file and provides detailed information that you may not think you need, but do. For example, the traffic log can have information about an application used (web: HTTP.Image), and whether or not the packet was SNAT or DNAT translated. The following is an example of a traffic log message.

```
2011-04-13
05:23:47
log_id=4
type=traffic
subtype=other
pri=notice
vd=root
status="start"
src="10.41.101.20"
srcname="10.41.101.20"
src_port=58115
dst="172.20.120.100"
dstname="172.20.120.100"
dst_country="N/A"
dst_port=137
tran_ip="N/A"
tran_port=0
tran_sip="10.31.101.41"
tran_sport=58115
service="137/udp"
proto=17
app_type="N/A"
duration=0
rule=1
policyid=1
sent=0
rcvd=0
shaper_drop_sent=0
shaper_drop_rcvd=0
perip_drop=0
```

```
src_int="internal"
dst_int="wan1"
SN=97404 app="N/A"
app_cat="N/A"
carrier_ep="N/A"
```

If you want to know more about logging, see the Logging and Reporting chapter in the FortiOS Handbook. If you want to know more about traffic log messages, see the FortiGate Log Message Reference.

## Quality of Service

The Quality of Service (QoS) feature allows the management of the level of service and preference given to the various types and sources of traffic going through the firewall so that the traffic that is important to the services and functions connecting through the firewall gets the treatment required to ensure the level of quality that is required.

QoS uses the following techniques:

---

<b>Traffic policing</b>	Packets are dropped that do not conform to bandwidth limitations
<b>Traffic Shaping</b>	Assigning minimum levels of bandwidth to be allocated to specific traffic flows to guarantee levels of servers or assigning maximum levels of bandwidth to be allocated to specific traffic flows so that they do not impede other flows of traffic.

---

This helps to ensure that the traffic may consume bandwidth at least at the guaranteed rate by assigning a greater priority queue if the guarantee is not being met. Traffic shaping also ensures that the traffic cannot consume bandwidth greater than the maximum at any given instant in time. Flows that are greater than the maximum rate are subject to traffic policing.

## Queuing

Assigning differing levels priority to different traffic flows so that traffic flows that are adversely effected by latency are prevented from being effected by traffic flows that are not subject to the effects of latency. All traffic in a higher priority traffic queue must be completely transmitted before traffic in lower priority queues will be transmitted.

An example of where you would want to use something like this is if you had competing traffic flows of Voice over IP traffic and email traffic. The VoIP traffic is highly susceptible to latency issues. If you have a delay of a few seconds it is quickly noticeable when it is occurring. Email on the other hand can have a time delay of much longer and it is highly unlikely that it will be noticed at all.



By default, the priority given to any traffic is high, so if you want to give one type of traffic priority over all other traffic you will need to lower the priority of all of the other traffic.

---

## Policy Monitor

Once policies have been configured and enabled it is useful to be able to monitor them. To get an overview about what sort of traffic the policies are processing go to Policy > Monitor > Policy Monitor.

The window is separated into two panes.

### Upper Pane

The upper pane displays a horizontal bar graph comparing the *Top Policy Usage* based on one of the following criteria:

- Active Sessions
- Bytes
- Packets

The criteria that the displayed graph is based on can be selected from the drop down menu in the upper right corner of the pane. The field name is *Report By:*.

The bars of the graph are interactive to an extent and can be used to drill down for more specific information. If you hover the cursor over the bar of the graph a small popup box will appear displaying more detailed information. If the bar of the graph is selected an entirely new window will be displayed using a vertical bar graph to divide the data that made up the first graph by IP address.

For example if the first graph was reporting usage by active sessions it would include a bar for each of the top policies with a number at the end showing how many sessions were currently going through that policy. If one of the bars of the graph was then selected the new bar graph would show the traffic of that policy separated by either *Source Address*, *Destination Address* or *Destination Port*. As in the other window, the selection for the reported criteria is in the upper right corner of the pane. If the parameter was by source address there would be a bar for each of the IP addresses sending a session through the policy and the end of the bar would show how many sessions.

To go back to the previous window of information in the graphs select the Return link in the upper left of the pane.

### Lower Pane

The lower pane contains a spreadsheet of the information that the bar graph will derive their information from. The column headings will include:

- Policy ID
- Source Interface/Zone
- Destination Interface/Zone
- Action
- Active Sessions
- Bytes
- Packets

# Network defense

This section describes in general terms the means by which attackers can attempt to compromise your network and steps you can take to protect it. The goal of an attack can be as complex as gaining access to your network and the privileged information it contains, or as simple as preventing customers from accessing your web server. Even allowing a virus onto your network can cause damage, so you need to protect against viruses and malware even if they are not specifically targeted at your network.

The following topics are included in this section:

- [Monitoring](#)
- [Blocking external probes](#)
- [Defending against DoS attacks](#)

## Monitoring

Monitoring, in the form of logging, alert email, and SNMP, does not directly protect your network. But monitoring allows you to review the progress of an attack, whether afterwards or while in progress. How the attack unfolds may reveal weaknesses in your preparations. The packet archive and sniffer policy logs can reveal more details about the attack. Depending on the detail in your logs, you may be able to determine the attackers location and identity.

While log information is valuable, you must balance the log information with the resources required to collect and store it.

## Blocking external probes

Protection against attacks is important, but attackers often use vulnerabilities and network tools to gather information about your network to plan an attack. It is often easier to prevent an attacker from learning important details about your network than to defend against an attack designed to exploit your particular network.

Attacks are often tailored to the hardware or operating system of the target, so reconnaissance is often the first step. The IP addresses of the hosts, the open ports, and the operating systems the hosts are running is invaluable information to an attacker. Probing your network can be as simple as an attacker performing an address sweep or port scan to a more involved operation like sending TCP packets with invalid combinations of flags to see how your firewall reacts.

### Address sweeps

An address sweep is a basic network scanning technique to determine which addresses in an address range have active hosts. A typical address sweep involves sending an ICMP ECHO request (a ping) to each address in an address range to attempt to get a response. A response signifies that there is a host at this address that responded to the ping. It then becomes a target for more detailed and potentially invasive attacks.

Address sweeps do not always reveal all the hosts in an address range because some systems may be configured to ignore ECHO requests and not respond, and some firewalls and gateways may be configured to prevent ECHO requests from being transmitted to the destination

network. Despite this shortcoming, Address sweeps are still used because they are simple to perform with software tools that automate the process.

Use the `icmp_sweep` anomaly in a DoS policy to protect against address sweeps.

There are a number of IPS signatures to detect the use of ICMP probes that can gather information about your network. These signatures include `AddressMask`, `Traceroute`, `ICMP.Invalid.Packet.Size`, and `ICMP.Oversized.Packet`. Include ICMP protocol signatures in your IPS sensors to protect against these probes/attacks.

## Port scans

Potential attackers may run a port scan on one or more of your hosts. This involves trying to establish a communication session to each port on a host. If the connection is successful, a service may be available that the attacker can exploit.

Use the DoS anomaly check for `tcp_port_scan` to limit the number of sessions (complete and incomplete) from a single source IP address to the configured threshold. If the number of sessions exceed the threshold, the configured action is taken.

Use the DoS anomaly check for `udp_scan` to limit UDP sessions in the same way.

## Probes using IP traffic options

Every TCP packet has space reserved for eight flags or control bits. They are used for communicating various control messages. Although space in the packet is reserved for all eight, there are various combinations of flags that should never happen in normal network operation. For example, the SYN flag, used to initiate a session, and the FIN flag, used to end a session, should never be set in the same packet.

Attackers may create packets with these invalid combinations to test how a host will react. Various operating systems and hardware react in different ways, giving a potential attackers clues about the components of your network.

The IPS signature `TCP.Bad.Flags` detects these invalid combinations. The default action is pass though you can override the default and set it to *Block* in your IPS sensor.

## Configure packet replay and TCP sequence checking

The anti-replay CLI command allows you to set the level of checking for packet replay and TCP sequence checking (or TCP Sequence (SYN) number checking). All TCP packets contain a Sequence Number (SYN) and an Acknowledgement Number (ACK). The TCP protocol uses these numbers for error free end-to-end communications. TCP sequence checking can also be used to validate individual packets.

FortiGate units use TCP sequence checking to make sure that a packet is part of a TCP session. By default, if a packet is received with sequence numbers that fall out of the expected range, the FortiGate unit drops the packet. This is normally a desired behavior, since it means that the packet is invalid. But in some cases you may want to configure different levels of anti-replay checking if some of your network equipment uses non-RFC methods when sending packets.

Configure the anti-replay CLI command:

```
config system global
 set anti-replay {disable | loose | strict}
end
```

You can set anti-replay protection to the following settings:

- `disable` — No anti-replay protection.
- `loose` — Perform packet sequence checking and ICMP anti-replay checking with the following criteria:
  - The SYN, FIN, and RST bit can not appear in the same packet.
  - The FortiGate unit does not allow more than one ICMP error packet through before it receives a normal TCP or UDP packet.
  - If the FortiGate unit receives an RST packet, and `check-reset-range` is set to `strict`, the FortiGate unit checks to determine if its sequence number in the RST is within the un-ACKed data and drops the packet if the sequence number is incorrect.
- `strict` — Performs all of the loose checking but for each new session also checks to determine if the TCP sequence number in a SYN packet has been calculated correctly and started from the correct value for each new session. Strict anti-replay checking can also help prevent SYN flooding.

If any packet fails a check it is dropped.

## Configure ICMP error message verification

Enable ICMP error message verification to ensure an attacker can not send an invalid ICMP error message.

```
config system global
 check-reset-range {disable | strict}
end
```

- `disable` — the FortiGate unit does not validate ICMP error messages.
- `strict` — enable ICMP error message checking.

If the FortiGate unit receives an ICMP error packet that contains an embedded IP(A,B) | TCP(C,D) header, then if FortiOS can locate the A:C->B:D session it checks to make sure that the sequence number in the TCP header is within the range recorded in the session. If the sequence number is not in range then the ICMP packet is dropped. Strict checking also affects how the anti-replay option checks packets.

## Protocol header checking

Select the level of checking performed on protocol headers.

```
config system global
 check-protocol-header {loose | strict}
end
```

- `loose` — the FortiGate unit performs basic header checking to verify that a packet is part of a session and should be processed. Basic header checking includes verifying that the layer-4 protocol header length, the IP header length, the IP version, the IP checksum, IP options are correct, etc.
- `strict` — the FortiGate unit does the same checking as above plus it verifies that ESP packets have the correct sequence number, SPI, and data length.

If the packet fails header checking it is dropped by the FortiGate unit.

## Evasion techniques

Attackers employ a wide range of tactics to try to disguise their techniques. If an attacker disguises a known attack in such a way that it is not recognized, the attack will evade your



security and possibly succeed. FortiGate security recognizes a wide variety of evasion techniques and normalizes data traffic before inspecting it.

### Packet fragmentation

Information sent across local networks and the Internet is encapsulated in packets. There is a maximum allowable size for packets and this maximum size varies depending on network configuration and equipment limitations. If a packet arrives at a switch or gateway and it is too large, the data it carries is divided among two or more smaller packets before being forwarded. This is called fragmentation.

When fragmented packets arrive at their destination, they are reassembled and read. If the fragments do not arrive together, they must be held until all of the fragments arrive. Reassembly of a packet requires all of the fragments.

The FortiGate unit automatically reassembles fragmented packets before processing them because fragmented packets can evade security measures. Both IP packets and TCP packets are reassembled by the IPS engine before examination.

For example, you have configured the FortiGate unit to block access to the example.org web site. Any checks for example.com will fail if a fragmented packet arrives and one fragment contains `http://www.exa` while the other contains `mple.com/`. Viruses and malware can be fragmented and avoid detection in the same way. The FortiGate unit will reassemble fragmented packets before examining network data to ensure that inadvertent or deliberate packet fragmentation does not hide threats in network traffic.

### Non-standard ports

Most traffic is sent on a standard port based on the traffic type. The FortiGate unit recognizes most traffic by packet content rather than the TCP/UDP port and uses the proper IPS signatures to examine it. Protocols recognized regardless of port include DHCP, DNP3, FTP, HTTP, IMAP, MS RPC, NNTP, POP3, RSTP, SIP, SMTP, and SSL, as well as the supported IM/P2P application protocols.

In this way, the FortiGate unit will recognize HTTP traffic being sent on port 25 as HTTP rather than SMTP, for example. Because the protocol is correctly identified, the FortiGate unit will examine the traffic for any enabled HTTP signatures.

### Negotiation codes

Telnet and FTP servers and clients support the use of negotiation information to allow the server to report what features it supports. This information has been used to exploit vulnerable servers. To avoid this problem, the FortiGate unit removes negotiation codes before IPS inspection.

### HTTP URL obfuscation

Attackers encode HTML links using various formats to evade detection and bypass security measures. For example, the URL `www.example.com/cgi.bin` could be encoded in a number of ways to avoid detection but still work properly, and be interpreted the same, in a web browser.

The FortiGate prevents the obfuscation by converting the URL to ASCII before inspection.

**Table 50:** HTTP URL obfuscation types

Encoding type	Example
No encoding	http://www.example.com/cgi.bin/
Decimal encoding	http://www.example.com/&#99;&#103;&#105;&#46;&#98;&#105;&#110;&#47;
URL encoding	http://www.example.com/%43%47%49%2E%42%49%4E%2F
ANSI encoding	http://www.example.com/%u0063%u0067%u0069%u002E%u0062%u0069%u006E/
Directory traversal	http://www.example.com/cgi.bin/test/..

### HTTP header obfuscation

The headers of HTTP requests or responses can be modified to make the discovery of patterns and attacks more difficult. To prevent this, the FortiGate unit will:

- remove junk header lines
- reassemble an HTTP header that's been folded onto multiple lines
- move request parameters to HTTP POST body from the URL

The message is scanned for any enabled HTTP IPS signatures once these problems are corrected.

### HTTP body obfuscation

The body content of HTTP traffic can be hidden in an attempt to circumvent security scanning. HTTP content can be GZipped or deflated to prevent security inspection. The FortiGate unit will uncompress the traffic before inspecting it.

Another way to hide the contents of HTTP traffic is to send the HTTP body in small pieces, splitting signature matches across two separate pieces of the HTTP body. The FortiGate unit reassembles these 'chunked bodies' before inspection.

### Microsoft RPC evasion

Because of its complexity, the Microsoft Remote Procedure Call protocol suite is subject to a number of known evasion techniques, including:

- SMB-level fragmentation
- DCERPC-level fragmentation
- DCERPC multi-part fragmentation
- DCERPC UDP fragmentation
- Multiple DCERPC fragments in one packet

The FortiGate unit reassembles the fragments into their original form before inspection.

## Defending against DoS attacks

A denial of service is the result of an attacker sending an abnormally large amount of network traffic to a target system. Having to deal with the traffic flood slows down or disables the target system so that legitimate users can not use it for the duration of the attack.


Any network traffic the target system receives has to be examined, and then accepted or rejected. TCP, UDP, and ICMP traffic is most commonly used, but a particular type of TCP traffic is the most effective. TCP packets with the SYN flag are the most efficient DoS attack tool because of how communication sessions are started between systems.

### The “three-way handshake”

Communication sessions between systems start with establishing a TCP/IP connection. This is a simple three step process, sometimes called a “three-way handshake,” initiated by the client attempting to open the connection.

1. The client sends a TCP packet with the SYN flag set. With the SYN packet, the client informs the server of its intention to establish a connection.
2. If the server is able to accept the connection to the client, it sends a packet with the SYN and the ACK flags set. This simultaneously acknowledges the SYN packet the server has received, and informs the client that the server intends to establish a connection.
3. To acknowledge receipt of the packet and establish the connection, the client sends an ACK packet.

**Figure 172:** Establishing a TCP/IP connection



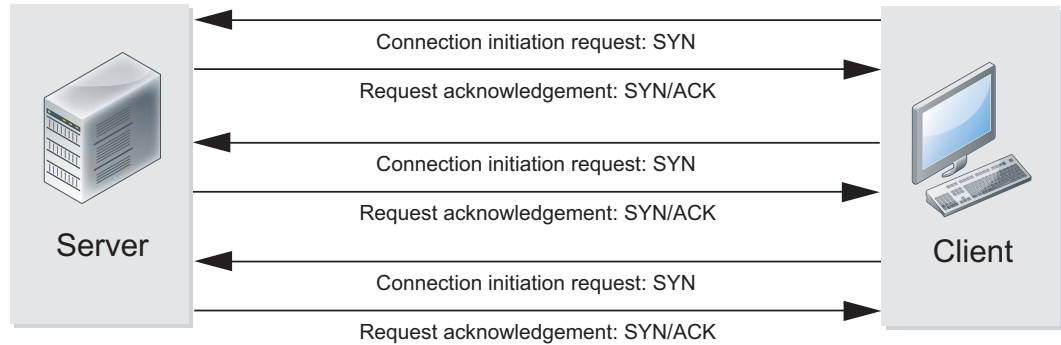
The three-way handshake is a simple way for the server and client to each agree to establish a connection and acknowledge the other party expressing its intent. Unfortunately, the three-way handshake can be used to interfere with communication rather than facilitate it.

### SYN flood

When a client sends a SYN packet to a server, the server creates an entry in its session table to keep track of the connection. The server then sends a SYN+ACK packet expecting an ACK reply and the establishment of a connection.

An attacker intending to disrupt a server with a denial of service (DoS) attack can send a flood of SYN packets and not respond to the SYN+ACK packets the server sends in response. Networks can be slow and packets can get lost so the server will continue to send SYN+ACK packets until it gives up, and removes the failed session from the session table. If an attacker sends enough SYN packets to the server, the session table will fill completely, and further connection attempts will be denied until the incomplete sessions time out. Until this happens, the server is unavailable to service legitimate connection requests.

**Figure 173:**A single client launches a SYN flood attack

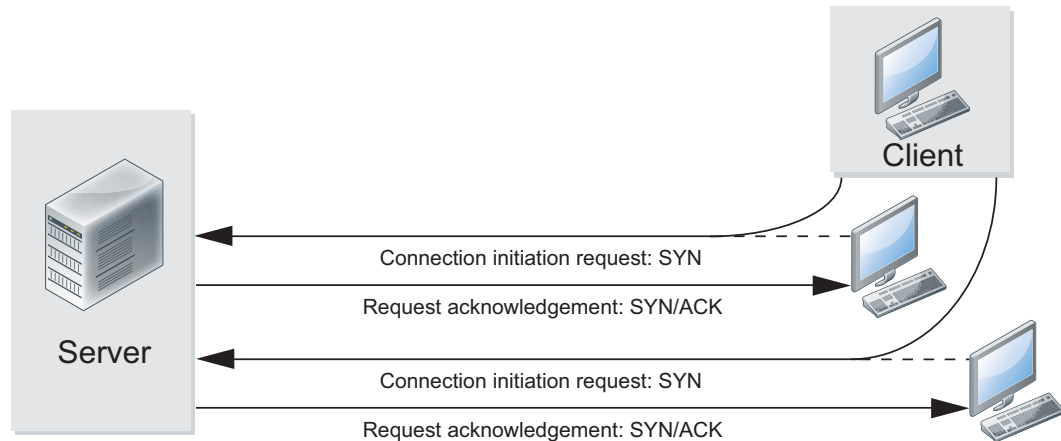


SYN floods are seldom launched from a single address so limiting the number of connection attempts from a single IP address is not usually effective.

## SYN spoofing

With a flood of SYN packets coming from a single attacker, you can limit the number of connection attempts from the source IP address or block the attacker entirely. To prevent this simple defense from working, or to disguise the source of the attack, the attacker may spoof the source address and use a number of IP addresses to give the appearance of a distributed denial of service (DDoS) attack. When the server receives the spoofed SYN packets, the SYN+ACK replies will go to the spoofed source IP addresses which will either be invalid, or the system receiving the reply will not know what to do with it.

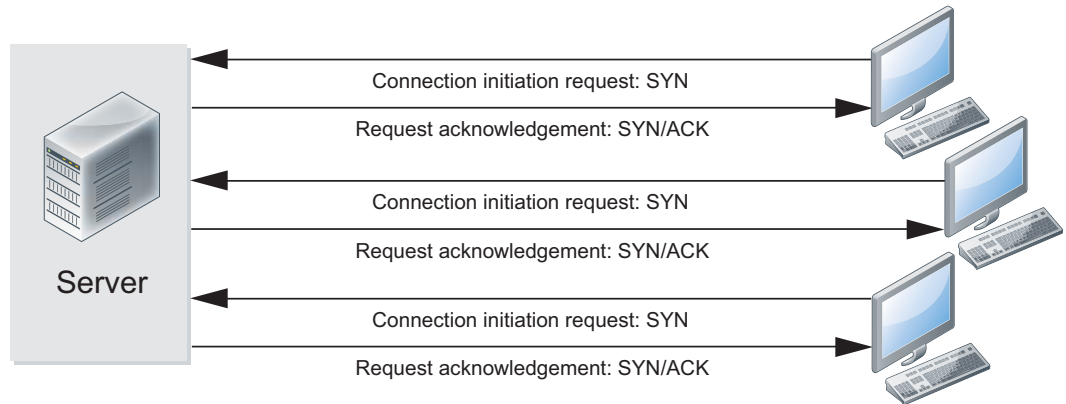
**Figure 174:**A client launches a SYN spoof attack



## DDoS SYN flood

The most severe form of SYN attack is the distributed SYN flood, one variety of distributed denial of service attack (DDoS). Like the SYN flood, the target receives a flood of SYN packets and the ACK+SYN replies are never answered. The attack is distributed across multiple sources sending SYN packets in a coordinated attack.

**Figure 175:** Multiple attackers launch a distributed SYN flood



The distributed SYN flood is more difficult to defend against because multiple clients are capable of creating a larger volume of SYN packets than a single client. Even if the server can cope, the volume of traffic may overwhelm a point in the network upstream of the targeted server. The only defence against this is more bandwidth to prevent any choke-points.

## Configuring the SYN threshold to prevent SYN floods

The preferred primary defence against any type of SYN flood is the DoS anomaly check for `tcp_syn_flood` threshold. The threshold value sets an upper limit on the number of new incomplete TCP connections allowed per second. If the number of incomplete connections exceeds the threshold value, and the action is set to *Pass*, the FortiGate unit will allow the SYN packets that exceed the threshold. If the action is set to *Block*, the FortiGate unit will block the SYN packets that exceed the threshold, but it will allow SYN packets from clients that send another SYN packet.

The tools attackers use to generate network traffic will not send a second SYN packet when a SYN+ACK response is not received from the server. These tools will not “retry.” Legitimate clients will retry when no response is received, and these retries are allowed even if they exceed the threshold with the action set to *Block*.

## SYN proxy

FortiGate units with network acceleration hardware, whether built-in or installed in the form of an add-on module, offer a third action for the `tcp_syn_flood` threshold. Instead of *Block* and *Pass*, you can choose to *Proxy* the incomplete connections that exceed the threshold value.

When the `tcp_syn_flood` threshold action is set to *f*, incomplete TCP connections are allowed as normal as long as the configured threshold is not exceeded. If the threshold is exceeded, the FortiGate unit will intercept incoming SYN packets from clients and respond with a SYN+ACK packet. If the FortiGate unit receives an ACK response as expected, it will “replay” this exchange to the server to establish a communication session between the client and the server, and allow the communication to proceed.

## Other flood types

UDP and ICMP packets can also be used for DoS attacks, though they are less common. TCP SYN packets are so effective because the target receives them and maintains a session table entry for each until they time out. Attacks using UDP or ICMP packets do not require the same level of attention from a target, rendering them less effective. The target will usually drop the offending packets immediately, closing the session.

Use the `udp_flood` and `icmp_flood` thresholds to defend against these DoS attacks.

## DoS policies

DDoS attacks vary in nature and intensity. Attacks aimed at saturating the available bandwidth upstream of your service can only be countered by adding more bandwidth. DoS policies can help protect against DDoS attacks that aim to overwhelm your server resources.

### DoS policy recommendations

- Use and configure DoS policies to appropriate levels based on your network traffic and topology. This will help drop traffic if an abnormal amount is received.
- It is important to set a good threshold. The threshold defines the maximum number of sessions/packets per second of normal traffic. If the threshold is exceeded, the action is triggered. Threshold defaults are general recommendations, although your network may require very different values.
- One way to find the correct values for your environment is to set the action to *Pass* and enable logging. Observe the logs and adjust the threshold values until you can determine the value at which normal traffic begins to generate attack reports. Set the threshold above this value with the margin you want. Note that the smaller the margin, the more protected your system will be from DoS attacks, but your system will also be more likely to generate false alarms.

# GUI & CLI - What You May Not Know

The Graphic User Interface (GUI) is designed to be as intuitive as possible but there are always a few things that are left out because to put all of that information on the interface would clutter it up to the point where it wouldn't be graphical and intuitive anymore.

This section is made up of knowledge that will make working with the both of the management interfaces easier because you won't have to find out about things like field limitations through trial and error. Some of it has to do with changing in how navigation in the GUI has changed.

The section includes the topics:

- [Mouse Tricks](#)
- [Changing the default column setting on the policy page](#)
- [Naming Rules and Restrictions](#)
- [Character Restrictions](#)
- [Length of Fields Restrictions](#)
- [Object Tagging and Coloring](#)
- [Numeric Values](#)
- [Selecting options from a list](#)
- [Enabling or disabling options](#)
- [To Enable or Disable Optionally Displayed Features](#)

## Mouse Tricks

In previous version of the firmware much of the navigation, editing or choosing of options in the Web-based Manager was carried out by using the mouse in combination with a number of icons visible on the interface. This version of the firmware makes more extensive use of the right or secondary mouse button as well as the "drag and drop" feature. If you are used to the old Web-based Manager interface you will notice that a number of the options at the top of the display window are not there anymore or there are fewer of them.

To get a feel for the new approach the Policy > Policy > Policy window is a noticeable place to see some of these changes in action.

The different view modes are still in the upper right-hand corner as they were before but now there is no column settings link to move or configure the columns of the window. Now if you wish to reposition a column just use the mouse to click on the column heading and drag it to its new position. If you wish to add a new column just right-click on one of the column headings and a drop down menu will appear with the option "Column Settings". Use the right pointing triangle to expand the "Column Settings" option to see a choice of possible columns for the window you are in. Those already selected will be at the top with a checked box and the available new ones will be at the bottom ready to be selected.

By right or secondary clicking the mouse cursor in the cells of the Policy window you will get a drop down menu that is contextual to the column and policy row where you made the click. For example if you right click in the "Schedule" column for the row that is for policy #5 you will get the option to select a schedule for policy #5 along with a number of other configuration options relating to that policy or its position in the sequence of policies.

You will find this approach used much more frequently through out the Web-based Manager, giving it a more modern and intuitive feel once you learn to use the right mouse button rather than finding a link displayed on the page.

## Changing the default column setting on the policy page

The Policy > Policy > Policy window is one of the more important ones in the Web based interface and has the capacity to display a lot of information, but displaying all of that information at the same time makes for a very busy screen. If all of the columns are displayed, depending on the screen size you may have to constantly use the scroll bars to see what you need to look at. The default installation shows some of the more commonly used columns but these list may not consist of the columns that you wish to look at or the order that you wish to view them in. For this reason it is possible, through the CLI to override these settings to establish a new default.

The syntax of the command starts with:

```
config system settings
 set gui-default-policy-columns
```

The rest of the command is a space delimited list that depends on the columns you wish to view and the order you wish to view them in. The possible selection is in the following table.

**Table 51:** Variables for the gui-default-policy-columns command

Variable Name	Column Heading
#	Sequence Number
policyid	Policy ID
srcintf	Source Interface
dstintf	Destination Interface
srcaddr	Source Addresses
dstaddr	Destination Addresses
schedule	Policy Schedule
service	Policy Services
action	Policy Action
logtraffic	Traffic Logging Status
nat	Policy NAT Status
status	Policy Status
authentication	Authentication Groups
count	Policy Traffic Counter
profile	Security Profiles
vpntunnel	VPN Tunnel
comments	Policy Comment



## Example:

If you wanted these columns in this order, Policy ID, Source Addresses, Destination Addresses, Security Profiles, Policy Comment. You would enter the command:

```
config system settings
 set gui-default-policy-columns policyid srcaddr dstaddr profile
 comments
```

## Naming Rules and Restrictions

The following are the specific rules that are obeyed by the FortiGate.

Duplicate Name Issues:

- A VLAN cannot have the same name as a physical interface.
- An Address must not have the same name as an Address Group.
- An Address or Address Group must not have the same name as a Virtual IP Address.
- A Service cannot have the same name as a Service Group.
- A VLAN must not have the same name as a VDOM.
- A VLAN or VDOM must not have the same name as a Zone.



Try to make each firewall object name as unique as possible so that it cannot be confused with another object.

---

## Character Restrictions

A name can contain numbers (0-9), uppercase and lowercase letters (A-Z, a-z), spaces, and the special characters - and \_. Other characters are not allowed

The special characters < > ( ) # " ' are allowed only in the following fields:

- Passwords
- Replacement message
- Firewall policy description
- IPS customized signature
- Antivirus blocked file pattern
- Web Filter banned word
- Spam filter banned word
- interface PPPoE client user name
- modem dialup account user name
- modem dialup telephone number



Where you would normally be tempted to use spaces in a name try using the "-" or "\_". There are a few name fields where it is not an issue but most of them will trigger serious and unpredictable issues if there is a space in the name field of the object.

---

## Length of Fields Restrictions

Most name fields accept 35 characters. The exceptions are:

**Table 52:** Characters allowed in some fields

Field	Characters allowed
VDOM names	12
VLAN name	15
RADIUS server secret	15
LDAP server common name identifier	15
Admin user password	32
Schedule names	32
Local certificate email	60
Modem dialup account user name, password, phone number fields	63
Firewall policy comments	63
RADIUS, LDAP server domain name	63
IPSec phase 1 local/peer ID	63
IPS custom signature name	63
Spam Filter MIME header name	63
Antivirus file block pattern	63
Local certificate organizational unit, organization, locality, state/province fields	127
IPSec phase 1 pre-shared key or certificate name	127
Web Filter banned word, URL, URL exempt, Pattern fields	127
Spam Filter RBL server name, email address, MIME header body	127
LDAP server distinguished name	128
IPS custom signature	511
Replacement message	1024

## Object Tagging and Coloring

For ease of administration and searching the features of object tagging and coloring are available. These features can make it easier to find similar objects in an administrative screen or to search based on common assigned criteria.

## Tags

The Tag Management menu provides a central location to view, search and manage tags that you created. Tags are keywords or a term that is assigned to a specific configuration that can be used for searching or filtering purposes. With the Tag Management feature you can:

- Search to find a specific tag
- View where a tag is referenced, for example, a single tag could be referenced in a security policy, predefined signature and application
- Go to where the tag is located, for example, a security policy
- View how many tags are currently unused
- Remove tags.

The Tag Management page also provides a way to easily locate a specific object, such as a security policy, because of how tags work. For example, an SSL VPN security policy is tagged with the keywords SSL VPN, remote, and SSL branch office;

1. From the Tag Management page, enter SSL and the tags for that security policy appear;
2. Select one of the tags and within the Object Usage window, select to go to the SSL VPN security policy.
3. You can view detailed information about what object is using a tag by selecting one of the tags in the rectangular area that contains a grey background. The Object Usage window appears, which displays similar information as when you select a number in the Ref. column.



In the Add Tags window, you can select to add existing tags to the security policy or address list; however, these tags belong to predefined signatures and applications as well as to other security policies and address lists so the tags may not be applicable. You should make sure that the tag is valid for its use when applied to a security policy or other object otherwise it becomes redundant.

---

## Coloring

You will sometimes be given the option of changing the color of the icon that will represent the firewall object you are configuring. You can do this by clicking on the [Change] link. You will be given the option of picking one of 32 colors.

## Numeric Values

Numeric values are used to configure various sizes, rates, numeric addresses, or other numeric values. For example, a static routing priority of 10, a port number of 8080, or an IP address of 10.10.10.1. Numeric values can be entered as a series of digits without spaces or commas (for example, 10 or 64400), in dotted decimal format (for example the IP address 10.10.10.1) or as in the case of MAC or IPv6 addresses separated by colons (for example, the MAC address 00:09:0F:B7:37:00). Most numeric values are standard base-10 numbers, but some fields (again such as MAC addresses) require hexadecimal numbers.

Most web-based manager numeric value fields make it easy to add the acceptable number of digits within the allowed range. CLI help includes information about allowed numeric value ranges. Both the web-based manager and the CLI prevent you from entering invalid numbers.

## Selecting options from a list

If a configuration field can only contain one of a number of selected options, the web-based manager and CLI present you a list of acceptable options and you can select one from the list. No other input is allowed. From the CLI you must spell the selection name correctly.

## Enabling or disabling options

If a configuration field can only be on or off (enabled or disabled) the web-based manager presents a check box or other control that can only be enabled or disabled. From the CLI you can set the option to enable or disable.

## To Enable or Disable Optionally Displayed Features

There are a number of features in the web-based manager that can be configured to either be displayed if you are likely to use them or disabled if you have no need to see them. The ones that may be relevant to the function of the Firewall are:

- Central NAT Table
- Dynamic Profile
- Explicit Proxy
- Implicit Firewall Policies
- IPv6
- Load Balance
- Local In Policy
- Object Tagging and Coloring

You can enable or disable these features by going to *System > Admin > Settings* or by using the following CLI options:

```
config system global
 set gui-ap-profile {disable | enable}
 set gui-central-nat-table {disable | enable}
 set gui-dns-database {disable | enable}
 set gui-dynamic-profile-display {disable | enable}
 set gui-icap {disable | enable}
 set gui-implicit-id-based-policy {disable | enable}
 set gui-implicit-policy {disable | enable}
 set gui-ipsec-manual-key {enable | disable}
 set gui-ipv6 {enable | disable}
 set gui-lines-per-page <gui_lines>
 set gui-load-balance {disable | enable}
 set gui-object-tags {disable | enable}
 set gui-policy-interface-pairs-view {enable | disable}
 set gui-voip-profile {disable | enable}
end
```

# Building firewall objects and policies

The other chapters in the Firewall book have so far been concerned primarily with concepts and abstract ideas that are designed to help you understand what is going on with the firewall and what it can do. Now that we have a good grounding in the “what” it is time to get into the “how”.

This section will provide the instructions for the web-based manager (when available) and the CLI for adding and or editing FortiGate firewall objects and then how to put them together when building a policy to govern the traffic flowing through your network. To give some context, scenarios have been included. The instructions here are concerned with the creation of the objects. The inclusion of these objects into firewall policies is not shown in these instructions.

This chapter includes the instructions for building the following:

- [IPv4 Firewall Addresses](#)
- [IPv6 Firewall Addresses](#)
- [FQDN address](#)
- [Changing the TTL of a FQDN address](#)
- [New Geography-based Address](#)
- [Wildcard Address](#)
- [IPv4 Address Group](#)
- [IPv6 Address Group](#)
- [Multicast Address](#)
- [Service Category](#)
- [TCP/UDP/SCTP Service](#)
- [ICMP Service](#)
- [ICMPv6 Service](#)
- [Service Group](#)
- [Virtual IP address](#)
- [IP Pool](#)
- [Central NAT Table](#)
- [Firewall Schedule - Recurring](#)
- [Firewall Schedule - One-time](#)
- [Schedule Group](#)
- [Proxy Option](#)
- [Firewall Address Policy](#)
- [Firewall User Identity Policy](#)
- [Firewall Device Identity Policy](#)
- [DoS Policy](#)

## IPv4 Firewall Addresses

### Scenario: Mail Server

You need to create an IPv4 address for the Mail Server on Port1 of your internal network.

- These server is on the network off of port1.
- The IP address is 192.168.1.27
- The subnet mask is 255.255.255.0
- There should be a tag for this address being for a server

Go to *Firewall Objects > Address > Addresses* and select *Create New > Address/FQDN*.

Fill out the fields with the following information:

<b>Address Name</b>	Mail_Server
<b>Comments</b>	<Input into this field is optional>
<b>Color</b>	<Changing this value is optional>
<b>Type</b>	Subnet / IP Range
<b>Subnet / IP Range</b>	192.168.1.27/255.255.255.0
<b>Interface</b>	port1
<b>Add Tags</b>	Server

Select *OK*.

Enter the following CLI command:

```
config firewall address
 edit Mail_Server
 set type ipmask
 set subnet 192.168.1.27 255.255.255.255
 set associated-interface port1
 set tags Server
 end
```

### Scenario: First Floor Network

You need to create an IPv4 address for the subnet of the internal network off of Port1.

- These computers are on the network off of port1.
- The subnet is the range from 192.168.1.1 to 192.168.1.255.
- The subnet mask is 255.255.255.0
- There should be a reference to this being the network for the 1st floor of the building.

Go to *Firewall Objects > Address > Addresses* and select *Create New > Address/FQDN*.

Fill out the fields with the following information

<b>Address Name</b>	Internal_Subnet_1
<b>Comments</b>	Network for 1st Floor

<b>Color</b>	<Changing this value is optional>
<b>Type</b>	Subnet / IP Range
<b>Subnet / IP Range</b>	192.168.1.0/24
<b>Interface</b>	port1
<b>Add Tags</b>	<Input into this field is optional>

Select *OK*.

Enter the following CLI command:

```
config firewall address
 edit Internal_Subnet_1
 set comment "Network for 1st Floor"
 set type ipmask
 set subnet 192.168.1.0/24
 set associated-interface port1
 end
```

## Scenario: Marketing Department

You need to create an IPv4 address for the address range for a group of computers used by the Marketing Department.

- These computers are on the network off of port1.
- The IP addresses for these computers range from 192.168.1.100 to 192.168.1.115

Go to *Firewall Objects > Address > Addresses* and select *Create New > Address/FQDN*.

Fill out the fields with the following information

<b>Address Name</b>	Marketing_computers
<b>Comments</b>	<Input into this field is optional>
<b>Color</b>	<Changing this value is optional>
<b>Type</b>	Subnet / IP Range
<b>Subnet / IP Range</b>	192.168.1.[100-115]
<b>Interface</b>	port1
<b>Add Tags</b>	<Input into this field is optional>

Select *OK*.

Enter the following CLI command:

```
config firewall address
 edit Internal_Subnet_1
 set type iprange
 set start-ip 192.168.1.100
 set end-ip 192.168.1.115
 set associated-interface port1
 end
```

## Verification

To verify that the addresses were added correctly:

Go to *Firewall Objects > Address > Addresses*. Check that the addresses have been added to the address list and that they are correct.

Enter the following CLI command:

```
config firewall address
 edit <the name of the address to verify>
 show full-configuration
```

## IPv6 Firewall Addresses

### Scenario: Mail Server

You need to create an IPv6 address for the Mail Server on Port1 of your internal network.

- These server is on the network off of port1.
- The IP address is 2001:db8:0:2::20/64
- There should be a tag for this address being for a server

Go to *Firewall Objects > Address > Addresses* and select *Create New > IPv6 Address*.

Fill out the fields with the following information

<b>Address Name</b>	Mail_Server
<b>Comments</b>	<Input into this field is optional>
<b>Color</b>	<Changing this value is optional>
<b>IPv6 Address</b>	2001:db8:0:2::20/64
<b>Add Tags</b>	Server

Select *OK*.



Enter the following CLI command:

```
config firewall address6
 edit Mail_Server
 set type ipmask
 set subnet 2001:db8:0:2::20/64
 set associated-interface port1
 set tags Server
 end
```

## Scenario: First Floor Network

You need to create an IPv4 address for the subnet of the internal network off of Port1.

- These computers are on the network off of port1.
- The Network uses the IPv6 addresses: fdde:5a7d:f40b:2e9d:xxxx:xxxx:xxxx:xxxx
- There should be a reference to this being the network for the 1st floor of the building.

Go to *Firewall Objects > Address > Addresses* and select *Create New > IPv6 Address*.

Fill out the fields with the following information

Field Name	Field Value
Address Name	Internal_Subnet_1
Comments	Network for 1st Floor
Color	<Changing this value is optional>
Type	Subnet / IP Range
Subnet / IP Range	2001:db8:0:2::/64
Interface	port1
Add Tags	<Input into this field is optional>

Select *OK*.

Enter the following CLI command:

```
config firewall address6
 edit Internal_Subnet_1
 set comment "Network for 1st Floor"
 set subnet 2001:db8:0:2::/64
 end
```

## Verification

To verify that the addresses were added correctly:

Go to *Firewall Objects > Address > Addresses*. Check that the addresses have been added to the address list and that they are correct.

Enter the following CLI command:

```
config firewall address6
 edit <the name of the address that you wish to verify>
 Show full-configuration
```

## FQDN address

You have to create a policy that will govern traffic that goes to a site that has a number of servers on the Internet. Depending on the traffic or the possibility that one of the servers is down network traffic can go to any one of those sites. The consistent factor is that they all use the same Fully Qualified Domain Name.

- The FQDN of the web site: example.com
- The number of ISP connections off of the FortiGate firewall: 2

Go to *Firewall Objects > Address > Addresses* and select *Create New > Address/FQDN*.

Fill out the fields with the following information

Field Name	Field Value
Address Name	BigWebsite.com
Comments	<Input into this field is optional>
Color	<Changing this value is optional>
Type	FQDN
FQDN	bigwebsite.com
Interface	any
Add Tags	<Input into this field is optional>

Select *OK*.

Enter the following CLI command:

```
config firewall address
 edit BigWebsite.com
 set type fqdn
 set associated-interface any
 set fqdn bigwebsite.com
 end
```

## Verification

To verify that the addresses were added correctly:

Go to *Firewall Objects > Address > Addresses*. Check that the addresses have been added to the address list and that they are correct.

Enter the following CLI command:

```
config firewall address
 edit <the name of the address that you wish to verify>
 Show full-configuration
```

## Changing the TTL of a FQDN address

To make sure that the FQDN resolves to the most recent active server you have been asked to make sure that the FortiGate has not cached the address for any longer than 10 minutes.

There is no field for the cached time-to-live in the web-based manager. It is only configurable in the CLI. Enter the following commands:

```
config firewall address
 edit BigWebsite.com
 set cache-ttl 600
 end
```

## New Geography-based Address

Your company is US based and has information on its web site that may be considered information that is not allowed to be sent to embargoed countries. In an effort to help reduce the possibility of sensitive information going to those countries you have been asked to set up addresses for those countries so that they can be blocked in the firewall policies.

- One of the countries you have been asked to block is Cuba
- You have been asked to Tag the addresses so that other administrators will know why they have been created

Go to *Firewall Objects > Address > Addresses* and select *Create New > Address/FQDN*.

Fill out the fields with the following information

Field Name	Field Value
Address Name	Cuba
Comments	<Input into this field is optional>
Color	<Changing this value is optional>
Type	Geography
Country	Cuba
Interface	any
Add Tags	Embargo

Select *OK*.

Enter the following CLI command:

```
config firewall address
edit Cuba
set type geography
set country CN
set interface wan1
end
```

## Wildcard Address

The company has a large network with multiple subnets. Each team has its own subnet in the 172.12.x.x range. To help keep things organized the IT department uses the same host address on each subnet for the servers. For instance the gateways are always .1 or .2. mail servers are always .5 and print servers are always .10.

In this case an address needs to be created for the mail servers for the entire company.

The addresses will be 172.12.0.5, 172.12.1.5, 172.12.2.5, etc.

Go to *Firewall Objects > Address > Addresses* and select *Create New > Address/FQDN*.

Fill out the fields with the following information

---

<b>Address Name</b>	Print_Servers
<b>Comments</b>	<Input into this field is optional>
<b>Color</b>	<Changing this value is optional>
<b>Type</b>	Subnet / IP Range
<b>Subnet / IP Range</b>	172.12.0.5 / 255.255.0.255
<b>Interface</b>	any

---

Select *OK*.



There will be a pop up window explaining that the address will automatically be converted to a wildcard address and ask if you would like to continue. When you see the address in the editing window you will notice that the type field shows wildcard even though that was not an option before.

Enter the following CLI command:

```
config firewall address
edit Print_Servers
Set type wildcard
Set wildcard 172.12.0.5 255.255.0.255
end
```

To verify that the addresses were added correctly:

Go to *Firewall Objects > Address > Addresses*. Check that the addresses have been added to the address list and that they are correct.

Enter the following CLI command:

```
config firewall address
 edit <the name of the address that you wish to verify>
 Show full-configuration
```

## IPv4 Address Group

Your company has a small division that is in Denmark that has a number of remote users that need to connect to a resource from either home, office or customer sites. The thing that they have in common is that there are all in Denmark. An address group needs to be created that will allow for this.

The preconfigured addresses to use will consist of:

- Denmark - a geography based address
- Denmark\_ISP1 - a IP range of address of an ISP that services Denmark
- Denmark\_ISP2 - a IP range of address of another ISP that services Denmark
- Denmark\_Division - the FQDN of the Denmark office that uses Dynamic DNS

Go to *Firewall Objects > Address > Groups* and select *Create New > Address Group*.

Fill out the fields with the following information

---

<b>Group Name</b>	Denmark_Users
<b>Comments</b>	<Input into this field is optional>
<b>Color</b>	<Changing this value is optional>
<b>Members</b>	Denmark Denmark_ISP1 Denmark_ISP2 Denmark_Division

---

Select *OK*.

Enter the following CLI command:

```
config firewall addrgrp
 edit Denmark_Users
 set member Denmark Denmark_ISP1 Denmark_ISP2 Denmark_Division
 end
```



If you need to edit out a member of an address group in the CLI you need to

To verify that the addresses were added correctly:

Go to *Firewall Objects > Address > Groups*. Check that the addresses have been added to the address list and that they are correct.

Enter the following CLI command:

```
config firewall addgrp
 edit <the name of the address that you wish to verify>
 Show full-configuration
```

## IPv6 Address Group

Create IPv6 address groups from existing IPv6 addresses

Your company has 3 internal servers with IPv6 addresses that it would like to group together for the purposes of a number of policies.

The preconfigured addresses to use will consist of:

- Web\_Server-1
- Web\_Server-2
- Web\_Server-3

Go to *Firewall Objects > Address > Groups* and select *Create New > IPv6 Address Group*.

Fill out the fields with the following information

<b>Group Name</b>	Web_Server_Cluster
<b>Comments</b>	<Input into this field is optional>
<b>Color</b>	<Changing this value is optional>
<b>Members</b>	Web_Server-1 Web_Server-2 Web_Server-3

Select *OK*.

Enter the following CLI command:

```
config firewall addrgrp6
 edit Web_Server_Cluster
 set member Web_Server-1 Web_Server-2 Web_Server-3
 end
```

To verify that the addresses were added correctly:

Go to *Firewall Objects > Address > Groups*. Check that the addresses have been added to the address list and that they are correct.

Enter the following CLI command:

```
config firewall addgrp6
 edit <the name of the address that you wish to verify>
 Show full-configuration
```

## Multicast Address

The company has a large high tech campus that has monitors in many of its meeting rooms. It is common practice for company wide notifications of importance to be done in a streaming video

format with the CEO of the company addressing everyone at once. The video is High Definition quality so takes up a lot of bandwidth. To minimize the impact on the network the network administrators have set things up to allow the use of multicasting to the monitors for these notifications. Now it has to be set up on the FortiGate firewall to allow the traffic.

- The range being used for the multicast is 239.5.0.0 to 239.5.255.255
- The interface on this FortiGate firewall will be on port 4

Go to *Firewall Objects > Address > Addresses* and select *Create New > Address/FQDN*.

Fill out the fields with the following information:

Field	Value
Category	Multicast Address
Name	Meeting_Room_Displays
Color	<optional>
Show in address list	enabled
Multicast IP Range	239.5.0.0-239.5.255.255
Interface	port4
Comments	<optional>

Select *OK*.

Enter the following CLI command:

```
config firewall multicast-address
 edit "meeting_room_display"
 set associated-interface "port9"
 set start-ip 239.5.0.0
 set end-ip 239.5.255.255
 set visibility enable
 next
end
```

To verify that the address range was added correctly:

Go to *Firewall Objects > Address > Groups*. Check that the addresses have been added to the address list and that they are correct.

Enter the following CLI command:

```
config firewall multicast-address
 edit <the name of the address that you wish to verify>
 Show full-configuration
```

## Service Category

Add a new category to the list of Service Categories

You plan on adding a number of devices such as web cameras that will allow the monitoring of the physical security of your datacenter. A number of non-standard services will have to be created and you would like to keep them grouped together under the heading of “Surveillance”

Go to *Firewall Objects > Service > Services* and select *Create New > Category*.

Fill out the fields with the following information

<b>Name</b>	Surveillance
<b>Comments</b>	For DataCenter Surveillance Devices

Select *OK*.

Enter the following CLI command:

```
config firewall service category
 Edit Surveillance
 Set comment "For DataCenter Surveillance Devices"
end
```

To verify that the category was added correctly:

Go to *Firewall Objects > Service > Services*. Select *Create New > Custom Service*. Go to the *Category* Field and use the drop down menu arrow to open the drop down menu. The new category should be displayed.

Enter the following CLI command:

```
config firewall service category
 show
```

This should bring up all of the categories. Check to see that the new one is displayed.

## TCP/UDP/SCTP Service

To create and configure a TCP/UDP/SCTP protocol type service.

You have set up some new web cams at work that send a constant live feed to a security service. Not only do these cameras have a feed that can be sent offsite they can be remotely managed from a browser or an application.

The ports that need to be opened to use all of the features of the web cams are:

- Management by browser - TCP on port 8000
- Real time video feed - UDP on port 4000
- Connection through vendor application - SCTP on port 1600

The IP address of the offsite service is 256.25.56.12 (Not a valid IP address. Used for example only)

- One service will be needed for the incoming connections
- One service will be needed for the outgoing connections

The IT manager would like the service for the outgoing data stream to be tied to the destination of the Surveillance service site so that service can only be used for that one vendor.

### To add the incoming service

Go to *Firewall Objects > Service > Services* and select *Create New > Custom Service*.



Fill out the fields with the following information

<b>Name</b>	WebCam_Connection-incoming
<b>Comments</b>	<Input into this field is optional>
<b>Service Type</b>	Firewall
<b>Color</b>	<Changing this value is optional>
<b>Show in Service List</b>	Check in check box
<b>Category</b>	Surveillance
<b>Protocol Type</b>	TCP/UDP/SCTP
<b>IP/FQDN</b>	<Leave blank>

Protocol	Destination Port		Source Port	
	Low	High	Low	High
<b>TCP</b>	8000	8000	1	65535
<b>SCTP</b>	16000	16000	1	65535



The source port range can be left blank as the default is 1 to 65535.

Select *OK*.

Enter the following CLI command:

```
config firewall service custom
 edit WebCam_Connection-incoming
 set protocol TCP/UDP/SCTP
 set tcp-portrange 8000
 set sctp-portrange 16000
 set visibility enable
 end
```

#### To add the outgoing service

Go to *Firewall Objects > Service > Services* and select *Create New > Custom Service*.

Fill out the fields with the following information

<b>Name</b>	WebCam_Connection-outgoing
<b>Comments</b>	<Input into this field is optional>
<b>Service Type</b>	Firewall
<b>Color</b>	<Changing this value is optional>

<b>Show in Service List</b> Check in check box				
<b>Category</b>	Surveillance			
<b>Protocol Type</b>	TCP/UDP/SCTP			
<b>IP/FQDN</b>	256.25.56.12			
	<b>Destination Port</b>		<b>Source Port</b>	
<b>Protocol</b>	Low	High	Low	High
<b>TCP</b>	4000	4000	1	65535

Select *OK*.

Enter the following CLI command:

```
config firewall service custom
edit WebCam_Connection-incoming
Set protocol TCP/UDP/SCTP
Set category Surveillance
Set udp-portrange 4000
Set iprange 256.25.56.12
Set visibility enable
end
```

To verify that the category was added correctly:

Go to *Firewall Objects > Service > Services*. Check that the services have been added to the service list and that they are correct.

Enter the following CLI command:

```
config firewall service custom
edit <the name of the service that you wish to verify>
Show full-configuration
```

This should bring up all of the details of the service.

## ICMP Service

The Security Officer would like to block the use of the traceroute utility through the network. The IT manager insists that ping and other ICMP utility must be allows for the task of diagnosing connectivity, so it is agreed that only traceroute functionality will be blocked.

The ICMP type for traceroute is 30. There is no codes with the type.

Web-based Manager Instructions

Go to *Firewall Objects > Service > Services* and select *Create New > Custom Service*.

Fill out the fields with the following information

Field Name	Field Value
<b>Name</b>	traceroute

<b>Comments</b>	<Input into this field is optional>
<b>Service Type</b>	Firewall
<b>Color</b>	<Changing this value is optional>
<b>Show in Service List</b>	Check in check box
<b>Category</b>	Uncategorized
<b>Protocol Type</b>	ICMP
<b>Type</b>	30
<b>Code</b>	<Leave blank>

Select *OK*.

Enter the following CLI command:

```
config firewall service custom
edit traceroute
set protocol ICMP
set icmp-type 30
set visibility enable
end
```

To verify that the category was added correctly:

Go to *Firewall Objects > Service > Services*. Check that the services have been added to the services list and that they are correct.

Enter the following CLI command:

```
config firewall service custom
edit <the name of the service that you wish to verify>
show full-configuration
```

## ICMPv6 Service

The IT Manager is doing some diagnostics and would like to temporarily block the successful replies of ICMP Node information Responses between 2 IPv6 networks.

The ICMP type for ICMP Node informations responses is 140. The codes for a successful response is 0.

Web-based Manager Instructions

Go to *Firewall Objects > Service > Services* and select *Create New > Custom Service*.

Fill out the fields with the following information

Field Name	Field Value
<b>Name</b>	diagnostic-test1
<b>Comments</b>	<Input into this field is optional>
<b>Service Type</b>	Firewall

<b>Color</b>	<Changing this value is optional>
<b>Show in Service List</b>	Check in check box
<b>Category</b>	Uncategorized
<b>Protocol Type</b>	ICMP6
<b>Type</b>	140
<b>Code</b>	0

Select *OK*.

Enter the following CLI command:

```
config firewall service custom
edit diagnostic-test1
set protocol ICMP6
set icmptype 140
set icmpcode 0
set visibility enable
end
```

To verify that the category was added correctly:

Go to *Firewall Objects > Service > Services*. Check that the services have been added to the services list and that they are correct.

Enter the following CLI command:

```
config firewall service custom
edit <the name of the service that you wish to verify>
show full-configuration
```

## Service Group

The company provide email services for a number of different companies. They have a standard list of services that they like to keep open to their customer's email servers, including webmail services. The company prides itself on getting a customer up and going the same day so they use standard templates for everything to make sure nothing is forgotten including the services that are available.

The services include:

- IMAP
- IMAPS
- POP3
- POP3S
- SMTP
- SMTPS
- HTTP
- HTTPS
- Email\_Admin - a custom service for administration of the servers

Go to Firewall Objects > Service > Groups and select *Create New*.

Fill out the fields with the following information:

Field	Value
Group Name	Cust_Email_Serv_Template
Comments	(Optional)
Color	(Optional)
Type	Firewall
Members	(click to add...choose from the drop down <ul style="list-style-type: none"><li>• IMAP</li><li>• IMAPS</li><li>• POP3</li><li>• POP3S</li><li>• SMTP</li><li>• SMTPS</li><li>• HTTP</li><li>• HTTPS</li><li>• Email_Admin</li></ul>

Select *OK*.

Enter the following CLI command:

```
config firewall service group
 edit Cust_Email_Serv_Template
 set member "IMAP" "IMAPS" "POP3" "POP3S" "SMTP" "SMTPS" "HTTP"
 "HTTPS" "Email_Admin"
 next
end
```

To verify that the category was added correctly:

Go to *Firewall Objects > Service > Groups*. Check that the service group has been added to the services list and that it is correct.

Enter the following CLI command:

```
config firewall service group
 edit <the name of the service that you wish to verify>
 show full-configuration
```

## Virtual IP address

The company has an web server on the internal network that needs to be accessed from the Internet.

- The internal IP address is 192.168.50.37
- The external IP address is 256.85.94.60 (for example use only. Not a valid IP address)
- The external IP address is assigned by ISP "A" on WAN1
- The port that needs to be mapped is 80

Go to *Firewall Objects > Virtual IP > Virtual IP* and select *Create New*.

Fill out the fields with the following information.

Field	Value
Name	Web1-VIP
Comments	Virtual IP for the Forum Webserver
Color	<optional>
External Interface	wan1
Type	(This field is only changeable in the CLI)
Source Address Filter	<leave blank or default setting>
External IP Address/Range	256.85.94.60
Mapped IP Address/Range	192.168.50.37
Port Forward	enabled
Protocol	TCP
External Service Port	80
Map to Port	80

Select *OK*.

Enter the following CLI command:

```
config firewall vip
 edit Web1-VIP
 set comment "Virtual IP for the Forum Webserver"
 set extintf wan1
 set extip 256.85.94.60
 set mappedip 192.168.50.37
 set portforward enable
 set protocol tcp
 set extport 80
 set mapped port 80
 end
end
```

To verify that the category was added correctly:

Go to *Firewall Objects > Virtual IP> Virtual IP*. Check that the virtual IP address has been added to the list and that it is correct.

Enter the following CLI command:

```
config firewall vip
 edit <the name of the vip that you wish to verify>
 show full-configuration
```

## VIP Group

The company has only a single external IP address but multiple servers with different functions running on its internal LAN that need to be accessed from the Internet.

- The external IP address of the company on wan1 is 256.34.56.149 (for example use only. Not a valid IP address)
- The webserver is on the internal LAN on 192.168.100.86
- The webserver needs to answer on ports 80 443
- The administration of the FortiGate firewall connects on port 4443 instead of 443
- There is are also a separate email server, FTP server, and Terminal Server for specialised applications.
- 2 Virtual IPs have been created to map 256.34.56.149 to 192.168.100.86 on ports 80 and 443. The names are webserver\_80 and webserver\_443 respectively.

Go to *Firewall Objects > Virtual IP> Virtual IP* and select *Create New*.

Fill out the fields with the following information.

Field	Value
Group Name	WebServer_Grp
Comments	(Optional)
Color	(Optional)
Interface	wan1

Move the Following “Available VIPs:” to the “Members” field:

- “webserver\_80”
- “webserver\_443”

Enter the following CLI command:

```
config firewall vipgrp
 edit WebServer_Grp
 set member "webserver_80" "webserver_443"
 next
end
```

To verify that the category was added correctly:

Go to *Firewall Objects > Virtual IP> Group*. Check that the virtual IP address group has been added to the list and that it is correct.

Enter the following CLI command:

```
config firewall vipgrp
 edit <the name of the vip that you wish to verify>
 show full-configuration
```

## IP Pool

Your company has an application server on the internal network that sends out regular data updates to an offsite service. In order to make the service site more secure, they only accept connections from predefined IP address. If the external IP address of the FortiGate firewall interface were used that would mean that the service would be accepting sessions from just about any user in the network so a separate IP address need so be assigned for the Network Address Translation.

- The external address that will be used is one that has been assigned to the company by the ISP on WAN2
- The address is 256.100.42.129 (for example use only. Not a valid IP address)

*Note:* the ARP interface cannot be set in the Web-based Manager but as this is the only path that the traffic will be coming from the outside this should not be an issue. The setting has been included in the CLI instructions so that you will now how to set it in a situation where you want the ARP replies to be answered only on a specific interface.

Go to *Firewall Objects > Virtual IP > IP Pools*

Fill out the fields with the following information:

Field	Value
Name	App_Server1
Comments	Addresses assignment for this server only.
Type	One-to-One
External IP Range/Subnet2	256.100.42.129
ARP Reply	enabled

Select *OK*

Enter the following CLI command:

```
config firewall ippool
 edit App_Server1
 set comments "Addresses assignment for this server only."
 set type one-to-one
 set arp-reply enable
 set arp-intf wan2
 set startip 256.100.42.129
 set endip 256.100.42.129
 end
```

To verify that the category was added correctly:

Go to *Firewall Objects > Virtual IP > IP Pools*



Check that the IP Pool has been added to the list of IP Pools and that the listed settings are correct.

Enter the following CLI command:

```
config firewall ippool
 edit <the name of the IP Pool you wish to verify>
 show full-configuration
```

## Central NAT Table

The company has a server on the Development LAN that needs to communicate with a server at a remote site over the Internet. One of the restrictions on the communications between these systems is that the IP address and source port must be specific.

- The traffic going out on to the Internet must be NATed
- The traffic is coming from a server with the IP address 192.168.150.86
- An address called "app-server" has been created for the address 192.168.150.86 on the port1 interface
- The external interface must be 256.23.45.67
- An address called "app-server-ext" has been created for the address 256.23.45.67 on the wan1 interface
- The originating traffic from the server originates in the port range from 2000 to 3000
- The remote site requires that the source TCP port must be within the 12000 to 13000 range

The original address and Translated Address fields require values that are address names that are listed in the address section of Firewall Objects.

Go to *Policy > Policy > Central NAT Table*.

Create a new NAT table

Fill out the fields with the following information:

Field	Value
Source Address	app-server
Translated Address	app-server-ext
Original Source Port	2000
Translated Port	12000-13000

Select *OK*

Enter the following CLI command:

```
config firewall central-nat
 edit 0
 set orig-addr app-server
 set nat-ippool app-server-ext
 set orig-port 2000
 set nat-port 12000-13000
 next
end
```

To verify that the table was added correctly:

Go to *Policy > Policy > Central NAT Table*

Check that the table has been added to the list of Central NAT Tables and that the listed settings are correct.

Enter the following CLI command:

```
config firewall central-nat
 show full-configuration
```

Verify that the listing of tables includes the one that you have just configured, with the correct settings.



When configuring the Central NAT in the GUI you may notice that only those addresses which have been configured to be associated with *any* interface are displayed in the drop down menu for choosing a Source Address and yet the CLI will allow any address to be used, not just those associated with *any* interface. This is because by default the policies in the GUI use a function of cross referencing which addresses are allowed based on which interface is involved in the policy. When combined with the aspect of Central NAT that doesn't restrict to a specific interface. This means the only addresses will be allowed are those associated with the *any* interface. The CLI does not have this cross referencing function which is why the CLI seems less restrictive. However, more care must be taken when using the CLI to make sure that appropriate addresses are used.

---

## Firewall Schedule - Recurring

The Company wants to allow the use of Facebook by employees, but only during none business hours and the lunch break.

- The business hours are 9:00 p.m. to 6:00 p.m.
- The Lunch break is 12:00 p.m. to 1:00 p.m.
- The plan is to create a schedule to cover the morning business hours and the afternoon business hours and block access to the Facebook web site during that time.

Go to *Firewall Objects > Schedule > IP Recurring*.

Create a new schedule

Fill out the fields with the following information:

Field	Value
Name	Morning_Business_Hours
Day of the Week	
Sunday	disabled
Monday	enabled
Tuesday	enabled
Wednesday	enabled
Thursday	enabled
Friday	enabled
Saturday	disabled

Start Time: Hour: 09 Minute: 00

Stop Time: Hour: 12 Minute: 00

Select *OK*

Enter the following CLI command:

```
config firewall schedule recurring
edit Morning_Business_Hours
set day monday tuesday wednesday thursday friday
set start 09:00
set end 12:00
end
```

Create a second new schedule.

Field	Value
Name	Afternoon_Business_Hours
Day of the Week	
Sunday	disabled
Monday	enabled
Tuesday	enabled
Wednesday	enabled
Thursday	enabled

Field	Value		
Friday	enabled		
Saturday	disabled		
Start Time: Hour:	13	Minute:	00
Stop Time: Hour:	18	Minute:	00

Select *OK*

Enter the following CLI command:

```
config firewall schedule recurring
edit Afternoon_Business_Hours
set day monday tuesday wednesday thursday friday
set start 13:00
set end 18:00
end
```

To verify that the schedule was added correctly:

Go to *Firewall Objects > Schedule > Recurring*

Check that the schedule with the name you used has been added to the list of recurring schedules and that the listed settings are correct.

Enter the following CLI command:

```
config firewall schedule recurring
edit <the name of the schedule you wish to verify>
show full-configuration
```

## Firewall Schedule - One-time

The company wants to change over their web site image to reference the new year. They have decided to take this opportunity to do some hardware upgrades as well. Their web site is business oriented so they have determined that over New Year's Eve there will be very limited traffic.

- They are going to need a maintenance window of 2 hours bracketing midnight on New Year's Eve.

Go to *Firewall Objects > Schedule > One-time*.

Create a new schedule

Fill out the fields with the following information:

	Year	Month	Day	Hour	Minute
Start	2012	12	31	23	00
Stop	2013	1	1	1	00

Select *OK*

Enter the following CLI command:

```
config firewall schedule onetime
 edit maintenance_window
 set start 23:00 2012/12/31
 set end 01:00 2013/01/01
 next
end
```

To verify that the schedule was added correctly:

Go to *Firewall Objects > Schedule > One-time*

Check that the schedule with the name you used has been added to the list of recurring schedules and that the listed settings are correct.

Enter the following CLI command:

```
config firewall schedule onetime
 edit <the name of the schedule you wish to verify>
 show full-configuration
```

## Schedule Group

In order to make the administration of the policies easier a group needs to be created.

Go to *Firewall Objects > Schedule > Groups*.

Create a new group

Fill out the fields with the following information:

Field	Value
Group Name	Business_Hours
Available Schedules	Morning_Business_Hours Afternoon_Business_Hours

Use arrows to move schedules from the “Members” field to the “Available Schedules:” field

Select *OK*

Enter the following CLI command:

```
config firewall service group
 edit Business_Hours
 set member Morning_Business_Hours Afternoon_Business _ours
 end
```

To verify that the schedule was added correctly:

Go to *Firewall Objects > Schedule > Groups*

Check that the schedule group with the name you used has been added to the list of schedule groups.

Enter the following CLI command:

```
config firewall service group
 edit <the name of the schedule group you wish to verify>
 show full-configuration
```

## Proxy Option

The company will be using a number of the Security Profiles features on various policies but wants to use as few profiles as possible to make administration simpler. The decision has been made to have two profiles, the default one and a single customized one that will be a combination of the settings required to cover the situations that will not be covered by the default profile.

The company profile will have the following parameters:

- There are no FTP servers running on the site so there is no need for FTP.
- The company will use the Fortinet supplied default certificate called “Fortinet\_CA\_SSLProxy”
- The company will only be doing inspecting of SSL over HTTP, SMTP and IMAP.
- The company has a non-standard IMAP implementation the uses ports 1143 and 1993 for IMAP and IMAPS respectively.
- Deep Scanning is to be enabled on any SSL traffic regardless of port and the traffic logged but nothing is blocked.
- The Comfort Clients is to be used with a ratio of 1 byte for every 15 seconds.
- There is a lot of varied email traffic so there is to be no blocking of emails due to size beyond the settings on the mail servers.
- The Security Officer insists that invalid SSL certificates not be allowed.

Go to Policy > Policy > Proxy Options

Create a new profile

Fill out the fields with the following information:

Field	Value
Name	example_standard
Comments	<optional>

### Protocol Port Mapping:

Enable	Protocol	Inspection Ports
enabled	HTTP	Specify and <leave on default setting.>
enabled	SMTP	Specify and <leave on default setting.>
enabled	POP3	Specify and <leave on default setting.>
enabled	IMAP	Specify and 1143
not enabled	FTP	
enabled	NNTP	Specify and <leave on default setting.>

Enable	Protocol	Inspection Ports
enabled	MAPI	<leave on default setting.>
enabled	DNS	<leave on default setting.>

### SSL Inspection Options

CA Certificate: "Fortinet\_CA\_SSLProxy" from drop down menu

Inspect All Ports: Not enabled

Enable	Protocol	Inspection Port(s)
enabled	HTTPS	<leave as default>
enabled	SMTPS	<leave as default>
enabled	POP3S	<leave as default>
enabled	IMAPS	1993
	FTPS	

### SSH Inspection Options

Enable SSH Deep Scan: enabled

Protocol	Inspection Ports
SSH	Any
Exec	Block: not enabled   Log: enabled
Port-Forward	Block: not enabled   Log: enabled
SSH-Shell	Block: not enabled   Log: enabled
x11-Filter	Block: not enabled   Log: enabled

### Common Options

Field	Value
Comfort Clients	enabled
• Interval (Seconds)	15
• Amount(bytes)	1
Block Oversized File/Email	not enabled
• Threshold(MB)	not enabled
Allow Invalid SSL Certificates	not enabled

### Web Options

Field	Value
Enabled Chunked Bypass	not enabled
Add Fortinet Bar	not enabled
• Communication Port	<leave as default>

### Email Options

Field	Value
Allow Fragmented Messages	not enabled
Append Signature (SMTP)	not enabled
Email Signature Text	not enabled



Select *OK*

Enter the following CLI command:

```
config firewall profile-protocol-options
 edit example_standard
 config http
 set options clientcomfort no-content-summary
 set comfort-interval 15
 next
 config https
 set status enable
 set options clientcomfort no-content-summary
 set comfort-interval 15
 next
 config ftp
 set status disable
 set options clientcomfort no-content-summary splice
 set comfort-interval 15
 next
 config ftps
 set options clientcomfort no-content-summary splice
 set comfort-interval 15
 next
 config imap
 set ports "1143"
 set options fragmail no-content-summary
 next
 config imaps
 set ports "1993"
 set status enable
 set options fragmail no-content-summary
 next
 config mapi
 set options fragmail no-content-summary
 next
 config pop3
 set options fragmail no-content-summary
 next
 config pop3s
 set status enable
 set options fragmail no-content-summary
 next
 config smtp
 set options fragmail no-content-summary splice
 next
 config smtps
 set status enable
 set options fragmail no-content-summary splice
 next
```

```

config nntp
 set options no-content-summary splice
 next
config ssh
 set inspect-all enable
 set log x11-filter ssh-shell exec port-forward
 next
end

```

## Oversized Files

A couple of variations on the example could have to do with the processing of oversized files at a level other than the default setting. The ways that it can be approached are:

1. Set a non default threshold size and block the files
2. Set a non default threshold size and not scan the files over the threshold but allow them to pass through the FortiGate firewall.

In the following instructions:

- We will just use 2 MB as the new threshold.
- In the CLI instructions we will limit the configuration to just the HTTP settings for the purposes of brevity and simplicity.

### Option 1

Option 1 can be done in the GUI.

Go to Policy > Policy > Proxy Options

Create a new profile

Fill out the fields with the following information:

#### Common Options

Field	Value
Comfort Clients	enabled
• Interval (Seconds)	15
• Amount(bytes)	1
Block Oversized File/Email	enabled
• Threshold(MB)	2
Allow Invalid SSL Certificates	not enabled

Select *OK*

Enter the following CLI command:

```
config firewall profile-protocol-options
 edit example_standard
 config http
 set options clientcomfort no-content-summary oversize
 set oversize-limit 2
 set comfort-interval 15
 next
 end
```

## Option 2

Option 2 can only be done in the CLI.

Enter the following CLI command:

```
config firewall profile-protocol-options
 edit example_standard
 config http
 set options clientcomfort no-content-summary
 set oversize-limit 2
 set comfort-interval 15
 next
 end
```

## Firewall Address Policy

A policy must be created to all traffic from one of the subnets off of the FortiGate out to the Internet.

- The traffic needs to be NATed
- The Internal subnet that needs to be connected is connected to port 1
- The subnet is from 192.168.1.1 to 192.168.1.255
- There is a user defined address for this subnet named “lan\_192.168.1.x”
- The policy is for general Internet access over wan1.
- The IT manager would like to keep an eye on the kinds of sites that users are going to so for the time being even allowed traffic will be logged.
- To prevent the spread of malware even the outgoing traffic will be screened for viruses using the “basic\_outgoing\_av” profile that has already been defined by the IT Manager.
- There is no schedule restriction on access to the Internet so the schedule is the predefined “always”.
- There is no restrictions yet on what protocols are allowed so the predefined services should be used.
- In order to prevent any HR issues a webfilter has been defined by the IT manager that blocks access to inappropriate sites. The profile is called “basic\_web\_filter”

Go to *Policy > Policy > Policy*.

Create a new policy

Fill out the fields with the following information:

<b>Field</b>	<b>Value</b>
Policy Type	Firewall
Policy Subtype	Address
Incoming Interface	port1
Source Address	lan_192.168.1.x
Outgoing Interface	wan1
Destination Address	all
Schedule	always
Service	ALL
Action	Accept
Enable NAT	enable
Use Destination Interface Address	enable
Log Allowed Traffic	enable
Enable WebCache	not enabled
Enable WAN Optimization	not enabled
Traffic Shaping	enabled

### **Security Profiles**

<b>Security Profiles</b>	<b>Status</b>	<b>Profile Name</b>
AntiVirus	ON	basic_outgoing_av
Web Filter	ON	basic_web_filter
Application Control	OFF	(option should be greyed out)
IPS	OFF	(option should be greyed out)
Email Filter	OFF	(option should be greyed out)
DLP Sensor	OFF	(option should be greyed out)
VoIP	OFF	(option should be greyed out)
ICAP	OFF	(option should be greyed out)

Select *OK*

Enter the following CLI command:

```
config firewall policy
 edit 0
 set srcintf port1
 set dstintf wan1
 set srcaddr lan_192.168.1.x
 set dstaddr "all"
 set action accept
 set schedule always
 set service ALL
 set logtraffic enable
 set nat enable
 set natip 0.0.0.0 0.0.0.0
 set utm-status enable
 set av-profile basic_outgoing_av
 set webfilter-profile basic_web_filter
 next
end
```

## Firewall User Identity Policy

There is a Server on the Internal LAN that is used as a resource by sales people that are out on the road and at home. The sales people and their administrator are the only ones that are allowed access to the server "SalesServer". Because of the importance of the data on the server it needs to be backed up on a regular bases and to make sure that all of the files are available for backup no one is allowed access during the scheduled backup time.

- Firewall User Group has been created: "sales". This is a group defined for more than one purpose.
- Firewall User has been created: "jsmith", who is the administrator. "jsmith" is not a member of the sales group and can't be placed in it because of other resources that the group has access to.
- A Virtual IP has been defined for the server "SalesServer" mapped from an external address to an internal address on port9. Port forwarding is enabled on TCP port 443 for the application that the sales people use.
- The schedule "Non-Maintenance\_Window" has been created (between 2 a.m. and 3 a.m.) by starting the window at 3:00 and stopping at 2:00 (If the stop time is set earlier than the start time, the stop time will be during the next day).
- Because there is access from the Internet to the internal network the company's HTTPS Server IPS profile, called "Protect\_HTTPS\_Server" will be enabled.

Go to *Policy > Policy > Policy*.

Create a new policy

Fill out the fields with the following information:

Field	Value
Policy Type	Firewall
Policy Subtype	User Identity

<b>Field</b>	<b>Value</b>
Incoming Interface	wan1
Source Address	all
Outgoing Interface	port9
Enable NAT	not enabled
Enable Web cache	not enabled
Enable WAN Optimization	not enabled
Certificate	No Certificate
Firewall	enabled
Fortinet Single Sign-On	not enabled
WiFi Single Sign-on	not enabled
NTLM Authentications	not enabled
Customize Authentication Messages	(optional)
Add tag	(optional)
Comments	(optional)

In the Configure Authentication Rules section, create a new Authentication rule by filling out the fields with the following information.

<b>Field</b>	<b>Value</b>
Destination Address	SalesServer (virtual IP address)
Group(s)	sales
User(s)	jsmith
Schedule	Non-Maintenance_Window
Service	HTTPS
Action	(automatically assigned value of "ACCEPT")
Log Allowed Traffic	(optional)
Traffic Shaping	(optional)

In the Security Profiles Section:

Enable and select the security profiles as follows:

Security Profiles	Status	Profile Name
AntiVirus	OFF	(option should be greyed out)
Web Filter	OFF	(option should be greyed out)
Application Control	OFF	(option should be greyed out)
IPS	ON	Protect_HTTPS_Server
Email Filter	OFF	(option should be greyed out)
DLP Sensor	OFF	(option should be greyed out)
VoIP	OFF	(option should be greyed out)
ICAP	OFF	(option should be greyed out)

Select *OK*

Enter the following CLI command:

```
config firewall policy
 edit 0
 set srcintf "wan1"
 set dstintf "port9"
 set srcaddr "all"
 set action accept
 set identity-based enable
 config identity-based-policy
 edit 1
 set schedule "Non-Maintenance_Window"
 set utm-status enable
 set profile-type single
 set groups "sales"
 set users "jsmith"
 set dstaddr "SalesServer"
 set service "HTTPS"
 set ips-sensor "Protect_HTTPS_Server"
 next
 end
```

## Firewall Device Identity Policy

The company is instituting a BYOD pilot project. They are going to be letting employee's personal devices connect through a wireless network. Most will only be allowed only to the Internet but some will have access to the Internal LAN.

- The wireless interface the users will be connecting on is "WiFi".
- The source address has been defined as the DHCP scope assigned to the wireless network, "Internal\_wireless".
- The Internal Network is on the interface designated as "LAN"
- There is no need to NAT the traffic.
- The AntiVirus Profile to be used is "Internal-AV", already defined.
- The Application Profile to be used is "Internal-AC", already defined.
- The IPS Profile to be used is "Internal-IPS", already defined.
- The Device group for this policy is "IT\_Personnel\_phones".
- The IT team is a trusted group that will be accessing practically everything so the schedule will be the predefined "always" and the Service will be the predefined "ALL".
- While the IT team is trusted the company would like to verify that they are compliant with the Endpoint profile so the check for this must be enabled on the policy.

Go to *Policy > Policy > Policy*.

Create a new policy

Fill out the fields with the following information:

Field	Value
Policy Type	Firewall
Policy Subtype	Device Identity
Incoming Interface	wifi
Source Address	all
Outgoing Interface	port9
Enable NAT	not enabled
Customize Authentication Messages	(optional)
Add tag	(optional)
Comments	(optional)



In the Configure Authentication Rules section, create a new Authentication rule by filling out the fields with the following information.

<b>Field</b>	<b>Value</b>
Destination Address	Internal_Network
Device	IT_Personnel_Phones
Compliant with Endpoint Profile	enable
Schedule	always
Service	ALL
Action	(automatically assigned value of "ACCEPT")
Log Allowed Traffic	(optional)
Traffic Shaping	(optional)

In the Security Profile Section:

Enable and select the security profiles as follows:

<b>Security Profiles</b>	<b>Status</b>	<b>Profile Name</b>
AntiVirus	ON	Internal-AV
Web Filter	OFF	(option should be greyed out)
Application Control	ON	Internal-AC
IPS	ON	Internal-IPS
Email Filter	OFF	(option should be greyed out)
DLP Sensor	OFF	(option should be greyed out)
VoIP	OFF	(option should be greyed out)
ICAP	OFF	(option should be greyed out)

Select *OK*

Enter the following CLI command:

```
config firewall policy
 edit 0
 set srcintf "wifi"
 set dstintf "LAN"
 set srcaddr "Internal_wireless"
 set action accept
 set identity-based enable
 set identity-from device
 config identity-based-policy
 edit 1
 set schedule always
 set utm-status enable
 set dstaddr Internal_Network
 set service ALL
 set devices IT_Personnel_phones
 set endpoint-compliance enable
 set av-profile Internal-AV
 set ips-sensor Internal-IPS
 set application-list Internal-AC
 next
 end
```

## DoS Policy

The company wishes to protect against Denial of Service attacks. They have chosen some where they wish to block the attacks if the incidence goes above a certain threshold and for some others they are just trying to get a baseline of activity for those types of attacks so they are letting the traffic pass through without action.

- The interface to the Internet is on WAN1
- There is no requirement to specify which addresses are being protected or protected from.
- The protection is to extend to all services.
- The TCP attacks are to be blocked
- The UDP, ICMP, and IP attacks are to be recorded but not blocked.
- The tcp\_syn\_flood attack's threshold is to be changed from the default to 1000

Go to *Policy > Policy > DoS Policy*.

Create a new policy

Fill out the fields with the following information:

Field	Value
Incoming Interface	wan1
Source Address	all

Field	Value
Destination Addresses	all
Service	ALL

### Anomalies

Name	Status	Logging	Action	Threshold
tcp_syn_flood	enabled	enabled	Block	1000
tcp_port_scan	enabled	enabled	Block	<default value>
tcp_src_session	enabled	enabled	Block	<default value>
tcp_dst_session	enabled	enabled	Block	<default value>
udp_flood	enabled	enabled	Pass	<default value>
udp_scan	enabled	enabled	Pass	<default value>
udp_src_session	enabled	enabled	Pass	<default value>
udp_dst_session	enabled	enabled	Pass	<default value>
icmp_flood	enabled	enabled	Pass	<default value>
icmp_sweep	enabled	enabled	Pass	<default value>
icmp_src_session	enabled	enabled	Pass	<default value>
icmp_dst_session	enabled	enabled	Pass	<default value>
ip_src_session	enabled	enabled	Pass	<default value>
ip_dst_session	enabled	enabled	Pass	<default value>
sctp_flood	not enabled	not enabled	Pass	<default value>
sctp_scan	not enabled	not enabled	Pass	<default value>
sctp_src_session	not enabled	not enabled	Pass	<default value>
sctp_dst_session	not enabled	not enabled	Pass	<default value>

Select OK

Enter the following CLI command:

```
config firewall DoS-policy
edit 0
set status enable
set interface ''
config anomaly
edit "tcp_syn_flood"
set status enable
set log enable
set action block
set threshold 1000
next
edit "tcp_port_scan"
set status enable
set log enable
set action block
next
edit "tcp_src_session"
set status enable
set log enable
set action block
next
edit "tcp_dst_session"
set status enable
set log enable
set action block
next
edit "udp_flood"
set status enable
set log enable
next
edit "udp_scan"
set status disable
set status enable
set log enable
next
edit "udp_src_session"
set status enable
set log enable
next
edit "udp_dst_session"
set status enable
set log enable
next
edit "icmp_flood"
set status enable
set log enable
next
```

```
edit "icmp_sweep"
 set status enable
 set log enable
next
edit "icmp_src_session"
 set status enable
 set log enable
next
edit "icmp_dst_session"
 set status enable
 set log enable
next
edit "ip_src_session"
 set status enable
 set log enable
next
edit "ip_dst_session"
 set status enable
 set log enable
next
end
next
end
```

# Multicast forwarding

Multicasting (also called IP multicasting) consists of using a single multicast source to send data to many receivers. Multicasting can be used to send data to many receivers simultaneously while conserving bandwidth and reducing network traffic. Multicasting can be used for one-way delivery of media streams to multiple receivers and for one-way data transmission for news feeds, financial information, and so on.

Also RIPv2 uses multicasting to share routing table information, OSPF uses multicasting to send hello packets and routing updates, Enhanced Interior Gateway Routing Protocol (EIGRP) uses multicasting to send routing information to all EIGRP routers on a network segment and the Bonjour network service uses multicasting for DNS.

A FortiGate unit can operate as a Protocol Independent Multicast (PIM) version 2 router. FortiGate units support PIM sparse mode (RFC 4601) and PIM dense mode (RFC 3973) and can service multicast servers or receivers on the network segment to which a FortiGate unit interface is connected. Multicast routing is not supported in transparent mode (TP mode).



To support PIM communications, the sending/receiving applications and all connecting PIM routers in between must be enabled with PIM version 2. PIM can use static routes, RIP, OSPF, or BGP to forward multicast packets to their destinations. To enable source-to-destination packet delivery, either sparse mode or dense mode must be enabled on the PIM-router interfaces. Sparse mode routers cannot send multicast messages to dense mode routers. In addition, if a FortiGate unit is located between a source and a PIM router, two PIM routers, or is connected directly to a receiver, you must create a security policy manually to pass encapsulated (multicast) packets or decapsulated data (IP traffic) between the source and destination.

---

A PIM domain is a logical area comprising a number of contiguous networks. The domain contains at least one Boot Strap Router (BSR), and if sparse mode is enabled, a number of Rendezvous Points (RPs) and Designated Routers (DRs). When PIM is enabled on a FortiGate unit, the FortiGate unit can perform any of these functions at any time as configured.

## Sparse mode

Initially, all candidate BSRs in a PIM domain exchange bootstrap messages to select one BSR to which each RP sends the multicast address or addresses of the multicast group(s) that it can service. The selected BSR chooses one RP per multicast group and makes this information available to all of the PIM routers in the domain through bootstrap messages. PIM routers use the information to build packet distribution trees, which map each multicast group to a specific RP. Packet distribution trees may also contain information about the sources and receivers associated with particular multicast groups.



When a FortiGate unit interface is configured as a multicast interface, sparse mode is enabled on it by default to ensure that distribution trees are not built unless at least one downstream receiver requests multicast traffic from a specific source. If the sources of multicast traffic and their receivers are close to each other and the PIM domain contains a dense population of active receivers, you may choose to enable dense mode throughout the PIM domain instead.

---

An RP represents the root of a non-source-specific distribution tree to a multicast group. By joining and pruning the information contained in distribution trees, a single stream of multicast

packets (for example, a video feed) originating from the source can be forwarded to a certain RP to reach a multicast destination.

Each PIM router maintains a Multicast Routing Information Base (MRIB) that determines to which neighboring PIM router join and prune messages are sent. An MRIB contains reverse-path information that reveals the path of a multicast packet from its source to the PIM router that maintains the MRIB.

To send multicast traffic, a server application sends IP traffic to a multicast group address. The locally elected DR registers the sender with the RP that is associated with the target multicast group. The RP uses its MRIB to forward a single stream of IP packets from the source to the members of the multicast group. The IP packets are replicated only when necessary to distribute the data to branches of the RP's distribution tree.

To receive multicast traffic, a client application can use Internet Group Management Protocol (IGMP) version 1 (RFC 1112), 2 (RFC 2236), or 3 (RFC 3376) control messages to request the traffic for a particular multicast group. The locally elected DR receives the request and adds the host to the multicast group that is associated with the connected network segment by sending a join message towards the RP for the group. Afterward, the DR queries the hosts on the connected network segment continually to determine whether the hosts are active. When the DR no longer receives confirmation that at least one member of the multicast group is still active, the DR sends a prune message towards the RP for the group.

FortiOS supports PIM sparse mode multicast routing for IPv6 multicast (multicast6) traffic and is compliant with RFC 4601: Protocol Independent Multicast - Sparse Mode (PIM-SM). You can use the following command to configure IPv6 PIM sparse multicast routing.

```
config router multicast6
 set multicast-routing {enable | disable}
 config interface
 edit <interface-name>
 set hello-interval <1-65535 seconds>
 set hello-holdtime <1-65535 seconds>
 end
 config pim-sm-global
 config rp-address
 edit <index>
 set ipv6-address <ipv6-address>
 end
```

The following diagnose commands for IPv6 PIM sparse mode are also available:

```
diagnose ipv6 multicast status
diagnose ipv6 multicast vif
diagnose ipv6 multicast mroute
```

## Dense mode

The packet organization used in sparse mode is also used in dense mode. When a multicast source begins to send IP traffic and dense mode is enabled, the closest PIM router registers the IP traffic from the multicast source (S) and forwards multicast packets to the multicast group address (G). All PIM routers initially broadcast the multicast packets throughout the PIM domain to ensure that all receivers that have requested traffic for multicast group address G can access the information if needed.

To forward multicast packets to specific destinations afterward, the PIM routers build distribution trees based on the information in multicast packets. Upstream PIM routers depend

on prune/graft messages from downstream PIM routers to determine if receivers are actually present on directly connected network segments. The PIM routers exchange state refresh messages to update their distribution trees. FortiGate units store this state information in a Tree Information Base (TIB), which is used to build a multicast forwarding table. The information in the multicast forwarding table determines whether packets are forwarded downstream. The forwarding table is updated whenever the TIB is modified.

PIM routers receive data streams every few minutes and update their forwarding tables using the source (S) and multicast group (G) information in the data stream. Superfluous multicast traffic is stopped by PIM routers that do not have downstream receivers—PIM routers that do not manage multicast groups send prune messages to the upstream PIM routers. When a receiver requests traffic for multicast address G, the closest PIM router sends a graft message upstream to begin receiving multicast packets.

FortiGate units operating in NAT mode can also be configured as multicast routers. You can configure a FortiGate unit to be a Protocol Independent Multicast (PIM) router operating in Sparse Mode (SM) or Dense Mode (DM).

## Multicast IP addresses

Multicast uses the Class D address space. The 224.0.0.0 to 239.255.255.255 IP address range is reserved for multicast groups. The multicast address range applies to multicast groups, not to the originators of multicast packets. [Table 53](#) lists reserved multicast address ranges and describes what they are reserved for:

**Table 53:** Reserved Multicast address ranges

Reserved Address Range	Use	Notes
224.0.0.0 to 224.0.0.255	Used for network protocols on local networks. For more information, see RFC 1700.	In this range, packets are not forwarded by the router but remain on the local network. They have a Time to Live (TTL) of 1. These addresses are used for communicating routing information.
224.0.1.0 to 238.255.255.255	Global addresses used for multicasting data between organizations and across the Internet. For more information, see RFC 1700.	Some of these addresses are reserved, for example, 224.0.1.1 is used for Network Time Protocol (NTP).
239.0.0.0 to 239.255.255.255	Limited scope addresses used for local groups and organizations. For more information, see RFC 2365.	Routers are configured with filters to prevent multicasts to these addresses from leaving the local system.

Creating multicast security policies requires multicast firewall addresses. You can add multicast firewall addresses by going to *Firewall Objects > Address > Addresses* and selecting *Create New > Multicast Address*. The factory default configuration includes multicast addresses for Bonjour(224.0.0.251-224.0.0.251, EIGRP (224.0.0.10-224.0.0.100), OSPF (224.0.0.5-224.0.0.60), all\_hosts (224.0.0.1-224.0.0.1), and all\_routers (224.0.0.2-224.0.0.2).



## PIM Support

A FortiGate unit can be configured to support PIM by going to *Router > Dynamic > Multicast* and enabling multicast routing. You can also enable multicast routing using the `config router multicast` CLI command. When PIM is enabled, the FortiGate unit allocates memory to manage mapping information. The FortiGate unit communicates with neighboring PIM routers to acquire mapping information and if required, processes the multicast traffic associated with specific multicast groups.



The end-user multicast client-server applications must be installed and configured to initiate Internet connections and handle broadband content such as audio/video information.

Client applications send multicast data by registering IP traffic with a PIM-enabled router. An end-user could type in a class D multicast group address, an alias for the multicast group address, or a call-conference number to initiate the session.

Rather than sending multiple copies of generated IP traffic to more than one specific IP destination address, PIM-enabled routers encapsulate the data and use the one multicast group address to forward multicast packets to multiple destinations. Because one destination address is used, a single stream of data can be sent. Client applications receive multicast data by requesting that the traffic destined for a certain multicast group address be delivered to them — end-users may use phone books, a menu of ongoing or future sessions, or some other method through a user interface to select the address of interest.

A class D address in the 224.0.0.0 to 239.255.255.255 range may be used as a multicast group address, subject to the rules assigned by the Internet Assigned Numbers Authority (IANA). All class D addresses must be assigned in advance. Because there is no way to determine in advance if a certain multicast group address is in use, collisions may occur (to resolve this problem, end-users may switch to a different multicast address).

### To configure a PIM domain

1. If you will be using sparse mode, determine appropriate paths for multicast packets.
2. Make a note of the interfaces that will be PIM-enabled. These interfaces may run a unicast routing protocol.
3. If you will be using sparse mode and want multicast packets to be handled by specific (static) RPs, record the IP addresses of the PIM-enabled interfaces on those RPs.
4. Enable PIM version 2 on all participating routers between the source and receivers. On FortiGate units, use the `config router multicast` command to set global operating parameters.
5. Configure the PIM routers that have good connections throughout the PIM domain to be candidate BSRs.
6. If sparse mode is enabled, configure one or more of the PIM routers to be candidate RPs.
7. If required, adjust the default settings of PIM-enabled interface(s).

## Multicast forwarding and FortiGate units

In both transparent mode and NAT mode you can configure FortiGate units to forward multicast traffic.

For a FortiGate unit to forward multicast traffic you must add FortiGate multicast security policies. Basic multicast security policies accept any multicast packets at one FortiGate

interface and forward the packets out another FortiGate interface. You can also use multicast security policies to be selective about the multicast traffic that is accepted based on source and destination address, and to perform NAT on multicast packets.

In the example shown in [Figure 176](#), a multicast source on the Marketing network with IP address 192.168.5.18 sends multicast packets to the members of network 239.168.4.0. At the FortiGate unit, the source IP address for multicast packets originating from workstation 192.168.5.18 is translated to 192.168.18.10. In this example, the FortiGate unit is not acting as a multicast router.

## Multicast forwarding and RIPv2


RIPv2 uses multicast to share routing table information. If your FortiGate unit is installed on a network that includes RIPv2 routers, you must configure the FortiGate unit to forward multicast packets so that RIPv2 devices can share routing data through the FortiGate unit. No special FortiGate configuration is required to share RIPv2 data, you can simply use the information in the following sections to configure the FortiGate unit to forward multicast packets.



RIPv1 uses broadcasting to share routing table information. To allow RIPv1 packets through a FortiGate unit you can add standard security policies. Security policies to accept RIPv1 packets can use the ANY predefined firewall service or the RIP predefined firewall service.

---

**Figure 176:**Example multicast network including a FortiGate unit that forwards multicast packets



## Configuring FortiGate multicast forwarding

You configure FortiGate multicast forwarding from the Command Line Interface (CLI). Two steps are required:

- [Adding multicast security policies](#)
- [Enabling multicast forwarding](#)

This second step is only required if your FortiGate unit is operating in NAT mode. If your FortiGate unit is operating in transparent mode, adding a multicast policy enables multicast forwarding.



There is sometimes a confusion between the terms “forwarding” and “routing”. These two functions should not be taking place at the same time.

It is mentioned that multicast-forward should be enabled when the FortiGate unit is in NAT mode and that this will forward any multicast packet to all interfaces. However, this parameter should **NOT** be enabled when the FortiGate unit operates as a multicast router (i.e. with a routing protocol enabled). It should only be enabled when there is no routing protocols activated.

## Adding multicast security policies

You need to add security policies to allow packets to pass from one interface to another. Multicast packets require multicast security policies. You add multicast security policies from the CLI using the `config firewall multicast-policy` command. As with unicast security policies, you specify the source and destination interfaces and optionally the allowed address ranges for the source and destination addresses of the packets.

You can also use multicast security policies to configure source NAT and destination NAT for multicast packets.

Keep the following in mind when configuring multicast security policies:

- The matched forwarded (outgoing) IP multicast source IP address is changed to the configured IP address.
- Source and Destination interfaces are optional. If left blank, then the multicast will be forwarded to ALL interfaces.
- Source and Destination addresses are optional. If left un set, then it will mean ALL addresses.
- The `nat` keyword is optional. Use it when source address translation is needed.

## Enabling multicast forwarding

Multicast forwarding is enabled by default. In NAT mode you must use the `multicast-forward` keyword of the `system settings` CLI command to enable or disable multicast forwarding. When `multicast-forward` is enabled, the FortiGate unit forwards any multicast IP packets in which the TTL is 2 or higher to all interfaces and VLAN interfaces except the receiving interface. The TTL in the IP header will be reduced by 1. Even though the multicast packets are forwarded to all interfaces, you must add security policies to actually allow multicast packets through the FortiGate. In our example, the security policy allows multicast packets received by the internal interface to exit to the external interface.



Enabling multicast forwarding is only required if your FortiGate unit is operating in NAT mode. If your FortiGate unit is operating in transparent mode, adding a multicast policy enables multicast forwarding.

---

Enter the following CLI command to enable multicast forwarding:

```
config system settings
 set multicast-forward enable
end
```

If multicast forwarding is disabled and the FortiGate unit drops packets that have multicast source or destination addresses.

You can also use the `multicast-ttl-notchange` keyword of the `system settings` command so that the FortiGate unit does not increase the TTL value for forwarded multicast packets. You should use this option only if packets are expiring before reaching the multicast router.

```
config system settings
 set multicast-ttl-notchange enable
end
```

In transparent mode, the FortiGate unit does not forward frames with multicast destination addresses. Multicast traffic such as the one used by routing protocols or streaming media may need to traverse the FortiGate unit, and should not be interfere with the communication. To

avoid any issues during transmission, you can set up multicast security policies. These types of security policies can only be enabled using the CLI.



The CLI parameter `multicast-skip-policy` must be disabled when using multicast security policies. To disable enter the command

```
config system settings
 set multicast-skip-policy disable
end
```

In this simple example, no check is performed on the source or destination interfaces. A multicast packet received on an interface is flooded unconditionally to all interfaces on the forwarding domain, except the incoming interface.

#### To enable the multicast policy

```
config firewall multicast-policy
 edit 1
 set action accept
 end
```

In this example, the multicast policy only applies to the source port of WAN1 and the destination port of Internal.

#### To enable the restrictive multicast policy

```
config firewall multicast-policy
 edit 1
 set srcintf wan1
 set dstintf internal
 set action accept
 end
```

In this example, packets are allowed to flow from WAN1 to Internal, and sourced by the address 172.20.120.129. This address needs to be configured as an address object. For simplicity we will assume that this address is represented by the address object "example-addr\_A".

#### To enable the restrictive multicast policy

```
config firewall multicast-policy
 edit 1
 set srcintf wan1
 set srcaddr example-addr_A
 set dstintf internal
 set action accept
 end
```

This example shows how to configure the multicast security policy required for the configuration shown. This policy accepts multicast packets that are sent from a PC with IP address 192.168.5.18, which will be represented by the address object "example-addr\_X", to destination address range 239.168.4.0, which will be represented by the address object "example-addr\_Y". The policy allows the multicast packets to enter the internal interface and then exit the external

interface. When the packets leave the external interface their source address is translated to 192.168.18.10.

```
config firewall multicast-policy
 edit 5
 set srcaddr example-addr_X
 set srcintf internal
 set destaddr example-addr_Y
 set dstintf external
 set nat 192.168.18.10
 end
```

This example shows how to configure a multicast security policy so that the FortiGate unit forwards multicast packets from a multicast Server with an IP 10.10.10.10, represented by address object "example-addr\_C", is broadcasting to address 225.1.1.1, represented by address object "example-addr\_D". This Server is on the network connected to the FortiGate DMZ interface.

```
config firewall multicast-policy
 edit 1
 set srcintf DMZ
 set srcaddr example-addr_C
 set dstintf Internal
 set destaddr example-addr_D
 set action accept
 edit 2
 set action deny
 end
```

## Multicast routing examples

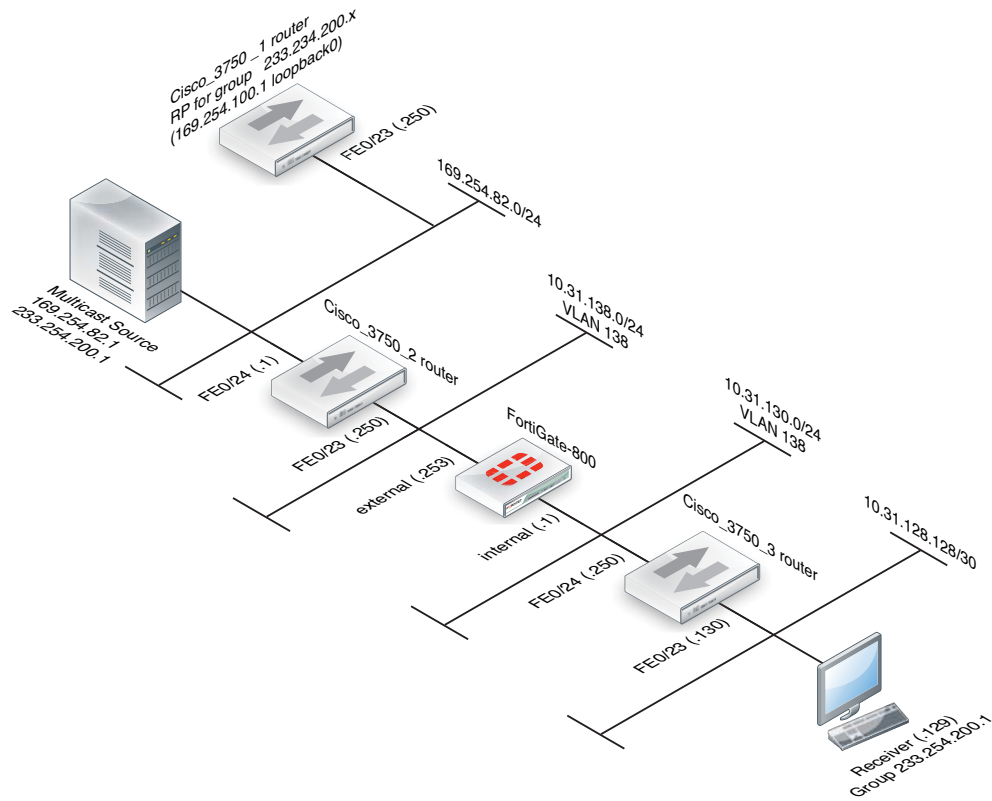
This section contains the following multicast routing configuration examples and information:

- [Example FortiGate PIM-SM configuration using a static RP](#)
- [FortiGate PIM-SM debugging examples](#)
- [Example multicast destination NAT \(DNAT\) configuration](#)
- [Example PIM configuration that uses BSR to find the RP](#)

### Example FortiGate PIM-SM configuration using a static RP

The example Protocol Independent Multicast Sparse Mode (PIM-SM) configuration shown in [Figure 177](#) has been tested for multicast interoperability using PIM-SM between Cisco 3750 switches running 12.2 and a FortiGate-800 running FortiOS v3.0 MR5 patch 1. In this configuration, the receiver receives the multicast stream when it joins the group 233.254.200.1.

**Figure 177:**Example FortiGate PIM-SM topology



The configuration uses a statically configured rendezvous point (RP) which resides on the Cisco\_3750\_1. Using a bootstrap router (BSR) was not tested in this example. See [“Example PIM configuration that uses BSR to find the RP”](#) on page 1055 for an example that uses a BSR.

### Configuration steps

The following procedures show how to configure the multicast configuration settings for the devices in the example configuration.

- [Cisco\\_3750\\_1 router configuration](#)
- [Cisco\\_3750\\_2 router configuration](#)
- [To configure the FortiGate-800 unit](#)
- [Cisco\\_3750\\_3 router configuration](#)

## Cisco\_3750\_1 router configuration

```
version 12.2
!
hostname Cisco-3750-1
!
switch 1 provision ws-c3750-24ts
ip subnet-zero
ip routing
!
ip multicast-routing distributed
!
spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
interface Loopback0
 ip address 169.254.100.1 255.255.255.255
!
interface FastEthernet1/0/23
 switchport access vlan 182
 switchport mode access
!
interface FastEthernet1/0/24
 switchport access vlan 172
 switchport mode access
!
interface Vlan172
 ip address 10.31.138.1 255.255.255.0
 ip pim sparse-mode
 ip igmp query-interval 125
 ip mroute-cache distributed
!
interface Vlan182
 ip address 169.254.82.250 255.255.255.0
 ip pim sparse-mode
 ip mroute-cache distributed
!
ip classless
ip route 0.0.0.0 0.0.0.0 169.254.82.1
ip http server
ip pim rp-address 169.254.100.1 Source-RP
!
ip access-list standard Source-RP
 permit 233.254.200.0 0.0.0.255
```



## Cisco\_3750\_2 router configuration

```
version 12.2
!
hostname Cisco-3750-2
!
switch 1 provision ws-c3750-24ts
ip subnet-zero
ip routing
!
ip multicast-routing distributed
!
spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
interface FastEthernet1/0/23
 switchport access vlan 138
 switchport mode access
!
interface FastEthernet1/0/24
 switchport access vlan 182
 witchport mode access
!
interface Vlan138
 ip address 10.31.138.250 255.255.255.0
 ip pim sparse-mode
 ip mroute-cache distributed
!
interface Vlan182
 ip address 169.254.82.1 255.255.255.0
 ip pim sparse-mode
 ip mroute-cache distributed
!
ip classless
ip route 0.0.0.0 0.0.0.0 10.31.138.253
ip route 169.254.100.1 255.255.255.255 169.254.82.250
ip http server
ip pim rp-address 169.254.100.1 Source-RP
!
!
ip access-list standard Source-RP
permit 233.254.200.0 0.0.0.255
```

## To configure the FortiGate-800 unit

### 1. Configure the internal and external interfaces.

- **Internal**

Go to *System > Network > Interfaces*.

Select the internal interface

Verify the following settings:

---

<b>Type:</b>	Physical Interface
<b>Addressing mode:</b>	Manual
<b>IP/Network Mask:</b>	10.31.138.253 255.255.255.0
<b>Administrative Access:</b>	PING

---

Select *OK*

- **External**

Go to *System > Network > Interfaces*.

Select the external interface

Verify the following settings:

---

<b>Type:</b>	Physical Interface
<b>Addressing mode:</b>	Manual
<b>IP/Network Mask:</b>	10.31.130.253 255.255.255.0
<b>Administrative Access:</b>	HTTPS and PING

---

Select *OK*.

2. Add a firewall addresses.

Go to *Firewall Objects > Address > Addresses*.

- RP

Select *Create New*.

Use the following settings:

---

<b>Category:</b>	Address
<b>Name:</b>	RP
<b>Type:</b>	Subnet
<b>Subnet/IP Range:</b>	169.254.100.1/32
<b>Interface:</b>	Any

---

Select *OK*.

- Multicast source subnet

Select *Create New*.

Use the following settings:

---

<b>Category:</b>	Address
<b>Name:</b>	multicast_source_subnet
<b>Type:</b>	Subnet
<b>Subnet/IP Range:</b>	169.254.82.0/24
<b>Interface:</b>	Any

---

Select *OK*.

3. Add destination multicast address

Go to *Firewall Objects > Address > Addresses*.

Select *Create New*.

Use the following settings:

---

<b>Category:</b>	Multicast Address
<b>Name:</b>	Multicast_stream
<b>Type:</b>	Broadcast Subnet
<b>Broadcast Subnet:</b>	233.254.200.0/24
<b>Interface:</b>	Any

---

Select *OK*.

4. Add standard security policies to allow traffic to reach the RP.

Go to Policy > Policy > Policy

- 1st policy

Select *Create New*

Use the following settings:

<b>Policy Type:</b>	Firewall
<b>Policy Subtype:</b>	Address
<b>Incoming Interfac:</b>	internal
<b>Source Address:</b>	all
<b>Outgoing Interface:</b>	external
<b>Destination Address:</b>	RP
<b>Schedule:</b>	always
<b>Service:</b>	ALL
<b>Action:</b>	ACCEPT

Select *OK*.

- 2nd policy

Select *Create New*

Use the following settings:

<b>Policy Type:</b>	Firewall
<b>Policy Subtype:</b>	Address
<b>Incoming Interfac:</b>	external
<b>Source Address:</b>	RP
<b>Outgoing Interface:</b>	internal
<b>Destination Address:</b>	all
<b>Schedule:</b>	always
<b>Service:</b>	ALL
<b>Action:</b>	ACCEPT

Select *OK*.

5. Add the multicast security policy.  
Go to Policy > Policy > Multicast Policy  
Select *Create New*  
Use the following settings:

<b>Incoming Interface:</b>	external
<b>Source Address:</b>	multicast_source_subnet
<b>Outgoing Interface:</b>	internal
<b>Destination Address:</b>	multicast_stream
<b>Protocol:</b>	Any
<b>Action:</b>	ACCEPT

Select *OK*.

6. Add an access list. (CLI only)

```
config router access-list
 edit Source-RP
 config rule
 edit 1
 set prefix 233.254.200.0 255.255.255.0
 set exact-match disable
 next
 end
```

7. Add some static routes.

Go to *Router > Static > Static Routes*.

- Route 1

Select *Create New*.

Use the following settings:

<b>Destination IP/Mask:</b>	0.0.0.0/0.0.0.0
<b>Device:</b>	internal
<b>Gateway:</b>	10.31.130.250

Select *OK*.

- Route 2

Select *Create New*.

Use the following settings:

<b>Destination IP/Mask:</b>	169.254.0.0/16
<b>Device:</b>	external
<b>Gateway:</b>	10.31.138.250

Select *OK*.

8. Configure multicast routing.

Go to *Router > Dynamic > Multicast*.

Add the following Static Rendezvous Point(s):

- 169.254.100.1
- Route 1

Select *Create New*.

Use the following settings:

---

<b>Interface:</b>	internal
<b>PIM Mode:</b>	Sparse Mode
<b>DR Priority:</b>	<not needed in this scenario>
<b>RP Candidate:</b>	<not needed in this scenario>
<b>RP Candidate Priority:</b>	<not needed in this scenario>

---

Select *OK*.

- Route 2

Select *Create New*.

Use the following settings:

---

<b>Interface:</b>	external
<b>PIM Mode:</b>	Sparse Mode
<b>DR Priority:</b>	
<b>RP Candidate:</b>	
<b>RP Candidate Priority:</b>	

---

Select *OK*.


### Cisco\_3750\_3 router configuration

```
version 12.2
!
hostname Cisco-3750-3
!
switch 1 provision ws-c3750-24ts
ip subnet-zero
ip routing
!
ip multicast-routing distributed
!
spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
interface FastEthernet1/0/23
 switchport access vlan 128
 switchport mode access
!
interface FastEthernet1/0/24
 switchport access vlan 130
 switchport mode access
!
interface Vlan128
 ip address 10.31.128.130 255.255.255.252
 ip pim sparse-mode
 ip mroute-cache distributed
!
interface Vlan130
 ip address 10.31.130.250 255.255.255.0
 ip pim sparse-mode
 ip mroute-cache distributed
!
ip classless
ip route 0.0.0.0 0.0.0.0 10.31.130.1
ip http server
ip pim rp-address 169.254.100.1 Source-RP
!
!
ip access-list standard Source-RP
 permit 233.254.200.0 0.0.0.255
```

### FortiGate PIM-SM debugging examples

Using the example topology shown in [Figure 178](#) you can trace the multicast streams and states within the three FortiGate units (FGT-1, FGT-2, and FGT-3) using the debug commands described in this section. The command output in this section is taken from FortiGate unit when the multicast stream is flowing correctly from source to receiver.

**Figure 178:**PIM-SM debugging topology



### Checking that the receiver has joined the required group

From the last hop router, FGT-3, you can use the following command to check that the receiver has correctly joined the required group.

```
FGT-3 # get router info multicast igmp groups
IGMP Connected Group Membership
Group Address Interface Uptime Expires Last Reporter
239.255.255.1 port3 00:31:15 00:04:02 10.167.0.62
```

Only 1 receiver is displayed for a particular group, this is the device that responded to the IGMP query request from the FGT-3. If a receiver is active the expire time should drop to approximately 2 minutes before being refreshed.

### Checking the PIM-SM neighbors

Next the PIM-SM neighbors should be checked. A PIM router becomes a neighbor when the PIM router receives a PIM hello. Use the following command to display the PIM-SM neighbors of FGT-3.

```
FGT-3 # get router info multicast pim sparse-mode neighbour
Neighbor Interface Uptime/Expires Ver DR
Address Priority/Mode
10.132.0.156 port2 01:57:12/00:01:33 v2 1 /
```



## Checking that the PIM router can reach the RP

The rendezvous point (RP) must be reachable for the PIM router (FGT-3) to be able to send the \*,G join to request the stream. This can be checked for FGT-3 using the following command:

```
FGT-3 # get router info multicast pim sparse-mode rp-mapping
PIM Group-to-RP Mappings
Group(s): 224.0.0.0/4, Static
RP: 192.168.1.1
Uptime: 07:23:00
```

## Viewing the multicast routing table (FGT-3)

The FGT-3 unicast routing table can be used to determine the path taken to reach the RP at 192.168.1.1. You can then check the stream state entries using the following commands:

```
FGT-3 # get router info multicast pim sparse-mode table
IP Multicast Routing Table

(*,*,RP) Entries: 0
(*,G) Entries: 1
(S,G) Entries: 1
(S,G,rpt) Entries: 1
FCR Entries: 0
```

(*,*,RP) Entries	This state may be reached by general joins for all groups served by a specified RP.
(*,G) Entries	State that maintains the RP tree for a given group.
(S,G) Entries	State that maintains a source-specific tree for source S and group G.
(S,G,rpt) Entries	State that maintains source-specific information about source S on the RP tree for G. For example, if a source is being received on the source-specific tree, it will normally have been pruned off the RP tree.
FCR	The FCR state entries are for tracking the sources in the <*, G> when <S, G> is not available for any reason, the stream would typically be flowing when this state exists.

Breaking down each entry in detail:

```
(*, 239.255.255.1)
RP: 192.168.1.1
RPF nbr: 10.132.0.156
RPF idx: port2
Upstream State: JOINED
Local:
 port3
Joined:
Asserted:
FCR:
```

The RP will always be listed in a \*,G entry, the RPF neighbor and interface index will also be shown. In this topology these are the same in all downstream PIM routers. The state is active so the upstream state is joined.

In this case FGT-3 is the last hop router so the IGMP join is received locally on port3. There is no PIM outgoing interface listed for this entry as it is used for the upstream PIM join.

```
(10.166.0.11, 239.255.255.1)
RPF nbr: 10.132.0.156
RPF idx: port2
SPT bit: 1
Upstream State: JOINED
Local:
Joined:
Asserted:
Outgoing:
 port3
```

This is the entry for the SPT, no RP IS listed. The S, G stream will be forwarded out of the stated outgoing interface.

```
(10.166.0.11, 239.255.255.1, rpt)
RP: 192.168.1.1
RPF nbr: 10.132.0.156
RPF idx: port2
Upstream State: NOT PRUNED
Local:
Pruned:
Outgoing:
```

The above S, G, RPT state is created for all streams that have both a S, G and a \*, G entry on the router. This is not pruned in this case because of the topology, the RP and source are reachable over the same interface.

Although not seen in this scenario, assert states may be seen when multiple PIM routers exist on the same LAN which can lead to more than one upstream router having a valid forwarding state. Assert messages are used to elect a single forwarder from the upstream devices.

## Viewing the PIM next-hop table

The PIM next-hop table is also very useful for checking the various states, it can be used to quickly identify the states of multiple multicast streams

```
FGT-3 # get router info multicast pim sparse-mode next-hop
Flags: N = New, R = RP, S = Source, U = Unreachable
Destination Type Nexthop Nexthop Nexthop Metric Pref
 Refcnt
 Num Addr

10.166.0.11 ..S. 1 10.132.0.156 9 21 110 3
192.168.1.1 .R.. 1 10.132.0.156 9 111 110 2
```

## Viewing the PIM multicast forwarding table

Also you can check the multicast forwarding table showing the ingress and egress ports of the multicast stream.

```
FGT-3 # get router info multicast table

IP Multicast Routing Table
Flags: I - Immediate Stat, T - Timed Stat, F - Forwarder installed
Timers: Uptime/Stat Expiry
Interface State: Interface (TTL threshold)

(10.166.0.11, 239.255.255.1), uptime 04:02:55, stat expires 00:02:25
Owner PIM-SM, Flags: TF
 Incoming interface: port2
 Outgoing interface list:
 port3 (TTL threshold 1)
```

## Viewing the kernel forwarding table

Also the kernel forwarding table can be verified, however this should give similar information to the above command:

```
FGT-3 # diag ip multicast mroute
grp=239.255.255.1 src=10.166.0.11 intf=9 flags=(0x10000000) []
 status=resolved
 last_assert=2615136 bytes=1192116 pkt=14538 wrong_if=0 num_ifs=1
 index(ttl)=[6(1),]
```

## Viewing the multicast routing table (FGT-2)

If you check the output on FGT-2 there are some small differences:

```
FGT-2 # get router info multicast pim sparse-mode table
IP Multicast Routing Table

(*,*,RP) Entries: 0
(*,G) Entries: 1
(S,G) Entries: 1
(S,G,rpt) Entries: 1
FCR Entries: 0

(*, 239.255.255.1)
RP: 192.168.1.1
RPF nbr: 0.0.0.0
RPF idx: None
Upstream State: JOINED
 Local:
 Joined:
 external
 Asserted:
FCR:
```

The `*,G` entry now has a joined interface rather than local because it has received a PIM join from FGT-3 rather than a local IGMP join.

```
(10.166.0.11, 239.255.255.1)
RPF nbr: 10.130.0.237
RPF idx: internal
SPT bit: 1
Upstream State: JOINED
Local:
Joined:
 external
Asserted:
Outgoing:
 external
```

The `S,G` entry shows that we have received a join on the external interface and the stream is being forwarded out of this interface.

```
(10.166.0.11, 239.255.255.1, rpt)
RP: 192.168.1.1
RPF nbr: 0.0.0.0
RPF idx: None
Upstream State: PRUNED
Local:
Pruned:
Outgoing:
 External
```

The `S,G,RPT` is different from FGT-3 because FGT-2 is the RP, it has pruned back the SPT for the RP to the first hop router.

### Viewing the multicast routing table (FGT-1)

FGT-1 again has some differences with regard to the PIM-SM states, there is no `*,G` entry because it is not in the path of a receiver and the RP.

```
FGT-1_master # get router info multicast pim sparse-mode table
IP Multicast Routing Table

(*,*,RP) Entries: 0
(*,G) Entries: 0
(S,G) Entries: 1
(S,G,rpt) Entries: 1
FCR Entries: 0
```

Below the `S,G` is the SPT termination because this FortiGate unit is the first hop router, the RPF neighbor always shows as 0.0.0.0 because the source is local to this device. Both the joined and outgoing fields show as external because the PIM join and the stream is egressing on this interface.

```
(10.166.0.11, 239.255.255.1)
RPF nbr: 0.0.0.0
RPF idx: None
SPT bit: 1
Upstream State: JOINED
```

```
Local:
Joined:
 external
Asserted:
Outgoing:
 external
```

The stream has been pruned back from the RP because the end-to-end SPT is flowing, there is no requirement for the stream to be sent to the RP in this case.

```
(10.166.0.11, 239.255.255.1, rpt)
RP: 0.0.0.0
RPF nbr: 10.130.0.156
RPF idx: external
Upstream State: RPT NOT JOINED
Local:
Pruned:
Outgoing:
```

## Example multicast destination NAT (DNAT) configuration

The example topology shown in [Figure 179](#) and described below shows how to configure destination NAT (DNAT) for two multicast streams. Both of these streams originate from the same source IP address, which is 10.166.0.11. The example configuration keeps the streams separate by creating 2 multicast NAT policies.

In this example the FortiGate units in [Figure 179](#) have the following roles:

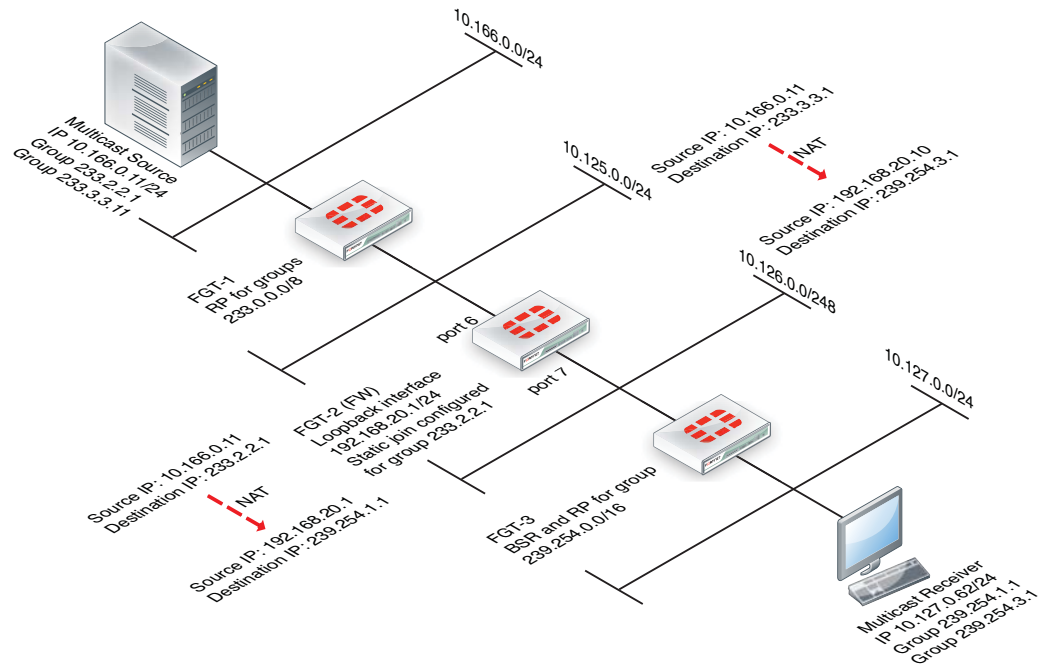
- FGT-1 is the RP for dirty networks, 233.0.0.0/8.
- FGT-2 performs all firewall and DNAT translations.
- FGT-3 is the RP for the clean networks, 239.254.0.0/16.
- FGT-1 and FGT-3 are functioning as PM enabled routers and could be replaced can be any PIM enabled router.

This example only describes the configuration of FGT-2.

FGT-2 performs NAT so that the receivers connected to FGT-3 receive the following translated multicast streams.

- If the multicast source sends multicast packets with a source and destination IP of 10.166.0.11 and 233.2.2.1; FGT-3 translates the source and destination IPs to 192.168.20.1 and 239.254.1.1
- If the multicast source sends multicast packets with a source and destination IP of 10.166.0.11 and 233.3.3.1; FGT-3 translates the source and destination IPs to 192.168.20.10 and 239.254.3.1

**Figure 179: Example multicast DNAT topology**



**To configure FGT-2 for DNAT multicast**

1. Add a loopback interface. In the example, the loopback interface is named loopback.

```
config system interface
 edit loopback
 set vdom root
 set ip 192.168.20.1 255.255.255.0
 set type loopback
 next
end
```

2. Add PIM and add a unicast routing protocol to the loopback interface as if it was a normal routed interface. Also add static joins to the loopback interface for any groups to be translated.

```
config router multicast
 config interface
 edit loopback
 set pim-mode sparse-mode
 config join-group
 edit 233.2.2.1
 next
 edit 233.3.3.1
 next
 end
 next
 next
```

3. In this example, to add firewall multicast policies, different source IP addresses are required so you must first add an IP pool:

```
config firewall ippool
 edit Multicast_source
 set endip 192.168.20.20
 set interface port6
 set startip 192.168.20.10
 next
end
```

4. Add the translation security policies.

Policy 2, which is the source NAT policy, uses the actual IP address of port6. Policy 1, the DNAT policy, uses an address from the IP pool.

```
config firewall multicast-policy
 edit 1
 set dnat 239.254.3.1
 set dstaddr example-addr_P
 set dstintf loopback
 set nat 192.168.20.10
 set srcaddr example-addr_Q
 set srcintf port6
 next
 edit 2
 set dnat 239.254.1.1
 set dstaddr example-addr_R
 set dstintf loopback
 set nat 192.168.20.1
 set srcaddr example-addr_Q
 set srcintf port6
 next
end
```

5. Add a firewall multicast policy to forward the stream from the loopback interface to the physical outbound interface.

This example is an any/any policy that makes sure traffic accepted by the other multicast policies can exit the FortiGate unit.

```
config firewall multicast-policy
 edit 3
 set dstintf port7
 set srcintf loopback
 next
end
```

## Example PIM configuration that uses BSR to find the RP

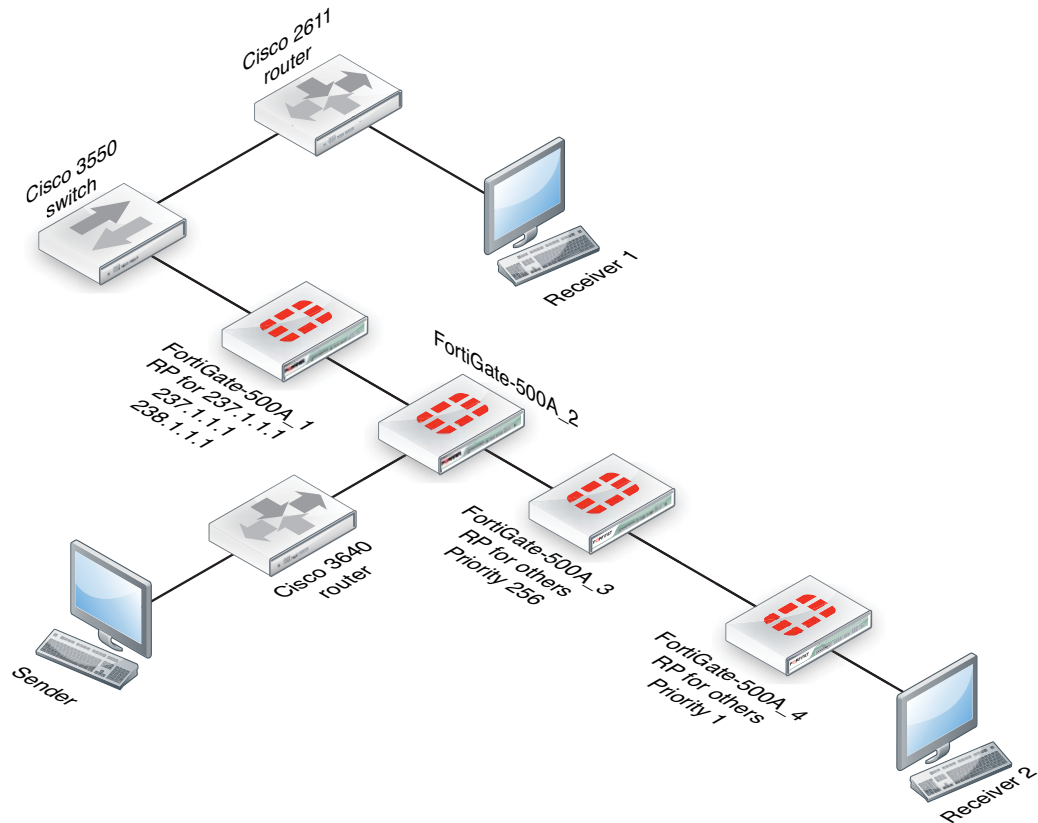
This example shows how to configure a multicast routing network for a network consisting of four FortiGate-500A units (FortiGate-500A\_1 to FortiGate-500A\_4, see [Figure 180](#)). A multicast sender is connected to FortiGate-500A\_2. FortiGate-500A\_2 forwards multicast packets in two directions to reach Receiver 1 and Receiver 2.

The configuration uses a Boot Start Router (BSR) to find the Rendezvous Points (RPs) instead of using static RPs. Under interface configuration, the loopback interface lo0 must join the 236.1.1.1 group (source).

This example describes:

- [Commands used in this example](#)
- [Configuration steps](#)
- [Example debug commands](#)

**Figure 180:** PIM network topology using BSR to find the RP



### Commands used in this example

This example uses CLI commands for the following configuration settings:

- [Adding a loopback interface \(lo0\)](#)
- [Defining the multicast routing](#)
- [Adding the NAT multicast policy](#)



## Adding a loopback interface (lo0)

Where required, the following command is used to define a loopback interface named lo0.

```
config system interface
 edit lo0
 set vdom root
 set ip 1.4.50.4 255.255.255.255
 set allowaccess ping https ssh snmp http telnet
 set type loopback
 next
end
```

## Defining the multicast routing

In this example, the following command syntax is used to define multicast routing. The example uses a Boot Start Router (BSR) to find the Rendezvous Points (RPs) instead of using static RPs. Under interface configuration, the loopback interface lo0 must join the 236.1.1.1 group (source).

```
config router multicast
 config interface
 edit port6
 set pim-mode sparse-mode
 next
 edit port1
 set pim-mode sparse-mode
 next
 edit lo0
 set pim-mode sparse-mode
 set rp-candidate enable
 config join-group
 edit 236.1.1.1
 next
 end
 set rp-candidate-priority 1
 next
 end
set multicast-routing enable
config pim-sm-global
 set bsr-allow-quick-refresh enable
 set bsr-candidate enable
 set bsr-interface lo0
 set bsr-priority 200
end
end
```

## Adding the NAT multicast policy

In this example, the incoming multicast policy does the address translation. The NAT address should be the same as the IP address of the of loopback interface. The DNAT address is the translated address, which should be a new group.

```
config firewall multicast-policy
 edit 1
 set dstintf port6
 set srcintf lo0
 next
 edit 2
 set dnat 238.1.1.1
 set dstintf lo0
 set nat 1.4.50.4
 set srcintf port1
 next
```

## Configuration steps

In this sample, FortiGate-500A\_1 is the RP for the group 228.1.1.1, 237.1.1.1, 238.1.1.1, and FortiGate-500A\_4 is the RP for the other group which has a priority of 1. OSPF is used in this example to distribute routes including the loopback interface. All firewalls have full mesh security policies to allow any to any.

- In the FortiGate-500A\_1 configuration, the NAT policy translates source address 236.1.1.1 to 237.1.1.1
- In the FortiGate-500A\_4, configuration, the NAT policy translates source 236.1.1.1 to 238.1.1.1
- Source 236.1.1.1 is injected into network as well.

The following procedures include the CLI commands for configuring each of the FortiGate units in the example configuration.

### To configure FortiGate-500A\_1

#### 1. Configure multicast routing.

```
config router multicast
 config interface
 edit port5
 set pim-mode sparse-mode
 next
 edit port4
 set pim-mode sparse-mode
 next
 edit lan
 set pim-mode sparse-mode
 next
 edit port1
 set pim-mode sparse-mode
 next
 edit lo999
 set pim-mode sparse-mode
 next
```

```

edit lo0
 set pim-mode sparse-mode
 set rp-candidate enable
 set rp-candidate-group 1
next
end
set multicast-routing enable
config pim-sm-global
 set bsr-candidate enable
 set bsr-interface lo0
end
end

```

**2. Add multicast security policies.**

```

config firewall multicast-policy
edit 1
 set dstintf port5
 set srcintf port4
next
edit 2
 set dstintf port4
 set srcintf port5
next
edit 3
next
end

```

**3. Add router access lists.**

```

config router access-list
edit 1
 config rule
 edit 1
 set prefix 228.1.1.1 255.255.255.255
 set exact-match enable
 next
 edit 2
 set prefix 237.1.1.1 255.255.255.255
 set exact-match enable
 next
 edit 3
 set prefix 238.1.1.1 255.255.255.255
 set exact-match enable
 next
end
next
end

```

## To configure FortiGate-500A\_2

### 1. Configure multicast routing.

```
config router multicast
 config interface
 edit "lan"
 set pim-mode sparse-mode
 next
 edit "port5"
 set pim-mode sparse-mode
 next
 edit "port2"
 set pim-mode sparse-mode
 next
 edit "port4"
 set pim-mode sparse-mode
 next
 edit "lo_5"
 set pim-mode sparse-mode
 config join-group
 edit 236.1.1.1
 next
 end
 next
 end
 set multicast-routing enable
end
```

### 2. Add multicast security policies.

```
config firewall multicast-policy
 edit 1
 set dstintf lan
 set srcintf port5
 next
 edit 2
 set dstintf port5
 set srcintf lan
 next
 edit 4
 set dstintf lan
 set srcintf port2
 next
 edit 5
 set dstintf port2
 set srcintf lan
 next
 edit 7
 set dstintf port1
 set srcintf port2
 next
 edit 8
```

```

 set dstintf port2
 set srcintf port1
 next
edit 9
 set dstintf port5
 set srcintf port2
next
edit 10
 set dstintf port2
 set srcintf port5
next
edit 11
 set dnat 237.1.1.1
 set dstintf lo_5
 set nat 5.5.5.5
 set srcintf port2
next
edit 12
 set dstintf lan
 set srcintf lo_5
next
edit 13
 set dstintf port1
 set srcintf lo_5
next
edit 14
 set dstintf port5
 set srcintf lo_5
next
edit 15
 set dstintf port2
 set srcintf lo_5
next
edit 16
next
end

```

### **To configure FortiGate-500A\_3**

#### **1. Configure multicast routing.**

```

config router multicast
config interface
 edit port5
 set pim-mode sparse-mode
 next
 edit port6
 set pim-mode sparse-mode
 next
 edit lo0
 set pim-mode sparse-mode

```

```

 set rp-candidate enable
 set rp-candidate-priority 255
 next
 edit lan
 set pim-mode sparse-mode
 next
end
set multicast-routing enable
config pim-sm-global
 set bsr-candidate enable
 set bsr-interface lo0
end
end

```

**2. Add multicast security policies.**

```

config firewall multicast-policy
 edit 1
 set dstintf port5
 set srcintf port6
 next
 edit 2
 set dstintf port6
 set srcintf port5
 next
 edit 3
 set dstintf port6
 set srcintf lan
 next
 edit 4
 set dstintf lan
 set srcintf port6
 next
 edit 5
 set dstintf port5
 set srcintf lan
 next
 edit 6
 set dstintf lan
 set srcintf port5
 next
end

```

## To configure FortiGate-500A\_4

### 1. Configure multicast routing.

```
config router multicast
 config interface
 edit port6
 set pim-mode sparse-mode
 next
 edit lan
 set pim-mode sparse-mode
 next
 edit port1
 set pim-mode sparse-mode
 next
 edit lo0
 set pim-mode sparse-mode
 set rp-candidate enable
 config join-group
 edit 236.1.1.1
 next
 end
 set rp-candidate-priority 1
 next
 end
set multicast-routing enable
config pim-sm-global
 set bsr-allow-quick-refresh enable
 set bsr-candidate enable
 set bsr-interface lo0
 set bsr-priority 1
end
end
```

### 2. Add multicast security policies.

```
config firewall policy
 edit 1
 set srcintf lan
 set dstintf port6
 set srcaddr all
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
 next
 edit 2
 set srcintf port6
 set dstintf lan
 set srcaddr all
 set dstaddr all
 set action accept
 set schedule always
```

```
 set service ANY
next
edit 3
 set srcintf port1
 set dstintf port6
 set srcaddr all
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
next
edit 4
 set srcintf port6
 set dstintf port1
 set srcaddr all
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
next
edit 5
 set srcintf port1
 set dstintf lan
 set srcaddr all
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
next
edit 6
 set srcintf lan
 set dstintf port1
 set srcaddr all
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
next
edit 7
 set srcintf port1
 set dstintf port1
 set srcaddr all
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
next
edit 8
 set srcintf port6
```



```

 set dstintf lo0
 set srcaddr all
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
 next
edit 9
 set srcintf port1
 set dstintf lo0
 set srcaddr all
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
next
edit 10
 set srcintf lan
 set dstintf lo0
 set srcaddr all
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
next
end

```

### Example debug commands

You can use the following CLI commands to view information about and status of the multicast configuration. This section includes `get` and `diagnose` commands and some sample output.

```
get router info multicast pim sparse-mode table 236.1.1.1
```

```
get router info multicast pim sparse-mode neighbor
```

Neighbor Address Mode	Interface	Uptime/Expires	Ver	DR	Priority/ DR
83.97.1.2	port6	02:22:01/00:01:44	v2	1	/ DR

```
diagnose ip multicast mroute
```

```
grp=236.1.1.1 src=19.2.1.1 intf=7 flags=(0x10000000) []
```

```
status=resolved
```

```
last_assert=171963 bytes=1766104 pkt=1718 wrong_if=1
```

```
num_ifs=2
```

```
index(ttl)=[6(1),10(1),]
```

```
grp=236.1.1.1 src=1.4.50.4 intf=10 flags=(0x10000000) []
```

```
status=resolved
```

```
last_assert=834864 bytes=4416 pkt=138 wrong_if=0 num_ifs=2
```

```
index(ttl)=[7(1),6(1),]
```

```
grp=238.1.1.1 src=1.4.50.4 intf=10 flags=(0x10000000) []
```

```
status=resolved
```

```
last_assert=834864 bytes=1765076 pkt=1717 wrong_if=0
```

```
num_ifs=1
```

```
index(ttl)=[7(1),]
```

```
get router info multicast igmp groups
```

```
IGMP Connected Group Membership
```

Group Address	Interface	Uptime	Expires	Last Reporter
236.1.1.1	lan	00:45:48	00:03:21	10.4.1.1
236.1.1.1	lo0	02:19:31	00:03:23	1.4.50.4

```
get router info multicast pim sparse-mode interface
```

Address	Interface	VIFindex	Ver/ Mode	Nbr Count	DR	DR Prior
10.4.1.2	lan	2	v2/S	0	1	10.4.1.2
83.97.1.1	port6	0	v2/S	1	1	83.97.1.2
1.4.50.4	lo0	3	v2/S	0	1	1.4.50.4

```
get router info multicast pim sparse-mode rp-mapping
```

```
PIM Group-to-RP Mappings
```

```
This system is the Bootstrap Router (v2)
```

```
Group(s): 224.0.0.0/4
```

```
RP: 1.4.50.4
```

```
Info source: 1.4.50.4, via bootstrap, priority 1
```

```
Uptime: 02:20:32, expires: 00:01:58
```

```
RP: 1.4.50.3
```

```
Info source: 1.4.50.3, via bootstrap, priority 255
```

```
Uptime: 02:20:07, expires: 00:02:24
```

```
Group(s): 228.1.1.1/32
```

```
RP: 1.4.50.1
```

```
Info source: 1.4.50.1, via bootstrap, priority 192
```

```
Uptime: 02:18:24, expires: 00:02:06
```

```
Group(s): 237.1.1.1/32
```

```
RP: 1.4.50.1
```

```
Info source: 1.4.50.1, via bootstrap, priority 192
```

```
Uptime: 02:18:24, expires: 00:02:06
```

```
Group(s): 238.1.1.1/32
```

```
RP: 1.4.50.1
```

```
Info source: 1.4.50.1, via bootstrap, priority 192
```

```
Uptime: 02:18:24, expires: 00:02:06
```

```
get router info multicast pim sparse-mode bsr-info
```

```
PIMv2 Bootstrap information
```

```
This system is the Bootstrap Router (BSR)
```

```
BSR address: 1.4.50.4
```

```
Uptime: 02:23:08, BSR Priority: 1, Hash mask length: 10
```

```
Next bootstrap message in 00:00:18
```

```
Role: Candidate BSR
```

```
State: Elected BSR
```

```
Candidate RP: 1.4.50.4(lo0)
```

```
Advertisement interval 60 seconds
```

```
Next Cand_RP_advertisement in 00:00:54
```

# Chapter 8 Hardware Acceleration

This FortiOS Handbook chapter contains the following sections:

[Hardware acceleration overview](#) describes the capabilities of FortiGate content processors (CPs), security processors (SPs) and network processors (NPs). This chapter also describes how to determine the hardware acceleration components installed in your FortiGate unit and contains some configuration details and examples.

[NP6 Acceleration](#) describes the FortiGate NP6 network processor.

[FortiGate NP6 architectures](#) contains details about the network processing architectures of FortiGate units that contain NP6 processors.

[NP4 Acceleration](#) describes the FortiGate NP4 network processor.

[FortiGate NP4 architectures](#) contains details about the network processing architectures of FortiGate units that contain NP4 processors.

# Hardware acceleration overview

All FortiGate models have specialized acceleration hardware that can offload resource intensive processing from main processing (CPU) resources. All FortiGate units include specialized content processors (CPs) that accelerate a wide range of important security processes such as virus scanning, attack detection, encryption and decryption. Many FortiGate models also contain security processors (SPs) that accelerate processing for specific security features such as IPS and network processors (NPs) that offload processing of high volume network traffic.

This chapter contains the following topics:

- [Content processors \(CP4, CP5, CP6 and CP8\)](#)
- [Security processors \(SPs\)](#)
- [Network processors \(NP1, NP2, NP3, NP4 and NP6\)](#)
- [Checking that traffic is offloaded by NP processors](#)
- [Controlling IPS NPx and CPx acceleration](#)

## Content processors (CP4, CP5, CP6 and CP8)

All FortiGate units contain FortiASIC Content Processors (CPs) that accelerate many common resource intensive security related processes. CPs work at the system level with tasks being offloaded to them as determined by the main CPU. Capabilities of the CPs vary by model. Newer FortiGate units include CP8 processors. Older CP versions still in use in currently operating FortiGate models include the CP4, CP5, and CP6.

### CP8 capabilities

The CP8 content processor provides the following services:

- IPS signature matching acceleration
- High performance VPN bulk data engine
  - IPSEC and SSL/TLS protocol processor
  - DES/3DES/AES in accordance with FIPS46-3/FIPS81/FIPS197
  - ARC4 in compliance with RC4
  - MD5/SHA-1/SHA256 with RFC1321 and FIPS180
  - HMAC in accordance with RFC2104/2403/2404 and FIPS198
- Key Exchange Processor support high performance IKE and RSA computation
  - Public key exponentiation engine with hardware CRT support
  - Primarily checking for RSA key generation
  - Handshake accelerator with automatic key material generation
  - Random Number generator compliance with ANSI X9.31
  - Sub public key engine (PKCE) to support up to 4094 bit operation directly
- Message authentication module offers high performance cryptographic engine for calculating SHA256/SHA1/MD5 of data up to 4G bytes (used by many applications)
- PCI express Gen 2 four lanes interface
- Cascade Interface for chip expansion

### CP6 capabilities

- Dual content processors
- FIPS-compliant DES/3DES/AES encryption and decryption
- SHA-1 and MD5 HMAC with RFC1321 and FIPS180
- HMAC in accordance with RFC2104/2403/2404 and FIPS198
- IPsec protocol processor
- High performance IPsec engine
- Random Number generator compliance with ANSI X9.31
- Key exchange processor for high performance IKE and RSA computation
- Script Processor
- SSL/TLS protocol processor for SSL content scanning and SSL acceleration

### CP5 capabilities

- FIPS-compliant DES/3DES/AES encryption and decryption
- SHA-1 and MD5 HMAC with RFC1321/2104/2403/2404 and FIPS180/FIPS198
- IPsec protocol processor
- High performance IPSEC Engine
- Random Number generator compliant with ANSI X9.31
- Public Key Crypto Engine supports high performance IKE and RSA computation
- Script Processor

### CP4 capabilities

- FIPS-compliant DES/3DES/AES encryption and decryption
- SHA-1 and MD5 HMAC
- IPSEC protocol processor
- Random Number generator
- Public Key Crypto Engine
- Content processing engine
- ANSI X9.31 and PKCS#1 certificate support

## Determining the content processor in your FortiGate unit

Use the `get hardware status` CLI command to determine which content processor your FortiGate unit contains. The output looks like this:

```
get hardware status
Model name: FortiGate-100D
ASIC version: CP8
ASIC SRAM: 64M
CPU: Intel(R) Atom(TM) CPU D525 @ 1.80GHz
Number of CPUs: 4
RAM: 1977 MB
Compact Flash: 15331 MB /dev/sda
Hard disk: 15272 MB /dev/sda
USB Flash: not available
Network Card chipset: Intel(R) PRO/1000 Network Connection (rev.0000)
Network Card chipset: bcm-sw Ethernet driver 1.0 (rev.)
```

The ASIC version line lists the content processor model number.

## Viewing SSL acceleration status

You can view the status of SSL acceleration using the following command:

```
get vpn status ssl hw-acceleration-status
Acceleration hardware detected: kxp=on cipher=on
```

## Disabling CP offloading

If you want to completely disable offloading to CP processors for test purposes or other reasons, you can do so in security policies. Here are some examples:

For IPv4 security policies.

```
config firewall policy
 edit 1
 set auto-asic-offload disable
 end
```

For IPv6 security policies.

```
config firewall policy6
 edit 1
 set auto-asic-offload disable
 end
```

For multicast security policies.

```
config firewall multicast-policy
 edit 1
 set auto-asic-offload disable
 end
```

## Security processors (SPs)

FortiGate Security Processing (SP) modules, such as the SP3 but also including the XG2, XE2, FE8, and CE4, work at both the interface and system level to increase overall system performance by accelerating specialized security processing. You can configure the SP to favor IPS over firewall processing in hostile high-traffic environments. The following security processors are available:

- The SP3 is built into the FortiGate-5101B and provides IPS acceleration. No special configuration is required. All IPS traffic is accelerated by the built-in SP3 processors.
- The FMC-XG2 is an FMC module with two 10Gb/s SPF+ interfaces that can be used on FortiGate-3950B and FortiGate-3951B units.
- The FortiGate-3140B also contains a built-in XG2 using ports 19 and 20.
- The ADM-XE2 is a dual-width AMC module with two 10Gb/s interfaces that can be used on FortiGate-3810A and FortiGate-5001A-DW systems.
- The ADM-FE8 is a dual-width AMC module with eight 1Gb/s interfaces that can be used with the FortiGate-3810A.
- The ASM-CE4 is a single-width AMC module with four 10/100/1000 Mb/s interfaces that can be used on FortiGate-3016B and FortiGate-3810A units.

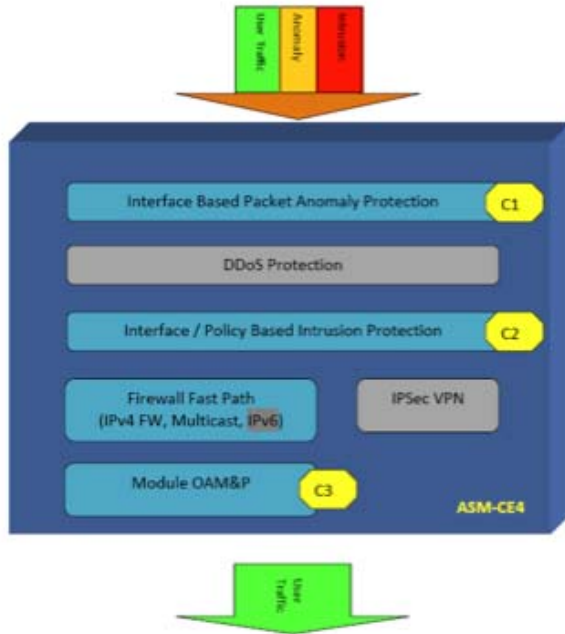
### SP Processing Flow

SP processors provide an integrated high performance fast path multilayer solution for both intrusion protection and firewall functions. The multilayered protection starts from anomaly checking at packet level to ensure each packet is sound and reasonable. Immediately after that, a sophisticated set of interface based packet anomaly protection, DDoS protection, policy based intrusion protection, firewall fast path, and behavior based methods are employed to prevent DDoS attacks from the rest of system.

Then the packets enter an interface/policy based intrusion protection system,

where each packet is evaluated against a set of signatures. The end result is streams of user packets that are free of anomaly and attacks, entering the fast path system for unicast or multicast fast path forwarding.

**Figure 181:**SP processing flow



## Displaying information about security processing modules

You can display information about installed SP modules using the CLI command

```
diagnose npu spm
```

For example, for the FortiGate-5101C:

```
FG-5101C # diagnose npu spm list
Available SP Modules:
```

ID	Model	Slot	Interface
0	xh0	built-in	port1, port2, port3, port4 base1, base2, fabric1, fabric2 eth10, eth11, eth12, eth13 eth14, eth15, eth16, eth17 eth18, eth19

You can also use this command to get more info about SP processing. This example shows how to display details about how the module is processing sessions using the syn proxy.

```
diagnose npu spm dos synproxy <sp_id>
```



This is a partial output of the command:

```
Number of proxied TCP connections : 0
Number of working proxied TCP connections : 0
Number of retired TCP connections : 0
Number of valid TCP connections : 0
Number of attacks, no ACK from client : 0
Number of no SYN-ACK from server : 0
Number of reset by server (service not supported): 0
Number of established session timeout : 0
Client timeout setting : 3 Seconds
Server timeout setting : 3 Seconds
```

## Network processors (NP1, NP2, NP3, NP4 and NP6)

FortiASIC network processors work at the interface level to accelerate traffic by offloading traffic from the main CPU. Current models contain NP4 and NP6 network processors. Older FortiGate models include NP1 network processors (also known as FortiAccel, or FA2) and NP2 network processors.

The traffic that can be offloaded, maximum throughput, and number of network interfaces supported by each varies by processor model:

- NP6 supports offloading of most IPv4 and IPv6 traffic, IPsec VPN encryption, CAPWAP traffic, and multicast traffic. The NP6 has a capacity of 40 Gbps through 4 x 10 Gbps interfaces or 3 x 10 Gbps and 16 x 1 Gbps interfaces. For details about the NP6 processor, see [“NP6 Acceleration” on page 1082](#) and for information about FortiGate models with NP6 processors, see [“FortiGate NP6 architectures” on page 1089](#).
- NP4 supports offloading of most IPv4 firewall traffic and IPsec VPN encryption. The NP4 has a capacity of 20 Gbps through 2 x 10 Gbps interfaces. For details about NP4 processors, see [“NP4 Acceleration” on page 1093](#) and for information about FortiGate models with NP4 processors, see [“FortiGate NP4 architectures” on page 1105](#).
- NP2 supports IPv4 firewall and IPsec VPN acceleration. The NP2 has a capacity of 2 Gbps through 2 x 10 Gbps interfaces or 4 x 1 Gbps interfaces.
- NP1 supports IPv4 firewall and IPsec VPN acceleration with 2 Gbps capacity. The NP1 has a capacity of 2 Gbps through 2 x 1 Gbps interfaces.
  - The NP1 does not support frames greater than 1500 bytes. If your network uses jumbo frames, you may need to adjust the MTU (Maximum Transmission Unit) of devices connected to NP1 ports. Maximum frame size for NP2, NP4, and NP6 processors is 9000 bytes.
  - For both NP1 and NP2 network processors, ports attached to a network processor cannot be used for firmware installation by TFTP.



Session that require proxy-based and flow based security features (for example, virus scanning, IPS, application control and so on) are not fast pathed and must be processed by the CPU.

---

## Determining the network processors installed on your FortiGate unit

Use the following command to list the NP6 processors in your FortiGate unit:

```
diagnose npu np6 port-list
```

To list other network processors on your FortiGate unit, use the following CLI command.

```
get hardware npu <model> list
```

<model> can be legacy, np1, np2 or np4.

The output lists the interfaces that have the specified processor. For example, for a FortiGate-5001B:

```
get hardware npu np4 list
```

ID	Model	Slot	Interface
0	On-board		port1 port2 port3 port4 fabric1 base1 npu0-vlink0 npu0-vlink1
1	On-board		port5 port6 port7 port8 fabric2 base2 npu1-vlink0 npu1-vlink1

The npu0-vlink0, npu1-vlink1 etc interfaces are used for accelerating inter-VDOM links.

## How NP hardware acceleration alters packet flow

NP hardware acceleration generally alters packet flow as follows:

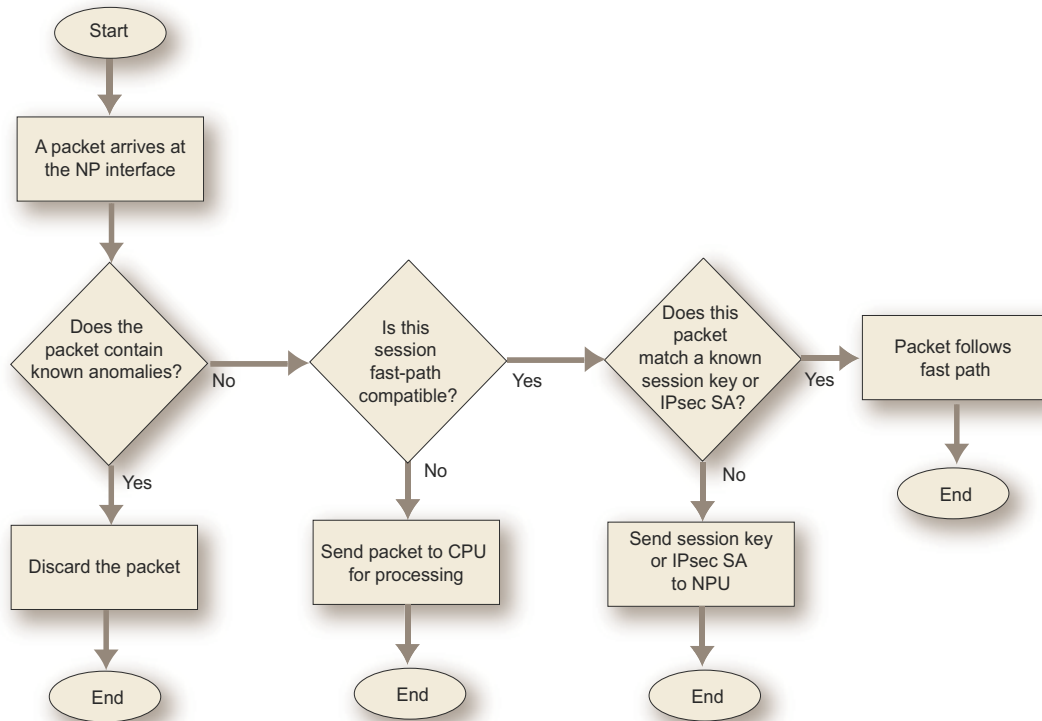
1. Packets initiating a session pass to the FortiGate unit's main processing resources (CPU).
2. The FortiGate unit assesses whether the session matches fast path (offload) requirements.

To be suitable for offloading, traffic must possess only characteristics that can be processed by the fast path. The list of requirements depends on the processor, see [“NP6 session fast path requirements” on page 1083](#) or [“NP4 session fast path requirements” on page 1094](#).

If the session can be fast pathed, the FortiGate unit sends the session key or IPsec security association (SA) and configured firewall processing action to the appropriate network processor.

3. Network processors continuously match packets arriving on their attached ports against the session keys and SAs they have received.
  - If a network processor's network interface is configured to perform hardware accelerated anomaly checks, the network processor drops or accepts packets which match the configured anomaly patterns. These checks are separate from and in advance of anomaly checks performed by IPS, which is not compatible with network processor offloading. See [“Offloading NP pre-IPS anomaly detection” on page 1076](#).
  - The network processor next checks for a matching session key or SA. If a matching session key or SA is found, and if the packet meets packet requirements, the network processor processes the packet according to the configured action and then sends the resulting packet. This is the actual offloading step. Performing this processing on the NP processor improves overall performance because the NP processor is optimized for this task. As well, overall FortiGate performance is improved because the CPU has fewer sessions to process.

**Figure 182:**NP network processor packet flow



- If a matching session key or SA is not found, or if the packet does not meet packet requirements, the packet cannot be offloaded. The network processor sends the data to the FortiGate unit's CPU, which processes the packet.

Encryption and decryption IPsec traffic originating from the FortiGate can utilize network processor encryption capabilities. See [“Configuring NP accelerated VPN encryption/decryption offloading”](#) on page 1078.

Packet forwarding rates vary by the percentage of offloadable processing and the type of network processing required by your configuration, but are independent of frame size. For optimal traffic types, network throughput can equal wire speed.

## NP processors and traffic logging and monitoring

Except for the NP6, network processors do not count offloaded packets, and offloaded packets are not logged by traffic logging and are not included in traffic statistics and traffic log reports.

NP6 processors support per-session traffic and byte counters, Ethernet MIB matching, and reporting through messages resulting in traffic statistics and traffic log reporting.

## NP session offloading in HA active-active configuration

Network processors can improve network performance in active-active (load balancing) high availability (HA) configurations, even though traffic deviates from general offloading patterns, involving more than one network processor, each in a separate FortiGate unit. No additional offloading requirements apply.

Once the primary FortiGate unit's main processing resources send a session key to its network processor(s), network processor(s) on the primary unit can redirect any subsequent session

traffic to other cluster members, reducing traffic redirection load on the primary unit's main processing resources.

As subordinate units receive redirected traffic, each network processor in the cluster assesses and processes session offloading independently from the primary unit. Session key states of each network processor are not part of synchronization traffic between HA members.

## Configuring NP HMAC check offloading

Hash-based Message Authentication Code (HMAC) checks offloaded to network processors by default. You can enter the following command to disable this feature:

```
configure system global
 set ipsec-hmac-offload disable
end
```

## Offloading NP pre-IPS anomaly detection

Network interfaces associated with a port attached to a network processor can be configured to offload anomaly checking. This anomaly checking happens before other offloading and separately from and in advance of DoS policy anomaly checking. Using the following command, each FortiGate interface can have a different anomaly checking configuration.

```
config system interface
 edit <port-name>
 set fp-anomaly <anomalies>
 end
```

where <anomalies> can be one, more than one or all of the following:

Anomaly	Description
drop_icmp_frag	Drop ICMP fragments to pass.
drop_icmpland	Drop ICMP Land.
drop_ipland	Drop IP Land.
drop_iplsrr	Drop IP with Loose Source Record Route option.
drop_iprr	Drop IP with Record Route option.
drop_ipsecurity	Drop IP with Security option.
drop_ipssrr	Drop IP with Strict Source Record Route option.
drop_ipstream	Drop IP with Stream option.
drop_iptimestamp	Drop IP with Timestamp option.
drop_ipunknown_option	Drop IP with malformed option.
drop_ipunknown_prot	Drop IP with Unknown protocol.
drop_tcp_fin_noack	Drop TCP FIN with no ACT flag set to pass.
drop_tcp_no_flag	Drop TCP with no flag set to pass.
drop_tcpland	Drop TCP Land.

Anomaly	Description
drop_udpland	Drop UDP Land.
drop_winnuke	Drop TCP WinNuke.
pass_icmp_frag	Allow ICMP fragments to pass.
pass_icmpland	Allow ICMP Land to pass.
pass_ipland	Allow IP land to pass.
pass_iplsrr	Allow IP with Loose Source Record Route option to pass.
pass_iprr	Allow IP with Record Route option to pass.
pass_ipsecurity	Allow IP with Security option to pass.
pass_ipssrr	Allow IP with Strict Source Record Route option to pass.
pass_ipstream	Allow IP with Stream option to pass.
pass_iptimestamp	Allow IP with Timestamp option to pass.
pass_ipunknown_option	Allow IP with malformed option to pass.
pass_ipunknown_prot	Allow IP with Unknown protocol to pass.
pass_tcp_fin_noack	Allow TCP FIN with no ACT flag set to pass.
pass_tcp_no_flag	Allow TCP with no flag set to pass.
pass_tcpland	Allow TCP Land to pass.
pass_udpland	Allow UDP Land to pass.
pass_winnuke	Allow TCP WinNuke to pass.

### Example

You might configure an NP4 to drop packets with TCP WinNuke or unknown IP protocol anomalies, but to pass packets with an IP time stamp, using hardware acceleration provided by the network processor.

```
config system interface
 edit port1
 set fp-anomaly drop_winnuke drop_ipunknown_prot pass_iptimestamp
 end
```

### Software switch interfaces and NP processors

FortiOS supports creating a software switch by grouping two or more FortiGate physical interfaces into a single virtual or software switch interface. All of the interfaces in this virtual switch act like interfaces in a hardware switch in that they all have the same IP address and can be connected to the same network. You create a software switch interface from the CLI using the command `config system switch-interface`.

The software switch is a bridge group of several interfaces, and the FortiGate CPU maintains the mac-port table for this bridge. As a result of this CPU involvement, traffic processed by a software switch interface is not offloaded to network processors.

## Configuring NP accelerated VPN encryption/decryption offloading

Network processing unit (npu) settings configure offloading behavior for IPsec VPN. Configured behavior applies to all network processors contained by the FortiGate unit itself or any installed AMC modules.

```
config system npu
 set enc-offload-antireplay {enable | disable}
 set dec-offload-antireplay {enable | disable}
 set offload-ipsec-host {enable | disable}
end
```

Variables	Description	Default
enc-offload-antireplay {enable   disable}	Enable or disable offloading of IPsec encryption. This option is used only when replay detection is enabled in Phase 2 configuration. If replay detection is disabled, encryption is always offloaded.	disable
dec-offload-antireplay {enable   disable}	Enable or disable offloading of IPsec decryption. This option is used only when replay detection is enabled in Phase 2 configuration. If replay detection is disabled, decryption is always offloaded.	enable
offload-ipsec-host {enable   disable}	Enable or disable offloading of IPsec encryption of traffic from local host (FortiGate unit). <b>Note:</b> For this option to take effect, the FortiGate unit must have previously sent the security association (SA) to the network processor.	disable

### Example

You could configure the offloading of encryption and decryption for an IPsec SA that was sent to the network processor.

```
config system npu
 set enc-offload-antireplay enable
 set dec-offload-antireplay enable
 set offload-ipsec-host enable
end
```

## Checking that traffic is offloaded by NP processors

A number of diagnose commands can be used to verify that traffic is being offloaded.

## Using the packet sniffer

Use the packet sniffer to verify that traffic is offloaded. Offloaded traffic is not picked up by the packet sniffer so if you are sending traffic through the FortiGate unit and it is not showing up on the packet sniffer you can conclude that it is offloaded.

```
diag sniffer packet port1 <option>
```

## Checking the firewall session offload tag

Use the `diagnose sys session list` command to display sessions. If the output for a session includes the `npu info` field you should see information about session being offloaded. If the output doesn't contain an `npu info` field then the session has not been offloaded.

```
diagnose sys session list
session info: proto=6 proto_state=01 duration=34 expire=3565
 timeout=3600 flags=00000000 sockflag=00000000 sockport=0
 av_idx=0 use=3
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 policy_dir=0 tunnel=/
state=may_dirty npu
statistic(bytes/packets/allow_err): org=295/3/1 reply=60/1/1
 tuples=2
origin->sink: org pre->post, reply pre->post dev=48->6/6->48
 gwy=10.1.100.11/11.11.11.1
hook=pre dir=org act=noop
 172.16.200.55:56453->10.1.100.11:80(0.0.0.0:0)
hook=post dir=reply act=noop
 10.1.100.11:80->172.16.200.55:56453(0.0.0.0:0)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=1 id_policy_id=0 auth_info=0 chk_client_info=0 vd=4
serial=0000091c tos=ff/ff ips_view=0 app_list=0 app=0
dd_type=0 dd_mode=0
per_ip_bandwidth meter: addr=172.16.200.55, bps=393
npu_state=00000000
npu info: flag=0x81/0x81, offload=4/4, ips_offload=0/0, epid=1/23,
 ipid=23/1, vlan=32779/0
```

## Verifying IPsec VPN traffic offloading

The following commands can be used to verify IPsec VPN traffic offloading to NP processors.

```
diagnose vpn ipsec status
```

```
NP1/NP2/NP4_0/sp_0_0:
```

```
 null: 0 0
 des: 0 0
 3des: 4075 4074
 aes: 0 0
 aria: 0 0
 seed: 0 0
 null: 0 0
 md5: 4075 4074
 sha1: 0 0
 sha256: 0 0
 sha384: 0 0
 sha512: 0 0
```

```
diagnose vpn tunnel list
```

```
list all ipsec tunnel in vd 3
```

```

name=p1-vdom1 ver=1 serial=5 11.11.11.1:0->11.11.11.2:0 lgwy=static
 tun=tunnel mode=auto bound_if=47
proxyid_num=1 child_num=0 refcnt=8 ilast=2 olast=2
stat: rxp=3076 txp=1667 rxb=4299623276 txb=66323
dpd: mode=active on=1 idle=5000ms retry=3 count=0 seqno=20
natt: mode=none draft=0 interval=0 remote_port=0
proxyid=p2-vdom1 proto=0 sa=1 ref=2 auto_negotiate=0 serial=1
 src: 0:0.0.0.0/0.0.0.0:0
 dst: 0:0.0.0.0/0.0.0.0:0
SA: ref=6 options=0000000e type=00 soft=0 mtu=1436 expire=1736
 replaywin=2048 seqno=680
life: type=01 bytes=0/0 timeout=1748/1800
dec: spi=ae01010c esp=3des key=24
 18e021bcace225347459189f292fbc2e4677563b07498a07
 ah=md5 key=16 b4f44368741632b4e33e5f5b794253d3
enc: spi=ae01010d esp=3des key=24
 42c94a8a2f72a44f9a3777f8e6aa3b24160b8af15f54a573
 ah=md5 key=16 6214155f76b63a93345dcc9ec02d6415
dec:pkts/bytes=3073/4299621477, enc:pkts/bytes=1667/66375
npu_flag=03 npu_rgwy=11.11.11.2 npu_lgwy=11.11.11.1 npu_selid=4
```



```

diagnose sys session list
session info: proto=6 proto_state=01 duration=34 expire=3565
 timeout=3600 flags=00000000 sockflag=00000000 sockport=0
 av_idx=0 use=3
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 policy_dir=0 tunnel=/p1-vdom2
state=re may_dirty npu
statistic(bytes/packets/allow_err): org=112/2/1 reply=112/2/1
 tuples=2
origin->sink: org pre->post, reply pre->post dev=57->7/7->57
 gwy=10.1.100.11/11.11.11.1
hook=pre dir=org act=noop
 172.16.200.55:35254->10.1.100.11:80(0.0.0.0:0)
hook=post dir=reply act=noop
 10.1.100.11:80->172.16.200.55:35254(0.0.0.0:0)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=1 id_policy_id=0 auth_info=0 chk_client_info=0 vd=4
serial=00002d29 tos=ff/ff ips_view=0 app_list=0 app=0
dd_type=0 dd_mode=0
per_ip_bandwidth meter: addr=172.16.200.55, bps=260
npu_state=00000000
npu info: flag=0x81/0x82, offload=7/7, ips_offload=0/0, epid=1/3,
 ipid=3/1, vlan=32779/0

```

## Controlling IPS NPx and CPx acceleration

You can use the following commands to enable or disable acceleration of IPS processing by NPx and CPx processors:

```

config ips global
 set np-accel-mode {none | basic}
 set cp-accel-mode {none | basic | advanced}
end

```

The network processor (NP) acceleration modes are:

none: Network Processor acceleration disabled  
basic: Basic Network Processor acceleration enabled

The content processor (CP) acceleration modes are:

none: Content Processor acceleration disabled  
basic: Basic Content Processor acceleration enabled  
advanced: Advanced Content Processor acceleration enabled

# NP6 Acceleration

NP6 network processors provide fastpath acceleration by offloading communication sessions from the FortiGate CPU. When the first packet of a new session is received by an interface connected to an NP6 processor, just like any session connecting with any FortiGate interface, the session is forwarded to the FortiGate CPU where it is matched with a security policy. If the session is accepted by a security policy and if the session can be offloaded its session key is copied to the NP6 processor that received the packet. All of the rest of the packets in the session are intercepted by the NP6 processor and fast-pathed out of the FortiGate unit to their destination without ever passing through the FortiGate CPU. The result is enhanced network performance provided by the NP6 processor plus the network processing load is removed from the CPU. In addition the NP6 processor can handle some CPU intensive tasks, like IPsec VPN encryption/decryption.

Session keys (and IPsec SA keys) are stored in the memory of the NP6 processor that is connected to the interface that received the packet that started the session. All sessions are fast-pathed and accelerated, even if they exit the FortiGate unit through an interface connected to another NP6. There is no dependence on getting the right pair of interfaces since the offloading is done by the receiving NP6. The key to making this possible is the Integrated Switch Fabric (ISF) that connects the NP6s and the FortiGate unit interfaces together. The ISF allows any port connectivity. All ports and NP6s can communicate with each other over the ISF.

There are no special ingress and egress fast path requirements as long as traffic enters and exits on interfaces connected to the same ISF and the NP6 processors. All FortiGate models with NP6 processors connect all interfaces and NP6 processors to the same ISF (except management interfaces) so this should not ever be a problem.

There are at least two limitations to keep in mind:

- The capacity of each NP6 processor. An individual NP6 processor can support between 10 and 16 million sessions. This number is limited by the amount of memory the processor has. Once an NP6 processor hits its session limit, sessions that are over the limit are sent to the CPU. You can avoid this problem by as much as possible distributing incoming sessions evenly among the NP6 processors. To be able to do this you need to be aware of which interfaces connect to which NP6 processors and distribute incoming traffic accordingly.
- Some FortiGate units may use some NP6 processors for special functions. For example, ports 25 to 32 of the FortiGate-3700D can be used for low latency offloading. See [“FortiGate-3700D fast path architecture” on page 1091](#) for more information.

This chapter contains the following topics:

- [NP6 session fast path requirements](#)
- [Viewing your FortiGate NP6 processor configuration](#)
- [Increasing NP6 offloading capacity using link aggregation groups \(LAGs\)](#)
- [Configuring Inter-VDOM link acceleration with NP6 processors](#)

## NP6 session fast path requirements

NP6 processors can offload the following traffic and services:

- IPv4 and IPv6 traffic and NAT64 and NAT46 traffic (as well as IPv4 and IPv6 versions of the following traffic types where appropriate)
- TCP, UDP, ICMP and SCTP traffic
- IPsec VPN traffic, and offloading of IPsec encryption/decryption (including SHA2-256 and SHA2-512)
- Anomaly-based intrusion prevention, checksum offload and packet defragmentation
- SIT and IPv6 Tunnelling sessions
- Multicast traffic (including Multicast over IPsec)
- CAPWAP and wireless bridge traffic tunnel encapsulation to enable line rate wireless forwarding from FortiAP devices
- Traffic shaping and priority queuing for both shared and per IP traffic shaping. An NP6 processor has 16 million queues for traffic shaping and statistics counting.
- Syn proxying
- Inter-VDOM link traffic

Sessions that are offloaded must be fast path ready. For a session to be fast path ready it must meet the following criteria:

- Layer 2 type/length must be 0x0800 for IPv4 or 0x86dd for IPv6 (IEEE 802.1q VLAN specification is supported)
- Link aggregation between any network interfaces sharing the same network processor(s) may be used (IEEE 802.3ad specification is supported)
- Layer 3 protocol can be IPv4 or IPv6
- Layer 4 protocol can be UDP, TCP, ICMP, or SCTP
- In most cases, Layer 3 / Layer 4 header or content modification sessions that require a session helper can be offloaded.
- Local host traffic (originated by the FortiGate unit) can be offloaded
- Application layer content modification is not supported (the firewall policy that accepts the session must not include virus scanning, web filtering, DLP, application control, IPS, email filtering, SSL/SSH inspection, VoIP or ICAP)



If you disable anomaly checks by Intrusion Prevention (IPS), you can still enable hardware accelerated anomaly checks using the `fp-anomaly` field of the `config system interface` CLI command. See [“Offloading NP pre-IPS anomaly detection” on page 1076](#).

---

If a session or is not fast path ready, the FortiGate unit will not send the session key or IPsec SA key to the NP6 processor. Without the session key, all session key lookup by a network processor for incoming packets of that session fails, causing all session packets to be sent to the FortiGate unit’s main processing resources, and processed at normal speeds.

If a session is fast path ready, the FortiGate unit will send the session key or IPsec SA key to the network processor. Session key or IPsec SA key lookups then succeed for subsequent packets from the known session or IPsec SA.

## Packet fast path requirements

Packets within the session must then also meet packet requirements.

- Incoming packets must not be fragmented.
- Outgoing packets must not require fragmentation to a size less than 385 bytes. Because of this requirement, the configured MTU (Maximum Transmission Unit) for network processors' network interfaces must also meet or exceed the network processors' supported minimum MTU of 385 bytes.

## Mixing fast path and non-fast path traffic

If packet requirements are not met, an individual packet will be processed by the FortiGate CPU regardless of whether other packets in the session are offloaded to the NP6.

Also, in some cases, a protocol's session(s) may receive a mixture of offloaded and non-offloaded processing. For example, VoIP control packets may not be offloaded but VoIP data packets (voice packets) may be offloaded.

## Viewing your FortiGate NP6 processor configuration

Use the following command to view the NP6 processor configuration of your FortiGate unit:

```
diagnose npu np6 port-list
```

For example output of this command for different FortiGate models, see [“FortiGate NP6 architectures” on page 1089](#).

## Increasing NP6 offloading capacity using link aggregation groups (LAGs)

NP6 processors can offload sessions received by interfaces in link aggregation (LAG) groups. You can use link aggregation groups to offload more traffic that would exceed the capacity of a single FortiGate interface. For example, if you want to offload sessions on a 30 GB link you can add three 10-GB interfaces to a LAG group and send 30 GB of traffic to the LAG group.

Just like with normal interfaces, traffic accepted by a LAG group is offloaded by the NP6 processor connected to the interfaces in the LAG group that receive the traffic to be offloaded. If all interfaces in a LAG group are connected to the same NP6 processor, traffic received by the LAG group is offloaded by that NP6 processor. The amount of traffic that can be offloaded is limited by the capacity of the NP6 processor.

LAG groups can include interfaces connected to more than one NP6 processor. For example, adding a second NP6 processor to a LAG group effectively doubles the offloading capacity of the LAG group. Adding a third further increases capacity. Using LAG groups allows you to increase offloading capacity for incoming traffic by sharing the traffic load across multiple NP6 processors. This increase in capacity is supported by the integrated switch fabric that allows the NP6 processors to share session information.

The increase in offloading capacity may not actually be doubled by adding a second NP6 processor to a LAG group. Traffic and load conditions and other factors may limit the actual achieved offloading result.

## Configuring Inter-VDOM link acceleration with NP6 processors

FortiGate units with NP6 processors include inter-VDOM links that can be used to accelerate inter-VDOM link traffic.

- For a FortiGate unit with two NP6 processors there are two accelerated inter-VDOM links, each with two interfaces:
  - npu0\_vlink
    - npu0\_vlink0
    - npu0\_vlink1
  - npu1\_vlink
    - npu1\_vlink0
    - npu1\_vlink1

These interfaces are visible from the GUI and CLI. For a FortiGate unit with NP6 interfaces, enter the following CLI command to display the NP6-accelerated inter-VDOM links:

```
get system interface
...
== [npu0_vlink0]
name: npu0_vlink0 mode: static ip: 0.0.0.0 0.0.0.0 status: down
netbios-forward: disable type: physical sflow-sampler: disable
explicit-web-proxy: disable explicit-ftp-proxy: disable
mtu-override: disable wccp: disable drop-overlapped-fragment:
disable drop-fragment: disable
== [npu0_vlink1]
name: npu0_vlink1 mode: static ip: 0.0.0.0 0.0.0.0 status: down
netbios-forward: disable type: physical sflow-sampler: disable
explicit-web-proxy: disable explicit-ftp-proxy: disable
mtu-override: disable wccp: disable drop-overlapped-fragment:
disable drop-fragment: disable
== [npu1_vlink0]
name: npu1_vlink0 mode: static ip: 0.0.0.0 0.0.0.0 status: down
netbios-forward: disable type: physical sflow-sampler: disable
explicit-web-proxy: disable explicit-ftp-proxy: disable
mtu-override: disable wccp: disable drop-overlapped-fragment:
disable drop-fragment: disable
== [npu1_vlink1]
name: npu1_vlink1 mode: static ip: 0.0.0.0 0.0.0.0 status: down
netbios-forward: disable type: physical sflow-sampler: disable
explicit-web-proxy: disable explicit-ftp-proxy: disable
mtu-override: disable wccp: disable drop-overlapped-fragment:
disable drop-fragment: disable
...
```

By default the interfaces in each inter-VDOM link are assigned to the root VDOM. To use these interfaces to accelerate inter-VDOM link traffic, assign each interface in the pair to the VDOMs that you want to offload traffic between. For example, if you have added a VDOM named *New-VDOM* to a FortiGate unit with NP4 processors, you can go to *System > Network > Interfaces* and edit the *npu0-vlink1* interface and set the *Virtual Domain* to *New-VDOM*. This results in an accelerated inter-VDOM link between root and *New-VDOM*. You can also do this from the CLI:

```

config system interface
 edit npu0-vlink1
 set vdom New-VDOM
 end

```

## Using VLANs to add more accelerated Inter-VDOM links

You can add VLAN interfaces to the accelerated inter-VDOM links to create inter-VDOM links between more VDOMs. For the links to work, the VLAN interfaces must be added to the same inter-VDOM link, must be on the same subnet, and must have the same VLAN ID.

For example, to accelerate inter-VDOM link traffic between VDOMs named Marketing and Engineering using VLANs with VLAN ID 100 go to *System > Network > Interfaces* and select *Create New* to create the VLAN interface associated with the Marketing VDOM:

<b>Name</b>	Marketing-link
<b>Type</b>	VLAN
<b>Interface</b>	npu0_vlink0
<b>VLAN ID</b>	100
<b>Virtual Domain</b>	Marketing
<b>IP/Network Mask</b>	172.20.120.12/24

Create the inter-VDOM link associated with Engineering VDOM:

<b>Name</b>	Engineering-link
<b>Type</b>	VLAN
<b>Interface</b>	npu0_vlink1
<b>VLAN ID</b>	100
<b>Virtual Domain</b>	Engineering
<b>IP/Network Mask</b>	172.20.120.22/24

Or do the same from the CLI:

```

config system interface
 edit Marketing-link
 set vdom Marketing
 set ip 172.20.120.12/24
 set interface npu0_vlink0
 set vlanid 100
 next
 edit Engineering-link
 set vdom Engineering
 set ip 172.20.120.22/24
 set interface npu0_vlink1
 set vlanid 100

```

## Confirm that the traffic is accelerated

Use the following CLI commands to obtain the interface index and then correlate them with the session entries. In the following example traffic was flowing between new accelerated inter-VDOM links and physical ports port1 and port 2 also attached to the NP6 processor.

### **diagnose ip address list**

```
IP=172.31.17.76->172.31.17.76/255.255.252.0 index=5 devname=port1
IP=10.74.1.76->10.74.1.76/255.255.252.0 index=6 devname=port2
IP=172.20.120.12->172.20.120.12/255.255.255.0 index=55
 devname=IVL-VLAN1_ROOT
IP=172.20.120.22->172.20.120.22/255.255.255.0 index=56
 devname=IVL-VLAN1_VDOM1
```

### **diagnose sys session list**

```
session info: proto=1 proto_state=00 duration=282 expire=24 timeout=0
 session info: proto=1 proto_state=00 duration=124 expire=59
 timeout=0 flags=00000000 sockflag=00000000 sockport=0 av_idx=0
 use=3
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 policy_dir=0 tunnel=/
state=may_dirty npu
statistic(bytes/packets/allow_err): org=180/3/1 reply=120/2/1
 tuples=2
origin->sink: org pre->post, reply pre->post dev=55->5/5->55
 gwy=172.31.19.254/172.20.120.22
hook=post dir=org act=snat
 10.74.2.87:768->10.2.2.2:8(172.31.17.76:62464)
hook=pre dir=reply act=dnat
 10.2.2.2:62464->172.31.17.76:0(10.74.2.87:768)
misc=0 policy_id=4 id_policy_id=0 auth_info=0 chk_client_info=0 vd=0
serial=0000004e tos=ff/ff ips_view=0 app_list=0 app=0
dd_type=0 dd_mode=0
per_ip_bandwidth meter: addr=10.74.2.87, bps=880
npu_state=00000000
npu info: flag=0x81/0x81, offload=8/8, ips_offload=0/0,
 epid=160/218, ipid=218/160, vlan=32769/0

session info: proto=1 proto_state=00 duration=124 expire=20 timeout=0
 flags=00000000 sockflag=00000000 sockport=0 av_idx=0 use=3
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 policy_dir=0 tunnel=/
state=may_dirty npu
statistic(bytes/packets/allow_err): org=180/3/1 reply=120/2/1
 tuples=2
origin->sink: org pre->post, reply pre->post dev=6->56/56->6
 gwy=172.20.120.12/10.74.2.87
```

```
hook=pre dir=org act=noop 10.74.2.87:768->10.2.2.2:8(0.0.0.0:0)
hook=post dir=reply act=noop 10.2.2.2:768->10.74.2.87:0(0.0.0.0:0)
misc=0 policy_id=3 id_policy_id=0 auth_info=0 chk_client_info=0 vd=1
serial=0000004d tos=ff/ff ips_view=0 app_list=0 app=0
dd_type=0 dd_mode=0
per_ip_bandwidth meter: addr=10.74.2.87, bps=880
npu_state=00000000
npu info: flag=0x81/0x81, offload=8/8, ips_offload=0/0,
epid=219/161, ipid=161/219, vlan=0/32769
total session 2
```



# FortiGate NP6 architectures

Many FortiGate models can offload some types of network traffic processing from main processing resources to specialized network processors. If your network has a significant volume of traffic that is suitable for offloading, this hardware acceleration can significantly improve your network throughput.

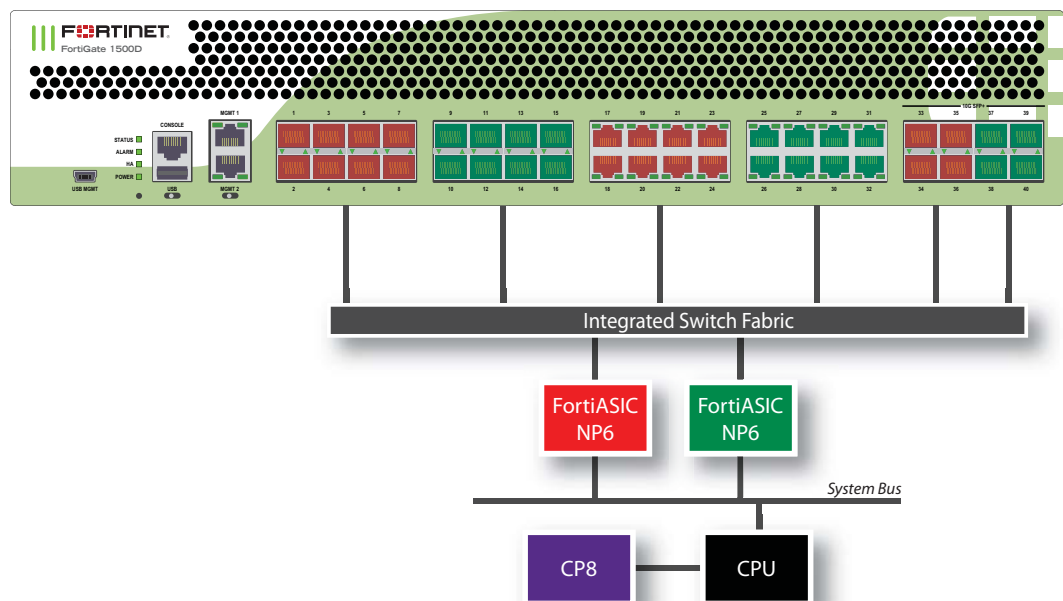
This chapter shows the fastpath architecture for the following FortiGate units:

- [FortiGate-1500D fast path architecture](#)
- [FortiGate-3700D fast path architecture](#)

## FortiGate-1500D fast path architecture

The FortiGate-1500D features two NP6 processors.

- Eight SFP 1Gb interfaces (port1-port8), eight RJ-45 Ethernet ports (port17-24) and four SFP+ 10Gb interfaces (port33-port36) share connections to the first NP6 processor.
- Eight SFP 1Gb interfaces (port9-port16), eight RJ-45 Ethernet ports (port25-32) and four SFP+ 10Gb interfaces (port37-port40) share connections to the second NP6 processor.



You can use the following command to display the FortiGate-1500D NP6 configuration. The command output shows two NP6s named NP6\_0 and NP6\_1. The output also shows the interfaces (ports) connected to each NP6.

```
diagnose npu np6 port-list
Chip XAUI Ports Max Cross-chip
 Speed offloading

np6_0 0 port1 10G Yes
 0 port5 10G Yes
 0 port17 10G Yes
 0 port21 10G Yes
 0 port33 10G Yes
 1 port2 10G Yes
 1 port6 10G Yes
 1 port18 10G Yes
 1 port22 10G Yes
 1 port34 10G Yes
 2 port3 10G Yes
 2 port7 10G Yes
 2 port19 10G Yes
 2 port23 10G Yes
 2 port35 10G Yes
 3 port4 10G Yes
 3 port8 10G Yes
 3 port20 10G Yes
 3 port24 10G Yes
 3 port36 10G Yes

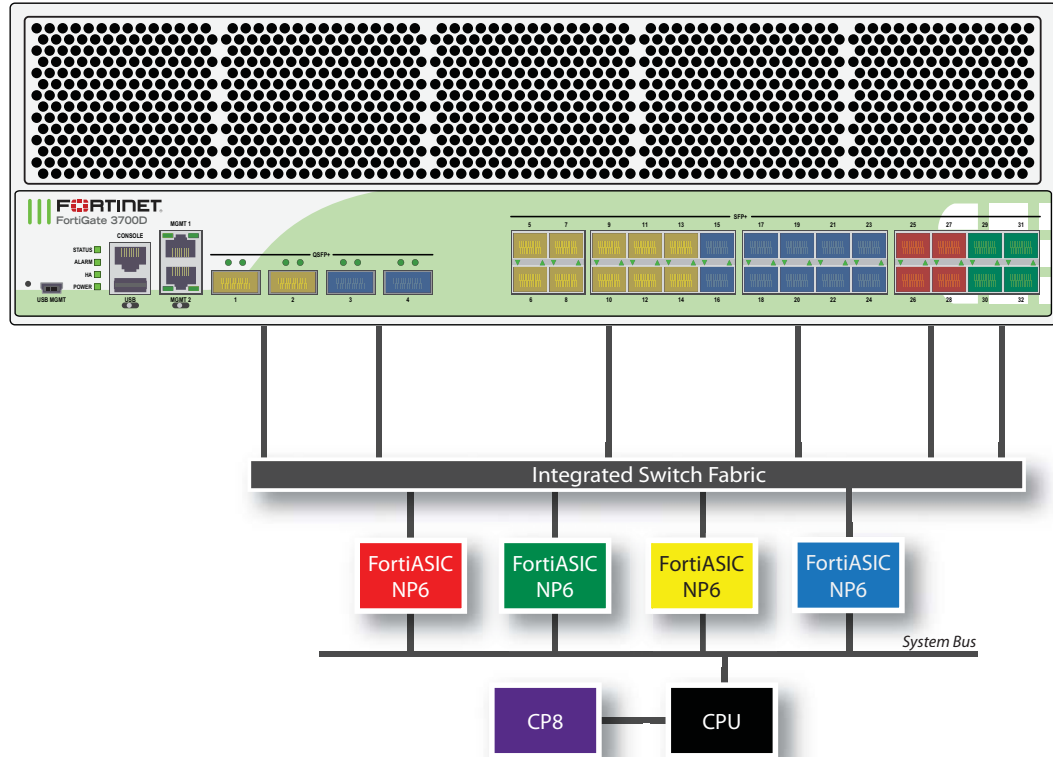
np6_1 0 port9 10G Yes
 0 port13 10G Yes
 0 port25 10G Yes
 0 port29 10G Yes
 0 port37 10G Yes
 1 port10 10G Yes
 1 port14 10G Yes
 1 port26 10G Yes
 1 port30 10G Yes
 1 port38 10G Yes
 2 port11 10G Yes
 2 port15 10G Yes
 2 port27 10G Yes
 2 port31 10G Yes
 2 port39 10G Yes
 3 port12 10G Yes
 3 port16 10G Yes
 3 port28 10G Yes
 3 port32 10G Yes
 3 port40 10G Yes

```

## FortiGate-3700D fast path architecture

The FortiGate-3700D features four NP6 processors.

- Port25 through port28, SFP+ 10Gb interfaces, share connections to the first NP6 processor.
- Port29 through port32, SFP+ 10Gb interfaces, share connections to the second NP6 processor.
- Ten SFP+ 10Gb interfaces, port5 through port14, and two 40Gb QSFP interfaces, port1 and port2 share connections to the third NP6 processor.
- Ten SFP+ 10Gb interfaces, port15 through port24, and two 40Gb QSFP interfaces, port3 and port4 share connections to the fourth NP6 processor.



Ports 25 to 32 can be used for low latency offloading. As long as traffic enters and exits the FortiGate-3700D through ports connected to the same NP6 processor and using these low latency ports the traffic will be offloaded and have lower latency than other NP6 offloaded traffic. Latency is reduced by bypassing the integrated switch fabric. Specifically:

- Port25 through port28, share connections to the first NP6 processor so sessions entering one of these ports and exiting through another will experience low latency
- Port29 through port32, share connections to the second NP6 processor so sessions entering one of these ports and exiting through another will experience low latency

You can use the following command to display the FortiGate-3700D NP6 configuration. The command output shows four NP6s named NP6\_0, NP6\_1, NP6\_2, and NP6\_3. The output also shows the interfaces (ports) connected to each NP6.

```
diag npu np6 port-list
Chip XAUI Ports Max Cross-chip
Speed offloading

np6_0 0 port25 10G Yes
1 port26 10G Yes
2 port27 10G Yes
3 port28 10G Yes

np6_1 0 port29 10G Yes
1 port30 10G Yes
2 port31 10G Yes
3 port32 10G Yes

np6_2 0 port5 10G Yes
0 port9 10G Yes
0 port13 10G Yes
1 port6 10G Yes
1 port10 10G Yes
1 port14 10G Yes
2 port7 10G Yes
2 port11 10G Yes
3 port8 10G Yes
3 port12 10G Yes
0-3 port1 40G Yes
0-3 port2 40G Yes

np6_3 0 port15 10G Yes
0 port19 10G Yes
0 port23 10G Yes
1 port16 10G Yes
1 port20 10G Yes
1 port24 10G Yes
2 port17 10G Yes
2 port21 10G Yes
3 port18 10G Yes
3 port22 10G Yes
0-3 port3 40G Yes
0-3 port4 40G Yes

```

# NP4 Acceleration

NP4 network processors provide fastpath acceleration by offloading communication sessions from the FortiGate CPU. When the first packet of a new session is received by an interface connected to an NP4 processor, just like any session connecting with any FortiGate interface, the session is forwarded to the FortiGate CPU where it is matched with a security policy. If the session is accepted by a security policy and if the session can be offloaded its session key is copied to the NP4 processor that received the packet. All of the rest of the packets in the session are intercepted by the NP4 processor and fast-pathed out of the FortiGate unit to their destination without ever passing through the FortiGate CPU. The result is enhanced network performance provided by the NP4 processor plus the network processing load is removed from the CPU. In addition, the NP4 processor can handle some CPU intensive tasks, like IPsec VPN encryption/decryption.

Session keys (and IPsec SA keys) are stored in the memory of the NP4 processor that is connected to the interface that received the packet that started the session. All sessions are fast-pathed and accelerated, even if they exit the FortiGate unit through an interface connected to another NP4. The key to making this possible is the Integrated Switch Fabric (ISF) that connects the NP4s and the FortiGate unit interfaces together. The ISF allows any port connectivity. All ports and NP4s can communicate with each other over the ISF.

There are no special ingress and egress fast path requirements because traffic enters and exits on interfaces connected to the same ISF. Most FortiGate models with multiple NP4 processors connect all interfaces and NP4 processors to the same ISF (except management interfaces) so this should not ever be a problem.

There is one limitation to keep in mind; the capacity of each NP4 processor. An individual NP4 processor has a capacity of 20 Gbps (10 Gbps ingress and 10 Gbps egress). Once an NP4 processor hits its limit, sessions that are over the limit are sent to the CPU. You can avoid this problem by as much as possible distributing incoming sessions evenly among the NP4 processors. To be able to do this you need to be aware of which interfaces connect to which NP4 processors and distribute incoming traffic accordingly.

Some FortiGate units contain one NP4 processor with all interfaces connected to it and to the ISF. As a result, offloading is supported for traffic between any pair of interfaces.

Some FortiGate units include NP4Lite processors. These network processors have the same functionality and limitations as NP4 processors but with about half the performance. NP4lite processors can be found in mid-range FortiGate models such as the FortiGate-200D and 240D.

This chapter contains the following topics:

- [Viewing your FortiGate's NP4 configuration](#)
- [Configuring NP4 traffic offloading](#)
- [NP4 traffic shaping offloading](#)
- [NP4 IPsec VPN offloading](#)
- [NP4 IPsec VPN offloading configuration example](#)
- [Configuring Inter-VDOM link acceleration with NP4 processors](#)

## Viewing your FortiGate's NP4 configuration

To list the NP4 network processors on your FortiGate unit, use the following CLI command.

```
get hardware npu np4 list
```

The output lists the interfaces that have NP4 processors. For example, for a FortiGate-5001C:

```
get hardware npu np4 list
```

ID	Model	Slot	Interface
0	On-board		port1 port2 port3 port4 fabric1 base1 npu0-vlink0 npu0-vlink1
1	On-board		port5 port6 port7 port8 fabric2 base2 npu1-vlink0 npu1-vlink1

Depending on the product, NP4 network processors may or may not be directly connected to each other on the circuit board through an EEI (Enhanced Extension Interface).

Directly connected network processors have an EEI, and can pass traffic between them without involving the FortiGate unit's main processing resources.

Indirectly connected network processors have no EEI, and cannot pass traffic between them without involving the FortiGate unit's main processing resources.

Sessions can only be offloaded if both the source and destination port are connected to the same network processor or directly (EEI) connected network processor pair.

For information about the network processors in specific FortiGate models, see [“FortiGate NP4 architectures” on page 1105](#).

### NP4lite CLI commands (disabling NP4Lite offloading)

If your FortiGate unit includes an NP4Lite processor the following commands will be available:

- Use the following command to disable or enable NP4Lite offloading. By default NP4lite offloading is enabled. If you want to disable NP4Lite offloading to diagnose a problem enter:

```
diagnose npu nplite fastpath disable
```

This command disables NP4Lite offloading until your FortiGate reboots. You can also re-enable offloading by entering the following command:

```
diagnose npu nplite fastpath enable
```

- NP4lite debug command. Use the following command to debug NP4Lite operation:

```
diagnose npl npl_debug {<parameters>}
```

## Configuring NP4 traffic offloading

Offloading traffic to a network processor requires that the FortiGate unit configuration and the traffic itself is suited to hardware acceleration. There are requirements for path the sessions and the individual packets.

### NP4 session fast path requirements

Sessions must be fast path ready. Fast path ready session characteristics are:

- Layer 2 type/length must be 0x0800 (IEEE 802.1q VLAN specification is supported); link aggregation between any network interfaces sharing the same network processor(s) may be used (IEEE 802.3ad specification is supported)
- Layer 3 protocol must be IPv4

- Layer 4 protocol must be UDP, TCP or ICMP
- Layer 3 / Layer 4 header or content modification must not require a session helper (for example, SNAT, DNAT, and TTL reduction are supported, but application layer content modification is not supported)
- Firewall policies must not include proxy-based or flow-based security features (antivirus, web filtering, email filtering, IPS, application control, or DLP)
- Origin must not be local host (the FortiGate unit)



If you disable anomaly checks by Intrusion Prevention (IPS), you can still enable hardware accelerated anomaly checks using the `fp-anomaly` field of the `config system interface` CLI command. See [“Offloading NP pre-IPS anomaly detection” on page 1076](#).

If a session is not fast path ready, the FortiGate unit will not send the session key to the network processor(s). Without the session key, all session key lookup by a network processor for incoming packets of that session fails, causing all session packets to be sent to the FortiGate unit’s main processing resources, and processed at normal speeds.

If a session is fast path ready, the FortiGate unit will send the session key to the network processor(s). Session key lookup then succeeds for subsequent packets from the known session.

## Packet fast path requirements

Packets within the session must then also meet packet requirements.

- Incoming packets must not be fragmented.
- Outgoing packets must not require fragmentation to a size less than 385 bytes. Because of this requirement, the configured MTU (Maximum Transmission Unit) for network processors’ network interfaces must also meet or exceed the network processors’ supported minimum MTU of 385 bytes.

If packet requirements are not met, an individual packet will use FortiGate unit main processing resources, regardless of whether other packets in the session are offloaded to the specialized network processor(s).

In some cases, due to these requirements, a protocol’s session(s) may receive a mixture of offloaded and non-offloaded processing.

For example, FTP uses two connections: a control connection and a data connection. The control connection requires a session helper, and cannot be offloaded, but the data connection does not require a session helper, and can be offloaded. Within the offloadable data session, fragmented packets will not be offloaded, but other packets will be offloaded.

Some traffic types differ from general offloading requirements, but still utilize some of the network processors’ encryption and other capabilities. Exceptions include IPsec traffic and active-active high availability (HA) load balanced traffic.

## Mixing fast path and non-fast path traffic

If packet requirements are not met, an individual packet will be processed by the FortiGate CPU regardless of whether other packets in the session are offloaded to the NP4.

Also, in some cases, a protocol’s session(s) may receive a mixture of offloaded and non-offloaded processing. For example, FTP control packets may not be offloaded but FTP data packets (voice packets) may be offloaded.

## NP4 traffic shaping offloading

Accelerated Traffic shaping is supported with the following limitations.

- NP4 processors support policy-based traffic shaping. However, fast path traffic and traffic handled by the FortiGate CPU (slow path) are controlled separately, which means the policy setting on fast path does not consider the traffic on the slow path.
- The port based traffic policing as defined by the inbandwidth and outbandwidth CLI commands is not supported.
- DSCP configurations are supported.
- 4QoS in general is not supported.

You can also use the traffic shaping features of the FortiGate unit's main processing resources by disabling the NP4 acceleration. See [“Disabling CP offloading” on page 1070](#).

## NP4 IPsec VPN offloading

NP4 processors improve IPsec tunnel performance by offloading IPsec encryption and decryption.

Requirements for hardware accelerated IPsec encryption or decryption are a modification of general offloading requirements. Differing characteristics are:

- Origin can be local host (the FortiGate unit)
- In Phase 1 configuration, Local Gateway IP must be specified as an IP address of a network interface for a port attached to a network processor
- SA must have been received by the network processor
- in Phase 2 configuration:
  - encryption algorithm must be DES, 3DES, AES-128, AES-192, AES-256, or null
  - authentication must be MD5, SHA1, or null
  - if encryption is null, authentication must not also be null
  - if replay detection is enabled, `enc-offload-antireplay` must also be `enable` in the CLI



If replay detection is enabled in the Phase 2 configuration, you can enable or disable IPsec encryption and decryption offloading from the CLI. Performance varies by those CLI options and the percentage of packets requiring encryption or decryption. For details, see [“Configuring NP accelerated VPN encryption/decryption offloading” on page 1078](#).

---

To apply hardware accelerated encryption and decryption, the FortiGate unit's main processing resources must first perform Phase 1 negotiations to establish the security association (SA). The SA includes cryptographic processing instructions required by the network processor, such as which encryption algorithms must be applied to the tunnel. After ISAKMP negotiations, the FortiGate unit's main processing resources send the SA to the network processor, enabling the network processor to apply the negotiated hardware accelerated encryption or decryption to tunnel traffic.

Possible accelerated cryptographic paths are:

- IPsec decryption offload
  - Ingress ESP packet > Offloaded decryption > Decrypted packet egress (fast path)
  - Ingress ESP packet > Offloaded decryption > Decrypted packet to FortiGate unit's main processing resources



- IPsec encryption offload
  - Ingress packet > Offloaded encryption > Encrypted (ESP) packet egress (fast path)
  - Packet from FortiGate unit's main processing resources > Offloaded encryption > Encrypted (ESP) packet egress

## NP4 IPsec VPN offloading configuration example

Hardware accelerated IPsec processing, involving either partial or full offloading, can be achieved in either tunnel or interface mode IPsec configurations.

To achieve offloading for both encryption and decryption:

- In Phase 1 configuration's Advanced section, Local Gateway IP must be specified as an IP address of a network interface associated with a port attached to a network processor. (In other words, if Phase 1's Local Gateway IP is Main Interface IP, or is specified as an IP address that is not associated with a network interface associated with a port attached to a network processor, IPsec network processing is not offloaded.)
- In Phase 2 configuration's P2 Proposal section, if the checkbox "Enable replay detection" is enabled, `enc-offload-antireplay` and `dec-offload-antireplay` must be set to `enable` in the CLI.
- `offload-ipsec-host` must be set to `enable` in the CLI.


This section contains example IPsec configurations whose IPsec encryption and decryption processing is hardware accelerated by an NP4 unit contained in a FortiGate-5001B at both ends of the VPN tunnel.



Hardware accelerated IPsec VPN does not require both tunnel endpoints to have the same network processor model. However, if hardware is not symmetrical, the packet forwarding rate is limited by the slower side.

---

**Figure 183:**Example network topology for offloaded IPsec processing



**Table 54:** Example ports and IP addresses for offloaded IPsec processing

	FortiGate_1		FortiGate_2	
	Port	IP	Port	IP
<b>IPsec tunnel</b>	FortiGate-5001B port 2	3.3.3.1/24	FortiGate-5001B port 2	3.3.3.2/24
<b>Protected network</b>	FortiGate-5001B port 1	1.1.1.0/24	FortiGate-5001B port 1	2.2.2.0/24

This section includes the following topics:

- [Accelerated policy mode IPsec configuration](#)
- [Accelerated interface mode IPsec configuration](#)

## Accelerated policy mode IPsec configuration

The following steps create a hardware accelerated policy mode IPsec tunnel between two FortiGate-5001B units, each containing two NP4 processors, the first of which will be used.

### To configure hardware accelerated policy mode IPsec

1. On FortiGate\_1, go to *VPN > IPsec > Auto Key (IKE)*.
2. Configure Phase 1.  
For tunnel mode IPsec and for hardware acceleration, specifying the Local Gateway IP is required.  
Select *Advanced*. In the Local Gateway IP section, select *Specify* and type the VPN IP address 3.3.3.2, which is the IP address of FortiGate\_2's FortiGate-ASM-FB4 module port 2.
3. Configure Phase 2.
4. Select *Enable replay detection*.
5. Use the following command to enable offloading antireplay packets:  

```
config system npu
 set enc-offload-antireplay enable
end
```

For details on encryption and decryption offloading options available in the CLI, see [“Configuring NP accelerated VPN encryption/decryption offloading”](#) on page 1078.

6. Go to *Policy > Policy > Policy*.
7. Configure a policy to apply the Phase 1 IPsec tunnel you configured in step 2 to traffic between FortiGate-5001B ports 1 and 2.
8. Go to *Router > Static > Static Route*.
9. Configure a static route to route traffic destined for FortiGate\_2's protected network to VPN IP address of FortiGate\_2's VPN gateway, 3.3.3.2, through the FortiGate-5001B port2.

You can also configure the static route using the following CLI command:

```
config router static
 edit 2
 set device "AMC-SW1/2"
 set dst 2.2.2.0 255.255.255.0
 set gateway 3.3.3.2
 end
```

10. On FortiGate\_2, go to *VPN > IPsec > Auto Key (IKE)*.

11. Configure Phase 1.

For tunnel mode IPsec and for hardware acceleration, specifying the Local Gateway IP is required.

Select Advanced. In the Local Gateway IP section, select Specify and type the VPN IP address 3.3.3.1, which is the IP address of FortiGate\_1's port2.

12. Configure Phase 2.

13. Select *Enable replay detection*.

14. Use the following command to enable offloading antireplay packets:

```
config system npu
 set enc-offload-antireplay enable
end
```

For details on encryption and decryption offloading options available in the CLI, see [“Configuring NP accelerated VPN encryption/decryption offloading” on page 1078](#).

15. Go to *Policy > Policy > Policy*.

16. Configure a policy to apply the Phase 1 IPsec tunnel you configured in step 9 to traffic between FortiGate-5001B ports 1 and 2.

17. Go to *Router > Static > Static Route*.

18. Configure a static route to route traffic destined for FortiGate\_1's protected network to VPN IP address of FortiGate\_1's VPN gateway, 3.3.3.1, through the FortiGate-5001B port2.

You can also configure the static route using the following CLI commands:

```
config router static
 edit 2
 set device "AMC-SW1/2"
 set dst 1.1.1.0 255.255.255.0
 set gateway 3.3.3.1
 end
```

19. Activate the IPsec tunnel by sending traffic between the two protected networks.

To verify tunnel activation, go to *VPN > Monitor > IPsec Monitor*.

## Accelerated interface mode IPsec configuration

The following steps create a hardware accelerated interface mode IPsec tunnel between two FortiGate units, each containing a FortiGate-ASM-FB4 module.

## To configure hardware accelerated interface mode IPsec

1. On FortiGate\_1, go to *VPN > IPsec > Auto Key (IKE)*.

2. Configure Phase 1.

For interface mode IPsec and for hardware acceleration, the following settings are required.

- Select *Advanced*.
- Enable the checkbox “Enable IPsec Interface Mode.”
- In the Local Gateway IP section, select *Specify* and type the VPN IP address 3.3.3.2, which is the IP address of FortiGate\_2’s port 2.

3. Configure Phase 2.

4. Select *Enable replay detection*.

5. Use the following command to enable offloading antireplay packets:

```
config system npu
 set enc-offload-antireplay enable
end
```

For details on encryption and decryption offloading options available in the CLI, see [“Configuring NP accelerated VPN encryption/decryption offloading” on page 1078](#).

6. Go to *Policy > Policy > Policy*.

7. Configure two policies (one for each direction) to apply the Phase 1 IPsec configuration you configured in step 2 to traffic leaving from or arriving on FortiGate-ASM-FB4 module port 1.

8. Go to *Router > Static > Static Route*.

9. Configure a static route to route traffic destined for FortiGate\_2’s protected network to the Phase 1 IPsec device, FGT\_1\_IPsec.

You can also configure the static route using the following CLI commands:

```
config router static
 edit 2
 set device "FGT_1_IPsec"
 set dst 2.2.2.0 255.255.255.0
 end
```

10. On FortiGate\_2, go to *VPN > IPsec > Auto Key (IKE)*.

11. Configure Phase 1.

For interface mode IPsec and for hardware acceleration, the following settings are required.

- Enable the checkbox “Enable IPsec Interface Mode.”
- In the Local Gateway IP section, select *Specify* and type the VPN IP address 3.3.3.1, which is the IP address of FortiGate\_1’s FortiGate-5001B port 2.

12. Configure Phase 2.

13. Select *Enable replay detection*.

14. Use the following command to enable offloading antireplay packets:

```
config system npu
 set enc-offload-antireplay enable
end
```

For details on encryption and decryption offloading options available in the CLI, see [“Configuring NP accelerated VPN encryption/decryption offloading” on page 1078](#).

15. Go to *Policy > Policy > Policy*.

16. Configure two policies (one for each direction) to apply the Phase 1 IPsec configuration you configured in step 9 to traffic leaving from or arriving on FortiGate-5001B port 1.

17. Go to *Router > Static > Static Route*.

**18.** Configure a static route to route traffic destined for FortiGate\_1's protected network to the Phase 1 IPsec device, FGT\_2\_IPsec.

You can also configure the static route using the following CLI commands:

```
config router static
 edit 2
 set device "FGT_2_IPsec"
 set dst 1.1.1.0 255.255.255.0
 next
end
```

**19.** Activate the IPsec tunnel by sending traffic between the two protected networks.

To verify tunnel activation, go to *VPN > Monitor > IPsec Monitor*.

## Configuring Inter-VDOM link acceleration with NP4 processors

FortiGate units with NP4 processors include inter-VDOM links that can be used to accelerate inter-VDOM link traffic.

- For a FortiGate unit with two NP4 processors there are also two inter-VDOM links, each with two interfaces:
  - npu0-vlink
    - npu0-vlink0
    - npu0-vlink1
  - npu1-vlink
    - npu1-vlink0
    - npu1-vlink1

These interfaces are visible from the GUI and CLI. For a FortiGate unit with NP4 interfaces, enter the following CLI command (output shown for a FortiGate-5001B):

```
get hardware npu np4 list
ID Model Slot Interface
0 On-board 0 port1 port2 port3 port4
 fabric1 base1 npu0-vlink0 npu0-vlink1
1 On-board 1 port5 port6 port7 port8
 fabric2 base2 npu1-vlink0 npu1-vlink1
```

By default the interfaces in each inter-VDOM link are assigned to the root VDOM. To use these interfaces to accelerate inter-VDOM link traffic, assign each interface in a pair to the VDOMs that you want to offload traffic between. For example, if you have added a VDOM named New-VDOM to a FortiGate unit with NP4 processors, you can go to *System > Network > Interfaces* and edit the *npu0-vlink1* interface and set the *Virtual Domain* to *New-VDOM*. This results in an inter-VDOM link between root and New-VDOM. You can also do this from the CLI:

```
config system interface
 edit npu0-vlink1
 set vdom New-VDOM
 end
```

## Using VLANs to add more accelerated Inter-VDOM links

You can add VLAN interfaces to the accelerated inter-VDOM links to create inter-VDOM links between more VDOMs. For the links to work, the VLAN interfaces must be added to the same inter-VDOM link, must be on the same subnet, and must have the same VLAN ID.

For example, to accelerate inter-VDOM link traffic between VDOMs named Marketing and Engineering using VLANs with VLAN ID 100 go to *System > Network > Interfaces* and select *Create New* to create the VLAN interface associated with the Marketing VDOM:

<b>Name</b>	Marketing-link
<b>Type</b>	VLAN
<b>Interface</b>	npu0-vlink0
<b>VLAN ID</b>	100
<b>Virtual Domain</b>	Marketing
<b>IP/Network Mask</b>	172.20.120.12/24

Create the inter-VDOM link associated with Engineering VDOM:

<b>Name</b>	Engineering-link
<b>Type</b>	VLAN
<b>Interface</b>	npu0-vlink1
<b>VLAN ID</b>	100
<b>Virtual Domain</b>	Engineering
<b>IP/Network Mask</b>	172.20.120.22/24

Or do the same from the CLI:

```
config system interface
 edit Marketing-link
 set vdom Marketing
 set ip 172.20.120.12/24
 set interface npu0-vlink0
 set vlanid 100
 next
 edit Engineering-link
 set vdom Engineering
 set ip 172.20.120.22/24
 set interface npu0-vlink1
 set vlanid 100
```

## Confirm that the traffic is accelerated

Use the following CLI commands to obtain the interface index and then correlate them with the session entries. In the following example traffic was flowing between new accelerated inter-VDOM links and physical ports port1 and port 2 also attached to the NP4 processor.

```
diagnose ip address list
IP=172.31.17.76->172.31.17.76/255.255.252.0 index=5 devname=port1
IP=10.74.1.76->10.74.1.76/255.255.252.0 index=6 devname=port2
IP=172.20.120.12->172.20.120.12/255.255.255.0 index=55
devname=IVL-VLAN1_ROOT
```

IP=172.20.120.22->172.20.120.22/255.255.255.0 index=56  
devname=IVL-VLAN1\_VDOM1

**diagnose sys session list**

session info: proto=1 proto\_state=00 duration=282 expire=24 timeout=0  
session info: proto=1 proto\_state=00 duration=124 expire=59  
timeout=0 flags=00000000 sockflag=00000000 sockport=0 av\_idx=0  
use=3  
origin-shaper=  
reply-shaper=  
per\_ip\_shaper=  
ha\_id=0 policy\_dir=0 tunnel=/  
state=may\_dirty npu  
statistic(bytes/packets/allow\_err): org=180/3/1 reply=120/2/1  
tuples=2  
origin->sink: org pre->post, reply pre->post dev=55->5/5->55  
gwy=172.31.19.254/172.20.120.22  
hook=post dir=org act=snat  
10.74.2.87:768->10.2.2.2:8(172.31.17.76:62464)  
hook=pre dir=reply act=dnat  
10.2.2.2:62464->172.31.17.76:0(10.74.2.87:768)  
misc=0 policy\_id=4 id\_policy\_id=0 auth\_info=0 chk\_client\_info=0 vd=0  
serial=0000004e tos=ff/ff ips\_view=0 app\_list=0 app=0  
dd\_type=0 dd\_mode=0  
per\_ip\_bandwidth meter: addr=10.74.2.87, bps=880  
npu\_state=00000000  
**npu info: flag=0x81/0x81, offload=4/4, ips\_offload=0/0,  
epid=160/218, ipid=218/160, vlan=32769/0**

session info: proto=1 proto\_state=00 duration=124 expire=20 timeout=0  
flags=00000000 sockflag=00000000 sockport=0 av\_idx=0 use=3  
origin-shaper=  
reply-shaper=  
per\_ip\_shaper=  
ha\_id=0 policy\_dir=0 tunnel=/  
state=may\_dirty npu  
statistic(bytes/packets/allow\_err): org=180/3/1 reply=120/2/1  
tuples=2  
origin->sink: org pre->post, reply pre->post dev=6->56/56->6  
gwy=172.20.120.12/10.74.2.87  
hook=pre dir=org act=noop 10.74.2.87:768->10.2.2.2:8(0.0.0.0:0)  
hook=post dir=reply act=noop 10.2.2.2:768->10.74.2.87:0(0.0.0.0:0)  
misc=0 policy\_id=3 id\_policy\_id=0 auth\_info=0 chk\_client\_info=0 vd=1  
serial=0000004d tos=ff/ff ips\_view=0 app\_list=0 app=0  
dd\_type=0 dd\_mode=0  
per\_ip\_bandwidth meter: addr=10.74.2.87, bps=880  
npu\_state=00000000  
**npu info: flag=0x81/0x81, offload=4/4, ips\_offload=0/0,  
epid=219/161, ipid=161/219, vlan=0/32769**

total session 2



# FortiGate NP4 architectures

This chapter shows the NP4 architecture for the following FortiGate units:


- FortiGate-600C
- FortiGate-800C
- FortiGate-1000C
- FortiGate-1240B
- FortiGate-3040B
- FortiGate-3140B
- FortiGate-3140B – load balance mode
- FortiGate-3240C
- FortiGate-3600C
- FortiGate-3950B and FortiGate-3951B
- FortiGate-5001C
- FortiGate-5001B

And includes the following reference information:

- [Setting switch-mode mapping on the ADM-XD4](#)

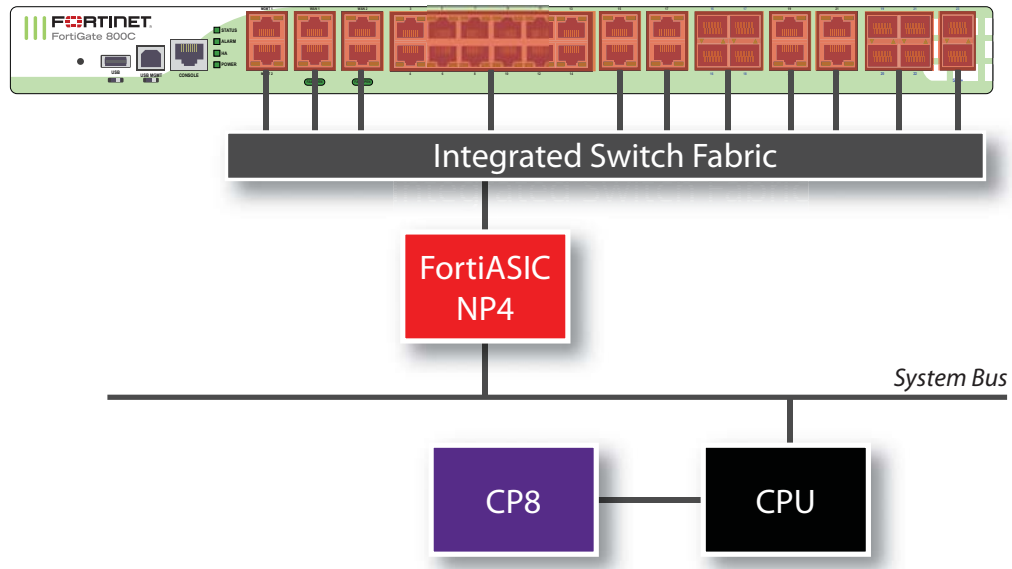
## FortiGate-600C

The FortiGate-600C features one NP4 processor. All the ports are connected to this NP4 over the Integrated Switch Fabric. Port1 and port2 are dual failopen redundant RJ-45 ports. Port3-port22 are RJ-45 ethernet ports, and there are four 1Gb SFP interface ports duplicating the port19-port22 connections.



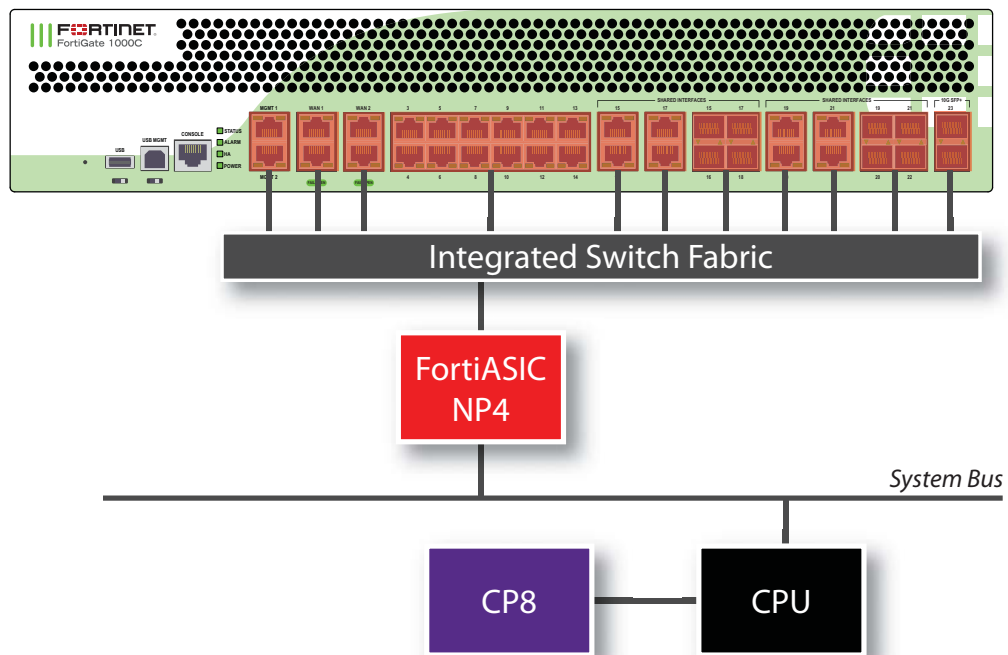
## FortiGate-800C

The FortiGate-800C features one NP4 processor. All the ports are connected to this NP4. Port1 and port2 are dual failopen redundant RJ-45 ports. Port3-port22 are RJ-45 ethernet ports, and there are eight 1Gb SFP interface ports duplicating the port15-18 and port19-port22 connections. There are also two 10Gb SFP+ ports, port23 and port24.



## FortiGate-1000C

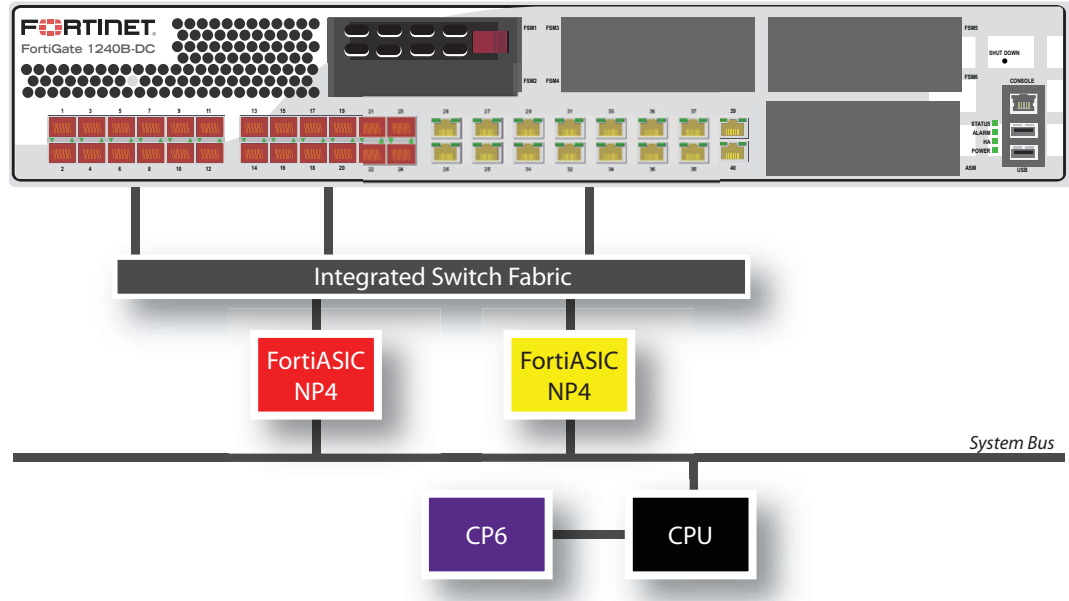
The FortiGate-1000C features one NP4 processor. All the ports are connected to this NP4. Port1 and port2 are dual failopen redundant RJ-45 ports. Port3-port22 are RJ-45 ethernet ports, and there are eight 1Gb SFP interface ports duplicating the port15-18 and port19-port22 connections. There are also two 10Gb SFP+ ports, port23 and port24.



## FortiGate-1240B

The FortiGate-1240B features two NP4 processors:

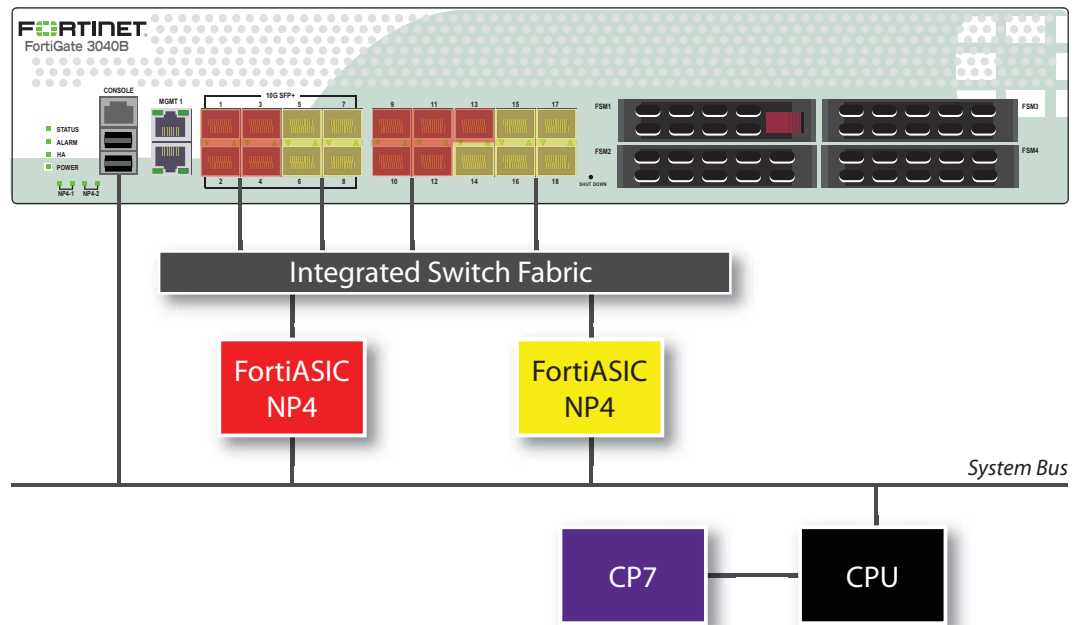
- Port1-port24 are 1Gb SFP interfaces connected to one NP4 processor.
- Port25-port40 are RJ-45 ethernet ports, connected to the other NP4 processor.



## FortiGate-3040B

The FortiGate-3040B features two NP4 processors:

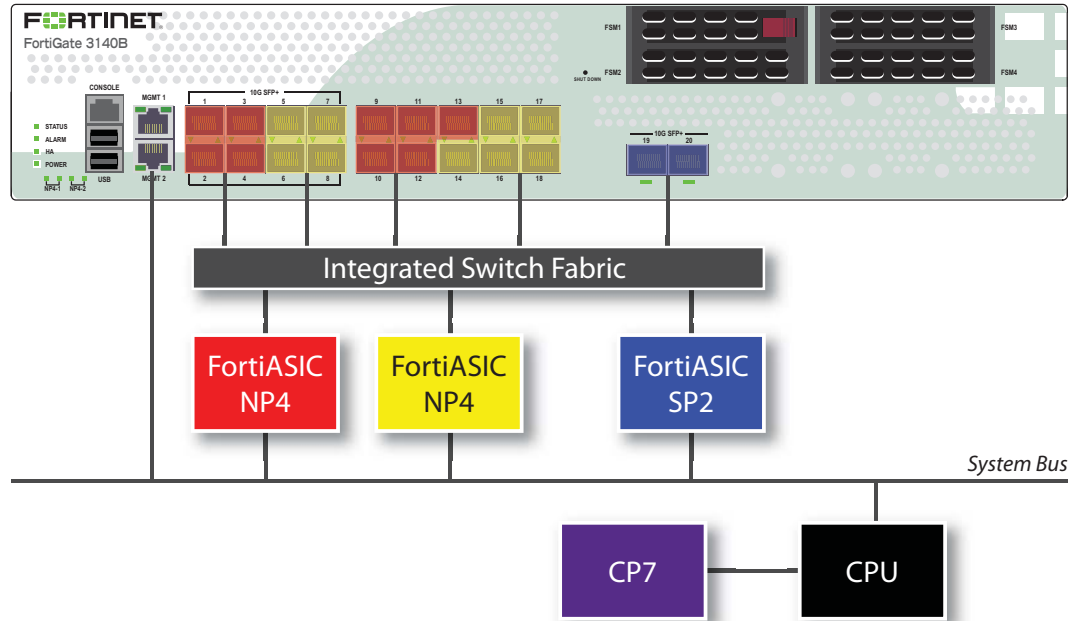
- The 10Gb interfaces, port1, port2, port3, port4, and the 1Gb interfaces, port9, port10, port11, port12, port13, share connections to one NP4 processor.
- The 10Gb interfaces, port5, port6, port7, port8, and the 1Gb interfaces, port14, port15, port16, port17, port18, share connections to the other NP4 processor.



# FortiGate-3140B

The FortiGate-3140B features two NP4 processors and one SP2 processor:

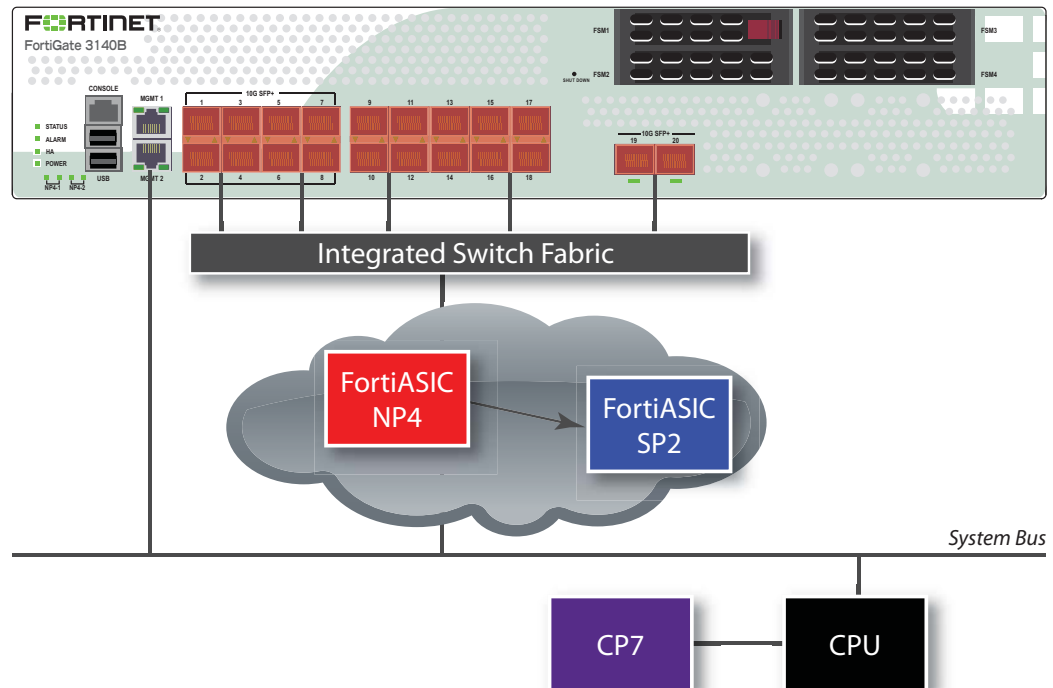
- The 10Gb interfaces, port1, port2, port3, port4, and the 1Gb interfaces, port9, port10, port11, port12, port13, share connections to one NP4 processor.
- The 10Gb interfaces, port5, port6, port7, port8, and the 1Gb interfaces, port14, port15, port16, port17, port18, share connections to the other NP4 processor.
- The 10Gb interfaces, port19 and port20, share connections to the SP2 processor.



## FortiGate-3140B — load balance mode

The FortiGate-3140B load balance mode allows you increased flexibility in how you use the interfaces on the FortiGate unit. When enabled, traffic between any two interfaces (excluding management and console) is accelerated. Traffic is not limited to entering and leaving the FortiGate unit in specific interface groupings to benefit from NP4 and SP2 acceleration. You can use any pair of interfaces.

Security acceleration in this mode is limited, however. Only IPS scanning is accelerated in load balance mode.



To enable this feature, issue this CLI command.

```
config system global
 set sp-load-balance enable
end
```

The FortiGate unit will then restart.

To return to the default mode, issue this CLI command.

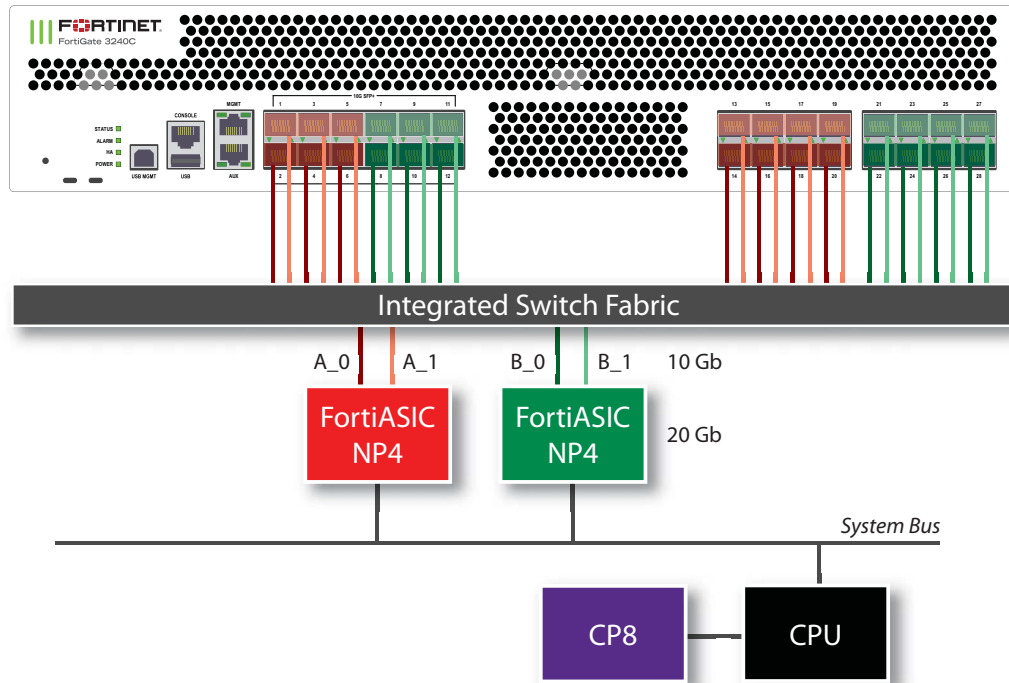
```
config system global
 set sp-load-balance disable
end
```

# FortiGate-3240C

The FortiGate-3240C features two NP4 processors:

- The 10Gb interfaces, port1 through port6, and the 1Gb interfaces, port13 through port20, share connections to one NP4 processor.
- The 10Gb interfaces, port7 through port12, and the 1Gb interfaces, port21 through port28, share connections to the other NP4 processor.


In addition to the ports being divided between the two NP4 processors, they are further divided between the two connections to each processor. Each NP4 can process 20 Gb of network traffic per second and each of two connections to each NP4 can move 10Gb of data to the processor per second, so the ideal configuration would have no more than 10 Gb of network traffic to each connection of each NP4 at any time.



## FortiGate-3600C

The FortiGate-3600C features three NP4 processors:

- The 10Gb interfaces, port1-port4, and the 1Gb interfaces, port13-port17, share connections to one NP4 processor.
- The 10Gb interfaces, port5-port8, and the 1Gb interfaces, port18-port22 share connections to the second NP4 processor.
- The 10Gb interfaces, port9-port12, and the 1Gb interfaces, port23-port28 share connections to the third NP4 processor.



## XAUI interfaces

Each NP4 processor connects to the integrated switch fabric through two XAUI interfaces: XAUI0 and XAUI1. On each NP4 processor all of the odd numbered interfaces use XAUI0 and all of the even numbered interfaces use XAUI1:

NPU1

XAUI0 = port1,port3,port13, port15, port17

XAUI1 = port2, port4, port14, port16

NPU2

XAUI0 = port5, port7, port18, port20, port22

XAUI1 = port6, port8, port19, port21

NPU3

XAUI0 = port9, port11, port23, port25, port27

XAUI1 = port10, port12, port24, port26, port28

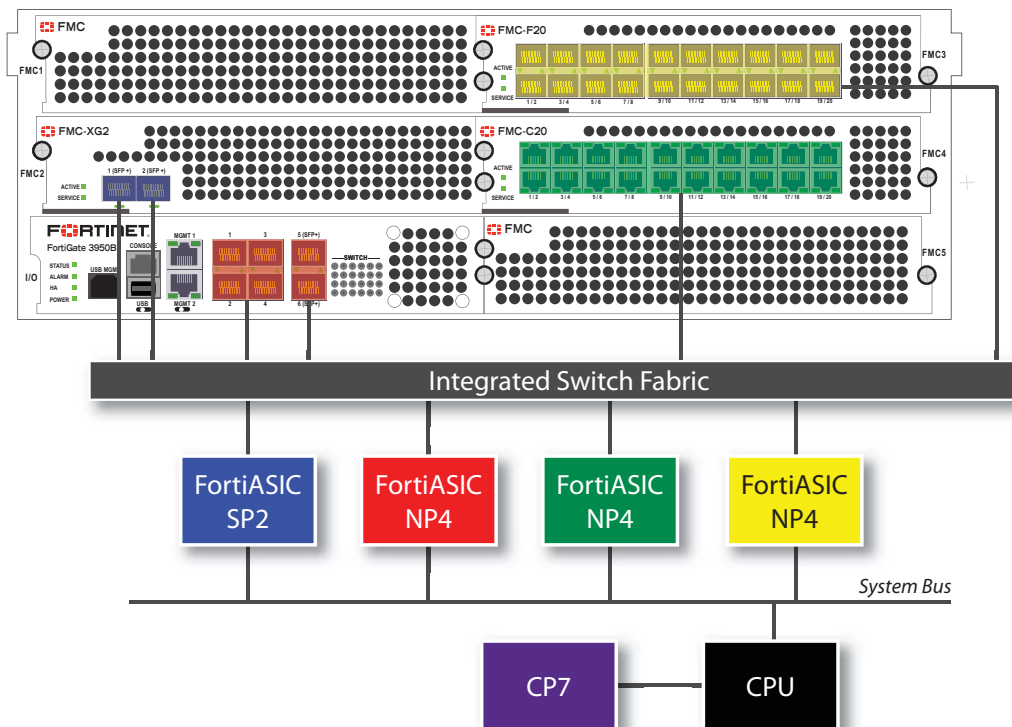
Usually you do not have to be concerned about the XAUI interface mapping. However, if an NP4 interface is processing a very high amount of traffic you should distribute that traffic among both of the XAUI interfaces connected to it. So if you have a very high volume of traffic flowing between two networks you should connect both networks to the same NP4 processor but to different XAUI links. So between even and an add numbered FortiGate-3600C ports. For example, you could connect one network to port5 and the other network to port6. In this configuration, the second NP4 processor would handle traffic acceleration and both XAUI interfaces would be processing traffic.

## FortiGate-3950B and FortiGate-3951B

The FortiGate-3950B features one NP4 processor. The 1Gb SPF interfaces, port1, port2, port3, port4, and the 10Gb SPF+ interfaces, port5, port6, share connections to one NP4 processor. The FortiGate-3951B is similar to the FortiGate-3950B, except it trades one FMC slot for four FSM slots. The network interfaces available on each model are identical.

You can add additional FMC interface modules. The diagram below shows a FortiGate-3950B with three modules installed: an FMC-XG2, an FMC-F20, and an FMC-C20.

- The FMC-XG2 has one SP2 processor. The 10Gb SPF+ interfaces, port1 and port2, share connections to the processor.
- The FMC-F20 has one NP4 processor and the twenty 1Gb SPF interfaces, port1 through port20, share connections to the NP4 processor.
- The FMC-C20 has one NP4 processor and the twenty 10/100/1000 interfaces, port1 through port20, share connections to the NP4 processor.





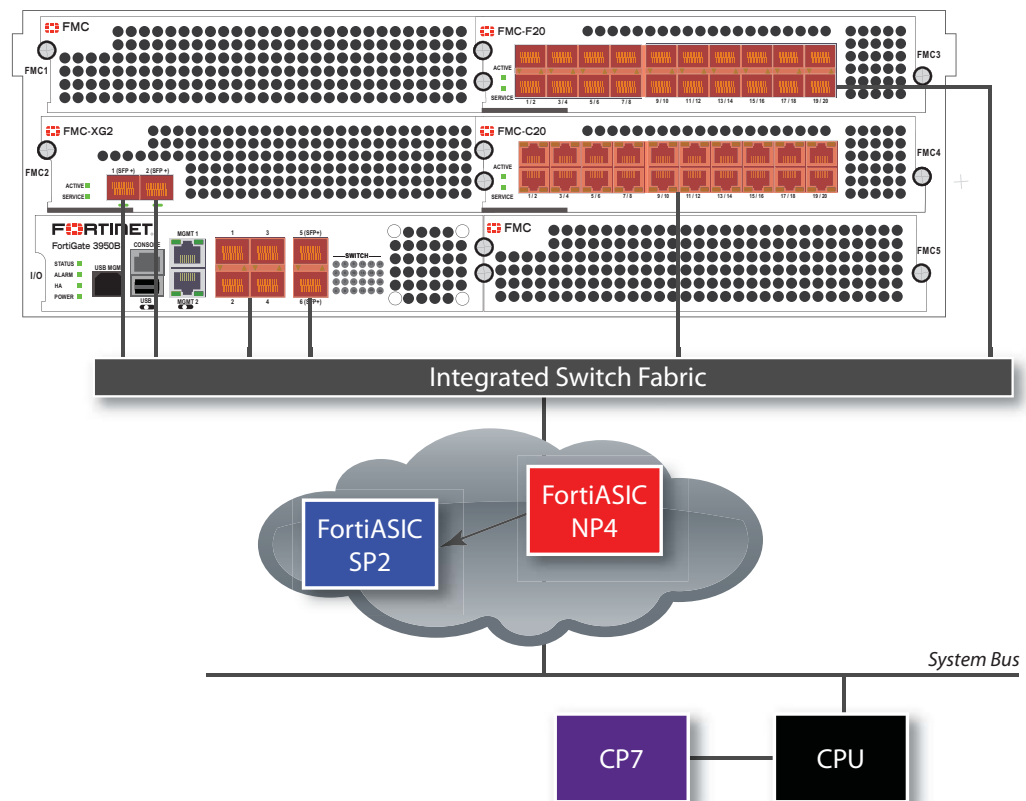
## FortiGate-3950B and FortiGate-3951B – load balance mode

Adding one or more FMC-XG2 modules to your FortiGate-3950B allows you to enable load balance mode. This feature allows you increased flexibility in how you use the interfaces on the FortiGate unit. The FortiGate-3951B is similar to the FortiGate-3950B, except it trades one FMC slot for four FSM slots. The network interfaces available on each model are identical.

When enabled, traffic between any two interfaces (excluding management and console) is accelerated whether they are the six interfaces on the FortiGate-3950B itself, or on any installed FMC modules. Traffic is not limited to entering and leaving the FortiGate unit in specific interface groupings to benefit from NP4 and SP2 acceleration. You can use any pair of interfaces.

Security acceleration in this mode is limited, however. Only IPS scanning is accelerated in load balance mode.

**Figure 184:**The FortiGate-3950B in load balance mode



To enable this feature, issue this CLI command.

```
config system global
 set sp-load-balance enable
end
```

The FortiGate unit will then restart.

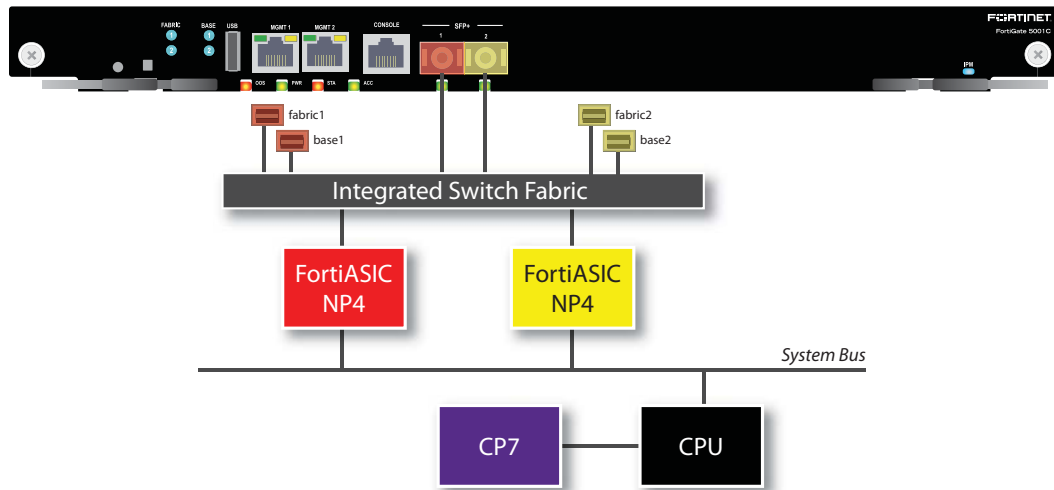
To return to the default mode, issue this CLI command.

```
config system global
 set sp-load-balance disable
end
```

## FortiGate-5001C

The FortiGate-5001C board includes two NP4 processors connected to an integrated switch fabric:

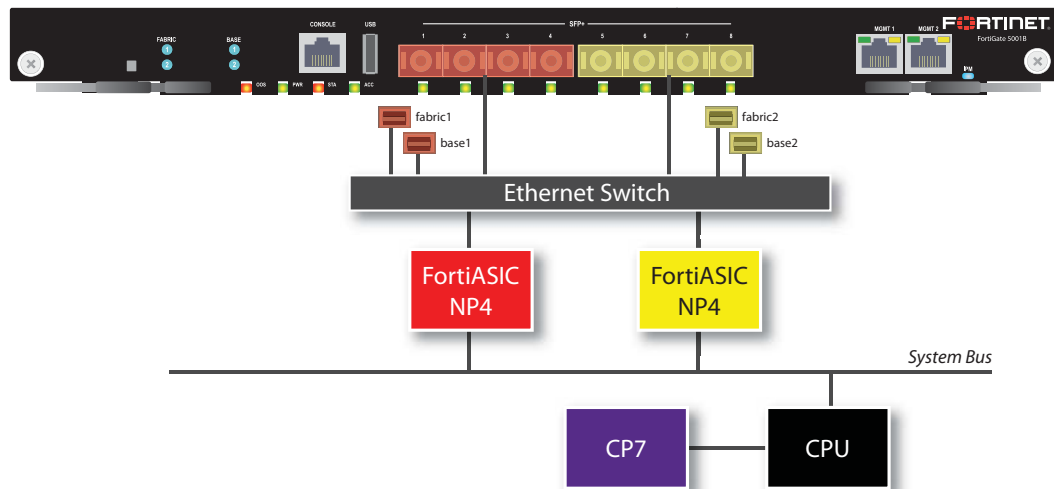
- The port1, fabric1, and base1 interfaces are connected to one NP4 processor.
- The port2, fabric2, and base2 interfaces are connected to the other NP4 processor.



## FortiGate-5001B

The FortiGate-5001B board includes two NP4 connected to an Ethernet switch. Traffic between interfaces that use the same NP4 processor experiences the highest acceleration since the FortiGate-5001B does not include an integrated switch fabric.

- The port1, port2, port3, port4, fabric1 and base1 interfaces are connected to one NP4 processor.
- The port5, port6, port7, port8, fabric2 and base2 interfaces are connected to the other NP4 processor.

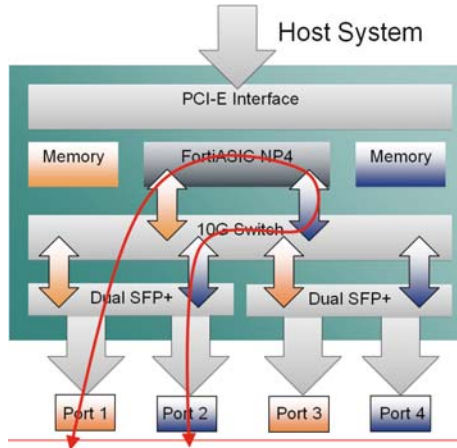


For example, for maximum NP4 acceleration of traffic received on port1 the traffic must exit the FortiGate-5001B board on port2, port3, port4, or fabric1. Also, for maximum acceleration of traffic received on port5 the traffic must exit the FortiGate-5001B board on port6, port7, port8, or fabric2.

## Setting switch-mode mapping on the ADM-XD4

The ADM-XD4 SP has four 10Gb/s ports, but the NP4 processor it contains has only two 10Gb/s ports. The external ports you use are important to optimize the SP for your application.

**Figure 185:**ADM-XD4 mapping mode



Ports 1 and 3 share one NP4 processor and ports 2 and 4 share the other. Performance ports sharing the same NP4 processor is far better than when forcing network data to move between NP4 processors by using one port from each, for example ports 1 and 2 or ports 3 and 4.

# Chapter 9 High Availability for FortiOS

## 5.0

This FortiOS Handbook chapter contains the following sections:

[Solving the High Availability problem](#) describes the high availability problem and introduces the FortiOS solutions described in this document (FGCP, VRRP, and standalone session synchronization).

[An introduction to the FGCP](#) introduces the FGCP clustering protocol and many of its features and terminology.

[Configuring and connecting HA clusters](#) describes configuring HA clusters and contains HA clustering configuration examples.

[Virtual clusters](#) describes configuring HA virtual clusters and contains virtual clustering configuration examples.

[Full mesh HA](#) describes configuring FortiGate Full mesh HA and contains a full mesh HA configuration example.

[Operating a cluster](#) describes how to operate a cluster and includes detailed information about how various FortiGate systems operate differently in a cluster.

[HA and failover protection](#) describes in detail how FortiGate HA device failover, link failover, and session failover work.

[HA and load balancing](#) describes in detail how FortiGate HA active-active load balancing load balances sessions.

[HA with FortiGate-VM and third-party products](#) describes how FortiGate units interact with third-party products.

[VRRP](#) describes FortiOS support of the Virtual Router Redundancy Protocol (VRRP) and its use for high availability.

[FortiGate Session Life Support Protocol \(FGSP\)](#) describes the FortiGate standalone session synchronization feature and its use for high availability.

[Configuring FRUP](#) describes how to set up a FortiGate Redundant UTM Protocol (FRUP) cluster consisting of two FortiGate-100D units.

-

# Solving the High Availability problem

The basic high availability (HA) problem for TCP/IP networks and security gateways is keeping network traffic flowing. Uninterrupted traffic flow is a critical component for online systems and media because critical business processes quickly come to a halt when the network is down.

The security gateway is a crucial component of most networks since all traffic passes through it. A standalone network security gateway is a single point of failure that is vulnerable to any number of software or hardware problems that could compromise the device and bring all traffic on the network to a halt.

A common solution to the high availability problem is to eliminate the security gateway as single point of failure by introducing redundancy. With two or more redundant security gateways, if one fails, the remaining one or more gateways keep the traffic flowing. FortiOS provides four redundancy solutions: industry standard VRRP as well as three proprietary solutions: FortiGate Cluster Protocol (FGCP) high availability, FortiGate Session Life Support Protocol (FGSP) high availability, and the Fortinet Redundant UTM protocol (FRUP) high availability.




You can combine more than one high availability solution into a single configuration. A common reason for doing this could be to add VRRP to an FGCP or FGSP configuration.

A strong and flexible High availability solution is required for many mission-critical firewall and UTM applications. Each FortiOS high availability solution can be fine tuned to fit into many different network scenarios.

## FortiGate Cluster Protocol (FGCP)

FGCP HA provides a solution for two key requirements of critical enterprise networking components: enhanced reliability and increased performance. Enhanced reliability is achieved through device failover protection, link failover protection, and remote link failover protection. Also contributing to enhanced reliability is session failover protection for most IPv4 and IPv6 sessions including TCP, UDP, ICMP, IPsec VPN, and NAT sessions. Increased performance is achieved through active-active HA load balancing. Extended FGCP features include full mesh HA and virtual clustering. You can also fine tune the performance of the FGCP to change how a cluster forms and shares information among cluster units and how the cluster responds to failures.

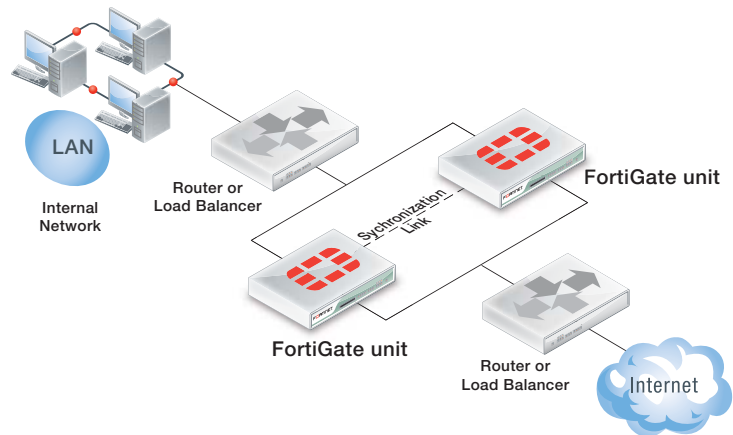


When configured onto your network an FGCP cluster appears to be a single FortiGate unit operating in NAT/Route or Transparent mode and configuration synchronization allows you to configure a cluster in the same way as a standalone FortiGate unit. If a failover occurs, the cluster recovers quickly and automatically and also sends administrator notifications so that the problem that caused the failure can be corrected and any failed equipment restored.

The FGCP is compatible with most network environments and most networking equipment. While initial configuration is relatively quick and easy, a large number of tools and configuration options are available to fine tune the cluster for most situations.

## FortiGate Session Life Support Protocol (FGSP)

In a network that already includes load balancing (either with load balancers or routers) for traffic redundancy, two identical FortiGate units can be integrated into the load balancing configuration using the FortiGate Session Life Support Protocol (FGSP). The external load balancers or routers can distribute sessions among the FortiGate units and the FGSP performs session synchronization of IPv4 and IPv6 TCP, UDP, ICMP, expectation, and NAT sessions to keep the session tables of both FortiGate units synchronized.



If one of the FortiGate units fails, session failover occurs and active sessions fail over to the unit that is still operating. This failover occurs without any loss of data. As well, the external load balancers or routers detect the failover and re-distribute all sessions to the unit that is still operating.

Load balancing and session failover is done by external routers or load balancers and not by the FGSP. The FortiGate units just perform session synchronization which allows session failover to occur without packet loss.


The FGSP also includes configuration synchronization, allowing you to make configuration changes once for both FortiGate units instead of requiring duplicate configuration changes on each unit. Settings that identify the FortiGate unit to the network, for example, interface IP addresses and BGP neighbor settings, are not synchronized so each FortiGate unit maintains its identity on the network. These settings must be configured separately for each FortiGate unit.



In previous versions of FortiOS the FGSP was called TCP session synchronization or standalone session synchronization. However, the FGSP has been expanded to include configuration synchronization and session synchronization of connectionless sessions, expectation sessions, and NAT sessions.

## VRRP

FortiGate units can function as master or backup Virtual Router Redundancy Protocol (VRRP) routers and can be quickly and easily integrated into a network that has already deployed VRRP. A FortiGate unit can be integrated into a VRRP group with any third-party VRRP devices and VRRP can provide redundancy between multiple FortiGate units.



In a VRRP configuration, when a FortiGate unit operating as the master unit fails, a backup unit takes its place and continues processing network traffic. If the backup unit is a FortiGate unit, the network continues to benefit from FortiOS security features. If the backup unit is a router, after a failure traffic will continue to flow, but FortiOS security features will be unavailable until the FortiGate unit is back on line. You can include different FortiGate models in the same VRRP group.

FortiOS supports VRRP between two or more FortiGate units and between FortiGate units and third-party routers that support VRRP. Using VRRP you can assign VRRP routers as master or backup routers. The master router processes traffic and the backup routers monitor the master router and can begin forwarding traffic if the master fails. Similar to the FGCP you can configuration VRRP between multiple FortiGate units to provide redundancy. You can also create a VRRP group with a FortiGate units and any routers that support VRRP.

In a VRRP configuration that consists of one FortiGate unit and one router, normally the FortiGate unit would be the master and all traffic would be processed by the FortiGate unit. If the FortiGate unit fails, all traffic switches to the router. Network connectivity is maintained even though FortiGate security features will be unavailable until the FortiGate unit can is back on line.

## Fortinet redundant UTM protocol (FRUP)


An extension to the FGCP combines switching HA and firewall HA into a single unified design. This feature is available on the FortiGate-100D and will be expanded to other models in future releases.

A FRUP setup consists of 2 (and only 2) identical FortiGate-100D units. The setup supports dual redundant HA links between the units for sharing session and configuration data.

FRUP requires redundant external routers where:

- One FortiGate unit has a primary connection to one of the routers and a backup connection to the other.
- The other FortiGate unit has the opposite configuration.

See “Configuring FRUP” on page 1380 for an example configuration.





# An introduction to the FGCP

A FortiGate HA cluster consists of two to four FortiGate units configured for HA operation. Each FortiGate unit in a cluster is called a cluster unit. All cluster units must be the same FortiGate model with the same FortiOS firmware build installed. All cluster units must also have the same hardware configuration (for example, the same AMC modules installed in the same slots, the same number of hard disks and so on) and be running in the same operating mode (NAT/Route mode or Transparent mode).



You can create an FGCP cluster of up to four FortiGate units.

---

In addition the cluster units must be able to communicate with each other through their heartbeat interfaces. This heartbeat communication is required for the cluster to be created and to continue operating. Without it, the cluster acts like a collection of standalone FortiGate units.

On startup, after configuring the cluster units with the same HA configuration and connecting their heartbeat interfaces, the cluster units use the FortiGate Clustering Protocol (FGCP) to find other FortiGate units configured for HA operation and to negotiate to create a cluster. During cluster operation, the FGCP shares communication and synchronization information among the cluster units over the heartbeat interface link. This communication and synchronization is called the FGCP heartbeat or the HA heartbeat. Often, this is shortened to just heartbeat.

The cluster uses the FGCP to select the primary unit, and to provide device, link and session failover. The FGCP also manages the two HA modes; active-passive (failover HA) and active-active (load balancing HA).


This chapter describes.

- [About the FGCP](#)
- [Synchronizing the configuration \(and settings that are not synchronized\)](#)
- [Configuring FortiGate units for FGCP HA operation](#)
- [Active-passive and active-active HA](#)
- [Identifying the cluster and cluster units](#)
- [Device failover, link failover, and session failover](#)
- [Primary unit selection](#)
- [HA override](#)
- [FortiGate HA compatibility with PPPoE and DHCP](#)
- [HA and distributed clustering](#)
- [Hard disk configuration and HA](#)
- [FGCP high availability best practices](#)
- [FGCP HA terminology](#)
- [HA web-based manager options](#)

## About the FGCP

FortiGate HA is implemented by configuring two or more FortiGate units to operate as an HA cluster. To the network, the HA cluster appears to function as a single FortiGate unit, processing network traffic and providing normal security services such as firewalling, security services, Unified Threat Management (UTM) and VPN services.

**Figure 186:** HA cluster installed between an internal network and the Internet



Inside the cluster the individual FortiGate units are called cluster units. These cluster units share state and configuration information. If one cluster unit fails, the other units in the cluster automatically replace that unit, taking over the work that the failed unit was doing. After the failure, the cluster continues to process network traffic and provide normal FortiGate services with virtually no interruption.

Every FortiGate cluster contains one primary unit (also called the master unit) and one or more subordinate units (also called slave or backup units). The primary unit controls how the cluster operates. The role that the subordinate units play depends on the mode in which the cluster operates: (Active-Passive (AP) or Active-Active (AA) (see [“Active-passive HA \(failover protection\)”](#) on page 1128 and [“Active-active HA \(load balancing and failover protection\)”](#) on page 1129).

The ability of an HA cluster to continue providing firewall services after a failure is called failover. FGCP failover means that your network does not have to rely on one FortiGate unit to continue functioning. You can install additional units and form an HA cluster.

A second HA feature, called load balancing, can be used to increase performance. A cluster of FortiGate units can increase overall network performance by sharing the load of processing network traffic and providing security services. The cluster appears to your network to be a single device, adding increased performance without changing your network configuration.

Virtual clustering extends HA features to provide failover protection and load balancing for Virtual Domains (VDOMs). See [“Virtual clusters”](#) on page 1217.

FortiGate models that support redundant interfaces can be configured to support full mesh HA. Full mesh HA is a method of reducing the number of single points of failure on a network that includes an HA cluster. For details about full mesh HA, see [“Full mesh HA”](#) on page 1239.

## FGCP failover protection

The FGCP provides IP/MAC takeover for failover protection by assigning virtual MAC addresses to the primary cluster unit and then sending gratuitous ARP packets from the primary unit interfaces to reprogram the network.

Failover times can be less than a second under optimal conditions. You can fine tune failover performance for your network by adjusting cluster status checking timers, routing table update timers, and wait timers.

An HA cluster fails over if the primary unit fails (a device failure) or experiences a link failure. The cluster can detect link failures for connections to the primary unit using port monitoring and for connections between downstream network components using remote IP monitoring. To compensate for a link failover, the cluster maintains active links to keep traffic flowing between high-priority networks. Port and remote IP monitoring can be fine tuned without disrupting cluster operation.

## Session Failover

FGCP session failover maintains TCP, SIP and IPsec VPN sessions after a failure. You can also configure session failover to maintain UDP and ICMP sessions. Session failover does not failover multicast, or SSL VPN sessions. Session failover may not be required for all networks because many TCP/IP, UDP, and ICMP protocols can resume sessions on their own. Supporting session failover adds extra overhead to cluster operations and can be disabled to improve cluster performance if it is not required.

## Load Balancing

Active-active HA load balances resource-intensive virus scanning, web filtering, intrusion protection, Application Control, email filtering and Data Leak Prevention operations among all cluster units to provide better performance than a standalone FortiGate unit. If network traffic consists of mainly TCP sessions, the FGCP can also load balance all TCP sessions to improve TCP performance in some network configurations. You can also load balance UDP sessions. You can use accelerated FortiGate interfaces to also accelerate HA load balancing and HA load balancing schedules can be adjusted to optimize performance for the traffic mix on your network. Weighted load balancing can be used to control the relative amount of sessions processed by each cluster unit.

## Virtual Clustering

Virtual clustering is an extension of the FGCP for a cluster of 2 FortiGate units operating with multiple VDOMS enabled. Not only does virtual clustering provide failover protection for a multiple VDOM configuration, but a virtual cluster can load balance traffic between the cluster units. Load balancing with virtual clustering is quite efficient and load balances all traffic (not just UTM and TCP traffic). Its possible to fine tune virtual clustering load balancing in real time to actively optimize load sharing between the cluster units without affecting the smooth operation of the cluster.

## Full Mesh HA

High availability improves the reliability of a network by replacing a single point of failure (a single FortiGate unit) with a cluster that can maintain network traffic if one of the cluster units fails. However, in a normal cluster configuration, single points of failure remain. Full mesh HA removes these single points of failure by allowing you to connect redundant switches to each cluster interface. Full mesh HA is achieved by configuring 802.3ad aggregate or redundant interfaces on the FortiGate unit and connecting redundant switches to these interfaces. Configuration is a relatively simple extension of the normal aggregate/redundant interface and HA configurations.

## Cluster Management

FortiOS HA provides a wide range of cluster management features:

- Automatic continuous configuration synchronization. You can get a cluster up and running almost as quickly as a standalone FortiGate unit by performing a few basic steps to configure HA settings and minimal network settings on each cluster unit. When the cluster is operating you can configure FortiGate features such as firewalling, content inspection, and VPN in the same way as for a standalone FortiGate unit. All configuration changes (even complex changes such as switching to multiple VDOM mode or from NAT/Route to Transparent mode) are synchronized among all cluster units.
- Firmware upgrades/downgrades. Upgrading or downgrading cluster firmware is similar to upgrading or downgrading standalone FortiGate firmware. The Firmware is uploaded once to the primary unit and the cluster automatically upgrades or downgrades all cluster units in one operation with minimal or no service interruption.
- Individual cluster unit management. In some cases you may want to manage individual cluster units. You can do so from cluster CLI by navigating to each cluster unit. You can also use the reserved management interface feature to give each cluster unit its own IP address and default route. You can use the reserved management interfaces and IP addresses to connect to the GUI and CLI of each cluster unit and configure an SNMP server to poll each cluster unit.
- Removing and adding cluster units. In one simple step any unit (even the primary unit) can be removed from a cluster and given a new IP address. The cluster keeps operating as it was; the transition happening without interrupting cluster operation. A new unit can also be added to an operating cluster without disrupting network traffic. All you have to do is connect the new unit and change its HA configuration to match the cluster's. The cluster automatically finds and adds the unit and synchronizes its configuration with the cluster.
- Debug and diagnose commands. An extensive range of debug and diagnose commands can be used to report on HA operation and find and fix problems.
- Logging and reporting. All cluster units can be configured to record all log messages. These message can be stored on the individual cluster units or sent to a FortiAnalyzer unit. You can view all cluster unit log messages by logging into any cluster unit.
- FortiManager support. FortiManager understands FortiOS HA and automatically recognizes when you add a FortiOS cluster to the FortiManager configuration.

## Synchronizing the configuration (and settings that are not synchronized)

The FGCP uses a combination of incremental and periodic synchronization to make sure that the configuration of all cluster units is synchronized to that of the primary unit. This means that in most cases you only have to make a configuration change once to have it synchronized to all cluster units.

Some configuration settings are not synchronized to support some aspects of FortiGate operation. The following settings are not synchronized among cluster units:

- The FortiGate unit host name. Allows you to identify cluster units.
- HA override (“HA override” on page 1138).
- HA device priority (“Primary unit selection and device priority” on page 1136).
- Virtual cluster 1 and Virtual cluster 2 device priorities (“Virtual clustering and load balancing or VDOM partitioning” on page 1218)
- The HA priority (`ha-priority`) setting for a ping server or dead gateway detection configuration (“Remote link failover” on page 1325).
- The system interface settings of the FortiGate interface that becomes the HA reserved management interface (“Managing individual cluster units using a reserved management interface” on page 1254).
- The default route for the reserved management interface, set using the `ha-mgt-interface-gateway` option of the `config system ha` command (“Managing individual cluster units using a reserved management interface” on page 1254).
- The dynamic weighted load balancing thresholds and high and low watermarks (“Dynamically optimizing weighted load balancing according to how busy cluster units are” on page 1350).

## Configuring FortiGate units for FGCP HA operation

Each FortiGate unit in the cluster must have the same HA configuration. Once the cluster is connected, you can configure it in the same way as you would configure a standalone FortiGate unit. The following procedures set the HA mode to active-passive and set the HA password to HA\_pass.



Make sure your FortiGate interfaces are configured with static IP addresses. If any interface gets its address using DHCP you will not be able to configure HA.

---

### To configure a FortiGate unit for HA operation - web-based manager

1. Power on the FortiGate unit to be configured.
2. Log into the web-based manager.
3. On the Dashboard *System Information* dashboard widget, beside *Host Name* select *Change*.
4. Enter a new Host Name for this FortiGate unit.

Changing the host name makes it easier to identify individual cluster units when the cluster is operating.

5. Go to *System > Config > HA* and change the following settings:

<b>Mode</b>	Active-Passive
<b>Group Name</b>	Example_cluster
<b>Password</b>	HA_pass The password must be the same for all FortiGate units in the cluster.

You can accept the default configuration for the remaining HA options and change them later, once the cluster is operating.

6. Select *OK*.

The FortiGate unit negotiates to establish an HA cluster. When you select *OK* you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 1300](#)). To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all ARP table entries). You may be able to delete the ARP table of your management PC from a command prompt using a command similar to `arp -d`.

7. Power off the FortiGate unit.
8. Repeat this procedure for all of the FortiGate units in the cluster.

Once all of the units are configured, continue with [“Connecting a FortiGate HA cluster” on page 1127](#).

### To configure a FortiGate unit for HA operation - CLI

1. Power on the FortiGate unit to be configured.
2. Log into the CLI.
3. Enter the following command to change the FortiGate unit host name.

```
config system global
 set hostname Example1_host
end
```

Changing the host name makes it easier to identify individual cluster units when the cluster is operating.

4. Enter the following command to enable HA:

```
config system ha
 set mode active-passive
 set group-name Example_cluster
 set password HA_pass
end
```

You can accept the default configuration for the remaining HA options and change them later, once the cluster is operating.

The FortiGate unit negotiates to establish an HA cluster. You may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and because the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 1300](#)). To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

5. Power off the FortiGate unit.

- Repeat this procedure for all of the FortiGate units in the cluster.  
Once all of the units are configured, continue with [“Connecting a FortiGate HA cluster”](#).

## Connecting a FortiGate HA cluster

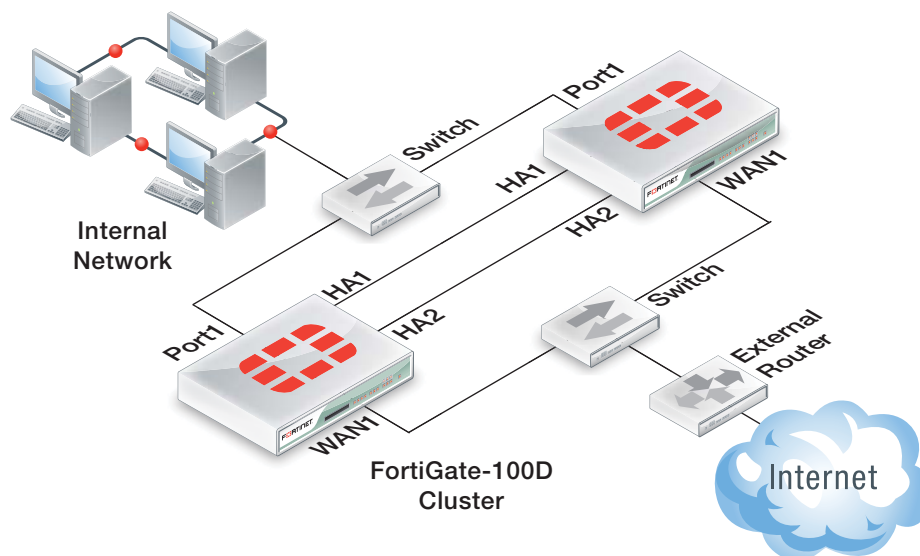
Use the following procedure to connect a cluster. Connect the cluster units to each other and to your network. You must connect all matching interfaces in the cluster to the same switch, then connect these interfaces to their networks using the same switch.

Although you can use hubs, Fortinet recommends using switches for all cluster connections for the best performance.

Connecting an HA cluster to your network temporarily interrupts communications on the network because new physical connections are being made to route traffic through the cluster. Also, starting the cluster interrupts network traffic until the individual cluster units are functioning and the cluster completes negotiation. Cluster negotiation is automatic and normally takes just a few seconds. During system startup and negotiation all network traffic is dropped.

This section describes how to connect the cluster shown in [Figure 187 on page 1127](#) that consists of two FortiGate-100D units to be connected between the Internet and a head office internal network. The wan1 interfaces of the FortiGate unit connect the cluster to the Internet and the internal interfaces connect the cluster to the internal network. The ha1 and ha2 interfaces are used for redundant HA heartbeat links.

**Figure 187:**Example cluster connections



### To connect a FortiGate HA cluster

- Connect the WAN1 interfaces of each cluster unit to a switch connected to the Internet.
- Connect the Port1 interfaces of each cluster unit to a switch connected to the internal network.
- Connect the HA1 interfaces of the cluster units together. You can use a crossover Ethernet cable or a regular Ethernet cable. (You can also connect the interfaces using Ethernet cables and a switch.)
- Connect the HA2 interfaces of the cluster units together. You can use a crossover Ethernet cable or a regular Ethernet cable. (You can also connect the interfaces using Ethernet cables and a switch.)

5. Power on both of the FortiGate units.

As the cluster units start, they negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally just takes a few seconds.

At least one heartbeat interface should be connected together for the cluster to operate.

Do not use a switch port for the HA heartbeat traffic. This configuration is not supported.

You could use one switch to connect all four heartbeat interfaces. However, this is not recommended because if the switch fails both heartbeat interfaces will become disconnected.

For more information about heartbeat interfaces, see [“HA heartbeat and communication between cluster units” on page 1293](#).

6. You can now configure the cluster as if it is a single FortiGate unit.

## Active-passive and active-active HA

The first decision to make when configuring FortiGate HA is whether to choose active-passive or active-active HA mode. To configure the HA mode, go to *System > Config > HA* and set Mode to *Active-Passive* or *Active-Active*.

From the CLI enter the following command to set the HA mode to active-passive:

```
config system ha
 set mode a-p
end
```

To form a cluster, all cluster units must be set to the same mode. You can also change the mode after the cluster is up and running. Changing the mode of a functioning cluster causes a slight delay while the cluster renegotiates to operate in the new mode and possibly select a new primary unit.

### Active-passive HA (failover protection)

An active-passive (A-P) HA cluster provides hot standby failover protection. An active-passive cluster consists of a primary unit that processes communication sessions, and one or more subordinate units. The subordinate units are connected to the network and to the primary unit but do not process communication sessions. Instead, the subordinate units run in a standby state. In this standby state, the configuration of the subordinate units is synchronized with the configuration of the primary unit and the subordinate units monitor the status of the primary unit.

Active-passive HA provides transparent device failover among cluster units. If a cluster unit fails, another immediately take its place. See [“Device failover” on page 1292](#).

Active-passive HA also provides transparent link failover among cluster units. If a cluster unit interface fails or is disconnected, this cluster unit updates the link state database and the cluster negotiates and may select a new primary unit. See [“Link failover \(port monitoring or interface monitoring\)” on page 1319](#) for more information.

If session failover (also called session pickup) is enabled, active-passive HA provides session failover for some communication sessions. See [“Session failover \(session pick-up\)” on page 1330](#) for information about session failover and its limitations.



The following example shows how to configure a FortiGate unit for active-passive HA operation. You would enter the exact same commands on every FortiGate unit in the cluster.

```
config system ha
 set mode a-p
 set group-name myname
 set password HApass
end
```

## Active-active HA (load balancing and failover protection)

Active-active (A-A) HA load balances resource-intensive content inspection processing among all cluster units. Content inspection processing applies protocol recognition, virus scanning, IPS, web filtering, email filtering, data leak prevention (DLP), application control, and VoIP content scanning and protection to HTTP, HTTPS, FTP, IMAP, IMAPS, POP3, POP3S, SMTP, SMTPS, IM, NNTP, SIP, SIMPLE, and SCCP sessions accepted by security policies. By load balancing this resource-intensive processing among all cluster units, an active-active HA cluster may provide better content inspection performance than a standalone FortiGate unit. Other features enabled in security policies such as Endpoint security, traffic shaping, user authentication, and device identification have no effect active-active load balancing.

Normally, sessions that don't include content inspection are not load balanced and are processed by the primary unit. You can configure active-active HA to load balance additional sessions. For more information see [“Load balancing UTM sessions, TCP sessions, and UDP sessions” on page 1347](#).

An active-active HA cluster consists of a primary unit that receives all communication sessions and load balances them among the primary unit and all of the subordinate units. In an active-active cluster the subordinate units are also considered active since they also process content processing sessions. In all other ways active-active HA operates the same as active-passive HA.

The following example shows how to configure a FortiGate unit for active-active HA operation. You would enter the exact same commands on every FortiGate unit in the cluster.

```
config system ha
 set mode a-a
 set group-name myname
 set password HApass
end
```

## Identifying the cluster and cluster units

You can use the cluster group name, group id, and password to identify a cluster and distinguish one cluster from another. If you have more than one cluster on the same network, each cluster must have a different group name, group id, and password.

### Group name

Use the group name to identify the cluster. The maximum length of the group name is 32 characters. The group name must be the same for all cluster units before the cluster units can form a cluster. After a cluster is operating, you can change the group name. The group name change is synchronized to all cluster units.

The default group name is *FGT-HA*. The group name appears on the FortiGate dashboard of a functioning cluster as the *Cluster Name*.

To change the group name from the web-based manager go to *Config > System > HA* and change the *Group Name*.

Enter the following CLI command to change the group name to *Cluster\_name*:

```
config system ha
 set group-name Cluster_name
end
```

## Password

Use the password to identify the cluster. You should always change the password when configuring a cluster. The password must be the same for all FortiGate units before they can form a cluster. When the cluster is operating you can change the password, if required. Two clusters on the same network cannot have the same password.

To change the password from the web-based manager go to *Config > System > HA* and change the *Password*.

Enter the following CLI command to change the password to *ha\_pwd*:

```
config system ha
 set password ha_pwd
end
```

## Group ID

Similar to the group name, the group ID is also identifies the cluster. In most cases you do not have to change the group ID. However, you should change the group ID if you have more than one cluster on the same network. All members of the HA cluster must have the same group ID. The group ID is a number from 0 to 255.

Changing the group ID changes the cluster virtual MAC address. See [“Cluster virtual MAC addresses” on page 1300](#).

Enter the following CLI command to change the group ID to 10:

```
config system ha
 set group-id 10
end
```

## Device failover, link failover, and session failover

The FGCP provides transparent device and link failover. You can also enable session pickup to provide session failover. A failover can be caused by a hardware failure, a software failure, or something as simple as a network cable being disconnected causing a link failover. When a failover occurs, the cluster detects and recognizes the failure and takes steps to respond so that the network can continue to operate without interruption. The internal operation of the cluster changes, but network components outside of the cluster notice little or no change.

If a failover occurs, the cluster also records log messages about the event and can be configured to send log messages to a syslog server and to a FortiAnalyzer unit. The cluster can also send SNMP traps and alert email messages. These alerts can notify network administrators of the failover and may contain information that the network administrators can use to find and fix the problem that caused the failure.

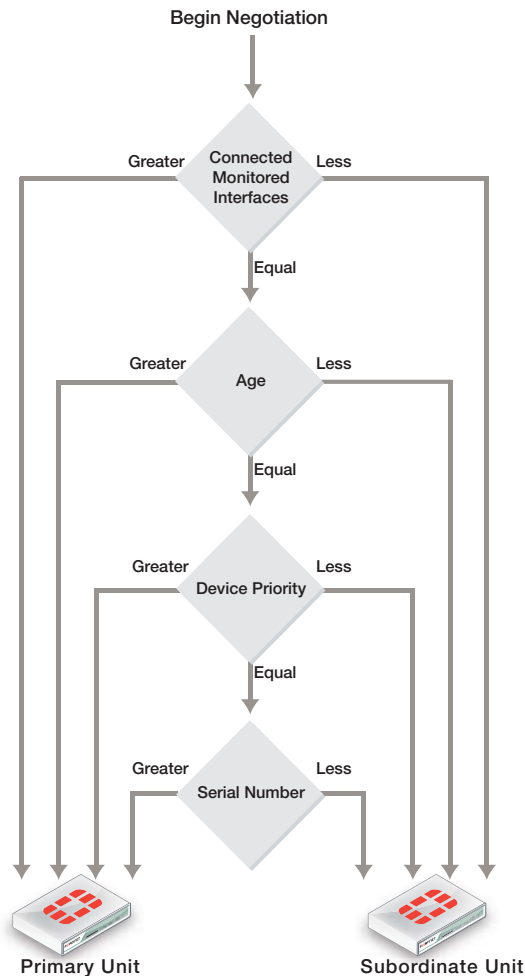
For a complete description of device failover, link failover, and session failover, how clusters support these types of failover, and how FortiGate HA clusters compensate for a failure to maintain network traffic flow see [“HA and failover protection” on page 1290](#).

## Primary unit selection

Once FortiGate units recognize that they can form a cluster, the cluster units negotiate to select a primary unit. Primary unit selection occurs automatically based on the criteria shown in [Figure 188](#). After the cluster selects the primary unit, all of the remaining cluster units become subordinate units.

Negotiation and primary unit selection also takes place if a primary unit fails (device failover) or if a monitored interface fails or is disconnected (link failover). During a device or link failover, the cluster renegotiates to select a new primary unit also using the criteria shown in [Figure 188](#).

**Figure 188:** Selecting the primary unit



For many basic HA configurations primary unit selection simply selects the cluster unit with the highest serial number to become the primary unit. A basic HA configuration involves setting the HA mode to active-passive or active-active and configuring the cluster group name and password. Using this configuration, the cluster unit with the highest serial number becomes the primary unit because primary unit selection disregards connected monitored interfaces (because interface monitoring is not configured), the age of the cluster units would usually always be the same, and all units would have the same device priority.

Using the serial number is a convenient way to differentiate cluster units; so basing primary unit selection on the serial number is predictable and easy to understand and interpret. Also the cluster unit with the highest serial number would usually be the newest FortiGate unit with the most recent hardware version. In many cases you may not need active control over primary unit selection, so basic primary unit selection based on serial number is sufficient.

In some situations you may want have control over which cluster unit becomes the primary unit. You can control primary unit selection by setting the device priority of one cluster unit to be higher than the device priority of all other cluster units. If you change one or more device priorities, during negotiation, the cluster unit with the highest device priority becomes the primary unit. As shown in [Figure 188](#) the FGCP selects the primary unit based on device priority before serial number. For more information about how to use device priorities, see [“Primary unit selection and device priority” on page 1136](#).

The only other way that you can influence primary unit selection is by configuring interface monitoring (also called port monitoring). Using interface monitoring you can make sure that cluster units with failed or disconnected monitored interfaces cannot become the primary unit. See [“Primary unit selection and monitored interfaces” on page 1132](#).

Finally, the age of a cluster unit is determined by a number of operating factors. Normally the age of all cluster units is the same so normally age has no effect on primary unit selection. Age does affect primary unit selection after a monitored interface failure. For more information about age, see [“Primary unit selection and age” on page 1133](#).

This section describes:

- [Primary unit selection and monitored interfaces](#)
- [Primary unit selection and age](#)
- [Primary unit selection and device priority](#)
- [Primary unit selection and the FortiGate unit serial number](#)
- [Points to remember about primary unit selection](#)

## Primary unit selection and monitored interfaces

If you have configured interface monitoring the cluster unit with the highest number of monitored interfaces that are connected to networks becomes the primary unit. Put another way, the cluster unit with the highest number of failed or disconnected monitored interfaces cannot become the primary unit.

Normally, when a cluster starts up, all monitored interfaces of all cluster units are connected and functioning normally. So monitored interfaces do not usually affect primary unit selection when the cluster first starts.

A cluster always renegotiates when a monitored interface fails or is disconnected (called link failover). A cluster also always renegotiates when a failed or disconnected monitored interface is restored.

If a primary unit monitored interface fails or is disconnected, the cluster renegotiates and if this is the only failed or disconnected monitored interface the cluster selects a new primary unit.

If a subordinate unit monitored interface fails or is disconnected, the cluster also renegotiates but will not necessarily select a new primary unit. However, the subordinate unit with the failed or disconnected monitored interface cannot become the primary unit.

Multiple monitored interfaces can fail or become disconnected on more than one cluster unit. Each time a monitored interface is disconnected or fails, the cluster negotiates to select the cluster unit with the most connected and operating monitored interfaces to become the primary unit. In fact, the intent of the link failover feature is just this, to make sure that the primary unit is always the cluster unit with the most connected and operating monitored interfaces. For information about monitored interfaces and link failover see [“Link failover \(port monitoring or interface monitoring\)” on page 1319](#).

## Primary unit selection and age

The cluster unit with the highest age value becomes the primary unit. The age of a cluster unit is the amount of time since a monitored interface failed or is disconnected. Age is also reset when a cluster unit starts (boots up). So, when all cluster units start up at about the same time, they all have the same age. Age does not affect primary unit selection when all cluster units start up at the same time. Age also takes precedence over priority for primary unit selection.

If a link failure of a monitored interface occurs, the age value for the cluster unit that experiences the link failure is reset. So, the cluster unit that experienced the link failure also has a lower age value than the other cluster units. This reduced age does not effect primary unit selection because the number of link failures takes precedence over the age.

If the failed monitored interface is restored the cluster unit that had the failed monitored interface cannot become the primary unit because its age is still lower than the age of the other cluster units.

In most cases, the way that age is handled by the cluster reduces the number of times the cluster selects a new primary unit, which results in a more stable cluster since selecting a new primary unit has the potential to disrupt traffic.

### Cluster age difference margin (grace period)

In any cluster, some of the cluster units may take longer to start up than others. This startup time difference can happen as a result of a number of issues and does not affect the normal operation of the cluster. To make sure that cluster units that start slower can still become primary units, by default the FGCP ignores age differences of up to 5 minutes (300 seconds).

In most cases, during normal operation this age difference margin or grace period helps clusters function as expected. However, the age difference margin can result in some unexpected behavior in some cases:

- During a cluster firmware upgrade with `uninterruptible-upgrade` enabled (the default configuration) the cluster should not select a new primary unit after the firmware of all cluster units has been updated. But since the age difference of the cluster units is most likely less than 300 seconds, age is not used to affect primary unit selection and the cluster may select a new primary unit. See [“Upgrading cluster firmware” on page 1274](#) for more information.
- During failover testing where cluster units are failed over repeatedly the age difference between the cluster units will most likely be less than 5 minutes. During normal operation, if a failover occurs, when the failed unit rejoins the cluster its age will be very different from the age of the still operating cluster units so the cluster will not select a new primary unit. However, if a unit fails and is restored in a very short time the age difference may be less than 5 minutes. As a result the cluster may select a new primary unit during some failover testing scenarios.

### Changing the cluster age difference margin

You can change the cluster age difference margin using the following command:

```
config system ha
 set ha-uptime-diff-margin 60
end
```

This command sets the cluster age difference margin to 60 seconds (1 minute). The age difference margin range 1 to 65535 seconds. The default is 300 seconds.

You may want to reduce the margin if during failover testing you don't want to wait the default age difference margin of 5 minutes. You may also want to reduce the margin to allow uninterruptible upgrades to work. See [“Upgrading cluster firmware” on page 1274](#).

You may want to increase the age margin if cluster unit startup time differences are larger than 5 minutes.

## Displaying cluster unit age differences

You can use the CLI command `diagnose sys ha dump-by all-vcluster` to display the age difference of the units in a cluster. This command also displays information about a number of HA-related parameters for each cluster unit. You can enter the command from the primary unit CLI or you can enter the command from a subordinate unit after using `execute ha manage` to log into a subordinate unit CLI. The information displayed by the command is relative to the unit that you enter the command from.

For example, a cluster of two FortiGate-5001C units with no changes to the default HA configuration except to enable port monitoring for port1. Entering the `diagnose sys ha dump-by all-vcluster` command from the primary unit CLI displays information similar to the following:

```
diagnose sys ha dump-by all-vcluster
 HA information.
vcluster id=1, nentry=2, state=work, digest=4.a5.60.11.cf.d4...
ventry idx=0, id=1, FG-5KC3E13800084, prio=128, -50, claimed=0,
 override=0, flag=1, time=0, mon=0
 mondev=port1, 50
ventry idx=1, id=1, FG-5KC3E13800046, prio=128, -50, claimed=0,
 override=0, flag=0, time=-98, mon=0
```

The command displays one `ventry` line for each cluster unit. The first `ventry` in the example contains information for the cluster unit that you are logged into (usually the primary unit). The other `ventry` lines contain information for the other units in the cluster (in the example there is only one other cluster unit). The command also includes a `mondev` entry that displays the interface monitoring configuration.

The `time` field is always 0 for the unit that you are logged into. The `time` field for the other cluster unit is the age difference between the unit that you are logged into and the other cluster unit. The age difference is in the form seconds/10.

In the example, the age of the primary unit is 12.9 seconds more than the age of the subordinate unit. The age difference is less than 5 minutes (less than 300 seconds) so age has no affect on primary unit selection. The cluster selected the unit with the highest serial number to be the primary unit.

If you use `execute ha manage 1` to log into the subordinate unit CLI and enter `diagnose sys ha dump 1` you get results similar to the following:

```
diagnose sys ha dump-by all-vcluster
 HA information.
vcluster id=1, nentry=2, state=standby, digest=4.a5.60.11.cf.d4...
ventry idx=1, id=1, FG-5KC3E13800046, prio=128, -50, claimed=0,
 override=0, flag=1, time=0, mon=0
 mondev=port1, 50
ventry idx=0, id=1, FG-5KC3E13800084, prio=128, -50, claimed=1,
 override=0, flag=0, time=98, mon=0
```

The `time` for the primary unit is 98, indicating that age of the subordinate unit age is 9.8 seconds higher than the primary unit age.

If port1 (the monitored interface) of the primary unit is disconnected, the cluster renegotiates and the former subordinate unit becomes the primary unit. When you log into the new primary unit CLI and enter `diagnose sys ha dump-by all-vcluster` you could get results similar to the following:

```
diagnose sys ha dump-by all-vcluster
 HA information.
vcluster id=1, nentry=2, state=work, digest=3.f8.d1.63.4d.d2...
ventry idx=0,id=1,FG-5KC3E13800046,prio=128,0,claimed=0,
 override=0,flag=1,time=0,mon=0
 mondev=port1,50
ventry idx=1,id=1,FG-5KC3E13800084,prio=128,-50,claimed=0,
 override=0,flag=0,time=1362,mon=0
```

The command results show that the age of the new primary unit is 136.2 seconds higher than the age of the new subordinate unit.

If port1 of the former primary unit is reconnected the cluster will once again make this the primary unit because the age difference will still be less than 300 seconds. When you log into the primary unit CLI and enter `diagnose sys ha dump-by all-vcluster` you get results similar to the following:

```
diagnose sys ha dump-by all-vcluster
 HA information.
vcluster id=1, nentry=2, state=work, digest=4.a5.60.11.cf.d4...
ventry idx=0,id=1,FG-5KC3E13800084,prio=128,0,claimed=0,
 override=0,flag=1,time=0,mon=0
 mondev=port1,50
ventry idx=1,id=1,FG-5KC3E13800046,prio=128,0,claimed=0,
 override=0,flag=0,time=-1362,mon=0
```

## Resetting the age of all cluster units

In some cases, age differences among cluster units can result in the wrong cluster unit or the wrong virtual cluster becoming the primary unit. For example, if a cluster unit set to a high priority reboots, that unit will have a lower age than other cluster units when it rejoins the cluster. Since age takes precedence over priority, the priority of this cluster unit will not be a factor in primary unit selection.

This problem also affects virtual cluster VDOM partitioning in a similar way. After a reboot of one of the units in a virtual cluster configuration, traffic for all VDOMs could continue to be processed by the cluster unit that did not reboot. This can happen because the age of both virtual clusters on the unit that did not reboot is greater than the age of both virtual clusters on the unit that rebooted.

One way to resolve this issue is to reboot all of the cluster units at the same time so that the age of all of the cluster units is reset. However, rebooting cluster units may interrupt or at least slow down traffic. If you would rather not reboot all of the cluster units you can instead use the following command to reset the age of individual cluster units.

```
diagnose sys ha reset-uptime
```

This command resets the age of a unit back to zero so that if no other unit in the cluster was reset at the same time, it will now have the lowest age. You would use this command to reset the age of the cluster unit that is currently the primary unit. Since it will have the lowest age, the other unit in the cluster will have the highest age and can then become the primary unit.



The `diagnose sys ha reset-uptime` command should only be used as a temporary solution. The command resets the HA age internally and does not affect the up time displayed for cluster units using the `diagnose sys ha dump-by all-vcluster` command or the up time displayed on the Dashboard or cluster members list. To make sure the actual up time for cluster units is the same as the HA age you should reboot the cluster units during a maintenance window.

## Primary unit selection and device priority

A cluster unit with the highest device priority becomes the primary unit when the cluster starts up or renegotiates. By default, the device priority for all cluster units is 128. You can change the device priority to control which FortiGate unit becomes the primary unit during cluster negotiation. All other factors that influence primary unit selection either cannot be configured (age and serial number) or are synchronized among all cluster units (interface monitoring). You can set a different device priority for each cluster unit. During negotiation, if all monitored interfaces are connected, and all cluster units enter the cluster at the same time (or have the same age), the cluster with the highest device priority becomes the primary unit.

A higher device priority does not affect primary unit selection for a cluster unit with the most failed monitored interfaces or with an age that is higher than all other cluster units because failed monitored interfaces and age are used to select a primary unit before device priority.

Increasing the device priority of a cluster unit does not always guarantee that this cluster unit will become the primary unit. During cluster operation, an event that may affect primary unit selection may not always result in the cluster renegotiating. For example, when a unit joins a functioning cluster, the cluster will not renegotiate. So if a unit with a higher device priority joins a cluster the new unit becomes a subordinate unit until the cluster renegotiates.



Enabling the `override` HA CLI keyword makes changes in device priority more effective by causing the cluster to negotiate more often to make sure that the primary unit is always the unit with the highest device priority. For more information about `override`, see [“HA override” on page 1138](#).

## Controlling primary unit selection by changing the device priority

You set a different device priority for each cluster unit to control the order in which cluster units become the primary unit when the primary unit fails.

To change the device priority from the web-based manager go to *Config > System > HA* and change the *Device Priority*.

Enter the following CLI command to change the device priority to 200:

```
config system ha
 set priority 200
end
```

The device priority is not synchronized among cluster units. In a functioning cluster you can change the device priority of any unit in the cluster. Whenever you change the device priority of a cluster unit, when the cluster negotiates, the unit with the highest device priority becomes the primary unit.



The following example shows how to change the device priority of a subordinate unit to 255 so that this subordinate unit becomes the primary unit. You can change the device priority of a subordinate unit by going to *Config > System > HA* and selecting the Edit icon for the subordinate unit. Or from the CLI you can use the `execute ha manage 0` command to connect to the highest priority subordinate unit. After you enter the following commands the cluster renegotiates and selects a new primary unit.

```
execute ha manage 1
 config system ha
 set priority 255
 end
```

If you have three units in a cluster you can set the device priorities as shown in [Table 55](#). When the cluster starts up, cluster unit A becomes the primary unit because it has the highest device priority. If unit A fails, unit B becomes the primary unit because unit B has a higher device priority than unit C.

**Table 55:** Example device priorities for a cluster of three FortiGate units

Cluster unit	Device priority
A	200
B	100
C	50

When configuring HA you do not have to change the device priority of any of the cluster units. If all cluster units have the same device priority, when the cluster first starts up the FGCP negotiates to select the cluster unit with the highest serial number to be the primary unit. Clusters also function normally if all units have the same device priority.

You can change the device priority if you want to control the roles that individual units play in the cluster. For example, if you want the same unit to always become the primary unit, set this unit device priority higher than the device priority of other cluster units. Also, if you want a cluster unit to always become a subordinate unit, set this cluster unit device priority lower than the device priority of other cluster units.

If you have a cluster of three units you can set a different priority for each unit to control which unit becomes the primary unit when all three cluster units are functioning and which will be the primary unit when two cluster units are functioning.

The device priority range is 0 to 255. The default device priority is 128.

If you are configuring a virtual cluster, if you have added virtual domains to both virtual clusters, you can set the device priority that the cluster unit has in virtual cluster 1 and virtual cluster 2. If a FortiGate unit has different device priorities in virtual cluster 1 and virtual cluster 2, the FortiGate unit may be the primary unit in one virtual cluster and the subordinate unit in the other. For more information, see [“Virtual clustering and load balancing or VDOM partitioning”](#) on [page 1218](#).

## Primary unit selection and the FortiGate unit serial number

The cluster unit with the highest serial number is more likely to become the primary unit. When first configuring FortiGate units to be added to a cluster, if you do not change the device priority of any cluster unit, then the cluster unit with the highest serial number always becomes the primary unit.

Age does take precedence over serial number, so if a cluster unit takes longer to join a cluster for some reason (for example if one cluster unit is powered on after the others), that cluster unit will not become the primary unit because the other units have been in the cluster longer.

Device priority and failed monitored interfaces also take precedence over serial number. A higher device priority means a higher priority. So if you set the device priority of one unit higher or if a monitored interface fails, the cluster will not use the FortiGate serial number to select the primary unit.

## Points to remember about primary unit selection

Some points to remember about primary unit selection:

- The FGCP compares primary unit selection criteria in the following order: Failed Monitored interfaces > Age > Device Priority > Serial number. The selection process stops at the first criteria that selects one cluster unit.
- Negotiation and primary unit selection is triggered if a cluster unit fails or if a monitored interface fails.
- If the HA age difference is more than 5 minutes (300 seconds), the cluster unit that is operating longer becomes the primary unit.
- If HA age difference is less than 5 minutes (300 seconds), the device priority and FortiGate serial number selects the cluster unit to become the primary unit.
- Every time a monitored interface fails the HA age of the cluster unit is reset to 0.
- Every time a cluster unit restarts the HA age of the cluster unit is reset to 0.

## HA override

The HA `override` CLI keyword is disabled by default. When `override` is disabled a cluster may not always renegotiate when an event occurs that affects primary unit selection. For example, when `override` is disabled a cluster will not renegotiate when you change a cluster unit device priority or when you add a new cluster unit to a cluster. This is true even if the unit added to the cluster has a higher device priority than any other unit in the cluster. Also, when `override` is disabled a cluster does not negotiate if the new unit added to the cluster has a failed or disconnected monitored interface.



For a virtual cluster configuration, `override` is enabled by default for both virtual clusters when you enable virtual cluster 2. For more information, see [“Virtual clustering and HA override” on page 1218](#).

In most cases you should keep `override` disabled to reduce how often the cluster negotiates. Frequent negotiations may cause frequent traffic interruptions.

However, if you want to make sure that the same cluster unit always operates as the primary unit and if you are less concerned about frequent cluster negotiation you can set its device priority higher than other cluster units and enable `override`.

To enable `override`, connect to each cluster unit CLI (using the `execute ha manage` command) and use the `config system ha` CLI command to enable `override`.

For `override` to be effective, you must also set the device priority highest on the cluster unit that you want to always be the primary unit. To increase the device priority, from the CLI use the `config system ha` command and increase the value of the `priority` keyword to a number higher than the default priority of 128.

You can also increase the device priority from the web-based manager by going to *System > Config > HA*. To increase the device priority of the primary unit select edit for the primary or subordinate unit and set the *Device Priority* to a number higher than 128.



The `override` setting and device priority value are not synchronized to all cluster units. You must enable `override` and adjust device priority manually and separately for each cluster unit.

---

With `override` enabled, the primary unit with the highest device priority will always become the primary unit. Whenever an event occurs that may affect primary unit selection, the cluster negotiates. For example, when `override` is enabled a cluster renegotiates when you change the device priority of any cluster unit or when you add a new unit to a cluster.

This section also describes:

- [Override and primary unit selection](#)
- [Controlling primary unit selection using device priority and override](#)
- [Points to remember about primary unit selection when override is enabled](#)
- [Configuration changes can be lost if override is enabled](#)
- [Override and disconnecting a unit from a cluster](#)

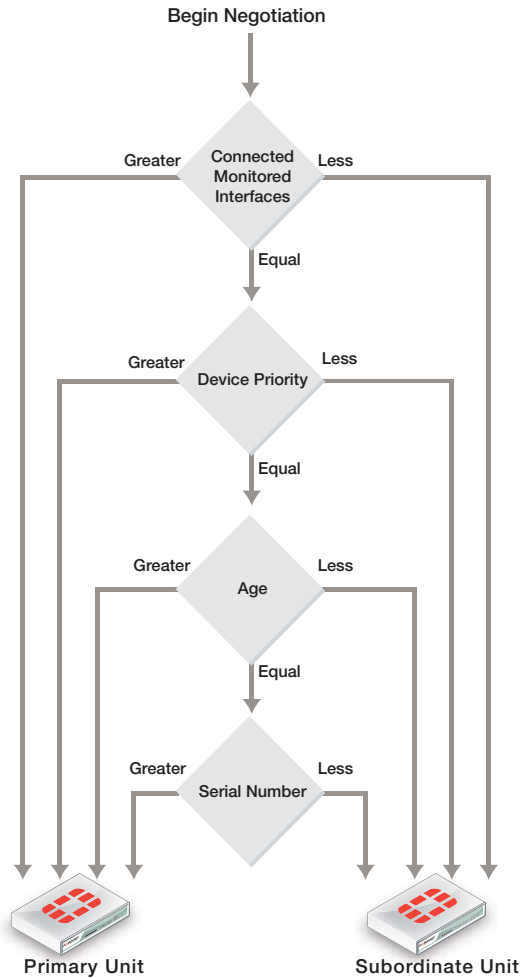
## Override and primary unit selection

Enabling `override` changes the order of primary unit selection. As shown in [Figure 189](#) if `override` is enabled, primary unit selection considers device priority before age and serial number. This means that if you set the device priority higher on one cluster unit, with `override` enabled this cluster unit becomes the primary unit even if its age and serial number are lower than other cluster units.

Similar to when `override` is disabled, when `override` is enabled primary unit selection checks for connected monitored interfaces first. So if interface monitoring is enabled, the cluster unit with the most disconnected monitored interfaces cannot become the primary unit, even if the unit has the highest device priority.

If all monitored interfaces are connected (or interface monitoring is not enabled) and the device priority of all cluster units is the same then age and serial number affect primary unit selection.

**Figure 189:** Selecting the primary unit with override enabled



## Controlling primary unit selection using device priority and override

To configure one cluster unit to always become the primary unit you should set its device priority to be higher than the device priorities of the other cluster units and you should enable `override` on all cluster units.

Using this configuration, when the cluster is operating normally the primary unit is always the unit with the highest device priority. If the primary unit fails the cluster renegotiates to select another cluster unit to be the primary unit. If the failed primary unit recovers, starts up again and rejoins the cluster, because `override` is enabled, the cluster renegotiates. Because the restarted primary unit has the highest device priority it once again becomes the primary unit.

In the same situation with `override` disabled, because the age of the failed primary unit is lower than the age of the other cluster units, when the failed primary unit rejoins the cluster it does not become the primary unit. Instead, even though the failed primary unit may have the highest device priority it becomes a subordinate unit because its age is lower than the age of all the other cluster units.

## Points to remember about primary unit selection when override is enabled

Some points to remember about primary unit selection when `override` is enabled:

- The FGCP compares primary unit selection criteria in the following order: Failed Monitored Interfaces > Device Priority > Age > Serial number. The selection process stops at the first criteria that selects one cluster unit.
- Negotiation and primary unit selection is triggered whenever an event occurs which may affect primary unit selection. For example negotiation occurs, when you change the device priority, when you add a new unit to a cluster, if a cluster unit fails, or if a monitored interface fails.
- Device priority is considered before age. Otherwise age is handled the same when `override` is enabled.

## Configuration changes can be lost if override is enabled

In some cases, when `override` is enabled and you make configuration changes to an HA cluster these changes can be lost. For example, consider the following sequence:

1. A cluster of two FortiGate units is operating with `override` enabled.
  - FGT-A: Primary unit with device priority 200 and with `override` enabled
  - FGT-B: Subordinate unit with device priority 100 and with `override` disabled
  - If both units are operating, FGT-A always becomes the primary unit because FGT-A has the highest device priority.
2. FGT-A fails and FGT-B becomes the new primary unit.
3. The administrator makes configuration changes to the cluster.

The configuration changes are made to FGT-B because FGT-B is operating as the primary unit. These configuration changes are not synchronized to FGT-A because FGT-A is not operating.
4. FGT-A is restored and starts up again.
5. The cluster renegotiates and FGT-A becomes the new primary unit.
6. The cluster recognizes that the configurations of FGT-A and FGT-B are not the same.
7. The configuration of FGT-A is synchronized to FGT-B.

The configuration is always synchronized from the primary unit to the subordinate units.
8. The cluster is now operating with the same configuration as FGT-A. The configuration changes made to FGT-B have been lost.

### The solution

When `override` is enabled, you can prevent configuration changes from being lost by doing the following:

- Verify that all cluster units are operating before making configuration changes (from the web-based manager go to *System > Config > HA* to view the cluster members list or from the FortiOS CLI enter `get system ha status`).
- Make sure the device priority of the primary unit is set higher than the device priorities of all other cluster units before making configuration changes.
- Disable `override` either permanently or until all configuration changes have been made and synchronized to all cluster units.

## Override and disconnecting a unit from a cluster

A similar scenario to that described in “[Configuration changes can be lost if override is enabled](#)” may occur when `override` is enabled and you use the Disconnect from Cluster option from the web-based manager or the `execute ha disconnect` command from the CLI to disconnect a cluster unit from a cluster.

Configuration changes made to the cluster can be lost when you reconnect the disconnected unit to the cluster. You should make sure that the device priority of the disconnected unit is lower than the device priority of the current primary unit. Otherwise, when the disconnected unit joins the cluster, if `override` is enabled, the cluster renegotiates and the disconnected unit may become the primary unit. If this happens, the configuration of the disconnected unit is synchronized to all other cluster units and any configuration changes made between when the unit was disconnected and reconnected are lost.

## FortiGate HA compatibility with PPPoE and DHCP

FortiGate HA is not compatible with PPP protocols such as PPPoE. FortiGate HA is also not compatible with DHCP. If one or more FortiGate unit interfaces is dynamically configured using DHCP or PPPoE you cannot switch to operate in HA mode. Also, you cannot switch to operate in HA mode if one or more FortiGate unit interfaces is configured as a PPTP or L2TP client.



Configuring an interface for DHCP or PPPoE is only supported in NAT/Route mode. So, usually when configuring HA in Transparent mode an interface being configured for DHCP or PPPoE should not affect HA operation. However, in some cases you may not be able to enable HA if you had configured an interface for DHCP or PPPoE before switching to Transparent mode. So, if you are blocked from operating a Transparent mode FortiGate unit in HA and cannot find another reason for the problem, try switching the FortiGate unit back to NAT/Route mode and setting all interface modes to static before switching to Transparent mode and enabling HA. You could also enable HA before switching to Transparent mode.

You can configure a cluster to act as a DHCP server or a DHCP relay agent. In both active-passive and active-active clusters DHCP relay sessions are always handled by the primary unit. It is possible that a DHCP relay session could be interrupted by a failover. If this occurs the DHCP relay session is not resumed after the failover and the DHCP client may have to repeat the DHCP request.

When a cluster is operating as a DHCP server the primary unit responds to all DHCP requests and maintains the DHCP server address lease database. The cluster also dynamically synchronizes the DHCP server address lease database to the subordinate units. If a failover occurs, the new primary unit will have an up-to-date DHCP server address lease database. Synchronizing the DHCP address lease database prevents the new primary unit from responding incorrectly to new DHCP requests after a failover.

Also, it is possible that when FortiGate units first negotiate to form a cluster that a unit that ends up as a subordinate unit in the cluster will have information in its DHCP address lease database that the cluster unit operating as the primary unit does not have. This can happen if a FortiGate unit responds to DHCP requests while operating as a standalone unit and then when the cluster is formed this unit becomes a subordinate unit. Because of this possibility, after a cluster is formed the DHCP address lease databases of all of the cluster units are merged into one database which is then synchronized to all cluster units.

## HA and distributed clustering

The FGCP supports widely separated cluster units installed in different physical locations. Distributed clusters can have cluster units in different rooms in the same building, different buildings in the same location, or even different geographical sites such as different cities, countries or continents.

Just like any cluster, distributed clusters require heartbeat communication between cluster units. In a distributed cluster this heartbeat communication can take place over the Internet or over other transmission methods including satellite linkups.

Because of the possible distance it may take a relatively long time for heartbeat packets to be transmitted between cluster units. To support a distributed cluster you may need to increase the heartbeat interval so that the cluster expects extra time between heartbeat packets. For information about changing the heartbeat interval and other heartbeat related settings, see [“Modifying heartbeat timing” on page 1298](#).

## Hard disk configuration and HA

If your cluster units include hard disks, all cluster units must have identical hard disk configurations. This means each cluster unit must have same number of hard disks (including AMC and FortiGate Storage Module (FSM) hard disks) and also means that matching hard disks in each cluster unit must be the same size, have the same hard disk format, and have the same number of partitions.

In most cases the default hard disk configuration of the cluster units will be compatible. However, a hard disk formatted by an older FortiGate firmware version may not be compatible with a hard disk formatted by a more recent firmware version. Problems may also arise if you have used the `execute scsi-dev` command to add or change hard disk protections.

If a cluster unit CLI displays hard disk compatibility messages, you may need to use the `execute scsi-dev delete` command to delete partitions. You can also use the `execute formatlogdisk` command to reformat hard disks. In some cases after deleting all partitions and reformatting the hard disks, you may still see hard disk incompatibility messages. If this happens, contact Fortinet Customer Support for assistance.

## FGCP high availability best practices

Fortinet suggests the following practices related to high availability:

- Use Active-Active HA to distribute TCP and UTM sessions among multiple cluster units. An active-active cluster may have higher throughput than a standalone FortiGate unit or than an active-passive cluster.
- Use a different host name on each FortiGate unit when configuring an HA cluster. Fewer steps are required to add host names to each cluster unit before configuring HA and forming a cluster.
- Consider adding an Alias to the interfaces used for the HA heartbeat so that you always get a reminder about what these interfaces are being used for.
- Enabling `load-balance-all` can increase device and network load since more traffic is load-balanced. This may be appropriate for use in a deployment using the firewall capabilities of the FortiGate unit and IPS but no other content inspection. See [“Load balancing UTM sessions, TCP sessions, and UDP sessions” on page 1347](#).
- An advantage of using session pickup is that non-content inspection sessions will be picked up by the new primary unit after a failover. The disadvantage is that the cluster generates more heartbeat traffic to support session pickup as a larger portion of the session table must

be synchronized. Session pickup should be configured only when required and is not recommended for use with SOHO FortiGate models. Session pickup should only be used if the primary heartbeat link is dedicated (otherwise the additional HA heartbeat traffic could affect network performance). See [“Session failover \(session pick-up\)” on page 1330](#).

- If session pickup is not selected, after a device or link failover all sessions are briefly interrupted and must be re-established at the application level after the cluster renegotiates. For example, after a failover, users browsing the web can just refresh their browsers to resume browsing. Users downloading large files may have to restart their download after a failover. Other protocols may experience data loss and some protocols may require sessions to be manually restarted. For example, a user downloading files with FTP may have to either restart downloads or restart their FTP client.
- If you need to enable session pickup, consider enabling session pickup delay to improve performance by reducing the number of sessions that are synchronized. If possible, also consider enabling session synchronization or multiple FortiGate Interfaces. See [“Improving session synchronization performance” on page 1331](#) for more information.
- To avoid unpredictable results, when you connect a switch to multiple redundant or aggregate interfaces in an active-passive cluster you should configure separate redundant or aggregate interfaces on the switch; one for each cluster unit. See [“HA MAC addresses and 802.3ad aggregation” on page 1188](#).
- Use SNMP, syslog, or email alerts to monitor a cluster for failover messages. Alert messages about cluster failovers may help find and diagnose network problems quickly and efficiently. See [“Operating a cluster” on page 1252](#).

## Heartbeat interfaces

Fortinet suggests the following practices related to heartbeat interfaces:



Do not use a FortiGate switch port for the HA heartbeat traffic. This configuration is not supported.

- 
- Isolate heartbeat interfaces from user networks. Heartbeat packets contain sensitive cluster configuration information and can consume a considerable amount of network bandwidth. If the cluster consists of two FortiGate units, connect the heartbeat interfaces directly using a crossover cable or a regular Ethernet cable. For clusters with more than two units, connect heartbeat interfaces to a separate switch that is not connected to any network.
  - If heartbeat traffic cannot be isolated from user networks, enable heartbeat message encryption and authentication to protect cluster information. See [“Enabling or disabling HA heartbeat encryption and authentication” on page 1299](#).
  - Configure and connect redundant heartbeat interfaces so that if one heartbeat interface fails or becomes disconnected, HA heartbeat traffic can continue to be transmitted using the backup heartbeat interface. If heartbeat communication fails, all cluster members will think they are the primary unit resulting in multiple devices on the network with the same IP addresses and MAC addresses (condition referred to as *Split Brain*) and communication will be disrupted until heartbeat communication can be reestablished.
  - Do not monitor dedicated heartbeat interfaces; monitor those interfaces whose failure should trigger a device failover.



## Interface monitoring (port monitoring)

Fortinet suggests the following practices related to interface monitoring (also called port monitoring):

- Wait until a cluster is up and running and all interfaces are connected before enabling interface monitoring. A monitored interface can easily become disconnected during initial setup and cause failovers to occur before the cluster is fully configured and tested.
- Monitor interfaces connected to networks that process high priority traffic so that the cluster maintains connections to these networks if a failure occurs.
- Avoid configuring interface monitoring for all interfaces.
- Supplement interface monitoring with remote link failover. Configure remote link failover to maintain packet flow if a link not directly connected to a cluster unit (for example, between a switch connected to a cluster interface and the network) fails. See [“Remote link failover” on page 1325](#).

## Troubleshooting

The following sections in this document contain troubleshooting information:

- [“Troubleshooting HA clusters” on page 1212](#)
- [“Troubleshooting virtual clustering” on page 1238](#)
- [“Troubleshooting full mesh HA” on page 1251](#)
- [“Troubleshooting layer-2 switches” on page 1361](#)

## FGCP HA terminology

The following HA-specific terms are used in this document.

### Cluster

A group of FortiGate units that act as a single virtual FortiGate unit to maintain connectivity even if one of the FortiGate units in the cluster fails.

### Cluster unit

A FortiGate unit operating in a FortiGate HA cluster.

### Device failover

Device failover is a basic requirement of any highly available system. Device failover means that if a device fails, a replacement device automatically takes the place of the failed device and continues operating in the same manner as the failed device. See also [“Device failover, link failover, and session failover” on page 1130](#).

### Failover

A FortiGate unit taking over processing network traffic in place of another unit in the cluster that suffered a device failure or a link failure.

### Failure

A hardware or software problem that causes a FortiGate unit or a monitored interface to stop processing network traffic.

## FGCP

The FortiGate clustering protocol (FGCP) that specifies how the FortiGate units in a cluster communicate to keep the cluster operating.

## Full mesh HA

Full mesh HA is a method of removing single points of failure on a network that includes an HA cluster. FortiGate models that support redundant interfaces can be used to create a cluster configuration called full mesh HA. Full mesh HA includes redundant connections between all network components. If any single component or any single connection fails, traffic switches to the redundant component or connection.

## HA virtual MAC address

When operating in HA mode, all of the interfaces of the primary unit acquire the same HA virtual MAC address. All communications with the cluster must use this MAC address. The HA virtual MAC address is set according to the group ID.

## Heartbeat

Also called FGCP heartbeat or HA heartbeat. The heartbeat constantly communicates HA status and synchronization information to make sure that the cluster is operating properly.

## Heartbeat device

An ethernet network interface in a cluster that is used by the FGCP for heartbeat communications among cluster units.

## Heartbeat failover

If an interface functioning as the heartbeat device fails, the heartbeat is transferred to another interface also configured as an HA heartbeat device.

## Hello state

In the hello state a cluster unit has powered on in HA mode, is using HA heartbeat interfaces to send hello packets, and is listening on its heartbeat interfaces for hello packets from other FortiGate units. Hello state may appear in HA log messages.

## High availability

The ability that a cluster has to maintain a connection when there is a device or link failure by having another unit in the cluster take over the connection, without any loss of connectivity. To achieve high availability, all FortiGate units in the cluster share session and configuration information.

## Interface monitoring

You can configure interface monitoring (also called port monitoring) to monitor FortiGate interfaces to verify that the monitored interfaces are functioning properly and connected to their networks. If a monitored interface fails or is disconnected from its network the interface leaves the cluster and a link failover occurs. For more information about interface monitoring, see [“Link failover \(port monitoring or interface monitoring\)” on page 1319](#).

## Link failover

Link failover means that if a monitored interface fails, the cluster reorganizes to re-establish a link to the network that the monitored interface was connected to and to continue operating with minimal or no disruption of network traffic. See also [“Device failover, link failover, and session failover” on page 1130](#).

## Load balancing

Also known as active-active HA. All units in the cluster process network traffic. The FGCP employs a technique similar to unicast load balancing. The primary unit interfaces are assigned virtual MAC addresses which are associated on the network with the cluster IP addresses. The primary unit is the only cluster unit to receive packets sent to the cluster. The primary unit can process packets itself, or propagate them to subordinate units according to a load balancing schedule. Communication between the cluster units uses the actual cluster unit MAC addresses.

## Monitored interface

An interface that is monitored by a cluster to make sure that it is connected and operating correctly. The cluster monitors the connectivity of this interface for all cluster units. If a monitored interface fails or becomes disconnected from its network, the cluster will compensate.

## Primary unit

Also called the primary cluster unit, this cluster unit controls how the cluster operates. The primary unit sends hello packets to all cluster units to synchronize session information, synchronize the cluster configuration, and to synchronize the cluster routing table. The hello packets also confirm for the subordinate units that the primary unit is still functioning.

The primary unit also tracks the status of all subordinate units. When you start a management connection to a cluster, you connect to the primary unit.

In an active-passive cluster, the primary unit processes all network traffic. If a subordinate unit fails, the primary unit updates the cluster configuration database.

In an active-active cluster, the primary unit receives all network traffic and re-directs this traffic to subordinate units. If a subordinate unit fails, the primary unit updates the cluster status and redistributes load balanced traffic to other subordinate units in the cluster.

The FortiGate firmware uses the term master to refer to the primary unit.

## Session failover

Session failover means that a cluster maintains active network sessions after a device or link failover. FortiGate HA does not support session failover by default. To enable session failover you must change the HA configuration to select Enable Session Pick-up. See also [“Device failover, link failover, and session failover” on page 1130](#).

## Session pickup

If you enable session pickup for a cluster, if the primary unit fails or a subordinate unit in an active-active cluster fails, all communication sessions with the cluster are maintained or picked up by the cluster after the cluster negotiates to select a new primary unit.

If session pickup is not a requirement of your HA installation, you can disable this option to save processing resources and reduce the network bandwidth used by HA session synchronization. In many cases interrupted sessions will resume on their own after a failover even if session pickup is not enabled. You can also enable session pickup delay to reduce the number of sessions that are synchronized by session pickup.

### Standby state

A subordinate unit in an active-passive HA cluster operates in the standby state. In a virtual cluster, a subordinate virtual domain also operates in the standby state. The standby state is actually a hot-standby state because the subordinate unit or subordinate virtual domain is not processing traffic but is monitoring the primary unit session table to take the place of the primary unit or primary virtual domain if a failure occurs.

In an active-active cluster all cluster units operate in a work state.

When standby state appears in HA log messages this usually means that a cluster unit has become a subordinate unit in an active-passive cluster or that a virtual domain has become a subordinate virtual domain.

### State synchronization

The part of the FGCP that maintains connections after failover.

### Subordinate unit

Also called the subordinate cluster unit, each cluster contains one or more cluster units that are not functioning as the primary unit. Subordinate units are always waiting to become the primary unit. If a subordinate unit does not receive hello packets from the primary unit, it attempts to become the primary unit.

In an active-active cluster, subordinate units keep track of cluster connections, keep their configurations and routing tables synchronized with the primary unit, and process network traffic assigned to them by the primary unit. In an active-passive cluster, subordinate units do not process network traffic. However, active-passive subordinate units do keep track of cluster connections and do keep their configurations and routing tables synchronized with the primary unit.

The FortiGate firmware uses the terms slave and subsidiary unit to refer to a subordinate unit.

### Virtual clustering

Virtual clustering is an extension of the FGCP for FortiGate units operating with multiple VDOMS enabled. Virtual clustering operates in active-passive mode to provide failover protection between two instances of a VDOM operating on two different cluster units. You can also operate virtual clustering in active-active mode to use HA load balancing to load balance sessions between cluster units. Alternatively, by distributing VDOM processing between the two cluster units you can also configure virtual clustering to provide load balancing by distributing sessions for different VDOMs to each cluster unit.

### Work state

The primary unit in an active-passive HA cluster, a primary virtual domain in a virtual cluster, and all cluster units in an active-active cluster operate in the work state. A cluster unit operating in the work state processes traffic, monitors the status of the other cluster units, and tracks the session table of the cluster.

When work state appears in HA log messages this usually means that a cluster unit has become the primary unit or that a virtual domain has become a primary virtual domain.

## HA web-based manager options

Go to *System > Config > HA* to change HA options. You can set the following options to put a FortiGate unit into HA mode. You can also change any of these options while the cluster is operating.

You can configure HA options for a FortiGate unit with virtual domains (VDOMs) enabled by logging into the web-based manager as the global admin administrator and going to *System > Config > HA*.

If already operating in HA mode, go to *System > Config > HA* to display the cluster members list (see “[Cluster members list](#)” on page 1268).

Go to *System > Config > HA* and select *View HA Statistics* to view statistics about cluster operation. See “[Viewing HA statistics](#)” on page 1271.



If your cluster uses virtual domains, you are configuring HA virtual clustering. Most virtual cluster HA options are the same as normal HA options. However, virtual clusters include VDOM partitioning options. Other differences between configuration options for regular HA and for virtual clustering HA are described below and see “[Virtual clusters](#)” on page 1217.



HA is not compatible with PPP protocols such as PPPoE. HA is also not compatible with DHCP. If one or more FortiGate interfaces is dynamically configured using DHCP or PPPoE, you cannot switch to operate in HA mode. You also cannot switch to operate in HA mode if one or more FortiGate interfaces is configured as a PPTP or L2TP client or if the FortiGate unit is configured for standalone session synchronization.

<b>Mode</b>	Select an HA mode for the cluster or return the FortiGate unit in the cluster to standalone mode. When configuring a cluster, you must set all members of the HA cluster to the same HA mode. You can select <i>Standalone</i> (to disable HA), <i>Active-Passive</i> , or <i>Active-Active</i> .  If virtual domains are enabled you can select <i>Active-Passive</i> or <i>Standalone</i> .
<b>Device Priority</b>	Optionally set the device priority of the cluster FortiGate unit. Each FortiGate unit in a cluster can have a different device priority. During HA negotiation, the FortiGate unit with the highest device priority usually becomes the primary unit. See “ <a href="#">Primary unit selection</a> ” on page 1131.  In a virtual cluster configuration, each cluster FortiGate unit can have two different device priorities, one for each virtual cluster. During HA negotiation, the FortiGate unit with the highest device priority in a virtual cluster becomes the primary FortiGate unit for that virtual cluster.  Changes to the device priority are not synchronized. You can accept the default device priority when first configuring a cluster.
<b>Reserve Management Port for Cluster Member</b>	You can provide direct management access to individual cluster units by reserving a management interface as part of the HA configuration. Once this management interface is reserved, you can configure a different IP address, administrative access and other interface settings for this interface for each cluster unit. Then by connecting this interface of each cluster unit to your network you can manage each cluster unit separately from a different IP address. See “ <a href="#">Managing individual cluster units using a reserved management interface</a> ” on page 1254.

---

<b>Group Name</b>	Enter a name to identify the cluster. The maximum length of the group name is 32 characters. The group name must be the same for all cluster units before the cluster units can form a cluster. After a cluster is operating, you can change the group name. The group name change is synchronized to all cluster units.
<b>Password</b>	Enter a password to identify the cluster. The password must be the same for all cluster FortiGate units before the cluster FortiGate units can form a cluster.  Two clusters on the same network must have different passwords.  The password is synchronized to all cluster units in an operating cluster. If you change the password of one cluster unit the change is synchronized to all cluster units.
<b>Enable Session pickup</b>	Select to enable session pickup so that if the primary unit fails, sessions are picked up by the cluster unit that becomes the new primary unit.  You must enable session pickup for session failover protection. If you do not require session failover protection, leaving session pickup disabled may reduce HA CPU usage and reduce HA heartbeat network bandwidth usage. See <a href="#">“Session failover (session pick-up)” on page 1330</a> .
<b>Port Monitor</b>	Select to enable or disable monitoring FortiGate interfaces to verify the monitored interfaces are functioning properly and are connected to their networks. See <a href="#">“Link failover (port monitoring or interface monitoring)” on page 1319</a> .  If a monitored interface fails or is disconnected from its network, the interface leaves the cluster and a link failover occurs. The link failover causes the cluster to reroute the traffic being processed by that interface to the same interface of another cluster FortiGate unit that still has a connection to the network. This other cluster FortiGate unit becomes the new primary unit.  Port monitoring (also called interface monitoring) is disabled by default. Leave port monitoring disabled until the cluster is operating and then only enable port monitoring for connected interfaces.  You can monitor up to 64 interfaces.

---

---

**Heartbeat Interface**

Select to enable or disable HA heartbeat communication for each interface in the cluster and set the heartbeat interface priority. The heartbeat interface with the highest priority processes all heartbeat traffic. If two or more heartbeat interfaces have the same priority, the heartbeat interface with the lowest hash map order value processes all heartbeat traffic. The web-based manager lists interfaces in alphanumeric order:

- port1
- port2 through 9
- port10

Hash map order sorts interfaces in the following order:

- port1
- port10
- port2 through port9

The default heartbeat interface configuration is different for each FortiGate model. This default configuration usually sets the priority of two heartbeat interfaces to 50. You can accept the default heartbeat interface configuration or change it as required.

The heartbeat interface priority range is 0 to 512. The default priority when you select a new heartbeat interface is 0.

You must select at least one heartbeat interface. If heartbeat communication is interrupted, the cluster stops processing traffic. See [“HA heartbeat and communication between cluster units” on page 1293](#).

You can select up to 8 heartbeat interfaces. This limit only applies to units with more than 8 physical interfaces.

---

**VDOM partitioning**

If you are configuring virtual clustering, you can set the virtual domains to be in virtual cluster 1 and the virtual domains to be in virtual cluster 2. The root virtual domain must always be in virtual cluster 1. See [“Virtual clusters” on page 1217](#).

---

# Configuring and connecting HA clusters

This chapter contains general procedures and descriptions as well as detailed configuration examples that describe how to configure FortiGate HA clusters.

The examples in this chapter include example values only. In most cases you will substitute your own values. The examples in this chapter also do not contain detailed descriptions of configuration parameters.

This chapter contains the following sections:

- [About the procedures in this chapter](#)
- [Example: NAT/Route mode active-passive HA configuration](#)
- [Example: Transparent mode active-active HA configuration](#)
- [Example: advanced Transparent mode active-active HA configuration](#)
- [Example: converting a standalone FortiGate unit to a cluster](#)
- [Example: adding a new unit to an operating cluster](#)
- [Example: replacing a failed cluster unit](#)
- [Example: HA and 802.3ad aggregated interfaces](#)
- [Example: HA and redundant interfaces](#)
- [Troubleshooting HA clusters](#)

## About the procedures in this chapter

The procedures in this chapter describe some of many possible sequences of steps for configuring HA clustering. As you become more experienced with FortiOS HA you may choose to use a different sequence of configuration steps.

For simplicity, many of these procedures assume that you are starting with new FortiGate units set to the factory default configuration. However, starting from the default configuration is not a requirement for a successful HA deployment. FortiGate HA is flexible enough to support a successful configuration from many different starting points.

## Example: NAT/Route mode active-passive HA configuration

This section describes a simple HA network topology that includes an HA cluster of two FortiGate-620B units in NAT/Route mode installed between an internal network and the Internet.


- [Example NAT/Route mode HA network topology](#)
- [General configuration steps](#)
- [Configuring a NAT/Route mode active-passive cluster of two FortiGate-620B units - web-based manager](#)
- [Configuring a NAT/Route mode active-passive cluster of two FortiGate-620B units - CLI](#)



## Example NAT/Route mode HA network topology

Figure 190 shows a typical FortiGate-620B HA cluster consisting of two FortiGate-620B units (620\_ha\_1 and 620\_ha\_2) connected to the same internal (port2) and external (port1) networks.

**Figure 190:**Example NAT/Route mode HA network topology



Port3 and port4 are the default FortiGate-620B heartbeat interfaces. Because the cluster consists of two FortiGate units, you can make the connections between the heartbeat interfaces using crossover cables. You could also use switches and regular ethernet cables.

## General configuration steps

The section includes web-based manager and CLI procedures. These procedures assume that the FortiGate-620B units are running the same FortiOS firmware build and are set to the factory default configuration.

### General configuration steps

1. Configure the FortiGate units for HA operation.
  - Optionally change each unit's host name.
  - Configure HA.
2. Connect the cluster to the network.
3. Confirm that the cluster units are operating as a cluster and add basic configuration settings to the cluster.
  - View cluster status from the web-based manager or CLI.
  - Add a password for the admin administrative account.
  - Change the IP addresses and netmasks of the internal and external interfaces.
  - Add a default route.

## Configuring a NAT/Route mode active-passive cluster of two FortiGate-620B units - web-based manager

Use the following procedures to configure two FortiGate-620B units for NAT/Route HA operation using the FortiGate web-based manager. These procedures assume you are starting with two FortiGate-620B units with factory default settings.



Give each cluster unit a unique host name to make the individual units easier to identify when they are part of a functioning cluster. The default FortiGate unit host name is the FortiGate serial number. You may want to change this host name to something more meaningful for your network.



Make sure your FortiGate interfaces are configured with static IP addresses. If any interface gets its address using DHCP you will not be able to configure HA.

### To configure the first FortiGate-620B unit (host name 620\_ha\_1)

1. Power on the first FortiGate unit.
2. On your management computer with an Ethernet connection, set the static IP address to 192.168.1.2 and the netmask to 255.255.255.0.
3. On a management computer, start a web browser and browse to the address <https://192.168.1.99> (remember to include the “s” in https://).  
The FortiGate login is displayed.
4. Type *admin* in the *Name* field and select *Login*.  
The FortiGate dashboard is displayed.
5. On the *System Information* dashboard widget beside *Host Name*, select *Change*.
6. Enter a new Host Name for this FortiGate unit.

<b>New Name</b>	620_ha_1
-----------------	----------

7. Select OK.
8. Go to *System > Config > HA* and change the following settings:

<b>Mode</b>	Active-Passive
-------------	----------------

<b>Group Name</b>	example1.com
-------------------	--------------

<b>Password</b>	HA_pass_1
-----------------	-----------



This is the minimum recommended configuration for an active-active HA cluster. You can also configure other HA options, but if you wait until after the cluster is operating you will only have to configure these options once for the cluster instead of separately for each unit in the cluster.

9. Select OK.

The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC](#)

addresses” on page 1300). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

To confirm these MAC address changes, you can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (MAC) and the port1 permanent MAC address (Permanent\_HWaddr):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

**10.** Power off the first FortiGate unit (620\_ha\_1).

**To configure the second FortiGate-620B unit (host name 620\_ha\_2)**

1. Power on the second FortiGate unit.
2. On a management computer, start a web browser and browse to the address <https://192.168.1.99> (remember to include the “s” in https://).

The FortiGate login is displayed.

3. Type *admin* in the *Name* field and select *Login*.  
The FortiGate dashboard is displayed.
4. On the *System Information* dashboard widget, beside *Host Name* select *Change*.
5. Enter a new Host Name for this FortiGate unit.

---

<b>New Name</b>	620_ha_2
-----------------	----------

---

6. Select OK.
7. Go to *System > Config > HA* and change the following settings:

---

<b>Mode</b>	Active-Passive
-------------	----------------

---

<b>Group Name</b>	example1.com
-------------------	--------------

---

<b>Password</b>	HA_pass_1
-----------------	-----------

---

8. Select OK.

The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and because the FGCP changes the MAC address of the FortiGate unit interfaces.

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

9. Power off the second FortiGate unit.

#### To connect the cluster to the network

1. Connect the port1 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the Internet.
2. Connect the port2 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the internal network.
3. Connect the port3 interfaces of 620\_ha\_1 and 620\_ha\_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
4. Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
5. Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

#### To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.



Once the cluster is operating, because configuration changes are synchronized to all cluster units, configuring the cluster is the same as configuring an individual FortiGate unit. You could have performed the following configuration steps separately on each FortiGate unit before you connected them to form a cluster.

---

1. Start Internet Explorer and browse to the address <https://192.168.1.99> (remember to include the “s” in https://).  
The FortiGate Login is displayed.
2. Type `admin` in the *Name* field and select Login.  
The FortiGate dashboard is displayed.  
The System Information dashboard widget shows the *Cluster Name* (`example1.com`) and the host names and serial numbers of the *Cluster Members*. The Unit Operation widget shows multiple cluster units.

**Figure 191:**Sample FortiGate-620B System Information dashboard widget

System Information		
Cluster Name	example1.com	
Cluster Members	620_ha_2/FG600B3908600825	(Master)
	620_ha_1/FG600B3908600705	(Slave)
Serial Number	FG600B3908600825	
Operation Mode	NAT [Change]	
HA Status	Active-Passive [Configure]	
System Time	Wed Feb 9 14:35:11 2011 [Change]	
Firmware Version	v4.0,build0415,110126 (Interim) [Update]	
System Configuration	Last Backup: N/A [Backup] [Restore]	
Current Administrator	admin [Change Password] /4 in Total [Details]	
Uptime	13 day(s) 7 hour(s) 34 min(s)	
Virtual Domain	Disabled [Enable]	

3. Go to *System > Config > HA* to view the cluster members list.  
The list shows both cluster units, their host names, their roles in the cluster, and their priorities. You can use this list to confirm that the cluster is operating normally. For example, if the list shows only one cluster unit then the other unit has left the cluster for some reason.

**Figure 192:**Sample FortiGate-620B cluster members list

HA Cluster		View HA Statistics		
	Cluster Member	Hostname	Role	Priority
		620_ha_2	MASTER	128
		620_ha_1	SLAVE	128

### To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units, the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 1212](#) to troubleshoot the cluster.

### To add basic configuration settings to the cluster

Use the following steps to configure the cluster to connect to its network. The following are example configuration steps only and do not represent all of the steps required to configure the cluster for a given network.

1. Log into the cluster web-based manager.
2. Go to *System > Admin > Administrators*.
3. For `admin`, select the *Change Password* icon
4. Enter and confirm a new password.

5. Select OK.
6. Go to *System > Network > Interfaces*.
7. Edit the *port2* interface and change *IP/Netmask* to 10.11.101.100/24.
8. Select OK.



After changing the IP address of the port1 interface you may have to change the IP address of your management computer and then reconnect to the port1 interface using the 172.20.120.141 IP address.

9. Edit the *port1* interface and change *IP/Netmask* to 172.20.120.141/24.
10. Select OK.
11. Go to *Router > Static > Static Routes*.
12. Change the default route.

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Gateway</b>	172.20.120.2
<b>Device</b>	port1
<b>Distance</b>	10

13. Select OK.

## Configuring a NAT/Route mode active-passive cluster of two FortiGate-620B units - CLI

Use the following procedures to configure two FortiGate-620B units for NAT/Route HA operation using the FortiGate CLI. These procedures assume you are starting with two FortiGate-620B units with factory default settings.



Make sure your FortiGate interfaces are configured with static IP addresses. If any interface gets its address using DHCP you will not be able to configure HA.

### To configure the first FortiGate-620B unit (host name 620\_ha\_1)

1. Power on the FortiGate unit.
2. Connect a null modem cable to the communications port of the management computer and to the FortiGate Console port.
3. Start HyperTerminal (or any terminal emulation program), enter a name for the connection, and select OK.
4. Configure HyperTerminal to connect directly to the communications port on the computer to which you have connected the null modem cable and select OK.
5. Select the following port settings and select OK.

<b>Bits per second</b>	9600
<b>Data bits</b>	8

<b>Parity</b>	None
<b>Stop bits</b>	1
<b>Flow control</b>	None

6. Press Enter to connect to the FortiGate CLI.

The FortiGate unit CLI login prompt appears.

If the prompt does not appear, press Enter. If it still does not appear, power off your FortiGate unit and power it back on. If you are connected, at this stage you will see startup messages that will confirm you are connected. The login prompt will appear after the startup has completed.

7. Type `admin` and press Enter twice.
8. Change the host name for this FortiGate unit.

```
config system global
 set hostname 620_ha_1
end
```

9. Configure HA settings.

```
config system ha
 set mode a-p
 set group-name example1.com
 set password HA_pass_1
end
```

The FortiGate unit negotiates to establish an HA cluster. You may temporarily lose network connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 1300](#)). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12

- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

To confirm these MAC address changes, you can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (MAC) and the port1 permanent MAC address (Permanent\_HWaddr):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

#### 10. Display the HA configuration (optional).

```
get system ha
 group-id : 0
 group-name : example1.com
 mode : a-p
 password : *
 hbdev : "port3" 50 "port4" 50
 session-sync-dev :
 route-ttl : 10
 route-wait : 0
 route-hold : 10
 sync-config : enable
 encryption : disable
 authentication : disable
 hb-interval : 2
 hb-lost-threshold : 6
 helo-holddown : 20
 arps : 5
 arps-interval : 8
 session-pickup : disable
 link-failed-signal : disable
 uninterruptible-upgrade : enable
 ha-mgmt-status : disable
 ha-eth-type : 8890
 hc-eth-type : 8891
 l2ep-eth-type : 8893
 subsecond : disable
 vcluster2 : disable
 vcluster-id : 1
 override : disable
```



```
priority : 128
monitor :
pingserver-monitor-interface:
pingserver-failover-threshold: 0
pingserver-flip-timeout: 60
vdom : "root"
```

**11. Power off the FortiGate unit.**

**To configure the second FortiGate-620B unit (host name 620\_ha\_2)**

1. Power on the FortiGate unit.
2. Connect a null modem cable to the communications port of the management computer and to the FortiGate Console port.
3. Start HyperTerminal, enter a name for the connection, and select OK.
4. Configure HyperTerminal to connect directly to the communications port on the computer to which you have connected the null modem cable and select OK.
5. Select the following port settings and select OK.

<b>Bits per second</b>	9600
<b>Data bits</b>	8
<b>Parity</b>	None
<b>Stop bits</b>	1
<b>Flow control</b>	None

6. Press Enter to connect to the FortiGate CLI.  
The FortiGate unit CLI login prompt appears.
7. Type `admin` and press Enter twice.
8. Change the host name for this FortiGate unit.

```
config system global
 set hostname 620_ha_2
end
```

9. Configure HA settings.

```
config system ha
 set mode a-p
 set group-name example1.com
 set password HA_pass_1
end
```

The FortiGate unit negotiates to establish an HA cluster. You may temporarily lose network connectivity with the FortiGate unit as the HA cluster negotiates and because the FGCP changes the MAC address of the FortiGate unit interfaces.

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

## 10. Display the HA configuration (optional).

```
get system ha
 group-id : 0
 group-name : example1.com
 mode : a-p
 password : *
 hbdev : "port3" 50 "port4" 50
 session-sync-dev :
 route-ttl : 10
 route-wait : 0
 route-hold : 10
 sync-config : enable
 encryption : disable
 authentication : disable
 hb-interval : 2
 hb-lost-threshold : 6
 helo-holddown : 20
 arps : 5
 arps-interval : 8
 session-pickup : disable
 link-failed-signal : disable
 uninterruptible-upgrade : enable
 ha-mgmt-status : disable
 ha-eth-type : 8890
 hc-eth-type : 8891
 l2ep-eth-type : 8893
 subsecond : disable
 vcluster2 : disable
 vcluster-id : 1
 override : disable
 priority : 128
 monitor :
 pingserver-monitor-interface :
 pingserver-failover-threshold : 0
 pingserver-flip-timeout : 60
 vdom : "root"
```

## 11. Power off the FortiGate unit.

### To connect the cluster to the network

1. Connect the port1 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the Internet.
2. Connect the port2 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the internal network.
3. Connect the port3 interfaces of 620\_ha\_1 and 620\_ha\_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
4. Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.

## 5. Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

### To view cluster status

Use the following steps to view cluster status from the CLI.

1. Determine which cluster unit is the primary unit.
  - Use the null-modem cable and serial connection to re-connect to the CLI of one of the cluster units.
  - Enter the command `get system status`.
  - If the command output includes `Current HA mode: a-a, master`, the cluster units are operating as a cluster and you have connected to the primary unit. Continue with Step 2.
  - If the command output includes `Current HA mode: a-a, backup`, you have connected to a subordinate unit. Connect the null-modem cable to the other cluster unit, which should be the primary unit and continue with Step 2.



If the command output includes `Current HA mode: standalone`, the cluster unit is not operating in HA mode and you should review your HA configuration.

---

2. Enter the following command to confirm the HA configuration of the cluster:

```
get system ha status
 Model: 620
 Mode: a-a
 Group: 0
 Debug: 0
 ses_pickup: disable
 Master:128 620_ha_2 FG600B3908600825 0
 Slave :128 620_ha_1 FG600B3908600705 1
 number of vcluster: 1
 vcluster 1: work 169.254.0.1
 Master:0 FG600B3908600825
 Slave :1 FG600B3908600705
```

The command output shows both cluster units, their host names, their roles in the cluster, and their priorities. You can use this command to confirm that the cluster is operating normally. For example, if the command shows only one cluster unit then the other unit has left the cluster for some reason.

### To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 1212](#) to troubleshoot the cluster.

### To add basic configuration settings to the cluster

Use the following steps to add some basic settings to the cluster so that it can connect to the network.

1. Log into the primary unit CLI.
2. Add a password for the admin administrative account.

```
config system admin
 edit admin
 set password <password_str>
 end
```

3. Configure the port1 and port2 interfaces.

```
config system interface
 edit port1
 set ip 172.20.120.141/24
 next
 edit port2
 set ip 10.11.101.100/24
 end
```

4. Add a default route.

```
config router static
 edit 1
 set dst 0.0.0.0 0.0.0.0
 set gateway 172.20.120.2
 set device port1
 end
```

## Example: Transparent mode active-active HA configuration

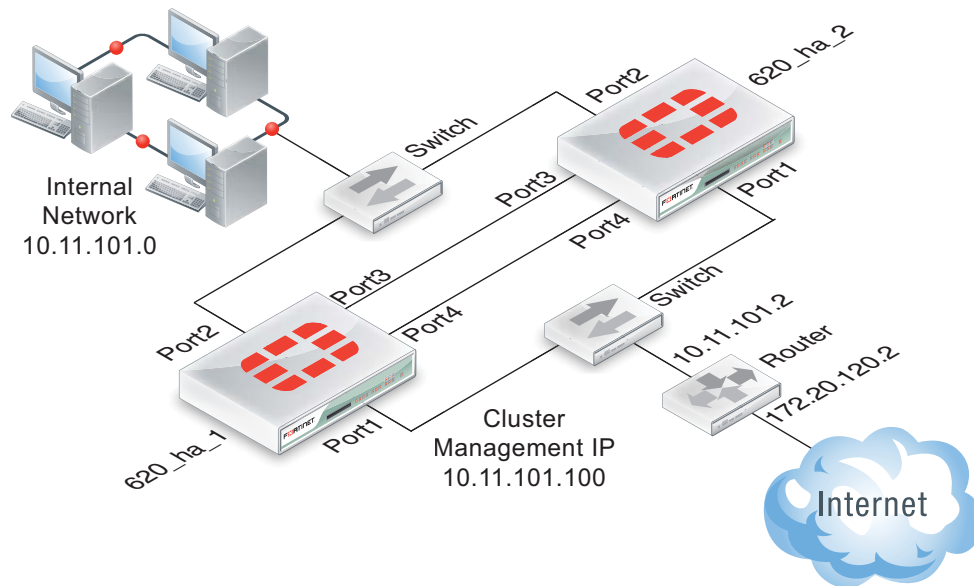
This section describes a simple HA network topology that includes an HA cluster of two FortiGate-620B units installed between an internal network and the Internet and running in Transparent mode.

- [Example Transparent mode HA network topology](#)
- [General configuration steps](#)

### Example Transparent mode HA network topology

Figure 193 shows a Transparent mode FortiGate-620B HA cluster consisting of two FortiGate-620B units (620\_ha\_1 and 620\_ha\_2) installed between the Internet and internal network. The topology includes a router that performs NAT between the internal network and the Internet. The cluster management IP address is 10.11.101.100.

**Figure 193:**Transparent mode HA network topology



Port3 and port4 are the default FortiGate-620B heartbeat interfaces. Because the cluster consists of two FortiGate units, you can make the connections between the heartbeat interfaces using crossover cables. You could also use switches and regular ethernet cables.

## General configuration steps

This section includes web-based manager and CLI procedures. These procedures assume that the FortiGate-620B units are running the same FortiOS firmware build and are set to the factory default configuration.

In this example, the configuration steps are identical to the NAT/Route mode configuration steps until the cluster is operating. When the cluster is operating, you can switch to Transparent mode and add basic configuration settings to cluster.

### General configuration steps

1. Configure the FortiGate units for HA operation.
  - Optionally change each unit's host name.
  - Configure HA.
2. Connect the cluster to the network.
3. Confirm that the cluster units are operating as a cluster.
4. Switch the cluster to Transparent mode and add basic configuration settings to the cluster.
  - Switch to Transparent mode, add the management IP address and a default route.
  - Add a password for the admin administrative account.
  - View cluster status from the web-based manager or CLI.

## Configuring a Transparent mode active-active cluster of two FortiGate-620B units - web-based manager

Use the following procedures to configure the FortiGate-620B units for HA operation using the FortiGate web-based manager. These procedures assume you are starting with two FortiGate-620B units with factory default settings.



Waiting until you have established the cluster to switch to Transparent mode means fewer configuration steps because you can switch the mode of the cluster in one step.

### To configure the first FortiGate-620B unit (host name 620\_ha\_1)

1. Power on the first FortiGate unit.
2. Set the IP address of a management computer with an Ethernet connection to the static IP address 192.168.1.2 and a netmask of 255.255.255.0.
3. On a management computer, start a web browser and browse to the address `https://192.168.1.99` (remember to include the “s” in `https://`).  
The FortiGate login is displayed.
4. Type `admin` in the *Name* field and select *Login*.  
The FortiGate dashboard is displayed.
5. On the *System Information* dashboard widget, beside *Host Name* select *Change*.
6. Enter a new Host Name for this FortiGate unit.

<b>New Name</b>	620_ha_1
-----------------	----------

7. Select OK.
8. Go to *System > Config > HA* and change the following settings:

<b>Mode</b>	Active-Active
-------------	---------------

<b>Group Name</b>	example2.com
-------------------	--------------

<b>Password</b>	HA_pass_2
-----------------	-----------



This is the minimum recommended configuration for an active-active HA cluster. You can configure other HA options at this point, but if you wait until the cluster is operating you will only have to configure these options once for the cluster instead of separately for each cluster unit.

9. Select OK.

The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 1300](#)). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02

- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

To confirm these MAC address changes, you can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (MAC) and the port1 permanent MAC address (Permanent\_HWaddr):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

**10.** Power off the first FortiGate unit.

**To configure the second FortiGate-620B unit (host name 620\_ha\_2)**

1. Power on second FortiGate unit.
2. On a management computer, start Internet Explorer and browse to the address `https://192.168.1.99` (remember to include the “s” in `https://`).  
The FortiGate login is displayed.
3. Type `admin` in the *Name* field and select *Login*.  
The FortiGate dashboard is displayed.
4. On the *System Information* dashboard widget, beside *Host Name* select *Change*.

5. Enter a new Host Name for this FortiGate unit.

---

<b>New Name</b>	620_ha_2
-----------------	----------

---

6. Select OK.
7. Go to *System > Config > HA* and change the following settings:

---

<b>Mode</b>	Active-Active
<b>Group Name</b>	example2.com
<b>Password</b>	HA_pass_2

---

8. Select OK.

The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and because the FGCP changes the MAC address of the FortiGate unit interfaces.

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

9. Power off the second FortiGate unit.

#### **To connect the cluster to the network**

1. Connect the port1 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the Internet.
2. Connect the port2 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the internal network.
3. Connect the port3 interfaces of 620\_ha\_1 and 620\_ha\_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
4. Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
5. Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

#### **To switch the cluster to Transparent mode**

Switching from NAT/Route to Transparent mode involves adding the Transparent mode management IP address and default route.



Since configuration changes are synchronized to all cluster units, switching the cluster to operate in Transparent mode once the cluster is operating is the same as switching an individual FortiGate unit to Transparent mode. You could have performed the following configuration steps separately on each FortiGate unit before you connected them to form a cluster.

1. Start a web browser and browse to the address `https://192.168.1.99` (remember to include the “s” in `https://`).  
The FortiGate Login is displayed.
2. Type admin in the Name field and select Login.



3. Under System Information, beside *Operation Mode* select *Change*.
4. Set Operation Mode to Transparent.
5. Configure basic Transparent mode settings.

---

<b>Operation Mode</b>	Transparent
<b>Management IP/Mask</b>	10.11.101.100/24
<b>Default Gateway</b>	10.11.101.2

---

6. Select Apply.

The cluster switches to operating in Transparent mode. The virtual MAC addresses assigned to the cluster interfaces do not change.

#### To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.



Once the cluster is operating, because configuration changes are synchronized to all cluster units, configuring the cluster is the same as configuring an individual FortiGate unit. You could have performed the following configuration steps separately on each FortiGate unit before you connected them to form a cluster.

1. Start Internet Explorer and browse to the address <https://10.11.101.100> (remember to include the “s” in https://).

The FortiGate Login is displayed.

2. Type admin in the Name field and select Login.

The FortiGate dashboard is displayed.

The System Information dashboard widget shows the *Cluster Name* (example2.com) and the host names and serial numbers of the *Cluster Members*. The Unit Operation widget shows multiple cluster units.

3. Go to *System > Config > HA* to view the cluster members list.

The list shows both cluster units, their host names, their roles in the cluster, and their priorities. You can use this list to confirm that the cluster is operating normally. For example, if the list shows only one cluster unit then the other unit has left the cluster for some reason.

#### To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units, the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 1212](#) to troubleshoot the cluster.

#### To add basic configuration settings to the cluster

Use the following steps to configure the cluster. Note that the following are example configuration steps only and do not represent all of the steps required to configure the cluster for a given network.

1. Log into the cluster web-based manager.
2. Go to *System > Admin > Administrators*.
3. For *admin*, select the *Change Password* icon
4. Enter and confirm a new password.

5. Select OK.



You added a default gateway when you switched to Transparent mode so you don't need to add a default route as part of the basic configuration of the cluster at this point.

## Configuring a Transparent mode active-active cluster of two FortiGate-620B units - CLI

Use the following procedures to configure the FortiGate-620B units for Transparent mode HA operation using the FortiGate CLI.

### To configure each FortiGate unit for HA operation

1. Power on the FortiGate unit.
2. Connect a null modem cable to the communications port of the management computer and to the FortiGate Console port.
3. Start HyperTerminal, enter a name for the connection, and select OK.
4. Configure HyperTerminal to connect directly to the communications port on the computer to which you have connected the null modem cable and select OK.
5. Select the following port settings and select OK.

<b>Bits per second</b>	9600
------------------------	------

<b>Data bits</b>	8
------------------	---

<b>Parity</b>	None
---------------	------

<b>Stop bits</b>	1
------------------	---

<b>Flow control</b>	None
---------------------	------

6. Press Enter to connect to the FortiGate CLI.

The FortiGate unit CLI login prompt appears. If the prompt does not appear, press Enter. If it still does not appear, power off your FortiGate unit and power it back on. If you are connected, at this stage you will see startup messages that will confirm you are connected. The login prompt will appear after the startup has completed.

7. Type `admin` and press Enter twice.
8. Change the host name for this FortiGate unit. For example:

```
config system global
 set hostname 620_ha_1
end
```

## 9. Configure HA settings.

```
config system ha
 set mode a-a
 set group-name example2.com
 set password HA_pass_2
end
```



This is the minimum recommended configuration for an active-active HA cluster. You can also configure other HA options, but if you wait until after the cluster is operating you will only have to configure these options once for the cluster instead of separately for each cluster unit.

The FortiGate unit negotiates to establish an HA cluster. You may temporarily lose network connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 1300](#)). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

To confirm these MAC address changes, you can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1

interface virtual MAC address (MAC) and the port1 permanent MAC address (Permanent\_HWaddr):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

**10. Display the HA configuration (optional).**

```
get system ha
 group-id : 0
 group-name : example2.com
 mode : a-a
 password : *
 hbdev : "port3" 50 "port4" 50
 session-sync-dev :
 route-ttl : 10
 route-wait : 0
 route-hold : 10
 sync-config : enable
 encryption : disable
 authentication : disable
 hb-interval : 2
 hb-lost-threshold : 6
 helo-holddown : 20
 arps : 5
 arps-interval : 8
 session-pickup : disable
 link-failed-signal : disable
 uninterruptible-upgrade: enable
 ha-mgmt-status : disable
 ha-eth-type : 8890
 hc-eth-type : 8891
 l2ep-eth-type : 8893
 subsecond : disable
 vcluster2 : disable
 vcluster-id : 1
 override : disable
 priority : 128
 monitor :
 pingserver-monitor-interface:
 pingserver-failover-threshold: 0
 pingserver-flip-timeout: 60
 vdom : "root"
```

**11. Power off the FortiGate unit.**

## To configure the second FortiGate-620B unit (host name 620\_ha\_2)

1. Power on the FortiGate unit.
2. Connect a null modem cable to the communications port of the management computer and to the FortiGate Console port.
3. Start HyperTerminal, enter a name for the connection, and select OK.
4. Configure HyperTerminal to connect directly to the communications port on the computer to which you have connected the null modem cable and select OK.
5. Select the following port settings and select OK.

<b>Bits per second</b>	9600
<b>Data bits</b>	8
<b>Parity</b>	None
<b>Stop bits</b>	1
<b>Flow control</b>	None

6. Press Enter to connect to the FortiGate CLI.

The FortiGate unit CLI login prompt appears. If the prompt does not appear, press Enter. If it still does not appear, power off your FortiGate unit and power it back on. If you are connected, at this stage you will see startup messages that will confirm you are connected. The login prompt will appear after the startup has completed.

7. Type `admin` and press Enter twice.
8. Change the host name for this FortiGate unit.

```
config system global
 set hostname 620_ha_2
end
```

9. Configure HA settings.

```
config system ha
 set mode a-a
 set group-name example2.com
 set password HA_pass_2
end
```

The FortiGate unit negotiates to establish an HA cluster. You may temporarily lose network connectivity with the FortiGate unit as the HA cluster negotiates and because the FGCP changes the MAC address of the FortiGate unit interfaces.

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

## 10. Display the HA configuration (optional).

```
get system ha
 group-id : 0
 group-name : example2.com
 mode : a-a
 password : *
 hbdev : "port3" 50 "port4" 50
 session-sync-dev :
 route-ttl : 10
 route-wait : 0
 route-hold : 10
 sync-config : enable
 encryption : disable
 authentication : disable
 hb-interval : 2
 hb-lost-threshold : 6
 helo-holddown : 20
 arps : 5
 arps-interval : 8
 session-pickup : disable
 link-failed-signal : disable
 uninterruptible-upgrade : enable
 ha-mgmt-status : disable
 ha-eth-type : 8890
 hc-eth-type : 8891
 l2ep-eth-type : 8893
 subsecond : disable
 vcluster2 : disable
 vcluster-id : 1
 override : disable
 priority : 128
 monitor :
 pingserver-monitor-interface :
 pingserver-failover-threshold : 0
 pingserver-flip-timeout : 60
 vdom : "root"
```

## 11. Power off the FortiGate unit.

### To connect the cluster to the network

1. Connect the port1 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the Internet.
2. Connect the port2 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the internal network.
3. Connect the port3 interfaces of 620\_ha\_1 and 620\_ha\_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
4. Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.

## 5. Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

### To connect to the cluster CLI and switch the cluster to Transparent mode

#### 1. Determine which cluster unit is the primary unit.

- Use the null-modem cable and serial connection to re-connect to the CLI of one of the cluster units.
- Enter the command `get system status`.
- If the command output includes `Current HA mode: a-a, master`, the cluster units are operating as a cluster and you have connected to the primary unit. Continue with Step 2.
- If the command output includes `Current HA mode: a-a, backup`, you have connected to a subordinate unit. Connect to the other cluster unit, which should be the primary unit and continue with Step 2.



If the command output includes `Current HA mode: standalone`, the cluster unit is not operating in HA mode. See [“Troubleshooting the initial cluster configuration” on page 1213](#).

---

#### 2. Change to transparent mode.

```
config system settings
 set opmode transparent
 set manageip 192.168.20.3/24
 set gateway 192.168.20.1
end
```

The cluster switches to Transparent Mode, and your administration session is disconnected.

You can now connect to the cluster CLI using SSH to connect to the cluster internal interface using the management IP address (192.168.20.3).

### To view cluster status

Use the following steps to view cluster status from the CLI.

#### 1. Determine which cluster unit is the primary unit.

- Use the null-modem cable and serial connection to re-connect to the CLI of one of the cluster units.
- Enter the command `get system status`.
- If the command output includes `Current HA mode: a-a, master`, the cluster units are operating as a cluster and you have connected to the primary unit. Continue with Step 2.

- If the command output includes `Current HA mode: a-a, backup`, you have connected to a subordinate unit. Connect the null-modem cable to the other cluster unit, which should be the primary unit and continue with Step 2.



If the command output includes `Current HA mode: standalone`, the cluster unit is not operating in HA mode and you should review your HA configuration.

2. Enter the following command to confirm the HA configuration of the cluster:

```
get system ha status
 Model: 620
 Mode: a-p
 Group: 0
 Debug: 0
 ses_pickup: disable
 Master:128 620_ha_2 FG600B3908600825 0
 Slave :128 620_ha_1 FG600B3908600705 1
 number of vcluster: 1
 vcluster 1: work 169.254.0.1
 Master:0 FG600B3908600825
 Slave :1 FG600B3908600705
```

The command output shows both cluster units, their host names, their roles in the cluster, and their priorities. You can use this command to confirm that the cluster is operating normally. For example, if the command shows only one cluster unit then the other unit has left the cluster for some reason.

### To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 1212](#) to troubleshoot the cluster.

### To add a password for the admin administrative account

1. Add a password for the admin administrative account.

```
config system admin
 edit admin
 set password <psswr>
 end
```

## Example: advanced Transparent mode active-active HA configuration

This section describes a more complex HA network topology that includes an HA cluster of three FortiGate-5002FA2 units running in Transparent mode and installed between an internal network and an engineering network.

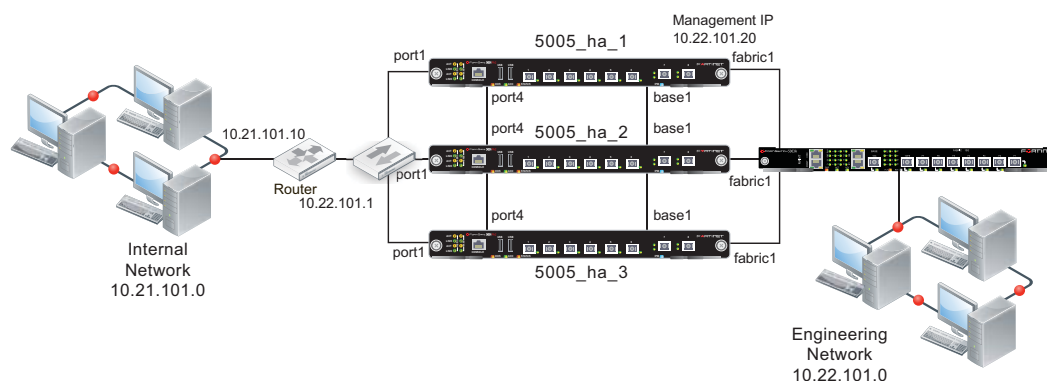
- [Example Transparent mode HA network topology](#)
- [General configuration steps](#)



## Example Transparent mode HA network topology

Figure 194 shows a Transparent mode FortiGate-5005FA2 HA cluster consisting of three FortiGate-5005FA2 units (5005\_ha\_1, 5005\_ha\_2, and 5005\_ha\_3) installed in a FortiGate-5000 series chassis with one FortiSwitch-5003A board. The cluster applies virus scanning to traffic passing between an engineering network and an internal network. The topology includes a router that performs NAT between the internal network and the engineering network. The cluster is connected to the engineering network with an management IP address of 10.22.101.20. This IP address is on the engineering network subnet.

Figure 194: Transparent mode HA network topology



By default fabric1 and fabric2 are the FortiGate-5005FA2 heartbeat interfaces. This example changes the heartbeat configuration to use the base1 and port4 interfaces for the heartbeat. The base1 connection is handled using the base backplane channel switched by the FortiSwitch-5003A board. The port4 connection is handled by connecting the port4 interfaces together using a switch.

The cluster connects to the engineering network using fabric1. The FortiSwitch-5003A board provides switching for the fabric1 interfaces and the fabric1 connection to the engineering network.

## Configuring a Transparent mode active-active cluster of three FortiGate-5005FA2 units - web-based manager

These procedures assume you are starting with three FortiGate-5005FA2 units with factory default settings but not installed in chassis slots and a FortiSwitch-5003A board installed in chassis slot 1. The chassis is powered on. This configuration works for a FortiGate-5050 chassis or for a FortiGate-5140 chassis. No configuration changes to the FortiSwitch-5003A board are required.

### To configure the FortiGate-5005FA2 units

1. Power on the first FortiGate unit by inserting it into chassis slot 5.
2. Connect port1 to the network and log into the web-based manager.
3. On the *System Information* dashboard widget, beside *Host Name* select *Change*.
4. Enter a new Host Name for this FortiGate unit.

<b>New Name</b>	5005_ha_1
-----------------	-----------

5. Select OK.
6. Go to *System > Network > Interfaces* and select *Show backplane interfaces*.

7. Make sure the administrative status and link status is for base1 and fabric1.

You can edit the interface to set the administrative status to up. The link status will be up if the administrative status is up and the FortiGate-5005FA2 board can connect to the FortiSwitch-5003A board.

8. Go to *System > Config > HA* and change the following settings:

<b>Mode</b>	Active-Active	
<b>Group Name</b>	example3.com	
<b>Password</b>	HA_pass_3	
<b>Heartbeat Interface</b>		
	<b>Enable</b>	<b>Priority</b>
<b>base1</b>	Select	50
<b>fabric1</b>	Clear check box	0
<b>fabric2</b>	Clear check box	0
<b>port4</b>	Select	50

9. Select OK.

The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 1300](#)). The MAC addresses of the FortiGate-5005FA2 interfaces change to the following virtual MAC addresses:

- base1 interface virtual MAC: 00-09-0f-09-00-00
- base2 interface virtual MAC: 00-09-0f-09-00-01
- fabric1 interface virtual MAC: 00-09-0f-09-00-02
- fabric2 interface virtual MAC: 00-09-0f-09-00-03
- port1 interface virtual MAC: 00-09-0f-09-00-04
- port2 interface virtual MAC: 00-09-0f-09-00-05
- port3 interface virtual MAC: 00-09-0f-09-00-06
- port4 interface virtual MAC: 00-09-0f-09-00-07
- port5 interface virtual MAC: 00-09-0f-09-00-08
- port6 interface virtual MAC: 00-09-0f-09-00-09
- port7 interface virtual MAC: 00-09-0f-09-00-0a
- port8 interface virtual MAC: 00-09-0f-09-00-0b

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use

the following command to view the port1 interface virtual MAC address (Current\_HWaddr) and the port1 permanent MAC address (Permanent\_HWaddr):

```
get hardware nic port1
.
.
.
Current_HWaddr 00:09:0f:09:00:04
Permanent_HWaddr 00:09:0f:71:0a:dc
.
.
.
```

10. Power off the first FortiGate unit.

11. Repeat these steps for the second and third FortiGate units, with the following difference.

Set the second FortiGate unit host name to:

<b>New Name</b>	5005_ha_2
-----------------	-----------

Set the third FortiGate unit host name to:

<b>New Name</b>	5005_ha_3
-----------------	-----------

As you insert and configure each FortiGate unit, they will negotiate and join the cluster using the base1 interface for HA heartbeat communication.

### To connect the cluster to the network

1. Connect the port1 interfaces of the cluster to a switch that can connect to the router and the internal network.
2. Connect the port4 interfaces of the cluster units together using a switch.  
These interfaces become the backup heartbeat interface.
3. Connect one of the FortiSwitch-5003A front panel fabric interfaces (for example, F3) to the engineering network.

### To switch the cluster to operate in Transparent mode

Switching from NAT/Route to Transparent mode also involves adding the Transparent mode management IP address and default route.

1. Log into the web-based manager.
2. Under System Information, beside *Operation Mode* select *Change*.
3. Set *Operation Mode* to *Transparent*.
4. Configure basic Transparent mode settings.

<b>Operation Mode</b>	Transparent
<b>Management IP/Mask</b>	10.22.101.20/24
<b>Default Gateway</b>	10.22.101.1

5. Select Apply.

The cluster switches to operating in Transparent mode. The virtual MAC addresses assigned to the cluster interfaces do not change. You must login again using the new TP address.

### To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.

1. View the system dashboard.

The System Information dashboard widget shows the *Cluster Name* (example3.com) and the host names and serial numbers of the *Cluster Members*. The Unit Operation widget shows multiple cluster units.

2. Go to *System > Config > HA* to view the cluster members list.

The list shows three cluster units, their host names, their roles in the cluster, and their priorities. You can use this list to confirm that the cluster is operating normally.

### To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units, the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 1212](#) to troubleshoot the cluster.

### To add basic configuration settings to the cluster

Use the following steps to configure the cluster. The following are example configuration steps only and do not represent all of the steps required to configure the cluster for a given network.

1. Log into the cluster web-based manager.
2. Go to *System > Admin > Administrators*.
3. For *admin*, select the *Change Password* icon
4. Enter and confirm a new password.
5. Select OK.

The default route was changed when you switched to Transparent mode.

## Configuring a Transparent mode active-active cluster of three FortiGate-5005FA2 units - CLI

Use the following procedures to configure the three FortiGate-5005FA2 units for Transparent mode HA operation using the FortiGate CLI.

### To configure the FortiGate-5005FA2 units

1. Power on the first FortiGate unit by inserting it into chassis slot 5.
2. Connect port1 to the network and log into the CLI.

You can also use a console connection.

3. Change the host name for this FortiGate unit. For example:

```
config system global
 set hostname 5005_ha_1
end
```

4. Enable showing backplane interfaces.

```
config system global
 set show-backplane-intf enable
end
```

5. Make sure the administrative status and link status is up for base1 and fabric1.

Enter `get system interface` to view the status of these interfaces.

You can use the following commands to set the administrative status to up for these interfaces.

```
config system interface
 edit base1
 set status up
 next
 edit fabricq
 set status up
end
```

6. Configure HA settings.

```
config system ha
 set mode a-a
 set group-name example3.com
 set password HA_pass_3
 set hbdev base1 50 port4 50
end
```



This is the minimum recommended configuration for an active-active HA cluster. You can also configure other HA options, but if you wait until after the cluster is operating you will only have to configure these options once for the cluster instead of separately for each cluster unit.

---

The FortiGate unit negotiates to establish an HA cluster. You may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 1300](#)). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- base1 interface virtual MAC: 00-09-0f-09-00-00
- base2 interface virtual MAC: 00-09-0f-09-00-01
- fabric1 interface virtual MAC: 00-09-0f-09-00-02
- fabric2 interface virtual MAC: 00-09-0f-09-00-03
- port1 interface virtual MAC: 00-09-0f-09-00-04
- port2 interface virtual MAC: 00-09-0f-09-00-05
- port3 interface virtual MAC: 00-09-0f-09-00-06
- port4 interface virtual MAC: 00-09-0f-09-00-07
- port5 interface virtual MAC: 00-09-0f-09-00-08
- port6 interface virtual MAC: 00-09-0f-09-00-09
- port7 interface virtual MAC: 00-09-0f-09-00-0a
- port8 interface virtual MAC: 00-09-0f-09-00-0b

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use

the following command to view the port1 interface virtual MAC address (Current\_HWaddr) and the port1 permanent MAC address (Permanent\_HWaddr):

```
get hardware nic port1
.
.
.
Current_HWaddr 00:09:0f:09:00:04
Permanent_HWaddr 00:09:0f:71:0a:dc
.
.
.
```

#### 7. Display the HA configuration (optional).

```
get system ha
group-id : 0
group-name : example3.com
mode : a-a
password : *
hbdev : "base1" 50 "port4" 50
session-sync-dev :
route-ttl : 10
route-wait : 0
route-hold : 10
sync-config : enable
encryption : disable
authentication : disable
hb-interval : 2
hb-lost-threshold : 20
helo-holddown : 20
arps : 5
arps-interval : 8
session-pickup : disable
link-failed-signal : disable
uninterruptible-upgrade: enable
ha-mgmt-status : disable
ha-eth-type : 8890
hc-eth-type : 8891
l2ep-eth-type : 8893
subsecond : disable
vcluster2 : disable
vcluster-id : 1
override : disable
priority : 128
schedule : round-robin
monitor :
pingserver-monitor-interface:
pingserver-failover-threshold: 0
pingserver-flip-timeout: 60
vdom : "root"
load-balance-all : disable
```

8. Repeat these steps for the second and third FortiGate units.

Set the second FortiGate unit host name to:

```
config system global
 set hostname 5005_ha_2
end
```

Set the third FortiGate unit host name to:

```
config system global
 set hostname 5005_ha_3
end
```

As you insert and configure each FortiGate unit they will negotiate and join the cluster using the base1 interface for HA heartbeat communication.

### To connect the cluster to the network

1. Connect the port1 interfaces of the cluster to a switch that can connect to the router and the internal network.
2. Connect the port4 interfaces of the cluster units together using a switch.  
These interfaces become the backup heartbeat interface.
3. Connect one of the FortiSwitch-5003A front panel fabric interfaces (for example, F3) to the engineering network.

### To switch the cluster to Transparent mode

1. Log into the cluster CLI.
2. Change to Transparent mode.

```
config system settings
 set opmode transparent
 set manageip 10.22.101.20/24
 set gateway 10.22.101.1
end
```

The cluster switches to Transparent Mode.

You can now connect to the cluster CLI using SSH to connect to the cluster internal interface using the management IP address (10.22.101.20).

### To view cluster status

Use the following steps to view cluster status from the CLI.

1. Log into the CLI.
2. To verify the HA status of the cluster unit that you logged into, enter the CLI command `get system status`. Look for the following information in the command output.

Current HA mode: a-a, master	The cluster units are operating as a cluster and you have connected to the primary unit.
Current HA mode: a-a, backup	The cluster units are operating as a cluster and you have connected to a subordinate unit.
Current HA mode: standalone	The cluster unit is not operating in HA mode

3. Enter the following command to confirm the HA configuration of the cluster:

```
get system ha status
 Model: 5005
 Mode: a-a
 Group: 0
 Debug: 0
 ses_pickup: disable
 load_balance: disable
 schedule: round robin
 Master:128 5005_ha_1 FG5A253E07600124 0
 Slave :128 5005_ha_2 FG5A253E06500088 1
 Slave :128 5005_ha_3 FG5A253E06500099 2
 number of vcluster: 1
 vcluster 1: work 169.254.0.1
 Master:0 FG5A253E07600124
 Slave :1 FG5A253E06500088
 Slave :2 FG5A253E06500099
```

The command output shows both cluster units, their host names, their roles in the cluster, and their priorities. You can use this command to confirm that the cluster is operating normally. For example, if the command shows only one cluster unit then the other unit has left the cluster for some reason.

### To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 1212](#) to troubleshoot the cluster.

### To add a password for the admin administrative account

1. Add a password for the admin administrative account.

```
config system admin
 edit admin
 set password <psswr>
end
```

## Example: converting a standalone FortiGate unit to a cluster

You can convert an already configured and installed FortiGate unit into a cluster by configuring this FortiGate unit to be a primary unit and then adding subordinate units.

General configuration steps:

- Configure the original FortiGate unit for HA operation.
- Set the HA Device Priority of the original FortiGate unit to 255 to make sure that this FortiGate unit becomes the primary unit after cluster negotiation and synchronization.
- Back up the configuration of the original FortiGate unit.
- Configure one or more new FortiGate units with the same HA configuration as the original FortiGate unit with one exception. Keep the Unit Priority at the default setting, which is 128.
- Connect the FortiGate units to form a cluster and connect the cluster to your network.

When you power on all of the FortiGate units in the cluster, the original FortiGate unit becomes the primary unit. Its configuration is synchronized to all of the subordinate units. The entire



cluster now operates with the original FortiGate unit configuration. No further configuration changes are required.

The new FortiGate units must:

- Have the same hardware configuration as the original FortiGate unit. Including the same hard disk configuration and the same AMC cards installed in the same slots.
- Have the same firmware build as the original FortiGate unit.
- Be set to the same operating mode (NAT or Transparent) as the original FortiGate unit.
- Be operating in single VDOM mode.

In addition to one or more new FortiGate units, you need sufficient switches to connect all of the FortiGate interfaces in the cluster. Generally you will need one switch per interface, as it will have to connect that same interface on all cluster units. That is, all port1 interfaces use the port1 switch, port2 interfaces use the port2 switch, and so on. Intelligent switches that can be partitioned can reduce your switch requirements.

Converting a FortiGate unit to a primary unit and adding in the subordinate unit or units results in a brief service interruption as you disconnect and reconnect FortiGate interfaces and as the cluster negotiates. Therefore, conversion should only be done during off peak hours.



Make sure your FortiGate interfaces are configured with static IP addresses. If any interface gets its address using DHCP you will not be able to configure HA.

#### To configure the original FortiGate unit for HA operation

1. Connect to the FortiGate unit web-based manager.
2. Go to *System > Config > HA*.
3. Configure the FortiGate unit for HA operation.

<b>Mode</b>	Active-Active
<b>Device Priority</b>	255
<b>Group Name</b>	example4.com
<b>Password</b>	HA_pass_4

You can make other HA configuration changes after the cluster is operating.

4. Select OK.

The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and because the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 1300](#)).

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

5. Configure the new FortiGate units with the same HA configuration as the original FortiGate unit. The one exception is to keep the device priorities of the new FortiGate units at 128 to ensure the original FortiGate unit will become the primary unit in the new cluster.

<b>Mode</b>	Active-Active
<b>Device Priority</b>	128
<b>Group Name</b>	example4.com
<b>Password</b>	HA_pass_4

6. Configure the other FortiGate units to the same operation mode as the original FortiGate unit.  
There is no need to make any other configuration changes (including network configuration changes) to the other FortiGate units.
7. Optionally power off all of the cluster units.  
If you don't power off all of the units they may not negotiate to form a cluster when they are connected together.
8. Connect the cluster to your network.  
For example, for a configuration similar to the FortiGate-620B cluster configuration described in this chapter, see [“To connect the cluster to the network” on page 1156](#).
9. Power on all of the cluster units.  
As the units start they change their MAC addresses and then negotiate to choose the primary unit and the subordinate units. This negotiation occurs with no user intervention and normally takes less than a minute.  
The original the FortiGate unit becomes the primary unit because the device priority of the original FortiGate unit is higher than the device priority of the other FortiGate units. The configuration of the original FortiGate unit is synchronized to all the cluster units. As a result, the cluster is quickly up and running and configured for your network. No further configuration changes are required.

## Example: adding a new unit to an operating cluster

This procedure describes how to add a new FortiGate unit to a functioning cluster. Adding a new unit to a cluster does not interrupt the operation of the cluster unless you have to change how the cluster is connected to the network to accommodate the new cluster unit.

You can use this procedure to add as many units as required to the cluster.

### To add a new unit to a functioning cluster

1. Install the same firmware build on the new cluster unit as is running on the cluster.
2. Configure the new cluster unit for HA operation with the same HA configuration as the other units in the cluster.
3. If the cluster is running in Transparent mode, change the operating mode of the new cluster unit to Transparent mode.
4. Connect the new cluster unit to the cluster.
5. For example, for a configuration similar to the FortiGate-620B cluster configuration described in this chapter, see [“To connect the cluster to the network” on page 1156](#).

**6. Power on the new cluster unit.**

When the unit starts it negotiates to join the cluster. After it joins the cluster, the cluster synchronizes the new unit configuration with the configuration of the primary unit.

You can add a new unit to a functioning cluster at any time. The new cluster unit must:

- Have the same hardware configuration as the cluster units. Including the same hard disk configuration and the same AMC cards installed in the same slots.
- Have the same firmware build as the cluster.
- Be set to the same operating mode (NAT or Transparent) as the cluster.
- Be operating in single VDOM mode.

## Example: replacing a failed cluster unit

This procedure describes how to remove a failed cluster unit from a cluster and add a new one to replace it. You can also use this procedure to remove a failed unit from a cluster, repair it and add it back to the cluster. Replacing a failed does not interrupt the operation of the cluster unless you have to change how the cluster is connected to the network to accommodate the replacement unit.

You can use this procedure to replace more than one cluster unit.

### To replace a failed cluster unit

**1. Disconnect the failed unit from the cluster and the network.**

If you maintain other connections between the network and the still functioning cluster unit or units and between remaining cluster units network traffic will continue to be processed.

**2. Repair the failed cluster unit, or obtain a replacement unit with the exact same hardware configuration as the failed cluster unit.**

**3. Install the same firmware build on the repaired or replacement unit as is running on the cluster.**

**4. Configure the repaired or replacement unit for HA operation with the same HA configuration as the cluster.**

**5. If the cluster is running in Transparent mode, change the operating mode of the repaired or replacement cluster unit to Transparent mode.**

**6. Connect the repaired or replacement cluster unit to the cluster.**

For example, for a configuration similar to the FortiGate-620B cluster configuration described in this chapter, see [“To connect the cluster to the network” on page 1156.](#)

**7. Power on the repaired or replacement cluster unit.**

When the unit starts it negotiates to join the cluster. After it joins the cluster, the cluster synchronizes the repaired or replacement unit configuration with the configuration of the primary unit.

You can add a repaired or replacement unit to a functioning cluster at any time. The repaired or replacement cluster unit must:

- Have the same hardware configuration as the cluster units. Including the same hard disk configuration and the same AMC cards installed in the same slots.
- Have the same firmware build as the cluster.
- Be set to the same operating mode (NAT or Transparent) as the cluster.
- Be operating in single VDOM mode.


## Example: HA and 802.3ad aggregated interfaces

On FortiGate models that support it you can use 802.3ad link aggregation to combine two or more interfaces into a single aggregated interface. 802.3ad Link Aggregation and its management protocol, Link Aggregation Control Protocol (LACP) are a method for combining multiple physical links into a single logical link. This increases both potential throughput and network resiliency. Using LACP, traffic is distributed among the physical interfaces in the link, potentially resulting in increased performance.

This example describes how to configure an HA cluster consisting of two FortiGate-620B units with two aggregated 1000 Mb connections to the Internet using port1 and port2 and two aggregated 1000 Mb connections to the internal network using port3 and port4. The aggregated interfaces are also configured as HA monitored interfaces.

Each of the aggregate links connects to a different switch. Each switch is configured for link aggregation (2x1000Mb).

**Figure 195:**Example cluster with aggregate interfaces



### HA interface monitoring, link failover, and 802.3ad aggregation

When monitoring the aggregated interface, HA interface monitoring treats the aggregated link as a single interface and does not monitor the individual physical interfaces in the link. HA interface monitoring registers the link to have failed only if all the physical interfaces in the link have failed. If only some of the physical interfaces in the link fail or become disconnected, HA considers the link to be operating normally.

### HA MAC addresses and 802.3ad aggregation

If a configuration uses the Link Aggregate Control Protocol (LACP) (either passive or active), LACP is negotiated over all of the interfaces in any link. For a standalone FortiGate unit, the FortiGate LACP implementation uses the MAC address of the first interface in the link to uniquely identify that link. For example, a link consisting of port1 and port2 interfaces would have the MAC address of port1.

In an HA cluster, HA changes the MAC addresses of the cluster interfaces to virtual MAC addresses. An aggregate interface in a cluster acquires the virtual MAC address that would have been acquired by the first interface in the aggregate.

## Link aggregation, HA failover performance, and HA mode

To operate an active-active or active-passive cluster with aggregated interfaces and for best performance of a cluster with aggregated interfaces, the switches used to connect the cluster unit aggregated interfaces together should support configuring multiple Link Aggregation (LAG) groups.

For example, the cluster shown in [Figure 195](#) should be configured into two LAG groups on the external switch: one for the port1 and port2 aggregated interface of 620\_ha\_1 and a second one for the port1 and port2 aggregate interface of 620\_ha\_2. You should also be able to do the same on the internal switch for the port3 and port4 aggregated interfaces of each cluster unit.

As a result, the subordinate unit aggregated interfaces would participate in LACP negotiation while the cluster is operating. In an active-active mode cluster, packets could be redirected to the subordinate unit interfaces. As well, in active-active or active-passive mode, after a failover the subordinate unit can become a primary unit without having to perform LACP negotiation before it can process traffic. Performing LACP negotiation causes a minor failover delay.

However if you cannot configure multiple LAG groups on the switches, due to the primary and subordinate unit interfaces having the same MAC address, the switch will put all of the interfaces into the same LAG group which would disrupt the functioning of the cluster. To prevent this from happening, you must change the FortiGate aggregated interface configuration to prevent subordinate units from participating in LACP negotiation.

For example, use the following command to prevent subordinate units from participating in LACP negotiation with an aggregate interface named Port1\_Port2:

```
config system interface
 edit Port1_Port2
 set lacp-ha-slave disable
 end
```

As a result of this setting, subordinate unit aggregated interfaces cannot accept packets. This means that you cannot operate the cluster in active-active mode because in active-active mode the subordinate units must be able to receive and process packets. Also, failover may take longer because after a failover the subordinate unit has to perform LACP negotiation before being able to process network traffic.

Also, it may also be necessary to configure the switch to use Passive or even Static mode for LACP to prevent the switch from sending packets to the subordinate unit interfaces, which won't be able to process them.

Finally, in some cases depending on the LACP configuration of the switches, you may experience delayed failover if the FortiGate LACP configuration is not compatible with the switch LACP configuration. For example, in some cases setting the FortiGate LACP mode to static reduces the failover delay because the FortiGate unit does not perform LACP negotiation. However there is a potential problem with this configuration because static LACP does not send periodic LAC Protocol Data Unit (LACPDU) packets to test the connections. So a non-physical failure (for example, if a device is not responding because its too busy) may not be detected and packets could be lost or delayed.

## General configuration steps

The section includes web-based manager and CLI procedures. These procedures assume that the FortiGate-620B units are running the same FortiOS firmware build and are set to the factory default configuration.

### General configuration steps

1. Configure the FortiGate units for HA operation.

- Change each unit's host name.
  - Configure HA.
2. Connect the cluster to the network.
  3. View cluster status.
  4. Add basic configuration settings and configure the aggregated interfaces.
    - Add a password for the admin administrative account.
    - Add the aggregated interfaces.
    - Disable `lACP-ha-slave` so that the subordinate unit does not send LACP packets.
    - Add a default route.

You could also configure aggregated interfaces in each FortiGate unit before the units form a cluster.

5. Configure HA port monitoring for the aggregated interfaces.

## Configuring active-passive HA cluster that includes aggregated interfaces - web-based manager

These procedures assume you are starting with two FortiGate-620B units with factory default settings.

### To configure the FortiGate-620B units for HA operation

1. Power on the first FortiGate-620B unit and log into the web-based manager.
2. On the *System Information* dashboard widget, beside *Host Name* select *Change*.
3. Enter a new Host Name for this FortiGate unit.

<b>New Name</b>	620_ha_1
-----------------	----------

4. Select OK.
5. Go to *System > Config > HA* and change the following settings.

<b>Mode</b>	Active-Passive	
<b>Group Name</b>	example5.com	
<b>Password</b>	HA_pass_5	
<b>Heartbeat Interface</b>		
	<b>Enable</b>	<b>Priority</b>
<b>port5</b>	Select	50
<b>port6</b>	Select	50

Since port3 and port4 will be used for a aggregated interface, you must change the HA heartbeat configuration to not use those interfaces.

6. Select OK.

The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 1300](#)). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

7. Power off the first FortiGate unit.
8. Repeat these steps for the second FortiGate unit.

Set the second FortiGate unit host name to:

---

<b>New Name</b>	620_ha_2
-----------------	----------

---

### To connect the cluster to the network

1. Connect the port1 and port2 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the Internet.

Configure the switch so that the port1 and port2 of 620\_ha\_1 make up an aggregated interface and port1 and port2 of 620\_ha\_2 make up a second aggregated interface.

2. Connect the port3 and port4 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the internal network.

Configure the switch so that the port3 and port4 of 620\_ha\_1 make up an aggregated interface and port3 and port4 of 620\_ha\_2 make up another aggregated interface.

3. Connect the port5 interfaces of 620\_ha\_1 and 620\_ha\_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
4. Connect the port5 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.

5. Power on the cluster units.

The units negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete, the cluster is ready to be configured for your network.

### To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.

1. View the system dashboard.

The System Information dashboard widget shows the *Cluster Name* (example5.com) and the host names and serial numbers of the *Cluster Members*. The Unit Operation widget shows multiple cluster units.

2. Go to *System > Config > HA* to view the cluster members list.

The list shows two cluster units, their host names, their roles in the cluster, and their priorities. You can use this list to confirm that the cluster is operating normally.

### To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units, the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 1212](#) to troubleshoot the cluster.

### To add basic configuration settings and the aggregate interfaces

Use the following steps to add a few basic configuration settings.

1. Log into the cluster web-based manager.
2. Go to *System > Admin > Administrators*.
3. For *admin*, select the *Change Password* icon.
4. Enter and confirm a new password.
5. Select OK.
6. Go to *Router > Static > Static Routes* and temporarily delete the default route.  
You cannot add an interface to a aggregated interface if any settings (such as the default route) are configured for it.
7. Go to *System > Network > Interfaces* and select *Create New* to add the aggregate interface to connect to the Internet.



8. Set *Type* to *802.3ad Aggregate* and configure the aggregate interface to be connected to the Internet:

<b>Name</b>	Port1_Port2
<b>Physical Interface Members</b>	
<b>Selected Interfaces</b>	port1, port2
<b>IP/Netmask</b>	172.20.120.141/24

9. Select OK.

10. Select *Create New* to add the aggregate interface to connect to the internal network.

11. Set *Type* to *802.3ad Aggregate* and configure the aggregate interface to be connected to the Internet:

<b>Name</b>	Port3_Port4
<b>Physical Interface Members</b>	
<b>Selected Interfaces</b>	port3, port4
<b>IP/Netmask</b>	10.11.101.100/24
<b>Administrative Access</b>	HTTPS, PING, SSH

12. Select OK.

The virtual MAC addresses of the FortiGate-620B interfaces change to the following. Note that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12

- port9 interface virtual MAC: 00-09-0f-09-00-13

13. Connect to the CLI and enter the following command to disable sending LACP packets from the subordinate unit:

```
config system interface
 edit Port1_Port2
 set lacp-ha-slave disable
 next
 edit Port3_Port4
 set lacp-ha-slave disable
end
```

14. Go to *Router > Static > Static Routes*.

15. Add the default route.

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Gateway</b>	172.20.120.2
<b>Device</b>	Port1_Port2
<b>Distance</b>	10

16. Select OK.

#### To configure HA port monitoring for the aggregate interfaces

1. Go to *System > Config > HA*.
2. In the cluster members list, edit the primary unit.
3. Configure the following port monitoring for the aggregate interfaces:

Port Monitor	
<b>Port1_Port2</b>	Select
<b>Port3_Port4</b>	Select

4. Select OK.

## Configuring active-passive HA cluster that includes aggregate interfaces - CLI

These procedures assume you are starting with two FortiGate-620B units with factory default settings.

#### To configure the FortiGate-620B units for HA operation

1. Power on the first FortiGate-620B unit and log into the CLI.
2. Change the host name for this FortiGate unit:

```
config system global
 set hostname 620_ha_1
end
```

### 3. Configure HA settings.

```
config system ha
 set mode a-p
 set group-name example5.com
 set password HA_pass_5
 set hbdev port5 50 port6 50
end
```

Since port3 and port4 will be used for an aggregated interface, you must change the HA heartbeat configuration.

The FortiGate unit negotiates to establish an HA cluster. You may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 1300](#)). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use

the following command to view the port1 interface virtual MAC address (Current\_HWaddr) and the port1 permanent MAC address (Permanent\_HWaddr):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

#### 4. Display the HA configuration (optional).

```
get system ha
group-id : 0
group-name : example5.com
mode : a-p
password : *
hbdev : "port5" 50 "port6" 50
session-sync-dev :
route-ttl : 10
route-wait : 0
route-hold : 10
sync-config : enable
encryption : disable
authentication : disable
hb-interval : 2
hb-lost-threshold : 20
helo-holddown : 20
arps : 5
arps-interval : 8
session-pickup : disable
link-failed-signal : disable
uninterruptible-upgrade: enable
ha-mgmt-status : disable
ha-eth-type : 8890
hc-eth-type : 8891
l2ep-eth-type : 8893
subsecond : disable
vcluster2 : disable
vcluster-id : 1
override : disable
priority : 128
monitor :
pingserver-monitor-interface:
pingserver-failover-threshold: 0
pingserver-flip-timeout: 60
vdom : "root"
```

5. Repeat these steps for the other FortiGate unit.

Set the other FortiGate unit host name to:

```
config system global
 set hostname 620_ha_2
end
```

### To connect the cluster to the network

1. Connect the port1 and port2 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the Internet.

Configure the switch so that the port1 and port2 of 620\_ha\_1 make up an aggregated interface and port1 and port2 of 620\_ha\_2 make up another aggregated interface.

2. Connect the port3 and port4 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the internal network.

Configure the switch so that the port3 and port4 of 620\_ha\_1 make up an interfaced and port3 and port4 of 620\_ha\_2 make up another aggregated interface.

3. Connect the port5 interfaces of 620\_ha\_1 and 620\_ha\_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
4. Connect the port5 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
5. Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

### To view cluster status

Use the following steps to view cluster status from the CLI.

1. Log into the CLI.
2. Enter `get system status` to verify the HA status of the cluster unit that you logged into. Look for the following information in the command output.

---

Current HA mode: a-a, master	The cluster units are operating as a cluster and you have connected to the primary unit.
------------------------------	------------------------------------------------------------------------------------------

---

Current HA mode: a-a, backup	The cluster units are operating as a cluster and you have connected to a subordinate unit.
------------------------------	--------------------------------------------------------------------------------------------

---

Current HA mode: standalone	The cluster unit is not operating in HA mode
-----------------------------	----------------------------------------------

---

3. Enter the following command to confirm the HA configuration of the cluster:

```
get system ha status
 Model: 620
 Mode: a-a
 Group: 0
 Debug: 0
 ses_pickup: disable
 Master:128 620_ha_2 FG600B3908600825 0
 Slave :128 620_ha_1 FG600B3908600705 1
 number of vcluster: 1
 vcluster 1: work 169.254.0.1
 Master:0 FG600B3908600825
 Slave :1 FG600B3908600705
```

The command output shows both cluster units, their host names, their roles in the cluster, and their priorities. You can use this command to confirm that the cluster is operating normally. For example, if the command shows only one cluster unit then the other unit has left the cluster for some reason.

### To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 1212](#) to troubleshoot the cluster.

### To add basic configuration settings and the aggregate interfaces

Use the following steps to add a few basic configuration settings and the aggregate interfaces.

1. Add a password for the admin administrative account.

```
config system admin
 edit admin
 set password <psswr>
 end
```

2. Temporarily delete the default route.

You cannot add an interface to an aggregate interface if any settings (such as the default route) are configured for it. In this example the index of the default route is 1.

```
config router static
 delete 1
end
```

### 3. Add the aggregate interfaces:

```
config system interface
 edit Port1_Port2
 set type aggregate
 set lacp-ha-slave disable
 set member port1 port2
 set ip 172.20.120.141/24
 set vdom root
 next
 edit Port3_Port4
 set type aggregate
 set lacp-ha-slave disable
 set member port3 port4
 set ip 10.11.101.100/24
 set vdom root
end
```

The virtual MAC addresses of the FortiGate-620B interfaces change to the following. Note that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

### 4. Add the default route.

```
config router static
 edit 1
 set dst 0.0.0.0 0.0.0.0
 set gateway 172.20.120.2
 set device Port1_Port2
end
```

## To configure HA port monitoring for the aggregate interfaces


1. Configure HA port monitoring for the aggregate interfaces.

```
config system ha
 set monitor Port1_Port2 Port3_Port4
end
```

## Example: HA and redundant interfaces

On FortiGate models that support it you can combine two or more interfaces into a single redundant interface. A redundant interface consists of two or more physical interfaces. Traffic is processed by the first physical interface in the redundant interface. If that physical interface fails, traffic fails over to the next physical interface. Redundant interfaces don't have the benefit of improved performance that aggregate interfaces can have, but they do provide failover if a physical interface fails or is disconnected.

**Figure 196:**Example cluster with a redundant interfaces



This example describes how to configure an HA cluster consisting of two FortiGate-620B units with a redundant interface connection to the Internet and to an internal network. The connection to the Internet uses port1 and port2. The connection to the internal network uses port3 and port4. The HA heartbeat uses port5 and port6.

The redundant interfaces are also configured as HA monitored interfaces.

## HA interface monitoring, link failover, and redundant interfaces

HA interface monitoring monitors the redundant interface as a single interface and does not monitor the individual physical interfaces in the redundant interface. HA interface monitoring registers the redundant interface to have failed only if all the physical interfaces in the redundant interface have failed. If only some of the physical interfaces in the redundant interface fail or become disconnected, HA considers the redundant interface to be operating normally.



## HA MAC addresses and redundant interfaces

For a standalone FortiGate unit a redundant interface has the MAC address of the first physical interface added to the redundant interface configuration. A redundant interface consisting of port1 and port2 would have the MAC address of port1.

In an HA cluster, HA changes the MAC addresses of the cluster interfaces to virtual MAC addresses. A redundant interface in a cluster acquires the virtual MAC address that would have been acquired by the first physical interface added to the redundant interface configuration.

## Connecting multiple redundant interfaces to one switch while operating in active-passive HA mode

HA assigns the same virtual MAC addresses to the subordinate unit interfaces as are assigned to the corresponding primary unit interfaces. Consider a cluster of two FortiGate units operating in active-passive mode with a redundant interface consisting of port1 and port2. You can connect multiple redundant interfaces to the same switch if you configure the switch so that it defines multiple separate redundant interfaces and puts the redundant interfaces of each cluster unit into separate redundant interfaces. In this configuration, each cluster unit forms a separate redundant interface with the switch.

However, if the switch is configured with a single four-port redundant interface configuration, because the same MAC addresses are being used by both cluster units, the switch adds all four interfaces (port1 and port2 from the primary unit and port1 and port2 from the subordinate unit) to the same redundant interface.

To avoid unpredictable results, when you connect a switch to multiple redundant interfaces in an active-passive cluster you should configure separate redundant interfaces on the switch; one for each cluster unit.

## Connecting multiple redundant interfaces to one switch while operating in active-active HA mode

In an active-active cluster, all cluster units send and receive packets. To operate a cluster with redundant interfaces in active-active mode, with multiple redundant interfaces connected to the same switch, you must separate the redundant interfaces of each cluster unit into different redundant interfaces on the connecting switch.

## General configuration steps

The section includes web-based manager and CLI procedures. These procedures assume that the FortiGate-620B units are running the same FortiOS firmware build and are set to the factory default configuration.

### General configuration steps

1. Configure the FortiGate units for HA operation.
  - Change each unit's host name.
  - Configure HA.
2. Connect the cluster to the network.
3. View cluster status.
4. Add basic configuration settings and configure the redundant interfaces.
  - Add a password for the admin administrative account.
  - Add the redundant interfaces.

- Add a default route.

You could also configure redundant interfaces in each FortiGate unit before they form a cluster.

5. Configure HA port monitoring for the redundant interfaces.

## Configuring active-passive HA cluster that includes redundant interfaces - web-based manager

These procedures assume you are starting with two FortiGate-620B units with factory default settings.

### To configure the FortiGate-620B units for HA operation

1. Power on the first FortiGate-620B unit and log into the web-based manager.
2. On the *System Information* dashboard widget, beside *Host Name* select *Change*.
3. Enter a new Host Name for this FortiGate unit.

<b>New Name</b>	620_ha_1
-----------------	----------

4. Select OK.
5. Go to *System > Config > HA* and change the following settings.

<b>Mode</b>	Active-Passive
-------------	----------------

<b>Group Name</b>	example6.com
-------------------	--------------

<b>Password</b>	HA_pass_6
-----------------	-----------

#### Heartbeat Interface

	Enable	Priority
<b>port5</b>	Select	50
<b>port6</b>	Select	50

Since port3 and port4 will be used for a redundant interface, you must change the HA heartbeat configuration.

6. Select OK.

The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 1300](#)). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07

- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

7. Power off the first FortiGate unit.
8. Repeat these steps for the second FortiGate unit.  
Set the second FortiGate unit host name to:

---

<b>New Name</b>	620_ha_2
-----------------	----------

---

### To connect the cluster to the network

1. Connect the port1 and port2 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the Internet.  
Configure the switch so that the port1 and port2 of 620\_ha\_1 make up a redundant interface and port1 and port2 of 620\_ha\_2 make up another redundant interface.
2. Connect the port3 and port4 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the internal network.  
Configure the switch so that the port3 and port4 of 620\_ha\_1 make up a redundant interface and port3 and port4 of 620\_ha\_2 make up another redundant interface.
3. Connect the port5 interfaces of 620\_ha\_1 and 620\_ha\_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.

4. Connect the port5 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
5. Power on the cluster units.
 

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

### To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.

1. View the system dashboard.
 

The System Information dashboard widget shows the *Cluster Name* (example5.com) and the host names and serial numbers of the *Cluster Members*. The Unit Operation widget shows multiple cluster units.
2. Go to *System > Config > HA* to view the cluster members list.
 

The list shows two cluster units, their host names, their roles in the cluster, and their priorities. You can use this list to confirm that the cluster is operating normally.

### To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 1212](#) to troubleshoot the cluster.

### To add basic configuration settings and the redundant interfaces

Use the following steps to add a few basic configuration settings.

1. Log into the cluster web-based manager.
2. Go to *System > Admin > Administrators*.
3. For *admin*, select the *Change Password* icon
4. Enter and confirm a new password.
5. Select OK.
6. Go to *Router > Static > Static Routes* and temporarily delete the default route.
 

You cannot add an interface to a redundant interface if any settings (such as the default route) are configured for it.
7. Go to *System > Network > Interfaces* and select *Create New* to add the redundant interface to connect to the Internet.
8. Set *Type* to *Redundant Interface* and configure the redundant interface to be connected to the Internet:

<b>Name</b>	Port1_Port2
<b>Physical Interface Members</b>	
<b>Selected Interfaces</b>	port1, port2
<b>IP/Netmask</b>	172.20.120.141/24

9. Select OK.
10. Select *Create New* to add the redundant interface to connect to the internal network.

11. Set *Type* to *Redundant Interface* and configure the redundant interface to be connected to the Internet:

<b>Name</b>	Port3_Port4
<b>Physical Interface Members</b>	
<b>Selected Interfaces</b>	port3, port4
<b>IP/Netmask</b>	10.11.101.100/24
<b>Administrative Access</b>	HTTPS, PING, SSH

12. Select OK.

The virtual MAC addresses of the FortiGate-620B interfaces change to the following. Note that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

13. Go to *Router > Static > Static Routes*.

14. Add the default route.

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Gateway</b>	172.20.120.2
<b>Device</b>	Port1_Port2
<b>Distance</b>	10

15. Select OK.

### To configure HA port monitoring for the redundant interfaces

1. Go to *System > Config > HA*.
2. In the cluster members list, edit the primary unit.
3. Configure the following port monitoring for the redundant interfaces:

Port Monitor	
Port1_Port2	Select
Port3_Port4	Select

4. Select OK.

## Configuring active-passive HA cluster that includes redundant interfaces - CLI

These procedures assume you are starting with two FortiGate-620B units with factory default settings.

### To configure the FortiGate-620B units for HA operation

1. Power on the first FortiGate-620B unit and log into the CLI.
2. Change the host name for this FortiGate unit:

```
config system global
 set hostname 620_ha_1
end
```

3. Configure HA settings.

```
config system ha
 set mode a-p
 set group-name example6.com
 set password HA_pass_6
 set hbdev port5 50 port6 50
end
```

Since port3 and port4 will be used for a redundant interface, you must change the HA heartbeat configuration.

The FortiGate unit negotiates to establish an HA cluster. You may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 1300](#)). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a

- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

**4. Display the HA configuration (optional).**

```
get system ha
group-id : 0
group-name : example6.com
mode : a-p
password : *
hbdev : "port5" 50 "port6" 50
session-sync-dev :
route-ttl : 10
route-wait : 0
route-hold : 10
sync-config : enable
encryption : disable
authentication : disable
hb-interval : 2
hb-lost-threshold : 20
helo-holddown : 20
arps : 5
arps-interval : 8
session-pickup : disable
link-failed-signal : disable
uninterruptible-upgrade: enable
ha-mgmt-status : disable
ha-eth-type : 8890
hc-eth-type : 8891
l2ep-eth-type : 8893
subsecond : disable
vcluster2 : disable
vcluster-id : 1
override : disable
priority : 128
monitor :
pingserver-monitor-interface:
pingserver-failover-threshold: 0
pingserver-flip-timeout: 60
vdom : "root"
```

**5. Repeat these steps for the other FortiGate unit.**

Set the other FortiGate unit host name to:

```
config system global
 set hostname 620_ha_2
end
```

**To connect the cluster to the network**

1. Connect the port1 and port2 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the Internet.

Configure the switch so that the port1 and port2 of 620\_ha\_1 make up a redundant interface and port1 and port2 of 620\_ha\_2 make up another redundant interface.



2. Connect the port3 and port4 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the internal network.  
Configure the switch so that the port3 and port4 of 620\_ha\_1 make up a redundant interface and port3 and port4 of 620\_ha\_2 make up another redundant interface.
3. Connect the port5 interfaces of 620\_ha\_1 and 620\_ha\_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
4. Connect the port5 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
5. Power on the cluster units.  
The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.  
When negotiation is complete the cluster is ready to be configured for your network.

### To view cluster status

Use the following steps to view cluster status from the CLI.

1. Log into the CLI.
2. Enter `get system status` to verify the HA status of the cluster unit that you logged into. Look for the following information in the command output.

Current HA mode: a-a, master	The cluster units are operating as a cluster and you have connected to the primary unit.
Current HA mode: a-a, backup	The cluster units are operating as a cluster and you have connected to a subordinate unit.
Current HA mode: standalone	The cluster unit is not operating in HA mode

3. Enter the following command to confirm the HA configuration of the cluster:

```
get system ha status
Model: 620
Mode: a-a
Group: 0
Debug: 0
ses_pickup: disable
Master:128 620_ha_2 FG600B3908600825 0
Slave :128 620_ha_1 FG600B3908600705 1
number of vcluster: 1
vcluster 1: work 169.254.0.1
Master:0 FG600B3908600825
Slave :1 FG600B3908600705
```

The command output shows both cluster units, their host names, their roles in the cluster, and their priorities. You can use this command to confirm that the cluster is operating normally. For example, if the command shows only one cluster unit then the other unit has left the cluster for some reason.

### To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 1212](#) to troubleshoot the cluster.

## To add basic configuration settings and the redundant interfaces

Use the following steps to add a few basic configuration settings and the redundant interfaces.

1. Add a password for the admin administrative account.

```
config system admin
 edit admin
 set password <psswr>
 end
```

2. Temporarily delete the default route.

You cannot add an interface to a redundant interface if any settings (such as the default route) are configured for it. In this example the index of the default route is 1.

```
config router static
 delete 1
end
```

### 3. Add the redundant interfaces:

```
config system interface
 edit Port1_Port2
 set type redundant
 set member port1 port2
 set ip 172.20.120.141/24
 set vdom root
 next
 edit Port3_Port4
 set type redundant
 set member port3 port4
 set ip 10.11.101.100/24
 set vdom root
end
```

The virtual MAC addresses of the FortiGate-620B interfaces change to the following. Note that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

### 4. Add the default route.

```
config router static
 edit 1
 set dst 0.0.0.0 0.0.0.0
 set gateway 172.20.120.2
 set device Port1_Port2
end
```

## To configure HA port monitoring for the redundant interfaces

1. Configure HA port monitoring for the redundant interfaces.

```
config system ha
 set monitor Port1_Port2 Port3_Port4
end
```

## Troubleshooting HA clusters

This section describes some HA clustering troubleshooting techniques.

### Ignoring hardware revisions

Some FortiGate platforms have gone through multiple hardware versions. In some cases the hardware changes between versions have meant that by default you cannot form a cluster if the FortiGate units in the cluster have different hardware versions. If you run into this problem you can use the following command on each FortiGate unit to cause the cluster to ignore different hardware versions:

```
execute ha ignore-hardware-revision {disable | enable | status}
```

This command is only available on FortiGate units that have had multiple hardware revisions. By default the command is set to prevent FortiOS from forming clusters between FortiGate units with different hardware revisions. You can enable this command to be able to create a cluster consisting of FortiGate units with different hardware revisions. Use the `status` option to verify the whether ignoring hardware revisions is enabled or disabled.

Affected hardware models include:

- FortiGate-100D
- FortiGate-300C
- FortiGate-80C and FortiWiFi-80C
- FortiGate-60C

### Before you set up a cluster

Before you set up a cluster ask yourself the following questions about the FortiGate units that you are planning to use to create a cluster.

1. Do all the FortiGate units have the same hardware configuration? Including the same hard disk configuration and the same AMC cards installed in the same slots?
2. Do all FortiGate units have the same firmware build?
3. Are all FortiGate units set to the same operating mode (NAT or Transparent)?
4. Are all the FortiGate units operating in single VDOM mode?
5. If the FortiGate units are operating in multiple VDOM mode do they all have the same VDOM configuration?



In some cases you may be able to form a cluster if different FortiGate units have different firmware builds, different VDOM configurations, and are in different operating modes. However, if you encounter problems they may be resolved by installing the same firmware build on each unit, and give them the same VDOM configuration and operating mode.

---

## Troubleshooting the initial cluster configuration

This section describes how to check a cluster when it first starts up to make sure that it is configured and operating correctly. This section assumes you have already configured your HA cluster.

### To verify that a cluster can process traffic and react to a failure

1. Add a basic security policy configuration and send network traffic through the cluster to confirm connectivity.

For example, if the cluster is installed between the Internet and an internal network, set up a basic internal to external security policy that accepts all traffic. Then from a PC on the internal network, browse to a website on the Internet or ping a server on the Internet to confirm connectivity.

2. From your management PC, set ping to continuously ping the cluster, and then start a large download, or in some other way establish ongoing traffic through the cluster.
3. While traffic is going through the cluster, disconnect the power from one of the cluster units. You could also shut down or restart a cluster unit. Traffic should continue with minimal interruption.
4. Start up the cluster unit that you disconnected. The unit should re-join the cluster with little or no affect on traffic.
5. Disconnect a cable for one of the HA heartbeat interfaces. The cluster should keep functioning, using the other HA heartbeat interface.
6. If you have port monitoring enabled, disconnect a network cable from a monitored interface. Traffic should continue with minimal interruption.

### To verify the cluster configuration - web-based manager

1. Log into the cluster web-based manager.
2. Check the system dashboard to verify that the System Information widget displays all of the cluster units.
3. Check the cluster member graphic to verify that the correct cluster unit interfaces are connected.
4. Go to *System > Config > HA* and verify that all of the cluster units are displayed on the cluster members list.
5. From the cluster members list, edit the primary unit (master) and verify the cluster configuration is as expected.

### To troubleshoot the cluster configuration - web-based manager

1. Connect to each cluster unit web-based manager and verify that the HA configurations are the same.
2. To connect to each web-based manager, you may need to disconnect some units from the network to connect to the other if the units have the same IP address.
3. If the configurations are the same, try re-entering the cluster *Password* on each cluster unit in case you made an error typing the password when configuring one of the cluster units.
4. Check that the correct interfaces of each cluster unit are connected.

Check the cables and interface LEDs.

Use the Unit Operation dashboard widget, system network interface list, or cluster members list to verify that each interface that should be connected actually is connected.

If Link is down re-verify the physical connection. Try replacing network cables or switches as required.

### To verify the cluster configuration - CLI

1. Log into each cluster unit CLI.

You can use the console connection if you need to avoid the problem of units having the same IP address.

2. Enter the command `get system status`.

Look for the following information in the command output.

Current HA mode: a-a, master	The cluster units are operating as a cluster and you have connected to the primary unit.
Current HA mode: a-a, backup	The cluster units are operating as a cluster and you have connected to a subordinate unit.
Current HA mode: standalone	The cluster unit is not operating in HA mode

3. Verify that the `get system ha status` command displays all of the cluster units.
4. Enter the `get system ha` command to verify that the HA configuration is correct and the same for each cluster unit.

### To troubleshoot the cluster configuration - CLI

1. Try using the following command to re-enter the cluster password on each cluster unit in case you made an error typing the password when configuring one of the cluster units.

```
config system ha
 set password <password>
end
```

2. Check that the correct interfaces of each cluster unit are connected.

Check the cables and interface LEDs.

Use `get hardware nic <interface_name>` command to confirm that each interface is connected. If the interface is connected the command output should contain a `Link: up` entry similar to the following:

```
get hardware nic port1
.
.
.
Link: up
.
.
.
```

If Link is down, re-verify the physical connection. Try replacing network cables or switches as required.

## More troubleshooting information

Much of the information in this HA guide can be useful for troubleshooting HA clusters. Here are some links to sections with more information.

- If sessions are lost after a failover you may need to change route-ttl to keep synchronized routes active longer. See [“Change how long routes stay in a cluster unit routing table” on page 1317](#).
- To control which cluster unit becomes the primary unit, you can change the device priority and enable override. See [“Controlling primary unit selection using device priority and override” on page 1140](#).
- Changes made to a cluster can be lost if override is enabled. See [“Configuration changes can be lost if override is enabled” on page 1141](#).
- In some cases, age differences among cluster units result in the wrong cluster unit becoming the primary unit. For example, if a cluster unit set to a high priority reboots, that unit will have a lower age than other cluster units. You can resolve this problem by resetting the age of one or more cluster units. See [“Resetting the age of all cluster units” on page 1135](#). You can also adjust how sensitive the cluster is to age differences. This can be useful if large age differences cause problems. See [“Cluster age difference margin \(grace period\)” on page 1133](#) and [“Changing the cluster age difference margin” on page 1133](#).
- If one of the cluster units needs to be serviced or removed from the cluster for other reasons, you can do so without affecting the operation of the cluster. See [“Disconnecting a cluster unit from a cluster” on page 1285](#).
- The web-based manager and CLI will not allow you to configure HA if:
  - You have configured a FortiGate interface to get its IP address using DHCP or PPPoE. See [“FortiGate HA compatibility with PPPoE and DHCP” on page 1142](#).
  - You have enabled VRRP. See [“VRRP” on page 1364](#).
  - You have enabled TCP session synchronization. See [“FortiGate Session Life Support Protocol \(FGSP\)” on page 1370](#).
- Some third-party network equipment may prevent HA heartbeat communication, resulting in a failure of the cluster or the creation of a split brain scenario. For example, some switches use packets with the same Ethertype as HA heartbeat packets use for internal functions and when used for HA heartbeat communication the switch generates CRC errors and the packets are not forwarded. See [“Heartbeat packet Ethertypes” on page 1297](#).
- Very busy clusters may not be able to send HA heartbeat packets quickly enough, also resulting in a split brain scenario. You may be able to resolve this problem by modifying HA heartbeat timing. See [“Modifying heartbeat timing” on page 1298](#).
- Very busy clusters may suffer performance reductions if session pickup is enabled. If possible you can disable this feature to improve performance. If you require session pickup for your cluster, several options are available for improving session pickup performance. See [“Improving session synchronization performance” on page 1331](#).
- If it takes longer than expected for a cluster to failover you can try changing how the primary unit sends gratuitous ARP packets. See [“Changing how the primary unit sends gratuitous ARP packets after a failover” on page 1301](#).
- You can also improve failover times by configuring the cluster for subsecond failover. See [“Subsecond failover” on page 1324](#) and [“Failover performance” on page 1340](#).
- When you first put a FortiGate unit in HA mode you may lose connectivity to the unit. This occurs because HA changes the MAC addresses of all FortiGate unit interfaces, including the one that you are connecting to. The cluster MAC addresses also change if you change the some HA settings such as the cluster group ID. The connection will be restored in a short time as your network and PC updates to the new MAC address. To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the

FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

- Since HA changes all cluster unit MAC addresses, if your network uses MAC address filtering you may have to make configuration changes to account for the HA MAC addresses.
- A network may experience packet loss when two FortiGate HA clusters have been deployed in the same broadcast domain. Deploying two HA clusters in the same broadcast domain can result in packet loss because of MAC address conflicts. The packet loss can be diagnosed by pinging from one cluster to the other or by pinging both of the clusters from a device within the broadcast domain. You can resolve the MAC address conflict by changing the HA Group ID configuration of the two clusters. The HA Group ID is sometimes also called the Cluster ID. See [“Diagnosing packet loss with two FortiGate HA clusters in the same broadcast domain” on page 1305](#).
- The cluster CLI displays `slave is not in sync` messages if there is a synchronization problem between the primary unit and one or more subordinate units. See [“How to diagnose HA out of sync messages” on page 1313](#).
- If you have configured dynamic routing and the new primary unit takes too long to update its routing table after a failover you can configure graceful restart and also optimize how routing updates are synchronized. See [“Configuring graceful restart for dynamic routing failover” on page 1315](#) and [“Controlling how the FGCP synchronizes kernel routing table updates” on page 1316](#).
- Some switches may not be able to detect that the primary unit has become a subordinate unit and will keep sending packets to the former primary unit. This can occur after a link failover if the switch does not detect the failure and does not clear its MAC forwarding table. See [“Updating MAC forwarding tables when a link failover occurs” on page 1323](#).
- If a link not directly connected to a cluster unit (for example, between a switch connected to a cluster interface and the network) fails you can enable remote link failover to maintain communication. See [“Remote link failover” on page 1325](#).
- If you find that some cluster units are not running the same firmware build you can reinstall the correct firmware build on the cluster to upgrade all cluster units to the same firmware build. See [“Synchronizing the firmware build running on a new cluster unit” on page 1275](#).



# Virtual clusters

This chapter provides an introduction to virtual clustering and also contains general procedures and configuration examples that describe how to configure FortiGate HA virtual clustering.

This chapter contains the following sections:

- [Virtual clustering overview](#)
- [Configuring HA for virtual clustering](#)
- [Example: virtual clustering with two VDOMs and VDOM partitioning](#)
- [Example: inter-VDOM links in a virtual clustering configuration](#)
- [Troubleshooting virtual clustering](#)

## Virtual clustering overview

Virtual clustering is an extension of the FGCP for a cluster of 2 FortiGate units operating with multiple VDOMS enabled. Virtual clustering operates in active-passive mode to provide failover protection between two instances of a VDOM operating on two different cluster units. You can also operate virtual clustering in active-active mode to use HA load balancing to load balance sessions between cluster units. Alternatively, by distributing VDOM processing between the two cluster units you can also configure virtual clustering to provide load balancing by distributing sessions for different VDOMs to each cluster unit.

**Figure** shows an example virtual cluster configuration consisting of two FortiGate-620B units. The virtual cluster has two virtual domains, root and Eng\_vdm.

The root virtual domain includes the port1 and port2 interfaces. The Eng\_vdm virtual domain includes the port5 and port6 interfaces. The port3 and port4 interfaces (not shown in the diagram) are the HA heartbeat interfaces.



FortiGate virtual clustering is limited to a cluster of 2 FortiGate units with multiple VDOMs enabled. If you want to create a cluster of more than 2 FortiGate units operating with multiple VDOMS you could consider other solutions that either do not include multiple VDOMs in one cluster or employ a feature such as standalone session synchronization. See [“FortiGate Session Life Support Protocol \(FGSP\)”](#) on page 1370.

---


## Virtual clustering and failover protection

Virtual clustering operates on a cluster of two (and only two) FortiGate units with VDOMs enabled. Each VDOM creates a cluster between instances of the VDOMs on the two FortiGate units in the virtual cluster. All traffic to and from the VDOM stays within the VDOM and is processed by the VDOM. One cluster unit is the primary unit for each VDOM and one cluster unit is the subordinate unit for each VDOM. The primary unit processes all traffic for the VDOM. The subordinate unit does not process traffic for the VDOM. If a cluster unit fails, all traffic fails over to the cluster unit that is still operating.

## Virtual clustering and heartbeat interfaces

The HA heartbeat provides the same HA services in a virtual clustering configuration as in a standard HA configuration. One set of HA heartbeat interfaces provides HA heartbeat services for all of the VDOMs in the cluster. You do not have to add a heartbeat interface for each VDOM.

**Figure 197:**Example virtual cluster of two FortiGate-620B units



## Virtual clustering and HA override

For a virtual cluster configuration, override is enabled by default for both virtual clusters when you:

- Enable VDOM partitioning from the web-based manager by moving virtual domains to virtual cluster 2
- Enter `set vcluster2 enable` from the CLI `config system ha` command to enable virtual cluster 2.

Usually you would enable virtual cluster 2 and expect one cluster unit to be the primary unit for virtual cluster 1 and the other cluster unit to be the primary unit for virtual cluster 2. For this distribution to occur override must be enabled for both virtual clusters. Otherwise you will need to restart the cluster to force it to renegotiate.



If override is enabled the cluster may renegotiate too often. You can choose to disable override at any time. If you decide to disable override, for best results, you should disable it for both cluster units.

For more information about HA override see [“HA override” on page 1138](#).

## Virtual clustering and load balancing or VDOM partitioning

There are two ways to configure load balancing for virtual clustering. The first is to set the HA mode to active-active. The second is to configure VDOM partitioning. For virtual clustering, setting the HA Mode to active-active has the same result as active-active HA for a cluster without virtual domains. The primary unit receives all sessions and load balances them among the cluster units according to the load balancing schedule. All cluster units process traffic for all virtual domains.

In a VDOM partitioning virtual clustering configuration, the HA mode is set to active-passive. Even though virtual clustering operates in active-passive mode you can configure a form of load balancing by using VDOM partitioning to distribute traffic between both cluster units. To configure VDOM partitioning you set one cluster unit as the primary unit for some virtual

domains and you set the other cluster unit as the primary unit for other virtual domains. All traffic for a virtual domain is processed by the primary unit for that virtual domain. You can control the distribution of traffic between the cluster units by adjusting which cluster unit is the primary unit for each virtual domain.

For example, you could have 4 VDOMs, two of which have a high traffic volume and two of which have a low traffic volume. You can configure each cluster unit to be the primary unit for one of the high volume VDOMs and one of the low volume VDOMs. As a result each cluster unit will be processing traffic for a high volume VDOM and a low volume VDOM, resulting in an even distribution of traffic between the cluster units. You can adjust the distribution at any time. For example, if a low volume VDOM becomes a high volume VDOM you can move it from one cluster unit to another until the best balance is achieved.

From the web-based manager you configure VDOM partitioning by setting the HA mode to active-passive and distributing virtual domains between Virtual Cluster 1 and Virtual Cluster 2. You can also configure different device priorities, port monitoring, and remote link failover, for Virtual Cluster 1 and Virtual Cluster 2.



The device priorities for virtual cluster 1 and virtual cluster 2 are not synchronized between the FortiGate units in the virtual cluster. You must configure these device priorities separately for each cluster unit.

---

From the CLI you configure VDOM partitioning by setting the HA mode to `a-p`. Then you configure device priority, port monitoring, and remote link failover and specify the VDOMs to include in virtual cluster 1. You do the same for virtual cluster 2 by entering the `config secondary-vcluster` command.

Failover protection does not change. If one cluster unit fails, all sessions are processed by the remaining cluster unit. No traffic interruption occurs for the virtual domains for which the still functioning cluster unit was the primary unit. Traffic may be interrupted temporarily for virtual domains for which the failed unit was the primary unit while processing fails over to the still functioning cluster unit.

If the failed cluster unit restarts and rejoins the virtual cluster, VDOM partitioning load balancing is restored.

## Configuring HA for virtual clustering

If your cluster uses VDOMs, you are configuring virtual clustering. Most virtual cluster HA options are the same as normal HA options. However, virtual clusters include VDOM partitioning options. Other differences between configuration options for regular HA and for virtual clustering HA are described below.

To configure HA options for a cluster with VDOMs enabled:

- Log into the global web-based manager and go to *System > Config > HA*.
- From the CLI, log into the Global Configuration:

The following example shows how to configure active-active virtual clustering:

```
config global
 config system ha
 set mode a-a
 set group-name vexample1.com
 set password vHA_pass_1
 end
end
```

The following example shows how to configure active-passive virtual clustering:

```
config global
 config system ha
 set mode a-p
 set group-name vexample1.com
 set password vHA_pass_1
 end
end
```

The following example shows how to configure VDOM partitioning for virtual clustering. In the example, the FortiGate unit is configured with three VDOMs (domain\_1, domain\_2, and domain\_3) in addition to the root VDOM. The example shows how to set up a basic HA configuration that sets the device priority of virtual cluster 1 to 200. The example also shows how to enable `vcluster2`, how to set the device priority of virtual cluster 2 to 100 and how to add the virtual domains `domain_2` and `domain_3` to virtual cluster 2.

When you enable multiple VDOMs, `vcluster2` is enabled by default. Even so the command to enable `vcluster2` is included in this example in case for some reason it has been disabled. When `vcluster2` is enabled, `override` is also enabled.

The result of this configuration would be that the cluster unit that you are logged into becomes the primary unit for virtual cluster 1. This cluster unit processes all traffic for the root and `domain_1` virtual domains.

```
config global
 config system ha
 set mode a-p
 set group-name vexample1.com
 set password vHA_pass_1
 set priority 200
 set vcluster2 enable
 config secondary-vcluster
 set vdom domain_2 domain_3
 set priority 100
 end
 end
end
```

The following example shows how to use the `execute ha manage` command to change the device priorities for virtual cluster 1 and virtual cluster 2 for the other unit in the cluster. The commands set the device priority of virtual cluster 1 to 100 and virtual cluster 2 to 200.

The result of this configuration would be that the other cluster unit becomes the primary unit for virtual cluster 2. This other cluster unit would process all traffic for the domain\_2 and domain\_3 virtual domains.

```
config global
 execute ha manage 1
 config system ha
 set priority 100
 set vcluster2 enable
 config secondary-vcluster
 set priority 200
 end
 end
end
end
end
```

## Example: virtual clustering with two VDOMs and VDOM partitioning

This section describes how to configure the example virtual clustering configuration shown in [Figure 198](#). This configuration includes two virtual domains, root and Eng\_vdm and includes VDOM partitioning that sends all root VDOM traffic to 620\_ha\_1 and all Eng\_vdm VDOM traffic to 620\_ha\_2. The traffic from the internal network and the engineering network is distributed between the two FortiGate units in the virtual cluster. If one of the cluster units fails, the remaining unit will process traffic for both VDOMs.

The procedures in this example describe some of many possible sequences of steps for configuring virtual clustering. For simplicity many of these procedures assume that you are starting with new FortiGate units set to the factory default configuration. However, this is not a requirement for a successful HA deployment. FortiGate HA is flexible enough to support a successful configuration from many different starting points.

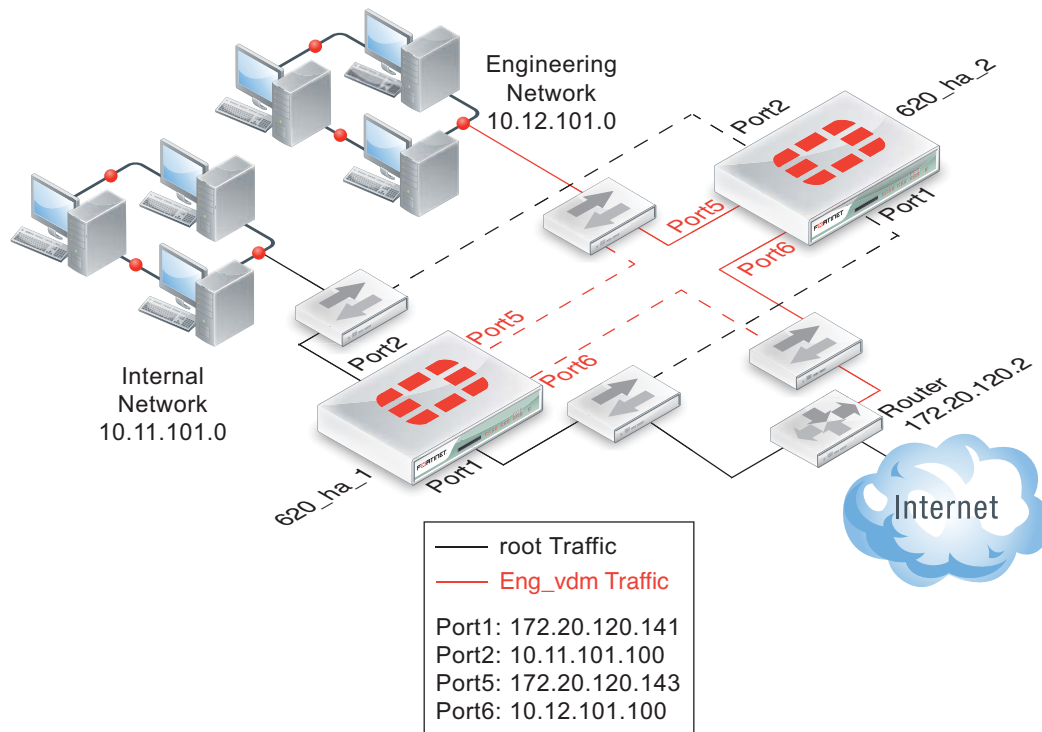
### Example virtual clustering network topology

[Figure 198](#) shows a typical FortiGate-620B HA virtual cluster consisting of two FortiGate-620B units (620\_ha\_1 and 620\_ha\_2) connected to an internal network, an engineering network and the Internet. To simplify the diagram the heartbeat connections are not shown.

The traffic from the internal network is processed by the root VDOM, which includes the port1 and port2 interfaces. The traffic from the engineering network is processed by the Eng\_vdm VDOM, which includes the port5 and port6 interfaces. VDOM partitioning is configured so that all traffic from the internal network is processed by 620\_ha\_1 and all traffic from the engineering network is processed by 620\_ha\_2.

This virtual cluster uses the default FortiGate-620B heartbeat interfaces (port3 and port4).

**Figure 198:**Example virtual cluster of two FortiGate-620B units showing VDOM partitioning



## General configuration steps

The section includes web-based manager and CLI procedures. These procedures assume that the FortiGate-620B units are running the same FortiOS firmware build and are set to the factory default configuration.

### General configuration steps

1. Configure the FortiGate units for HA operation.
  - Optionally change each unit's host name.
  - Configure HA.
2. Connect the cluster to the network.
3. Configure VDOM settings for the cluster:
  - Enable multiple VDOMs.
  - Add the Eng\_vdm VDOM.
  - Add port5 and port6 to the Eng\_vdm.
4. Configure VDOM partitioning.
5. Confirm that the cluster units are operating as a virtual cluster and add basic configuration settings to the cluster.
  - View cluster status from the web-based manager or CLI.
  - Add a password for the admin administrative account.
  - Change the IP addresses and netmasks of the port1, port2, port5, and port6 interfaces.
  - Add a default routes to each VDOM.

## Configuring virtual clustering with two VDOMs and VDOM partitioning - web-based manager

These procedures assume you are starting with two FortiGate-620B units with factory default settings.

### To configure the FortiGate-620B units for HA operation

1. Power on the first FortiGate-620B unit and log into the web-based manager.
2. On the *System Information* dashboard widget, beside *Host Name* select *Change*.
3. Enter a new Host Name for this FortiGate unit.

---

<b>New Name</b>	620_ha_1
-----------------	----------

---

4. Select OK.
5. Go to *System > Config > HA* and change the following settings.

---

<b>Mode</b>	Active-Passive
-------------	----------------

---

<b>Group Name</b>	vexample2.com
-------------------	---------------

---

<b>Password</b>	vHA_pass_2
-----------------	------------

---

6. Select OK.

The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC”](#))

addresses” on page 1300). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

7. Power off the first FortiGate unit.
8. Repeat these steps for the second FortiGate unit.  
Set the second FortiGate unit host name to:

---

<b>New Name</b>	620_ha_2
-----------------	----------

---



### To connect the cluster to the network

1. Connect the port1 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the Internet.
2. Connect the port5 interfaces of 620\_ha\_1 and 620\_ha\_2 to switch connected to the Internet. You could use the same switch for the port1 and port5 interfaces.
3. Connect the port2 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the internal network.
4. Connect the port6 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the engineering network.
5. Connect the port3 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
6. Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
7. Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention.

When negotiation is complete you can continue.

### To configure VDOM settings for the cluster

1. Log into the web-based manager.
2. On the *System Information* dashboard widget, beside *Virtual Domain* select *Enable*.
3. Select OK and then log back into the web-based manager.
4. Go to *System > VDOM* and select *Create New* to add a new VDOM.

---

<b>Name</b>	Eng_vdm
-------------	---------

---

5. Go to *System > Network > Interfaces*.
6. Edit the *port5* interface, add it to the Eng\_vdm VDOM and configure other interface settings:

---

<b>Alias</b>	Engineering_external
<b>Virtual Domain</b>	Eng_vdm
<b>IP/Netmask</b>	172.20.120.143/24

---

7. Select OK.
8. Edit the *port6* interface, add it to the Eng\_vdm VDOM and configure other interface settings:

---

<b>Alias</b>	Engineering_internal
<b>Virtual Domain</b>	Eng_vdm
<b>IP/Netmask</b>	10.120.101.100/24
<b>Administrative Access</b>	HTTPS, PING, SSH

---

9. Select OK.

### To add a default route to each VDOM

1. Go to *System > VDOM* and Enter the root VDOM.

2. Go to *Router > Static > Static Routes*.
3. Change the default route.

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Gateway</b>	172.20.120.2
<b>Device</b>	port1
<b>Distance</b>	10

4. Select *Global*.
5. Go to *System > VDOM* and Enter the Eng\_vdm VDOM.
6. Go to *Router > Static > Static Routes*.
7. Change the default route.

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Gateway</b>	172.20.120.2
<b>Device</b>	port5
<b>Distance</b>	10

#### To configure VDOM partitioning

1. Go to *System > Config > HA*.  
The cluster members shows two cluster units in Virtual Cluster 1.
2. Edit the cluster unit with the *Role of MASTER*.
3. Change *VDOM partitioning* to move the *Eng\_vdm* to the *Virtual Cluster 2* list.
4. Select OK.
5. Change the Virtual Cluster 1 and Virtual Cluster 2 device priorities for each cluster unit to the following:

<b>Host Name</b>	<b>Device Priority</b>	
	<b>Virtual Cluster 1</b>	<b>Virtual Cluster 2</b>
<b>620_ha_1</b>	200	100
<b>620_ha_2</b>	100	200

You can do this by editing the HA configurations of each cluster unit in the cluster members list and changing device priorities.

Since the device priority of Virtual Cluster 1 is highest for 620\_ha\_1 and since the root VDOM is in Virtual Cluster 1, all traffic for the root VDOM is processed by 620\_ha\_1.

Since the device priority of Virtual Cluster 2 is highest for 620\_ha\_2 and since the Eng\_vdm VDOM is in Virtual Cluster 2, all traffic for the Eng\_vdm VDOM is processed by 620\_ha\_2.

#### To view cluster status and verify the VDOM partitioning configuration

1. Log into the web-based manager.

2. Go to *System > Config > HA*.

The cluster members list should show the following:

- Virtual Cluster 1 contains the root VDOM.
- 620\_ha\_1 is the primary unit (master) for Virtual Cluster 1.
- Virtual Cluster 2 contains the Eng\_vdm VDOM.
- 620\_ha\_2 is the primary unit (master) for Virtual Cluster 2.

**Figure 199:**Example virtual clustering cluster members list

Virtual Cluster 1						<a href="#">View HA Statistics</a>
Virtual Domains: root						
	Cluster Member	Hostname	Role	Priority		
		620_ha_1	MASTER	128		
		620_ha_2	SLAVE	128		
Virtual Cluster 2						
Virtual Domains: Eng_vdm						
	Cluster Member	Hostname	Role	Priority		
		620_ha_2	MASTER	128		
		620_ha_1	SLAVE	128		

**To test the VDOM partitioning configuration**

You can do the following to confirm that traffic for the root VDOM is processed by 620\_ha\_1 and traffic for the Eng\_vdm is processed by 620\_ha\_2.

1. Log into the web-based manager by connecting to port2 using IP address 10.11.101.100.  
 You will log into 610\_ha\_1 because port2 is in the root VDOM and all traffic for this VDOM is processed by 610\_ha\_1. You can confirm that you have logged into 610\_ha\_1 by checking the HTML title displayed by your web browser. The title will include the 610\_ha\_1 host name. Also on the System Information dashboard widget displays the serial number of the 610\_ha\_1 FortiGate unit.
2. Log into the web-based manager by connecting to port6 using IP address 10.12.101.100.  
 You will log into 610\_ha\_2 because port6 is in the Eng\_vdm VDOM and all traffic for this VDOM is processed by 610\_ha\_2.
3. Add security policies to the root virtual domain that allows communication from the internal network to the Internet and connect to the Internet from the internal network.
4. Log into the web-based manager and go to *Config > System > HA* and select *View HA Statistics*.  
 The statistics display shows more active sessions, total packets, network utilization, and total bytes for the 620\_ha\_1 unit.
5. Add security policies to the Eng\_vdm virtual domain that allow communication from the engineering network to the Internet and connect to the Internet from the engineering network.

6. Log into the web-based manager and go to *Config > System > HA* and select *View HA Statistics*.

The statistics display shows more active sessions, total packets, network utilization, and total bytes for the 620\_ha\_2 unit.

## Configuring virtual clustering with two VDOMs and VDOM partitioning - CLI

These procedures assume you are starting with two FortiGate-620B units with factory default settings.

### To configure the FortiGate-620B units for HA operation

1. Power on the first FortiGate-620B unit and log into the CLI.
2. Change the host name for this FortiGate unit:

```
config system global
 set hostname 620_ha_1
end
```

3. Configure HA settings.

```
config system ha
 set mode a-p
 set group-name vexample2.com
 set password vHA_pass_2
end
```

The FortiGate unit negotiates to establish an HA cluster. You may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses”](#) on

page 1300). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

4. Display the HA configuration (optional).

```
get system ha
group-id : 0
group-name : vexample2.com
mode : a-p
password : *
hbdev : "port3" 50 "port4" 50
session-sync-dev :
route-ttl : 10
route-wait : 0
route-hold : 10
sync-config : enable
encryption : disable
authentication : disable
hb-interval : 2
hb-lost-threshold : 20
helo-holddown : 20
arps : 5
arps-interval : 8
session-pickup : disable
link-failed-signal : disable
uninterruptible-upgrade: enable
ha-mgmt-status : disable
ha-eth-type : 8890
hc-eth-type : 8891
l2ep-eth-type : 8893
subsecond : disable
vcluster2 : disable
vcluster-id : 1
override : disable
priority : 128
monitor :
pingserver-monitor-interface:
pingserver-failover-threshold: 0
pingserver-flip-timeout: 60
vdom : "root"
```

5. Power off the first FortiGate unit.

6. Repeat these steps for the second FortiGate unit.

Set the other FortiGate unit host name to:

```
config system global
 set hostname 620_ha_2
end
```

**To connect the cluster to the network**

1. Connect the port1 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the Internet.
2. Connect the port5 interfaces of 620\_ha\_1 and 620\_ha\_2 to switch connected to the Internet.  
You could use the same switch for port1 and port5.

3. Connect the port2 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the internal network.
4. Connect the port6 interfaces of 620\_ha\_1 and 620\_ha\_2 to a switch connected to the engineering network.
5. Connect the port3 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
6. Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
7. Power on the cluster units.  
 The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention.  
 When negotiation is complete you can continue.

### **To configure VDOM settings for the cluster**

1. Log into the CLI.
2. Enter the following command to enable multiple VDOMs for the cluster.  

```
config system global
 set vdom-admin enable
end
```
3. Log back into the CLI.
4. Enter the following command to add the Eng\_vdm VDOM:  

```
config vdom
 edit Eng_vdm
end
```
5. Edit the port5 interface, add it to the Eng\_vdm VDOM and configure other interface settings:  

```
config global
 config system interface
 edit port5
 set vdom Eng_vdm
 set alias Engineering_external
 set ip 172.20.12.143/24
 next
 edit port6
 set vdom Eng_vdm
 set alias Engineering_internal
 set ip 10.120.101.100/24
 end
 end
end
```

### To add a default route to each VDOM

1. Enter the following command to add default routes to the root and Eng\_vdm VDOMs.

```
config vdom
 edit root
 config router static
 edit 1
 set dst 0.0.0.0/0.0.0.0
 set gateway 172.20.120.2
 set device port1
 end
 next
 edit Eng_vdm
 config router static
 edit 1
 set dst 0.0.0.0/0.0.0.0
 set gateway 172.20.120.2
 set device port5
 end
 end
end
```

### To configure VDOM partitioning

1. Enter the `get system ha status` command to view cluster unit status:

For example, from the 620\_ha\_2 cluster unit CLI:

```
config global
 get system ha status
 Model: 620
 Mode: a-p
 Group: 0
 Debug: 0
 ses_pickup: disable
 Master:128 620_ha_2 FG600B3908600825 0
 Slave :128 620_ha_1 FG600B3908600705 1
 number of vcluster: 1
 vcluster 1: work 169.254.0.1
 Master:0 FG600B3908600825
 Slave :1 FG600B3908600705
```

This command output shows that VDOM partitioning has not been configured because only virtual cluster 1 is shown. The command output also shows that the 620\_ha\_2 is the primary unit for the cluster and for virtual cluster 1 because this cluster unit has the highest serial number

2. Enter the following commands to configure VDOM partitioning:

```
config global
 config system ha
 set vcluster2 enable
 config secondary-vcluster
 set vdom Eng_vdm
 end
 end
end
```



3. Enter the `get system ha status` command to view cluster unit status:

For example, from the `620_ha_2` cluster unit CLI:

```
config global
 get system ha status
 Model: 620
 Mode: a-p
 Group: 0
 Debug: 0
 ses_pickup: disable
 Master:128 620_ha_2 FG600B3908600825 0
 Slave :128 620_ha_1 FG600B3908600705 1
 number of vcluster: 2
 vcluster 1: work 169.254.0.1
 Master:0 FG600B3908600825
 Slave :1 FG600B3908600705
 vcluster 2: work 169.254.0.1
 Master:0 FG600B3908600825
 Slave :1 FG600B3908600705
```

This command output shows VDOM partitioning has been configured because both virtual cluster 1 and virtual cluster 2 are visible. However the configuration is not complete because `620_ha_2` is the primary unit for both virtual clusters. The command output shows this because under both `vcluster` entries the `Master` entry shows `FG600B3908600825`, which is the serial number of `620_ha_2`. As a result of this configuration, `620_ha_2` processes traffic for both VDOMs and `620_ha_1` does not process any traffic.

4. Change the Virtual Cluster 1 and Virtual Cluster 2 device priorities for each cluster unit so that `620_ha_1` processes virtual cluster 1 traffic and `620_ha_2` processes virtual cluster 2 traffic.

Since the root VDOM is in virtual cluster 1 and the `Eng_vdm` VDOM is in virtual cluster 2 the result of this configuration will be that `620_ha_1` will process all root VDOM traffic and

620\_ha\_2 will process all Eng\_vdm traffic. You make this happen by changing the cluster unit device priorities for each virtual cluster. You could use the following settings:

Device Priority		
Host Name	Virtual Cluster 1	Virtual Cluster 2
620_ha_1	200	100
620_ha_2	100	200

Since the device priority is not synchronized you can edit the device priorities of each virtual cluster on each FortiGate unit separately. To do this:

- Log into the CLI and note the FortiGate unit you have actually logged into (for example, by checking the host name displayed in the CLI prompt).
- Change the virtual cluster 1 and 2 device priorities for this cluster unit.
- Then use the `execute ha manage` command to log into the other cluster unit CLI and set its virtual cluster 1 and 2 device priorities.

Enter the following commands from the 620\_ha\_1 cluster unit CLI:

```
config global
 config system ha
 set priority 200
 config secondary-vcluster
 set priority 100
 end
 end
end
```

Enter the following commands from the 620\_ha\_2 cluster unit CLI:

```
config global
 config system ha
 set priority 100
 config secondary-vcluster
 set priority 200
 end
 end
end
```



The cluster may renegotiate during this step resulting in a temporary loss of connection to the CLI and a temporary service interruption.

Since the device priority of Virtual Cluster 1 is highest for 620\_ha\_1 and since the root VDOM is in Virtual Cluster 1, all traffic for the root VDOM is processed by 620\_ha\_1.

Since the device priority of Virtual Cluster 2 is highest for 620\_ha\_2 and since the Eng\_vdm VDOM is in Virtual Cluster 2, all traffic for the Eng\_vdm VDOM is processed by 620\_ha\_2.

## To verify the VDOM partitioning configuration

1. Log into the 620\_ha\_2 cluster unit CLI and enter the following command:

```
config global
 get system ha status
 Model: 620
 Mode: a-p
 Group: 0
 Debug: 0
 ses_pickup: disable
 Slave :100 620_ha_2 FG600B3908600825 0
 Master:200 620_ha_1 FG600B3908600705 1
 number of vcluster: 2
 vcluster 1: standby 169.254.0.2
 Slave :1 FG600B3908600825
 Master:0 FG600B3908600705
 vcluster 2: work 169.254.0.1
 Master:0 FG600B3908600825
 Slave :1 FG600B3908600705
```

The command output shows that 620\_ha\_1 is the primary unit for virtual cluster 1 (because the command output show the `Master` of virtual cluster 1 is the serial number of 620\_ha\_1) and that 620\_ha\_2 is the primary unit for virtual cluster 2.

If you enter the same command from the 620\_ha\_1 CLI the same information is displayed but in a different order. The command always displays the status of the cluster unit that you are logged into first.

```
config global
 get system ha status
 Model: 620
 Mode: a-p
 Group: 0
 Debug: 0
 ses_pickup: disable
 Master:200 620_ha_1 FG600B3908600705 1
 Slave :100 620_ha_2 FG600B3908600825 0
 number of vcluster: 2
 vcluster 1: work 169.254.0.2
 Master:0 FG600B3908600705
 Slave :1 FG600B3908600825
 vcluster 2: standby 169.254.0.1
 Slave :1 FG600B3908600705
 Master:0 FG600B3908600825
```

## To test the VDOM partitioning configuration

You can do the following to confirm that traffic for the root VDOM is processed by 620\_ha\_1 and traffic for the Eng\_vdm is processed by 620\_ha\_2. These steps assume the cluster is operating correctly.

1. Log into the CLI by connecting to port2 using IP address 10.11.101.100.

You will log into 610\_ha\_1 because port2 is in the root VDOM and all traffic for this VDOM is processed by 610\_ha\_1. You can confirm that you have logged into 610\_ha\_1 by checking

the host name in the CLI prompt. Also the `get system status` command displays the status of the 610\_ha\_1 cluster unit.

2. Log into the web-based manager or CLI by connecting to port6 using IP address 10.12.101.100.

You will log into 610\_ha\_2 because port6 is in the Eng\_vdm VDOM and all traffic for this VDOM is processed by 610\_ha\_2.

3. Add security policies to the root virtual domain that allow communication from the internal network to the Internet and connect to the Internet from the internal network.
4. Log into the web-based manager and go to *Config > System > HA* and select *View HA Statistics*.

The statistics display shows more active sessions, total packets, network utilization, and total bytes for the 620\_ha\_1 unit.

5. Add security policies to the Eng\_vdm virtual domain that allow communication from the engineering network to the Internet and connect to the Internet from the engineering network.

6. Log into the web-based manager and go to *Config > System > HA* and select *View HA Statistics*.

The statistics display shows more active sessions, total packets, network utilization, and total bytes for the 620\_ha\_2 unit.

## Example: inter-VDOM links in a virtual clustering configuration

In a virtual domain configuration you can use inter-VDOM links to route traffic between two virtual domains operating in a single FortiGate unit without using physical interfaces. Adding an inter-VDOM link has the affect of adding two interfaces to the FortiGate unit and routing traffic between the virtual domains using the inter-VDOM link interfaces.

In a virtual clustering configuration inter-VDOM links can only be made between virtual domains that are in the same virtual cluster. So, if you are planning on configuring inter-VDOM links in a virtual clustering configuration, you should make sure the virtual domains that you want to link are in the same virtual cluster.

For example, the following tables show an example virtual clustering configuration where each virtual cluster contains three virtual domains. In this configuration you can configure inter-VDOM links between root and vdom\_1 and between vdom\_2 and vdom\_3. But, you cannot configure inter-VDOM links between root and vdom\_2 or between vdom\_1 and vdom\_3 (and so on).

Virtual Domains	Hostname	
	FortiGate_A	FortiGate_B
root	Priority	Priority
vdom_1	200	100
	Role	Role
	Primary	Subordinate

Virtual Domains	Hostname	
	FortiGate_A	FortiGate_B
vdom_2	Priority	Priority
vdom_3	100	200
	Role	Role
	Subordinate	Primary

## Configuring inter-VDOM links in a virtual clustering configuration

Configuring inter-VDOM links in a virtual clustering configuration is very similar to configuring inter-VDOM links for a standalone FortiGate unit. The main difference the `config system vdom-link` command includes the `vcluster` keyword. The default setting for `vcluster` is `vcluster1`. So you only have to use the `vcluster` keyword if you are added an inter-VDOM link to virtual cluster 2.

### To add an inter-VDOM link to virtual cluster 1

This procedure describes how to create an inter-VDOM link to virtual cluster 1 that results in a link between the root and `vdom_1` virtual domains.



Inter-VDOM links are also called internal point-to-point interfaces.

- 1 Add an inter-VDOM link called `vc1link`.

```
config global
 config system vdom-link
 edit vc1link
 end
```

Adding the inter-VDOM link also adds two interfaces. In this example, these interfaces are called `vc1link0` and `vc1link1`. These interfaces appear in all CLI and web-based manager interface lists. These interfaces can only be added to virtual domains in virtual cluster 1.

2. Bind the `vc1link0` interface to the root virtual domain and bind the `vc1link1` interface to the `vdom_1` virtual domain.

```
config system interface
 edit vc1link0
 set vdom root
 next
 edit vc1link1
 set vdom vdom_1
 end
```

### To add an inter-VDOM link to virtual cluster 2

This procedure describes how to create an inter-VDOM link to virtual cluster 2 that results in a link between the `vdom_2` and `vdom_3` virtual domains.

- 1 Add an inter-VDOM link called `vc2link`.

```
config global
 config system vdom-link
 edit vc2link
 set vcluster vcluster2
 end
```

Adding the inter-VDOM link also adds two interfaces. In this example, these interfaces are called `vc2link0` and `vc2link1`. These interfaces appear in all CLI and web-based manager interface lists. These interfaces can only be added to virtual domains in virtual cluster 2.

2. Bind the `vc2link0` interface to the `vdom_2` virtual domain and bind the `vc2link1` interface to the `vdom_3` virtual domain.

```
config system interface
 edit vc2link0
 set vdom vdom_2
 next
 edit vc2link1
 set vdom vdom_3
 end
```

## Troubleshooting virtual clustering

Troubleshooting virtual clusters is similar to troubleshooting any cluster (see [“Troubleshooting HA clusters” on page 1212](#)). This section describes a few testing and troubleshooting techniques for virtual clustering.

### To test the VDOM partitioning configuration

You can do the following to confirm that traffic for different VDOMs will be distributed among both FortiGate units in the virtual cluster. These steps assume the cluster is otherwise operating correctly.

1. Log into the web-based manager or CLI using the IP addresses of interfaces in each VDOM. Confirm that you have logged into the FortiGate unit that should be processing traffic for that VDOM by checking the HTML title displayed by your web browser or the CLI prompt. Both of these should include the host name of the cluster unit that you have logged into. Also on the system Dashboard, the System Information widget displays the serial number of the FortiGate unit that you logged into. From the CLI the `get system status` command displays the status of the cluster unit that you logged into.
2. To verify that the correct cluster unit is processing traffic for a VDOM:
  - Add security policies to the VDOM that allow communication between the interfaces in the VDOM.
  - Optionally enable traffic logging and other monitoring for that VDOM and these security policies.
  - Start communication sessions that pass traffic through the VDOM.
  - Log into the web-based manager and go to *Config > System > HA* and select *View HA Statistics*. Verify that the statistics display shows more active sessions, total packets, network utilization, and total bytes for the unit that should be processing all traffic for the VDOM.
  - Optionally check traffic logging and the Top Sessions Widget for the FortiGate unit that should be processing traffic for that VDOM to verify that the traffic is being processed by this FortiGate unit.

# Full mesh HA

This chapter provides an introduction to full mesh HA and also contains general procedures and configuration examples that describe how to configure FortiGate full mesh HA.

The examples in this chapter include example values only. In most cases you will substitute your own values. The examples in this chapter also do not contain detailed descriptions of configuration parameters.

This chapter contains the following sections:

- [Full mesh HA overview](#)
- [Example: full mesh HA configuration](#)
- [Troubleshooting full mesh HA](#)

## Full mesh HA overview

When two or more FortiGate units are connected to a network in an HA cluster the reliability of the network is improved because the HA cluster replaces a single FortiGate unit as a single point of failure. With a cluster, a single FortiGate unit is replaced by a cluster of two or more FortiGate units.

However, even with a cluster, potential single points of failure remain. The interfaces of each cluster unit connect to a single switch and that switch provides a single connection to the network. If the switch fails or if the connection between the switch and the network fails service is interrupted to that network.


The HA cluster does improve the reliability of the network because switches are not as complex components as FortiGate units, so are less likely to fail. However, for even greater reliability, a configuration is required that includes redundant connections between the cluster the networks that it is connected to.

FortiGate models that support 802.3ad Aggregate or Redundant interfaces can be used to create a cluster configuration called full mesh HA. Full mesh HA is a method of reducing the number of single points of failure on a network that includes an HA cluster.

This redundant configuration can be achieved using FortiGate 802.3ad Aggregate or Redundant interfaces and a full mesh HA configuration. In a full mesh HA configuration, you connect an HA cluster consisting of two or more FortiGate units to the network using 802.3ad Aggregate or Redundant interfaces and redundant switches. Each 802.3ad Aggregate or Redundant interface is connected to two switches and both of these switches are connected to the network.

The resulting full mesh configuration, an example is shown in [Figure 200](#), includes redundant connections between all network components. If any single component or any single connection fails, traffic automatically switches to the redundant component and connection and traffic flow resumes.

**Figure 200:**Single points of failure in a standalone and HA network configuration



### Full mesh HA and redundant heartbeat interfaces

A full mesh HA configuration also includes redundant HA heartbeat interfaces. At least two heartbeat interfaces should be selected in the HA configuration and both sets of HA heartbeat interfaces should be connected. The HA heartbeat interfaces do not have to be configured as redundant interfaces because the FGCP handles failover between heartbeat interfaces.

### Full mesh HA, redundant interfaces and 802.3ad aggregate interfaces

Full mesh HA is supported for both redundant interfaces and 802.3ad aggregate interfaces. In most cases you would simply use redundant interfaces. However, if your switches support 802.3ad aggregate interfaces and split multi-trunking you can use aggregate interfaces in place of redundant interfaces for full mesh HA. One advantage of using aggregate interfaces is that all



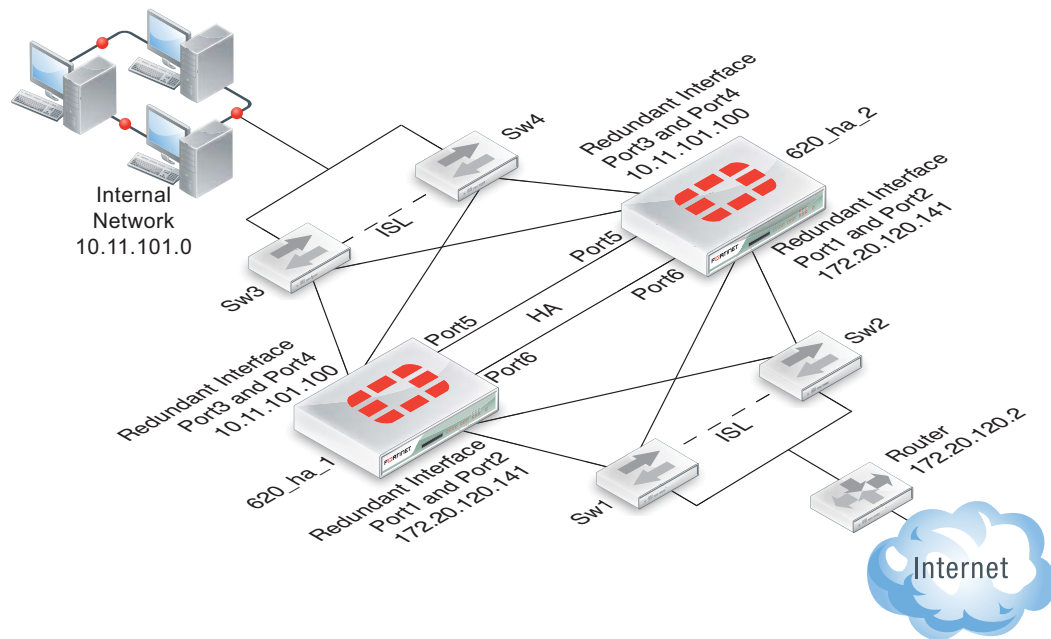
of the physical interfaces in the aggregate interface can send and receive packets. As a result, using aggregate interfaces may increase the bandwidth capacity of the cluster.

Usually redundant and aggregate interfaces consist of two physical interfaces. However, you can add more than two physical interfaces to a redundant or aggregate interface. Adding more interfaces can increase redundancy protection. Adding more interfaces can also increase bandwidth capacity if you are using 802.3ad aggregate interfaces.

## Example: full mesh HA configuration

Figure 200 shows a full mesh HA configuration with a cluster of two FortiGate-620b units. This section describes the FortiGate configuration settings and network components required for a full mesh HA configuration. This section also contains example steps for setting up this full mesh HA configuration. The procedures in this section describe one of many possible sequences of steps for configuring full mesh HA. As you become more experienced with FortiOS, HA, and full mesh HA you may choose to use a different sequence of configuration steps.

Figure 201: Full Mesh HA configuration



For simplicity these procedures assume that you are starting with two new FortiGate units set to the factory default configuration. However, starting from the default configuration is not a requirement for a successful HA deployment. FortiGate HA is flexible enough to support a successful configuration from many different starting points.

These procedures describe how to configure a cluster operating in NAT/Route mode because NAT/Route is the default FortiGate operating mode. However, the steps are the same if the cluster operates in Transparent mode. You can either switch the cluster units to operate in Transparent mode before beginning these procedures, or you can switch the cluster to operate in Transparent mode after HA is configured and the cluster is connected and operating.

## FortiGate-620B full mesh HA configuration

The two FortiGate-620B units (620\_ha\_1 and 620\_ha\_2) can be operating in NAT/Route or Transparent mode. Aside from the standard HA settings, the FortiGate-620B configuration includes the following:

- The port5 and port6 interfaces configured as heartbeat interfaces. A full mesh HA configuration also includes redundant HA heartbeat interfaces.
- The port1 and port2 interfaces added to a redundant interface. Port1 is the active physical interface in this redundant interface. To make the port1 interface the active physical interface it should appear above the port2 interface in the redundant interface configuration.
- The port3 and port4 interfaces added to a redundant interface. Port3 is the active physical interface in this redundant interface. To make the port3 interface the active physical interface it should appear above the port4 interface in the redundant interface configuration.

## Full mesh switch configuration

The following redundant switch configuration is required:

- Two redundant switches (Sw3 and Sw4) connected to the internal network. Establish an interswitch-link (ISL) between them.
- Two redundant switches (Sw1 and Sw2) connected to the Internet. Establish an interswitch-link (ISL) between them.

## Full mesh network connections

Make the following physical network connections for 620\_ha\_1:

- Port1 to Sw1 (active)
- Port2 to Sw2 (inactive)
- Port3 to Sw3 (active)
- Port4 to Sw4 (inactive)

Make the following physical network connections for 620\_ha\_2:

- Port1 to Sw2 (active)
- Port2 to Sw1 (inactive)
- Port3 to Sw4 (active)
- Port4 to Sw3 (inactive)

## How packets travel from the internal network through the full mesh cluster and to the Internet

If the cluster is operating in active-passive mode and 620\_ha\_2 is the primary unit, all packets take the following path from the internal network to the internet:

1. From the internal network to Sw4. Sw4 is the active connection to 620\_ha\_2; which is the primary unit. The primary unit receives all packets.
2. From Sw4 to the 620\_ha\_2 port3 interface. Active connection between Sw4 and 620\_ha\_2. Port3 is the active member of the redundant interface.
3. From 620\_ha\_2 port3 to 620\_ha\_2 port1. Active connection between 620\_ha\_2 and Sw2. Port1 is the active member of the redundant interface.
4. From Sw2 to the external router and the Internet.

## Configuring FortiGate-620B units for HA operation - web-based manager

Each FortiGate-620B unit in the cluster must have the same HA configuration.

### To configure the FortiGate-620B units for HA operation

1. Connect to the web-based manager of one of the FortiGate-620B units.
2. On the *System Information* dashboard widget, beside *Host Name* select *Change*.
3. Enter a new Host Name for this FortiGate unit.

---

<b>New Name</b>	620_ha_1
-----------------	----------

---

4. Go to *System > Config > HA* and change the following settings.

---

<b>Mode</b>	Active-Active
-------------	---------------

---

<b>Group Name</b>	Rexample1.com
-------------------	---------------

---

<b>Password</b>	RHA_pass_1
-----------------	------------

---

#### Heartbeat Interface

---

	<b>Enable</b>	<b>Priority</b>
<b>port5</b>	Select	50
<b>port6</b>	Select	50

---

5. Select OK.

The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 1300](#)). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10

- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

6. Power off the first FortiGate unit.
7. Repeat these steps for the second FortiGate unit.  
Set the second FortiGate unit host name to:

---

<b>New Name</b>	620_ha_2
-----------------	----------

---

### To connect the cluster to your network

1. Make the following physical network connections for 620\_ha\_1:
  - Port1 to Sw1 (active)
  - Port2 to Sw2 (inactive)
  - Port3 to Sw3 (active)
  - Port4 to Sw4 (inactive)
2. Make the following physical network connections for 620\_ha\_2:
  - Port1 to Sw2 (active)
  - Port2 to Sw1 (inactive)
  - Port3 to Sw4 (active)
  - Port4 to Sw3 (inactive)
3. Connect Sw3 and Sw4 to the internal network.
4. Connect Sw1 and Sw2 to the external router.
5. Enable ISL communication between Sw1 and Sw2 and between Sw3 and Sw4.
6. Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention.

When negotiation is complete the cluster is ready to be configured for your network.

## To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.

1. View the system dashboard.

The System Information dashboard widget shows the *Cluster Name* (Rexample1.com) and the host names and serial numbers of the *Cluster Members*. The Unit Operation widget shows multiple cluster units.

2. Go to *System > Config > HA* to view the cluster members list.

The list shows two cluster units, their host names, their roles in the cluster, and their priorities. You can use this list to confirm that the cluster is operating normally.

## To troubleshoot the cluster configuration

If the cluster members list and the dashboard does not display information for both cluster units the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 1212](#) to troubleshoot the cluster.

## To add basic configuration settings and the redundant interfaces

Use the following steps to add a few basic configuration settings.

1. Log into the cluster web-based manager.
2. Go to *System > Admin > Administrators*.
3. For *admin*, select the *Change Password* icon
4. Enter and confirm a new password.
5. Select OK.
6. Go to *Router > Static > Static Routes* and temporarily delete the default route.  
You cannot add an interface to a redundant interface if any settings (such as the default route) are configured for it.
7. Go to *System > Network > Interfaces* and select *Create New* and configure the redundant interface to connect to the Internet.

<b>Name</b>	Port1_Port2
<b>Type</b>	Redundant
<b>Physical Interface Members</b>	
<b>Selected Interfaces</b>	port1, port2
<b>IP/Netmask</b>	172.20.120.141/24

8. Select OK.
9. Select *Create New* and configure the redundant interface to connect to the internal network.

<b>Name</b>	Port3_Port4
<b>Type</b>	Redundant
<b>Physical Interface Members</b>	
<b>Selected Interfaces</b>	port3, port4

<b>IP/Netmask</b>	10.11.101.100/24
<b>Administrative Access</b>	HTTPS, PING, SSH

**10.**Select OK.

The virtual MAC addresses of the FortiGate-620B interfaces change to the following. Notice that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

**11.**Go to *Router > Static > Static Routes*.

**12.**Add the default route.

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Gateway</b>	172.20.120.2
<b>Device</b>	Port1_Port2
<b>Distance</b>	10

**13.**Select OK.

**To configure HA port monitoring for the redundant interfaces**

1. Go to *System > Config > HA*.
2. In the cluster members list, edit the primary unit.

3. Configure the following port monitoring for the redundant interfaces:

Port Monitor	
Port1_Port2	Select
Port3_Port4	Select

4. Select OK.

## Configuring FortiGate-620B units for HA operation - CLI

Each FortiGate-620B unit in the cluster must have the same HA configuration. Use the following procedure to configure the FortiGate-620B units for HA operation.

### To configure the FortiGate-620B units for HA operation

1. Connect to the CLI of one of the FortiGate-620B units.
2. Enter a new Host Name for this FortiGate unit.

```
config system global
 set hostname 620_ha_1
end
```

3. Configure HA settings.

```
config system ha
 set mode a-a
 set group-name Rexample1.com
 set password RHA_pass_1
 set hbdev port5 50 port6 50
end
```

The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 1300](#)). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e

- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

4. Power off the first FortiGate unit.
5. Repeat these steps for the second FortiGate unit.

Set the other FortiGate unit host name to:

```
config system global
 set hostname 620_ha_2
end
```

### To connect the cluster to your network

1. Make the following physical network connections for 620\_ha\_1:
  - Port1 to Sw1 (active)
  - Port2 to Sw2 (inactive)
  - Port3 to Sw3 (active)
  - Port4 to Sw4 (inactive)
2. Make the following physical network connections for 620\_ha\_2:
  - Port1 to Sw2 (active)
  - Port2 to Sw1 (inactive)
  - Port3 to Sw4 (active)
  - Port4 to Sw3 (inactive)
3. Connect Sw3 and Sw4 to the internal network.
4. Connect Sw1 and Sw2 to the external router.
5. Enable ISL communication between Sw1 and Sw2 and between Sw3 and Sw4.
6. Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention.

When negotiation is complete the cluster is ready to be configured for your network.



## To view cluster status

Use the following steps to view cluster status from the CLI.

1. Log into the CLI.
2. Enter `get system status` to verify the HA status of the cluster unit that you logged into.  
If the command output includes `Current HA mode: a-a, master`, the cluster units are operating as a cluster and you have connected to the primary unit.  
If the command output includes `Current HA mode: a-a, backup`, you have connected to a subordinate unit.  
If the command output includes `Current HA mode: standalone` the cluster unit is not operating in HA mode.
3. Enter the following command to confirm the HA configuration of the cluster:

```
get system ha status
 Model: 620
 Mode: a-a
 Group: 0
 Debug: 0
 ses_pickup: disable
 Master:128 620_ha_2 FG600B3908600825 0
 Slave :128 620_ha_1 FG600B3908600705 1
 number of vcluster: 1
 vcluster 1: work 169.254.0.1
 Master:0 FG600B3908600825
 Slave :1 FG600B3908600705
```

The command output shows both cluster units, their host names, their roles in the cluster, and their priorities. You can use this command to confirm that the cluster is operating normally. For example, if the command shows only one cluster unit then the other unit has left the cluster for some reason.

4. Use the `execute ha manage` command to connect to the other cluster unit's CLI and use these commands to verify cluster status.

## To troubleshoot the cluster configuration

If the cluster members list and the dashboard does not display information for both cluster units the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 1212](#) to troubleshoot the cluster.

## To add basic configuration settings and the redundant interfaces

Use the following steps to add a few basic configuration settings. Some steps use the CLI and some the web-based manager.

1. Log into the cluster CLI.
2. Add a password for the admin administrative account.

```
config system admin
 edit admin
 set password <password_str>
 end
```

3. Temporarily delete the default route.

You cannot add an interface to a redundant interface if any settings (such as the default route) are configured for it.

```
config router static
 delete 1
end
```

4. Go to *System > Network > Interface* and select *Create New* to add the redundant interface to connect to the Internet.

5. Add the redundant interface to connect to the Internet.

```
config sysetem interface
 edit Port1_Port2
 set type redundant
 set member port1 port2
 end
```

6. Add the redundant interface to connect to the internal network.

```
config sysetem interface
 edit Port3_Port4
 set type redundant
 set member port3 port4
 end
```

The virtual MAC addresses of the FortiGate-620B interfaces change to the following. Note that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

7. Go to *Router > Static*.

8. Add the default route.

```
config router static
 edit 1
 set dst 0.0.0.0 0.0.0.0
 set gateway 172.20.120.2
 set device Port1_Port2
 end
```

**To configure HA port monitoring for the redundant interfaces**

1. Enter the following command to configure port monitoring for the redundant interfaces:

```
config system ha
 set monitor Port1_Port2 Port3_Port4
end
```

## Troubleshooting full mesh HA

Troubleshooting full mesh HA clusters is similar to troubleshooting any cluster (see [“Troubleshooting HA clusters” on page 1212](#) or [“Troubleshooting virtual clustering” on page 1238](#)). The configuration and operation of a full mesh HA cluster is very similar to the configuration and operation of a standard cluster. The only differences relate to the configuration, connection, and operation of the redundant interfaces and redundant switches.

- Make sure the redundant interfaces and switches are connected correctly. With so many connections it is possible to make mistakes or for cables to become disconnected.
- Confirm that the configuration of the cluster unit 802.3ad Aggregate or Redundant interfaces is correct according to the configuration procedures in this chapter.
- In some configurations with some switch hardware, MAC-learning delays on the inter-switch links on the surrounding topologies may occur. The delays occur if the gratuitous ARP packets sent by the cluster after a failover are delayed by the switches before being sent across the inter-switch link. If this happens the surrounding topologies may be delayed in recognizing the failover and will keep sending packets to the MAC address of the failed primary unit resulting in lost traffic. Resolving this problem may require changing the configuration of the switch or replacing them with switch hardware that does not delay the gratuitous ARP packets.

# Operating a cluster

With some exceptions, you can operate a cluster in much the same way as you operate a standalone FortiGate unit. This chapter describes those exceptions and also the similarities involved in operating a cluster instead of a standalone FortiGate unit.

This chapter contains the following sections:

- [Operating a cluster](#)
- [Operating a virtual cluster](#)
- [Managing individual cluster units using a reserved management interface](#)
- [The primary unit acts as a router for subordinate unit management traffic](#)
- [Clusters and FortiGuard services](#)
- [Clusters and logging](#)
- [Clusters and SNMP](#)
- [Clusters and file quarantine](#)
- [Cluster members list](#)
- [Virtual cluster members list](#)
- [Viewing HA statistics](#)
- [Changing the HA configuration of an operating cluster](#)
- [Changing the HA configuration of an operating virtual cluster](#)
- [Changing the subordinate unit host name and device priority](#)
- [Upgrading cluster firmware](#)
- [Downgrading cluster firmware](#)
- [Backing up and restoring the cluster configuration](#)
- [Monitoring cluster units for failover](#)
- [Viewing cluster status from the CLI](#)
- [Disconnecting a cluster unit from a cluster](#)
- [Adding a disconnected FortiGate unit back to its cluster](#)

## Operating a cluster

The configurations of all of the FortiGate units in a cluster are synchronized so that the cluster units can simulate a single FortiGate unit. Because of this synchronization, you manage the HA cluster instead of managing the individual cluster units. You manage the cluster by connecting to the web-based manager using any cluster interface configured for HTTPS or HTTP administrative access. You can also manage the cluster by connecting to the CLI using any cluster interface configured for SSH or telnet administrative access.

The cluster web-based manager dashboard displays the cluster name, the host name and serial number of each cluster member, and also shows the role of each unit in the cluster. The roles can be master (primary unit) and slave (subordinate units). The dashboard also displays a cluster unit front panel illustration.

**Figure 202:**Example cluster web-based manager dashboard

System Information		
Cluster Name	FGT-HA	
Cluster Members	620_ha_2/FG600B3908600825 (Master)	
	620_ha_1/FG600B3908600705 (Slave)	
Serial Number	FG600B3908600825	
Operation Mode	NAT [Change]	
HA Status	Active-Passive [Configure]	
System Time	Wed Oct 13 11:43:58 2010 [Change]	
Firmware Version	v4.0.build0395.101005 (Interim) [Update]	
System Configuration	Last Backup: N/A [Backup] [Restore]	
Current Administrator	admin [Change Password] / 1 in Total [Details]	
Uptime	5 day(s) 22 hour(s) 55 min(s)	
Virtual Domain	Disabled [Enable]	

**Unit Operation**

FortiGate 620B

FortiAnalyzer FortiManager FortiClient

Reboot Shutdown

**Log and Archive Statistics (Since 2010-10-07 12:49:11)**

DLP Archive -- Average 0 B per day since last reset

You can also go to *System > Config > HA* to view the cluster members list. This includes status information for each cluster unit. You can also use the cluster members list for a number of cluster management functions including changing the HA configuration of an operating cluster, changing the host name and device priority of a subordinate unit, and disconnecting a cluster unit from a cluster. See [“Cluster members list” on page 1268](#).

You can use log messages to view information about the status of the cluster. See [“Viewing and managing log messages for individual cluster units” on page 1261](#). You can use SNMP to manage the cluster by configuring a cluster interface for SNMP administrative access. Using an SNMP manager you can get cluster configuration information and receive traps.

You can configure a reserved management interface to manage individual cluster units. You can use this interface to access the web-based manager or CLI and to configure SNMP management for individual cluster units. See [“Managing individual cluster units using a reserved management interface” on page 1254](#).

You can manage individual cluster units by using SSH, telnet, or the CLI console on the web-based manager dashboard to connect to the CLI of the cluster. From the CLI you can use the `execute ha manage` command to connect to the CLI of any unit in the cluster.

You can also manage individual cluster units by using a null-modem cable to connect to any cluster unit CLI. From there you can use the `execute ha manage` command to connect to the CLI of each unit in the cluster.

## Operating a virtual cluster

Managing a virtual cluster is very similar to managing a cluster that does not contain multiple virtual domains. Most of the information in this chapter applies to managing both kinds of clusters. This section describes what is different when managing a virtual cluster.

If virtual domains are enabled, the cluster web-based manager dashboard displays the cluster name and the role of each cluster unit in virtual cluster 1 and virtual cluster 2.

**Figure 203:**Example virtual clustering web-based manager dashboard

The screenshot displays the FortiGate web-based manager dashboard. The left panel, titled 'System Information', shows details for a FortiGate HA cluster. The right panel, titled 'Unit Operation', shows a graphical representation of the cluster units and their status. Below the graphical representation are buttons for 'Reboot' and 'Shutdown'. At the bottom, there is a 'Log and Archive Statistics' section showing 'DLP Archive' statistics.

System Information		
Cluster Name	FGT-HA	
Virtual Cluster 1	620_ha_2/FG600B3908600825	(Master)
Virtual Cluster 2	620_ha_2/FG600B3908600825	(Master)
Serial Number	FG600B3908600825	
HA Status	Active-Passive [Configure]	
System Time	Wed Oct 13 11:41:40 2010 [Change]	
Firmware Version	v4.0_build0395,101005 (Interim) [Update]	
System Configuration	Last Backup: N/A [Backup] [Restore]	
Current Administrator	admin [Change Password] /1 in Total [Details]	
Uptime	5 day(s) 22 hour(s) 53 min(s)	
Virtual Domain	Enabled [Disable]	

**Unit Operation**

FortiAnalyzer FortiManager

FortiGate 620B

FortiClient

Reboot Shutdown

**Log and Archive Statistics (Since 2010-10-07 12:49:11)**

DLP Archive -- Average 0 B per day since last reset

The configuration and maintenance options that you have when you connect to a virtual cluster web-based manager or CLI depend on the virtual domain that you connect to and the administrator account that you use to connect.

If you connect to a cluster as the administrator of a virtual domain, you connect directly to the virtual domain. Since HA virtual clustering is a global configuration, virtual domain administrators cannot see HA configuration options. However, virtual domain administrators see the host name of the cluster unit that they are connecting to on the web browser title bar or CLI prompt. This host name is the host name of the primary unit for the virtual domain. Also, when viewing log messages by going to *Log & Report > Log Access* virtual domain administrator can select to view log messages for either of the cluster units.

If you connect to a virtual cluster as the admin administrator you connect to the global web-based manager or CLI. Even so, you are connecting to an interface and to the virtual domain that the interface has been added to. The virtual domain that you connect to does not make a difference for most configuration and maintenance operations. However, there are a few exceptions. You connect to the FortiGate unit that functions as the primary unit for the virtual domain. So the host name displayed on the web browser title bar and on the CLI is the host name of this primary unit.

## Managing individual cluster units using a reserved management interface

You can provide direct management access to all cluster units by reserving a management interface as part of the HA configuration. Once this management interface is reserved, you can configure a different IP address, administrative access and other interface settings for this interface for each cluster unit. Then by connecting this interface of each cluster unit to your network you can manage each cluster unit separately from a different IP address. Configuration changes to the reserved management interface are not synchronized to other cluster units.

The reserved management interface provides direct management access to each cluster unit and gives each cluster unit a different identity on your network. This simplifies using external services, such as SNMP, to separately monitor and manage each cluster unit.



The reserved management interface is not assigned an HA virtual MAC address like other cluster interfaces. Instead the reserved management interface retains the permanent hardware address of the physical interface unless you change it using the `config system interface` command.

The reserved management interface and IP address should not be used for managing a cluster using FortiManager. To correctly manage a FortiGate HA cluster with FortiManager use the IP address of one of the cluster unit interfaces.

If you enable SNMP administrative access for the reserved management interface you can use SNMP to monitor each cluster unit using the reserved management interface IP address. To monitor each cluster unit using SNMP, just add the IP address of each cluster unit's reserved management interface to the SNMP server configuration. You must also enable direct management of cluster members in the cluster SNMP configuration.

If you enable HTTPS or HTTP administrative access for the reserved management interfaces you can connect to the web-based manager of each cluster unit. Any configuration changes made to any of the cluster units is automatically synchronized to all cluster units. From the subordinate units the web-based manager has the same features as the primary unit except that unit-specific information is displayed for the subordinate unit, for example:

- The *Dashboard System Information* widget displays the subordinate unit serial number but also displays the same information about the cluster as the primary unit
- On the Cluster members list (go to *System > Config > HA*) you can change the HA configuration of the subordinate unit that you are logged into. For the primary unit and other subordinate units you can change only the host name and device priority.
- Log Access displays the logs of the subordinate that you are logged into first, You use the HA Cluster list to view the log messages of other cluster units including the primary unit.


If you enable SSH or TELNET administrative access for the reserved management interfaces you can connect to the CLI of each cluster unit. The CLI prompt contains the host name of the cluster unit that you have connected to. Any configuration changes made to any of the cluster units is automatically synchronized to all cluster units. You can also use the `execute ha manage` command to connect to other cluster unit CLIs.

The reserved management interface is available in NAT/Route and in Transparent mode. It is also available if the cluster is operating with multiple VDOMs. In Transparent mode you cannot normally add an IP address to an interface. However, you can add an IP address to the reserved management interface.

## Configuring the reserved management interface and SNMP remote management of individual cluster units

This example describes how to configure SNMP remote management of individual cluster units using the HA reserved management interface. The configuration consists of two FortiGate-620B units already operating as a cluster. In the example, the port8 interface of each cluster unit is connected to the internal network using the switch and configured as the reserved management interface.

**Figure 204:**SNMP remote management of individual cluster units



#### To configure the reserved management interface - web-based manager

1. Go to *System > Config > HA*.
2. Edit the primary unit.
3. Select *Reserve Management Port for Cluster Member* and select port8.
4. Select OK.

#### To configure the reserved management interface - CLI

From the CLI you can also configure a default route that is only used by the reserved management interface.

1. Log into the CLI of any cluster unit.
2. Enter the following command to enable the reserved management interface, set port8 as the reserved interface, and add a default route of 10.11.101.100 for the reserved management interface.

```
config system ha
 set ha-mgmt-status enable
 set ha-mgmt-interface port8
 set ha-mgmt-interface-gateway 10.11.101.100
end
```

The reserved management interface default route is not synchronized to other cluster units.

#### To change the primary unit reserved management interface configuration - web-based manager

You can change the IP address of the primary unit reserved management interface from the primary unit web-based manager. Configuration changes to the reserved management interface are not synchronized to other cluster units.

1. From a PC on the internal network, browse to <http://10.11.101.100> and log into the cluster web-based manager.

This logs you into the primary unit web-based manager.

You can identify the primary unit from its serial number or host name that appears on the System Information dashboard widget.



2. Go to *System > Network > Interfaces* and edit the port8 interface as follows:

<b>Alias</b>	primary_reserved
<b>IP/Netmask</b>	10.11.101.101/24
<b>Administrative Access</b>	Ping, SSH, HTTPS, SNMP

3. Select OK.

You can now log into the primary unit web-based manager by browsing to <https://10.11.101.101>. You can also log into this primary unit CLI by using an SSH client to connect to 10.11.101.101.

### To change subordinate unit reserved management interface configuration - CLI

At this point you cannot connect to the subordinate unit reserved management interface because it does not have an IP address. Instead, this procedure describes connecting to the primary unit CLI and using the `execute ha manage` command to connect to subordinate unit CLI to change the port8 interface. You can also use a serial connection to the cluster unit CLI. Configuration changes to the reserved management interface are not synchronized to other cluster units.

1. Connect to the primary unit CLI and use the `execute ha manage` command to connect to a subordinate unit CLI.

You can identify the subordinate unit from its serial number or host name. The host name appears in the CLI prompt.

2. Enter the following command to change the port8 IP address to 10.11.101.102 and set management access to HTTPS, ping, SSH, and SNMP.

```
config system interface
 edit port8
 set ip 10.11.101.102/24
 set allowaccess https ping ssh snmp
 end
```

You can now log into the subordinate unit web-based manager by browsing to <https://10.11.101.102>. You can also log into this subordinate unit CLI by using an SSH client to connect to 10.11.101.102.

### To configure the cluster for SNMP management using the reserved management interfaces - CLI

This procedure describes how to configure the cluster to allow the SNMP server to get status information from the primary unit and the subordinate unit. The SNMP configuration is synchronized to all cluster units. To support using the reserved management interfaces, you must add at least one HA direct management host to an SNMP community. If your SNMP configuration includes SNMP users with user names and passwords you must also enable HA direct management for SNMP users.

1. Enter the following command to add an SNMP community called `Community` and add a host to the community for the reserved management interface of each cluster unit. The host includes the IP address of the SNMP server (10.11.101.20).

```
config system snmp community
 edit 1
 set name Community
 config hosts
 edit 1
 set ha-direct enable
 set ip 10.11.101.20
 end
 end
end
```

2. Enter the following command to add an SNMP user for the reserved management interface.

```
config system snmp user
 edit 1
 set ha-direct enable
 set notify-hosts 10.11.101.20
 end
```

Configure other settings as required.

### **To get CPU, memory, and network usage of each cluster unit using the reserved management IP addresses**

From the command line of an SNMP manager, you can use the following SNMP commands to get CPU, memory and network usage information for each cluster unit. In the examples, the community name is `Community`. The commands use the MIB field names and OIDs listed in [Table 57 on page 1265](#).

Enter the following commands to get CPU, memory and network usage information for the primary unit with reserved management IP address 10.11.101.101 using the MIB fields:

```
snmpget -v2c -c Community 10.11.101.101 fgHaStatsCpuUsage
snmpget -v2c -c Community 10.11.101.101 fgHaStatsMemUsage
snmpget -v2c -c Community 10.11.101.101 fgHaStatsNetUsage
```

Enter the following commands to get CPU, memory and network usage information for the primary unit with reserved management IP address 10.11.101.101 using the OIDs:

```
snmpget -v2c -c Community 10.11.101.101
1.3.6.1.4.1.12356.101.13.2.1.1.3.1
snmpget -v2c -c Community 10.11.101.101
1.3.6.1.4.1.12356.101.13.2.1.1.4.1
snmpget -v2c -c Community 10.11.101.101
1.3.6.1.4.1.12356.101.13.2.1.1.5.1
```

Enter the following commands to get CPU, memory and network usage information for the subordinate unit with reserved management IP address 10.11.101.102 using the MIB fields:

```
snmpget -v2c -c Community 10.11.101.102 fgHaStatsCpuUsage
snmpget -v2c -c Community 10.11.101.102 fgHaStatsMemUsage
snmpget -v2c -c Community 10.11.101.102 fgHaStatsNetUsage
```

Enter the following commands to get CPU, memory and network usage information for the subordinate unit with reserved management IP address 10.11.101.102 using the OIDs:

```
snmpget -v2c -c Community 10.11.101.102
1.3.6.1.4.1.12356.101.13.2.1.1.3.1
snmpget -v2c -c Community 10.11.101.102
1.3.6.1.4.1.12356.101.13.2.1.1.4.1
snmpget -v2c -c Community 10.11.101.102
1.3.6.1.4.1.12356.101.13.2.1.1.5.1
```

## The primary unit acts as a router for subordinate unit management traffic

HA uses routing and inter-VDOM links to route subordinate unit management traffic through the primary unit to the network. Similar to a standalone FortiGate unit, subordinate units may generate their own management traffic, including:

- DNS queries.
- FortiGuard Web Filtering rating requests.
- Log messages to be sent to a FortiAnalyzer unit, to a syslog server, or to the FortiGuard Analysis and Management Service.
- Log file uploads to a FortiAnalyzer unit.
- Quarantine file uploads to a FortiAnalyzer unit.
- SNMP traps.
- Communication with remote authentication servers (RADIUS, LDAP, TACACS+ and so on)

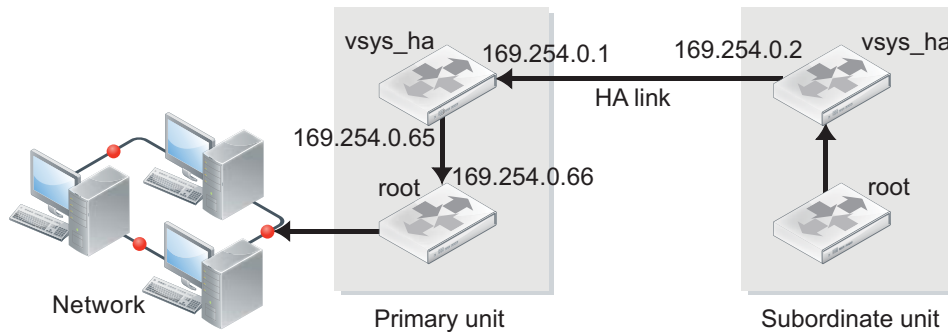
Subordinate units send this management traffic over the HA heartbeat link to the primary unit. The primary unit forwards the management traffic to its destination. The primary unit also routes replies back to the subordinate unit in the same way.

HA uses a hidden VDOM called `vsys_ha` for HA operations. The `vsys_ha` VDOM includes the HA heartbeat interfaces, and all communication over the HA heartbeat link goes through the `vsys_ha` VDOM. To provide communication from a subordinate unit to the network, HA adds hidden inter-VDOM links between the primary unit management VDOM and the primary unit `vsys_ha` VDOM. By default, `root` is the management VDOM.

Management traffic from the subordinate unit originates in the subordinate unit `vsys_ha` VDOM. The `vsys_ha` VDOM routes the management traffic over the HA heartbeat link to the primary unit `vsys_ha` VDOM. This management traffic is then routed to the primary unit management VDOM and from there out onto the network.

DNS queries and FortiGuard Web Filtering and Email Filter requests are still handled by the HA proxy so the primary unit and subordinate units share the same DNS query cache and the same FortiGuard Web Filtering and Email Filter cache. In a virtual clustering configuration, the cluster unit that is the primary unit for the management virtual domain maintains the FortiGuard Web Filtering, Email Filtering, and DNS query cache.

**Figure 205:**Subordinate unit management traffic path



## Cluster communication with RADIUS and LDAP servers

In an active-passive cluster, only the primary unit processes traffic, so the primary unit communicates with RADIUS or LDAP servers. In a cluster that is operating in active-active mode, subordinate units send RADIUS and LDAP requests to the primary unit over the HA heartbeat link and the primary unit routes them to their destination. The primary unit relays the responses back to the subordinate unit.

## Clusters and FortiGuard services

This section describes how various FortiGate HA clustering configurations communicate with the FDN.

In an operating cluster, the primary unit communicates directly with the FortiGuard Distribution Network (FDN). Subordinate units also communicate directly with the FDN but as described in [“The primary unit acts as a router for subordinate unit management traffic” on page 1259](#), all communication between subordinate units and the FDN is routed through the primary unit.

You must register and licence all of the units in a cluster for all required FortiGuard services, both because all cluster units communicate with the FDN and because any cluster unit could potentially become the primary unit.

### FortiGuard and active-passive clusters

For an active-passive cluster, only the primary unit processes traffic. Even so, all cluster units communicate with the FDN. Only the primary unit sends FortiGuard Web Filtering and Antispam requests to the FDN. All cluster units receive FortiGuard Antivirus, IPS, and application control updates from the FDN.

In an active-passive cluster the FortiGuard Web Filter and Email Filter caches are located on the primary unit in the same way as for a standalone FortiGate unit. The caches are not shared among cluster units so after a failover the new primary unit must build up new caches.

In an active-passive cluster all cluster units also communicate with the FortiGuard Analysis and Management Service (FAMS).

### FortiGuard and active-active clusters

For an active-active cluster, both the primary unit and the subordinate units process traffic. Communication between the cluster units and the FDN is the same as for active-passive clusters with the following exception.

Because the subordinate units process traffic, they may also be making FortiGuard Web Filtering and Email Filter requests. The primary unit receives all such requests from the subordinate units and relays them to the FDN and then relays the FDN responses back to the subordinate units. The FortiGuard Web Filtering and Email Filtering URL caches are maintained on the primary unit. The primary unit caches are used for primary and subordinate unit requests.

## FortiGuard and virtual clustering

For a virtual clustering configuration the management virtual domain of each cluster unit communicates with the FDN. The cluster unit that is the primary unit for the management virtual domain maintains the FortiGuard Web Filtering and Email Filtering caches. All FortiGuard Web Filtering and Email Filtering requests are proxied by the management VDOM of the cluster unit that is the primary unit for the management virtual domain.

## Clusters and logging

This section describes the log messages that provide information about how HA is functioning, how to view and manage logs for each unit in a cluster, and provides some example log messages that are recorded during specific cluster events.

You configure logging for a cluster in the same way as you configuring logging for a standalone FortiGate unit. Log configuration changes made to the cluster are synchronized to all cluster units.

All cluster units record log messages separately to the individual cluster unit's log disk, to the cluster unit's system memory, or both. You can view and manage log messages for each cluster unit from the cluster web-based manager Log Access page.

When remote logging is configured, all cluster units send log messages to remote FortiAnalyzer units or other remote servers as configured. HA uses routing and inter-VDOM links to route subordinate unit log traffic through the primary unit to the network. See [“The primary unit acts as a router for subordinate unit management traffic” on page 1259](#).

When you configure a FortiAnalyzer unit to receive log messages from a FortiGate cluster, you should add a cluster to the FortiAnalyzer unit configuration so that the FortiAnalyzer unit can receive log messages from all cluster units.

## Viewing and managing log messages for individual cluster units

This section describes how to view and manage log messages for an individual cluster unit.

### To view HA cluster log messages

1. Log into the cluster web-based manager.
2. Go to *Log&Report > Log Access* and select Memory or Disk.  
For each log display, the *HA Cluster* list displays the serial number of the cluster unit for which log messages are displayed. The serial numbers are displayed in order in the list.
3. Set *HA Cluster* to the serial number of one of the cluster units to display log messages for that unit.

You can view logs saved to memory or logs saved to the hard disk for the cluster unit.

### About HA event log messages

HA event log messages always include the host name and serial number of the cluster unit that recorded the message. HA event log messages also include the HA state of the unit and also

indicate when a cluster unit switches (or moves) from one HA state to another. Cluster units can operate in the HA states listed in [Table 56](#):

**Table 56:** HA states

<b>Hello</b>	A FortiGate unit configured for HA operation has started up and is looking for other FortiGate units with which to form a cluster.
<b>Work</b>	In an active-passive cluster a cluster unit is operating as the primary unit. In an active-active cluster unit is operating as the primary unit or a subordinate unit.
<b>Standby</b>	In an active-passive cluster the cluster unit is operating as a subordinate unit.

HA log Event log messages also indicate the virtual cluster that the cluster unit is operating in as well as the member number of the unit in the cluster. If virtual domains are not enabled, all clusters unit are always operating in virtual cluster 1. If virtual domains are enabled, a cluster unit may be operating in virtual cluster 1 or virtual cluster 2. The member number indicates the position of the cluster unit in the cluster members list. Member 0 is the primary unit. Member 1 is the first subordinate unit, member 2 is the second subordinate unit, and so on.

## HA log messages

See the [FortiGate Log Message Reference](#) for a listing of and descriptions of the HA log messages.

## Fortigate HA message "HA master heartbeat interface <intf\_name> lost neighbor information"

The following HA log messages may be recorded by an operating cluster:

```
2009-02-16 11:06:34 device_id=FG2001111111 log_id=0105035001 type=event
subtype=ha pri=critical vd=root msg="HA slave heartbeat interface internal lost neighbor
information"
```

```
2009-02-16 11:06:40 device_id=FG2001111111 log_id=0105035001 type=event
subtype=ha pri=notice vd=root msg="Virtual cluster 1 of group 0 detected new joined HA
member"
```

```
2009-02-16 11:06:40 device_id=FG2001111111 log_id=0105035001 type=event
subtype=ha pri=notice vd=root msg="HA master heartbeat interface internal get peer
information"
```

These log messages indicate that the cluster units could not connect to each other over the HA heartbeat link for the period of time that is given by `hb-interval x hb-lost-threshold`, which is 1.2 seconds with the default values.

### To diagnose this problem

1. Check all heartbeat interface connections including cables and switches to make sure they are connected and operating normally.
2. Use the following commands to display the status of the heartbeat interfaces.

```
get hardware nic <heartbeat_interface_name>
diagnose hardware deviceinfo nic <heartbeat_interface_name>
```

The status information may indicate the interface status and link status and also indicate if a large number of errors have been detected.

3. If the log message only appears during peak traffic times, increase the tolerance for missed HA heartbeat packets by using the following commands to increase the lost heartbeat threshold and heartbeat interval:

```
config system ha
 set hb-lost-threshold 12
 set hb-interval 4
end
```

These settings multiply by 4 the loss detection interval. You can use higher values as well.

This condition can also occur if the cluster units are located in different buildings or even different geographical locations. Called a distributed cluster, as a result of the separation it may take a relatively long time for heartbeat packets to be transmitted between cluster units. You can support a distributed cluster by increasing the heartbeat interval so that the cluster expects extra time between heartbeat packets.

4. Optionally disable session-pickup to reduce the processing load on the heartbeat interfaces.
5. Instead of disabling session-pickup you can enable `session-pickup-delay` to reduce the number of sessions that are synchronized. With this option enabled only sessions that are active for more than 30 seconds are synchronized.

It may be useful to monitor CPU and memory usage to check for low memory and high CPU usage. You can configure event logging to monitor CPU and memory usage. You can also enable the CPU over usage and memory low SNMP events.

Once this monitoring is in place, try and determine if there have been any changes in the network or an increase of traffic recently that could be the cause. Check to see if the problem happens frequently and if so what the pattern is.

To monitor the CPU of the cluster units and troubleshoot further, use the following procedure and commands:

```
get system performance status
get sys performance top 2
diagnose sys top 2
```

These commands repeated at frequent intervals will show the activity of the CPU and the number of sessions.

Search the [Fortinet Knowledge Base](#) for articles about monitoring CPU and Memory usage.

If the problem persists, gather the following information (a console connection might be necessary if connectivity is lost) and provide it to Technical Support when opening a ticket:

- Debug log from the web-based manager: *System > Config > Advanced > Download Debug Log*
- CLI command output:

```
diag sys top 2 (keep it running for 20 seconds)
get sys perf status (repeat this command multiple times to get good samples)
get sys ha status
diagnose sys ha status
diagnose sys ha dump-by {all options}
diagnose netlink device list
diagnose hardware deviceinfo nic <heartbeat-interface-name>
execute log filter category 1
execute log display
```

## Formatting cluster unit hard disks (log disks)

If you need to format the hard disk (also called log disk or disk storage) of one or more cluster units you should disconnect the unit from the cluster and use the `execute formatlogdisk` command to format the cluster unit hard disk then add the unit back to the cluster.

For information about how to remove a unit from a cluster and add it back, see [“Disconnecting a cluster unit from a cluster” on page 1285](#) and [“Adding a disconnected FortiGate unit back to its cluster” on page 1286](#).

Once you add the cluster unit with the formatted log disk back to the cluster you should make it the primary unit before removing other units from the cluster to format their log disks and then add them back to the cluster.

## Clusters and SNMP

You can use SNMP to manage a cluster by configuring a cluster interface for SNMP administrative access. Using an SNMP manager you can get cluster configuration and status information and receive traps.

You configure SNMP for a cluster in the same way as configuring SNMP for a standalone FortiGate unit. SNMP configuration changes made to the cluster are shared by all cluster units.

Each cluster unit sends its own traps and SNMP manager systems can use SNMP get commands to query each cluster unit separately. To set SNMP get queries to each cluster unit you must create a special get command that includes the serial number of the cluster unit.

Alternatively you can use the HA reserved management interface feature to give each cluster unit a different management IP address. Then you can create an SNMP get command for each cluster unit that just includes the management IP address and does not have to include the serial number. See [“Managing individual cluster units using a reserved management interface” on page 1254](#).

For a list of HA MIB fields and OIDs, see [“Fortinet MIBs” on page 1520](#).

## SNMP get command syntax for the primary unit

Normally, to get configuration and status information for a standalone FortiGate unit or for a primary unit, an SNMP manager would use an SNMP get commands to get the information in a MIB field. The SNMP get command syntax would be similar to the following:

```
snmpget -v2c -c <community_name> <address_ipv4> {<OID> | <MIB_field>}
```

where:

`<community_name>` is an SNMP community name added to the FortiGate configuration. You can add more than one community name to a FortiGate SNMP configuration. The most commonly used community name is `public`.

`<address_ipv4>` is the IP address of the FortiGate interface that the SNMP manager connects to.

`{<OID> | <MIB_field>}` is the object identifier (OID) for the MIB field or the MIB field name itself. The HA MIB fields and OIDs are listed in [Table 57](#).



**Table 57:** SNMP field names and OIDs

<b>MIB field</b>	<b>OID</b>	<b>Description</b>
fgHaSystemMode	.1.3.6.1.4.1.12356.101.13.1.1.0	HA mode (standalone, a-a, or a-p)
fgHaGroupId	.1.3.6.1.4.1.12356.101.13.1.2.0	The HA priority of the cluster unit. Default 128.
fgHaPriority	.1.3.6.1.4.1.12356.101.13.1.3.0	The HA priority of the cluster unit. Default 128.
fgHaOverride	.1.3.6.1.4.1.12356.101.13.1.4.0	Whether HA override is disabled or enabled for the cluster unit.
fgHaAutoSync	.1.3.6.1.4.1.12356.101.13.1.5.0	Whether automatic HA synchronization is disabled or enabled.
fgHaSchedule	.1.3.6.1.4.1.12356.101.13.1.6.0	The HA load balancing schedule. Set to none unless operating in a-p mode.
fgHaGroupName	.1.3.6.1.4.1.12356.101.13.1.7.0	The HA group name.
fgHaStatsIndex	.1.3.6.1.4.1.12356.101.13.2.1.1.1.1	The cluster index of the cluster unit. 1 for the primary unit, 2 to x for the subordinate units.
fgHaStatsSerial	.1.3.6.1.4.1.12356.101.13.2.1.1.2.1	The serial number of the cluster unit.
fgHaStatsCpuUsage	.1.3.6.1.4.1.12356.101.13.2.1.1.3.1	The cluster unit's current CPU usage.
fgHaStatsMemUsage	.1.3.6.1.4.1.12356.101.13.2.1.1.4.1	The cluster unit's current Memory usage.
fgHaStatsNetUsage	.1.3.6.1.4.1.12356.101.13.2.1.1.5.1	The cluster unit's current Network bandwidth usage.
fgHaStatsSesCount	.1.3.6.1.4.1.12356.101.13.2.1.1.6.1	The cluster unit's current session count.
fgHaStatsPktCount	.1.3.6.1.4.1.12356.101.13.2.1.1.7.1	The cluster unit's current packet count.
fgHaStatsByteCount	.1.3.6.1.4.1.12356.101.13.2.1.1.8.1	The cluster unit's current byte count.
fgHaStatsIdsCount	.1.3.6.1.4.1.12356.101.13.2.1.1.9.1	The number of attacks reported by the IPS for the cluster unit.

**Table 57:** SNMP field names and OIDs

MIB field	OID	Description
fgHaStatsAvCount	.1.3.6.1.4.1.12356.101.13.2.1.1.10.1	The number of viruses reported by the antivirus system for the cluster unit.
fgHaStatsHostname	.1.3.6.1.4.1.12356.101.13.2.1.1.11.1	The hostname of the cluster unit.

**To get the HA priority for the primary unit**

The following SNMP get command gets the HA priority for the primary unit. The community name is `public`. The IP address of the cluster interface configured for SNMP management access is `10.10.10.1`. The HA priority MIB field is `fgHaPriority` and the OID for this MIB field is `1.3.6.1.4.1.12356.101.13.1.3.0`. The first command uses the MIB field name and the second uses the OID:

```
snmpget -v2c -c public 10.10.10.1 fgHaPriority
snmpget -v2c -c public 10.10.10.1 1.3.6.1.4.1.12356.101.13.1.3.0
```

**SNMP get command syntax for any cluster unit**

To get configuration status information for a specific cluster unit (for the primary unit or for any subordinate unit), the SNMP manager must add the serial number of the cluster unit to the SNMP get command after the community name. The community name and the serial number are separated with a dash. The syntax for this SNMP get command would be:

```
snmpget -v2c -c <community_name>-<fgt_serial> <address_ipv4> {<OID> | <MIB_field>}
```

where:

`<community_name>` is an SNMP community name added to the FortiGate configuration. You can add more than one community name to a FortiGate SNMP configuration. All units in the cluster have the same community name. The most commonly used community name is `public`.

`<fgt_serial>` is the serial number of any cluster unit. For example, `FGT4002803033172`. You can specify the serial number of any cluster unit, including the primary unit, to get information for that unit.

`<address_ipv4>` is the IP address of the FortiGate interface that the SNMP manager connects to.

`{<OID> | <MIB_field>}` is the object identifier (OID) for the MIB field or the MIB field name itself. To find OIDs and MIB field names see [“Fortinet MIBs” on page 1520](#).

If the serial number matches the serial number of a subordinate unit, the SNMP get request is sent over the HA heartbeat link to the subordinate unit. After processing the request, the subordinate unit sends the reply back over the HA heartbeat link back to the primary unit. The primary unit then forwards the response back to the SNMP manager.

If the serial number matches the serial number of the primary unit, the SNMP get request is processed by the primary unit. You can actually add a serial number to the community name of any SNMP get request. But normally you only need to do this for getting information from a subordinate unit.

## To get the CPU usage for a subordinate unit

The following SNMP get command gets the CPU usage for a subordinate unit in a FortiGate-5001SX cluster. The subordinate unit has serial number FG50012205400050. The community name is `public`. The IP address of the FortiGate interface is 10.10.10.1. The HA status table MIB field is `fgHaStatsCpuUsage` and the OID for this MIB field is 1.3.6.1.4.1.12356.101.13.2.1.1.3.1. The first command uses the MIB field name and the second uses the OID for this table:

```
snmpget -v2c -c public-FG50012205400050 10.10.10.1 fgHaStatsCpuUsage
snmpget -v2c -c public-FG50012205400050 10.10.10.1
1.3.6.1.4.1.12356.101.13.2.1.1.3.1
```

FortiGate SNMP recognizes the community name with syntax `<community_name>-<fgt_serial>`. When the primary unit receives an SNMP get request that includes the community name followed by serial number, the FGCP extracts the serial number from the request. Then the primary unit redirects the SNMP get request to the cluster unit with that serial number. If the serial number matches the serial number of the primary unit, the SNMP get is processed by the primary unit.

## Getting serial numbers of cluster units

The following SNMP get commands use the MIB field name `fgHaStatsSerial.<index>` to get the serial number of each cluster unit. Where `<index>` is the cluster unit's cluster index and 1 is the cluster index of the primary unit, 2 is the cluster index of the first subordinate unit, and 3 is the cluster index of the second subordinate unit.

The OID for this MIB field is 1.3.6.1.4.1.12356.101.13.2.1.1.2.1. The community name is `public`. The IP address of the FortiGate interface is 10.10.10.1.

The first command uses the MIB field name and the second uses the OID for this table and gets the serial number of the primary unit:

```
snmpget -v2c -c public 10.10.10.1 fgHaStatsSerial.1
snmpget -v2c -c public 10.10.10.1 1.3.6.1.4.1.12356.101.13.2.1.1.2.1
```

The second command uses the MIB field name and the second uses the OID for this table and gets the serial number of the first subordinate unit:

```
snmpget -v2c -c public 10.10.10.1 fgHaStatsSerial.2
snmpget -v2c -c public 10.10.10.1 1.3.6.1.4.1.12356.101.13.2.2.2
```

## SNMP get command syntax - reserved management interface enabled

To get configuration and status information for any cluster unit where you have enabled the HA reserved management interface feature and assigned IP addresses to the management interface of each cluster unit, an SNMP manager would use the following get command syntax:

```
snmpget -v2c -c <community_name> <mgt_address_ipv4> {<OID> |
<MIB_field>}
```

where:

`<community_name>` is an SNMP community name added to the FortiGate configuration. You can add more than one community names to a FortiGate SNMP configuration. The most commonly used community name is `public`.

`<mgt_address_ipv4>` is the IP address of the FortiGate HA reserved management interface that the SNMP manager connects to.

`{<OID> | <MIB_field>}` is the object identifier (OID) for the MIB field or the MIB field name itself. To find OIDs and MIB field names see your FortiGate unit's online help.

See [“To get CPU, memory, and network usage of each cluster unit using the reserved management IP addresses”](#) on page 1258.

## Clusters and file quarantine

You can configure file quarantine for a cluster in the same way as configuring file quarantine for a standalone FortiGate unit. Quarantine configuration changes made to the cluster are shared by all cluster units.

In an active-active cluster, both the primary unit and the subordinate units accept antivirus sessions and may quarantine files. In an active-passive cluster, only the primary unit quarantines files. Multiple cluster units in an active-passive cluster may have quarantined files if different cluster units have been the primary unit.

All cluster units quarantine files separately to their own hard disk. You can go to *Log&Report > Archive Access > Quarantine* to view and manage the quarantine file list for each cluster unit.

All cluster units can also quarantine files to a FortiAnalyzer unit. When you configure a FortiAnalyzer unit to receive quarantine files from a cluster, you should add each cluster unit to the FortiAnalyzer device configuration so that the FortiAnalyzer unit can receive quarantine files from all cluster units.

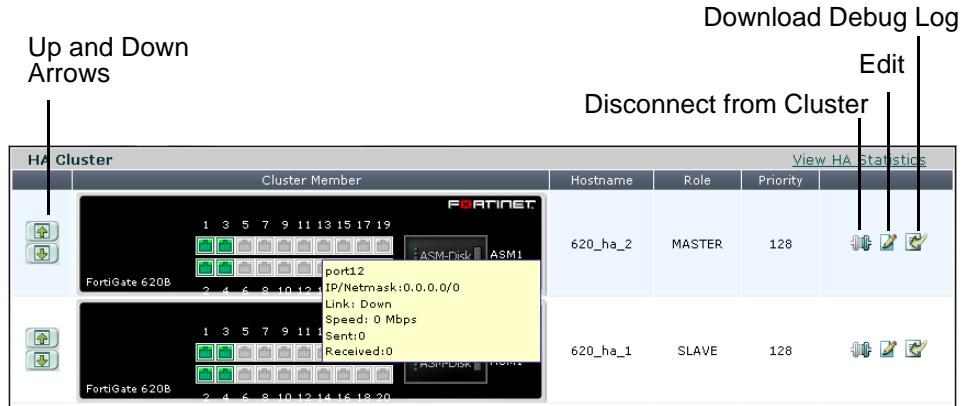
## Cluster members list

Display the cluster members list to view the status of the FortiGate units in an operating cluster. To display the cluster members list, go to *System > Config > HA*.

From the cluster members list you can also:

- View HA statistics (see [“Viewing HA statistics”](#) on page 1271).
- View and optionally change the HA configuration of the operating cluster (see [“Changing the HA configuration of an operating cluster”](#) on page 1273).
- View and optionally change the host name and device priority of a subordinate unit (see [“Changing the subordinate unit host name and device priority”](#) on page 1273).
- Disconnect a cluster unit from a cluster (see [“Disconnecting a cluster unit from a cluster”](#) on page 1285).
- Download the Debug log for any cluster unit. You can send this debug log file to Fortinet Technical Support to help diagnose problems with the cluster or with individual cluster units.

**Figure 206:**Example cluster members list



**View HA Statistics** Display the serial number, status, and monitor information for each cluster unit. See [“Viewing HA statistics” on page 1271](#).

**Up and down arrows** Change the order in which cluster members are listed. The operation of the cluster or of the units in the cluster are not affected. All that changes is the order in which cluster units are displayed on the cluster members list.

**Cluster member** Illustrations of the front panels of the cluster units. If the network jack for an interface is shaded green, the interface is connected. Pause the mouse pointer over each illustration to view the cluster unit host name, serial number, and how long the unit has been operating (up time). The list of monitored interfaces is also displayed.

**Hostname** The host name of the FortiGate unit. The default host name of the FortiGate unit is the FortiGate unit serial number.

- To change the primary unit host name, go to the system dashboard and select Change beside the current host name in System Information widget.
- To change a subordinate unit host name, from the cluster members list select the edit icon for a subordinate unit.

**Role** The status or role of the cluster unit in the cluster.

- Role is MASTER for the primary (or master) unit
- Role is SLAVE for all subordinate (or backup) cluster units

**Priority** The device priority of the cluster unit. Each cluster unit can have a different device priority. During HA negotiation, the unit with the highest device priority becomes the primary unit.

The device priority range is 0 to 255. The default device priority is 128.

**Disconnect from cluster** Disconnect the cluster unit from the cluster. See [“Disconnecting a cluster unit from a cluster” on page 1285](#).

---

**Edit**

Select Edit to change a cluster unit HA configuration.

- For a primary unit, select Edit to change the cluster HA configuration. You can also change the device priority of the primary unit.
- For a primary unit in a virtual cluster, select Edit to change the virtual cluster HA configuration. You can also change the virtual cluster 1 and virtual cluster 2 device priority of this cluster unit.
- For a subordinate unit, select Edit to change the subordinate unit host name and device priority. See [“Changing the subordinate unit host name and device priority”](#) on page 1273.
- For a subordinate unit in a virtual cluster, select Edit to change the subordinate unit host name. In addition you can change the device priority for the subordinate unit for the selected virtual cluster.

---

**Download debug log** Download an encrypted debug log to a file. You can send this debug log file to Fortinet Technical Support to help diagnose problems with the cluster or with individual cluster units.

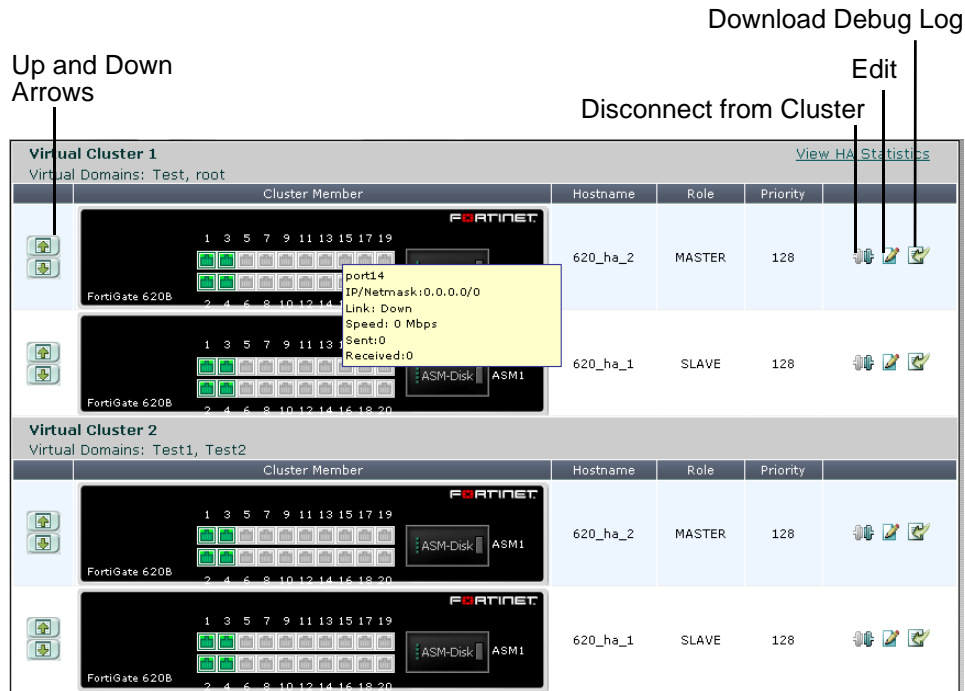
---

## Virtual cluster members list

If virtual domains are enabled, you can display the cluster members list to view the status of the operating virtual clusters. The virtual cluster members list shows the status of both virtual clusters including the virtual domains added to each virtual cluster.

To display the virtual cluster members list for an operating cluster log in as the admin administrator, select Global Configuration and go to *System > Config > HA*.

**Figure 207:**Example FortiGate-5001SX virtual cluster members list



The fields and functions of the virtual cluster members list are the same as the fields and functions described in “Cluster members list” on page 1268 with the following exceptions.

- When you select the edit icon for a primary unit in a virtual cluster, you can change the virtual cluster 1 and virtual cluster 2 device priority of this cluster unit and you can edit the VDOM partitioning configuration of the cluster.
- When you select the edit icon for a subordinate unit in a virtual cluster, you can change the device priority for the subordinate unit for the selected virtual cluster.

Also, the HA cluster members list changes depending on the cluster unit. For the virtual cluster described in the “Example: virtual clustering with two VDOMs and VDOM partitioning” on page 1221 if you connect to port5 using you are connecting to 620b\_ha\_2 (620b\_ha\_2 is displayed on the web browser title bar or in the CLI prompt).

If you connect to port1 you are connecting to 620b\_ha\_1 (620b\_ha\_2 is displayed on the web browser title bar or in the CLI prompt).

## Viewing HA statistics

From the cluster members list you can select View HA statistics to display the serial number, status, and monitor information for each cluster unit. To view HA statistics, go to *System > Config > HA* and select View HA Statistics.

**Figure 208:**Example HA statistics (active-passive cluster)

Refresh every		none		<a href="#">Back to HA monitor &gt;&gt;</a>		
Unit	Status	Up Time	Monitor			
620_ha_2 FG600B3908600825	✓	5 days	CPU Usage	Active Sessions	Total Packets	Virus Detected
		22 hours		42	74875	0
		57 minutes	Memory Usage	Network Utilization	Total Bytes	Intrusion Detected
		17 seconds		30 Kbps	26981277	0
620_ha_1 FG600B3908600705	✓	5 days	CPU Usage	Active Sessions	Total Packets	Virus Detected
		22 hours		21	12115	0
		48 minutes	Memory Usage	Network Utilization	Total Bytes	Intrusion Detected
		58 seconds		19 Kbps	930358	0

<b>Refresh every</b>	Select to control how often the web-based manager updates the HA statistics display.
<b>Back to HA monitor</b>	Close the HA statistics list and return to the cluster members list.
<b>Serial No.</b>	Use the serial number ID to identify each FortiGate unit in the cluster. The cluster ID matches the FortiGate unit serial number.
<b>Status</b>	Indicates the status of each cluster unit. A green check mark indicates that the cluster unit is operating normally. A red X indicates that the cluster unit cannot communicate with the primary unit.
<b>Up Time</b>	The time in days, hours, minutes, and seconds since the cluster unit was last started.
<b>Monitor</b>	Displays system status information for each cluster unit.
<b>CPU Usage</b>	The current CPU status of each cluster unit. The web-based manager displays CPU usage for core processes only. CPU usage for management processes (for example, for HTTPS connections to the web-based manager) is excluded.
<b>Memory Usage</b>	The current memory status of each cluster unit. The web-based manager displays memory usage for core processes only. Memory usage for management processes (for example, for HTTPS connections to the web-based manager) is excluded.
<b>Active Sessions</b>	The number of communications sessions being processed by the cluster unit.
<b>Total Packets</b>	The number of packets that have been processed by the cluster unit since it last started up.
<b>Virus Detected</b>	The number of viruses detected by the cluster unit.
<b>Network Utilization</b>	The total network bandwidth being used by all of the cluster unit interfaces.



---

<b>Total Bytes</b>	The number of bytes that have been processed by the cluster unit since it last started up.
--------------------	--------------------------------------------------------------------------------------------

---

<b>Intrusion Detected</b>	The number of intrusions or attacks detected by Intrusion Protection running on the cluster unit.
---------------------------	---------------------------------------------------------------------------------------------------

---

## Changing the HA configuration of an operating cluster

To change the configuration settings of an operating cluster, go to *System > Config > HA* to display the cluster members list. Select Edit for the master (or primary) unit in the cluster members list to display the HA configuration page for the cluster.

You can use the HA configuration page to check and fine tune the configuration of the cluster after the cluster is up and running. For example, if you connect or disconnect cluster interfaces you may want to change the Port Monitor configuration.

Any changes you make on this page, with the exception of changes to the device priority, are first made to the primary unit configuration and then synchronized to the subordinate units. Changing the device priority only affects the primary unit.

## Changing the HA configuration of an operating virtual cluster

To change the configuration settings of the primary unit in a functioning cluster with virtual domains enabled, log in as the admin administrator, select Global Configuration and go to *System > Config > HA* to display the cluster members list. Select Edit for the master (or primary) unit in virtual cluster 1 or virtual cluster 2 to display the HA configuration page for the virtual cluster.

You can use the virtual cluster HA configuration page to check and fine tune the configuration of both virtual clusters after the cluster is up and running. For example, you may want to change the Port Monitor configuration for virtual cluster 1 and virtual cluster 2 so that each virtual cluster monitors its own interfaces.

You can also use this configuration page to move virtual domains between virtual cluster 1 and virtual cluster 2. Usually you would distribute virtual domains between the two virtual clusters to balance the amount of traffic being processed by each virtual cluster.

Any changes you make on this page, with the exception of changes to the device priorities, are first made to the primary unit configuration and then synchronized to the subordinate unit.

You can also adjust device priorities to configure the role of this cluster unit in the virtual cluster. For example, to distribute traffic to both cluster units in the virtual cluster configuration, you would want one cluster unit to be the primary unit for virtual cluster 1 and the other cluster unit to be the primary unit for virtual cluster 2. You can create this configuration by setting the device priorities. The cluster unit with the highest device priority in virtual cluster 1 becomes the primary unit for virtual cluster 1. The cluster unit with the highest device priority in virtual cluster 2 becomes the primary unit in virtual cluster 2.

## Changing the subordinate unit host name and device priority

To change the host name and device priority of a subordinate unit in an operating cluster, go to *System > Config > HA* to display the cluster members list. Select Edit for any slave (subordinate) unit in the cluster members list.

To change the host name and device priority of a subordinate unit in an operating cluster with virtual domains enabled, log in as the admin administrator, select Global Configuration and go to *System > Config > HA* to display the cluster members list. Select Edit for any slave (subordinate) unit in the cluster members list.

You can change the host name (Peer) and device priority (Priority) of this subordinate unit. These changes only affect the configuration of the subordinate unit.

The device priority is not synchronized among cluster members. In a functioning cluster you can change device priority to change the priority of any unit in the cluster. The next time the cluster negotiates, the cluster unit with the highest device priority becomes the primary unit.

The device priority range is 0 to 255. The default device priority is 128.

## Upgrading cluster firmware

You can upgrade the FortiOS firmware running on an HA cluster in the same manner as upgrading the firmware running on a standalone FortiGate unit. During a normal firmware upgrade, the cluster upgrades the primary unit and all subordinate units to run the new firmware image. The firmware upgrade takes place without interrupting communication through the cluster.



Upgrading cluster firmware to a new major release (for example upgrading from 3.0 MRx to 4.0 MRx) is supported for clusters. Make sure you are taking an upgrade path described in the release notes. Even so you should back up your configuration and only perform such a firmware upgrade during a maintenance window.

---

To upgrade the firmware without interrupting communication through the cluster, the cluster goes through a series of steps that involve first upgrading the firmware running on the subordinate units, then making one of the subordinate units the primary unit, and finally upgrading the firmware on the former primary unit. These steps are transparent to the user and the network, but depending upon your HA configuration may result in the cluster selecting a new primary unit.

The following sequence describes in detail the steps the cluster goes through during a firmware upgrade and how different HA configuration settings may affect the outcome.

1. The administrator uploads a new firmware image from the web-based manager or CLI.
2. If the cluster is operating in active-active mode load balancing is turned off.
3. The cluster upgrades the firmware running on all of the subordinate units.
4. Once the subordinate units have been upgraded, a new primary unit is selected.

This primary unit will be running the new upgraded firmware.

5. The cluster now upgrades the firmware of the former primary unit.

If the age of the new primary unit is more than 300 seconds (5 minutes) greater than the age of all other cluster units, the new primary unit continues to operate as the primary unit.

This is the intended behavior but does not usually occur because the age difference of the cluster units is usually less than the cluster age difference margin of 300 seconds. So instead, the cluster negotiates again to select a primary unit as described in [“Primary unit selection” on page 1131](#).

You can keep the cluster from negotiating again by reducing the cluster age difference margin using the `ha-uptime-diff-margin` option. However, you should be cautious when reducing the age or other problems may occur. For information about the cluster age difference margin, see [“Cluster age difference margin \(grace period\)” on page 1133](#). For

more information about changing the cluster age margin, see [“Changing the cluster age difference margin” on page 1133](#).

6. If the cluster is operating in active-active mode, load balancing is turned back on.

## Changing how the cluster processes firmware upgrades

By default cluster firmware upgrades proceed as uninterruptible upgrades that do not interrupt traffic flow. If required, you can use the following CLI command to change how the cluster handles firmware upgrades. You might want to change this setting if you are finding uninterruptible upgrades take too much time.

```
config system ha
 set uninterruptible-upgrade disable
end
```

`uninterruptible-upgrade` is enabled by default. If you disable `uninterruptible-upgrade` the cluster still upgrades the firmware on all cluster units, but all cluster units are upgraded at once; which takes less time but interrupts communication through the cluster.

## Synchronizing the firmware build running on a new cluster unit

If the firmware build running on a FortiGate unit that you add to a cluster is older than the cluster firmware build, you may be able to use the following steps to synchronize the firmware running on the new cluster unit.

This procedure describes re-installing the same firmware build on a cluster to force the cluster to upgrade all cluster units to the same firmware build.

Due to firmware upgrade and synchronization issues, in some cases this procedure may not work. In all cases it will work to install the same firmware build on the new unit as the one that the cluster is running before adding the new unit to the cluster.

### To synchronize the firmware build running on a new cluster unit

1. Obtain a firmware image that is the same as build already running on the cluster.
2. Connect to the cluster using the web-based manager.
3. Go to the *System Information* dashboard widget.
4. Select *Update* beside *Firmware Version*.  
You can also install a newer firmware build.
5. Select OK.

After the firmware image is uploaded to the cluster, the primary unit upgrades all cluster units to this firmware build.

## Downgrading cluster firmware

For various reasons you may need to downgrade the firmware that a cluster is running. You can use the information in this section to downgrade the firmware version running on a cluster.

In most cases you can downgrade the firmware on an operating cluster using the same steps as for a firmware upgrade. A warning message appears during the downgrade but the downgrade usually works and after the downgrade the cluster continues operating normally with the older firmware image.

Downgrading between some firmware versions, especially if features have changed between the two versions, may not always work without the requirement to fix configuration issues after the downgrade.

Only perform firmware downgrades during maintenance windows and make sure you back up your cluster configuration before the downgrade.

If the firmware downgrade that you are planning may not work without configuration loss or other problems, you can use the following downgrade procedure to make sure your configuration is not lost after the downgrade.

### To downgrade cluster firmware

This example shows how to downgrade the cluster shown in [Figure 190 on page 1153](#). The cluster consists of two cluster units (620\_ha\_1 and 620\_ha\_2). The port1 and port2 interfaces are connected networks and the port3 and port4 interfaces are connected together for the HA heartbeat.

This example, describes separating each unit from the cluster and downgrading the firmware for the standalone FortiGate units. There are several ways you could disconnect units from the cluster. This example describes using the disconnect from cluster function described in [“Disconnecting a cluster unit from a cluster” on page 1285](#).

1. Go to the *System Information* dashboard widget and backup the cluster configuration.

From the CLI use `execute backup config`.

2. Go to *System > Config > HA* and for 620\_ha\_1 select the *Disconnect from cluster* icon.
3. Select the port2 interface and enter an IP address and netmask of 10.11.101.101/24 and select OK.

From the CLI you can enter the following command (FG600B3908600705 is the serial number of the cluster unit) to be able to manage the standalone FortiGate unit by connecting to the port2 interface with IP address and netmask 10.11.101.101/24.

```
execute ha disconnect FG600B3908600705 port2 10.11.101.101/24
```

After 620\_ha\_1 is disconnected, 620\_ha\_2 continues processing traffic.

4. Connect to the 620\_ha\_1 web-based manager or CLI using IP address 10.11.101.101/24 and follow normal procedures to downgrade standalone FortiGate unit firmware.
5. When the downgrade is complete confirm that the configuration of 620\_ha\_1 is correct.
6. Set the HA mode of 620\_ha\_2 to Standalone and follow normal procedures to downgrade standalone FortiGate unit firmware.
7. When the downgrade is complete confirm that the configuration of 620\_ha\_2 is correct.
8. Set the HA mode of 620\_ha\_2 to Active-Passive or the required HA mode.
9. Set the HA mode of 620\_ha\_1 to the same mode as 620\_ha\_2.

Network communication will be interrupted for a short time during the downgrade. If you have not otherwise changed the HA settings of the cluster units and if the firmware downgrades have not affected the configurations the units should negotiate and form cluster running the downgraded firmware.

## Backing up and restoring the cluster configuration

You can backup the configuration of the primary unit by logging into the web-based manager or CLI and following normal configuration backup procedures.

The following configuration settings are not synchronized to all cluster units:

- HA override and priority
- The interface configuration of the HA reserved management interface (`config system interface`)
- The HA reserved management interface default route (`ha-mgmt-interface-gateway`)
- The FortiGate unit host name.

To backup these configuration settings for each cluster unit you must log into each cluster unit and backup its configuration.

If you need to restore the configuration of the cluster including the configuration settings that are not synchronized you should first restore the configuration of the primary unit and then restore the configuration of each cluster unit. Alternatively you could log into each cluster unit and manually add the configuration settings that were not restored.

## Monitoring cluster units for failover

If the primary unit in the cluster fails, the units in the cluster renegotiate to select a new primary unit. Failure of the primary unit results in the following:

- If SNMP is enabled, the new primary unit sends HA trap messages. The messages indicate a cluster status change, HA heartbeat failure, and HA member down. For more info about HA and SNMP, see [“Clusters and SNMP” on page 1264](#).
- If event logging is enabled and HA activity event is selected, the new primary unit records log messages that show that the unit has become the primary unit.
- If alert email is configured to send email for HA activity events, the new primary unit sends an alert email containing the log message recorded by the event log.
- The cluster contains fewer FortiGate units. The failed primary unit no longer appears on the Cluster Members list.
- The host name and serial number of the primary unit changes. You can see these changes when you log into the web-based manager or CLI.
- The cluster info displayed on the dashboard, cluster members list or from the `get system ha status` command changes.

If a subordinate unit fails, the cluster continues to function normally. Failure of a subordinate unit results in the following:

- If event logging is enabled and HA activity event is selected, the primary unit records log messages that show that a subordinate has been removed from the cluster.
- If alert email is configured to send email for HA activity events, the new primary unit sends an alert email containing the log message recorded by the event log.
- The cluster contains fewer FortiGate units. The failed unit no longer appears on the Cluster Members list.

## Viewing cluster status from the CLI

Use the `get system ha status` command to display information about an HA cluster. The command displays general HA configuration settings. The command also displays information about how the cluster unit that you have logged into is operating in the cluster.

Usually you would log into the primary unit CLI using SSH or telnet. In this case the `get system ha status` command displays information about the primary unit first, and also displays the HA state of the primary unit (the primary unit operates in the work state). However,

if you log into the primary unit and then use the `execute ha manage` command to log into a subordinate unit, (or if you use a console connection to log into a subordinate unit) the `get system status` command displays information about this subordinate unit first, and also displays the HA state of this subordinate unit. The state of a subordinate unit is work for an active-active cluster and standby for an active-passive cluster.

For a virtual cluster configuration, the `get system ha status` command displays information about how the cluster unit that you have logged into is operating in virtual cluster 1 and virtual cluster 2. For example, if you connect to the cluster unit that is the primary unit for virtual cluster 1 and the subordinate unit for virtual cluster 2, the output of the `get system ha status` command shows virtual cluster 1 in the work state and virtual cluster 2 in the standby state. The `get system ha status` command also displays additional information about virtual cluster 1 and virtual cluster 2.

The command display includes the following fields.

Fields	Description
Model	The FortiGate model number.
Mode	The HA mode of the cluster: a-a or a-p.
Group	The group ID of the cluster.
Debug	The debug status of the cluster.
ses_pickup	The status of session pickup: enable or disable.
load balance	The status of the <code>load-balance-all</code> keyword: enable or disable. Relevant to active-active clusters only.
schedule	The active-active load balancing schedule. Relevant to active-active clusters only.
Master Slave	<p><code>Master</code> displays the device priority, host name, serial number, and cluster index of the primary (or master) unit.</p> <p><code>Slave</code> displays the device priority, host name, serial number, and cluster index of the subordinate (or slave, or backup) unit or units.</p> <p>The list of cluster units changes depending on how you log into the CLI. Usually you would use SSH or telnet to log into the primary unit CLI. In this case the primary unit would be at the top the list followed by the other cluster units.</p> <p>If you use <code>execute ha manage</code> or a console connection to log into a subordinate unit CLI, and then enter <code>get system ha status</code> the subordinate unit that you have logged into appears at the top of the list of cluster units.</p>
number of vcluster	The number of virtual clusters. If virtual domains are not enabled, the cluster has one virtual cluster. If virtual domains are enabled the cluster has two virtual clusters.

Fields	Description
vcluster 1 Master Slave	<p>The HA state (hello, work, or standby) and HA heartbeat IP address of the cluster unit that you have logged into in virtual cluster 1. If virtual domains are not enabled, <code>vcluster 1</code> displays information for the cluster. If virtual domains are enabled, <code>vcluster 1</code> displays information for virtual cluster 1.</p> <p>The HA heartbeat IP address is 169.254.0.2 if you are logged into the primary unit of virtual cluster 1 and 169.254.0.1 if you are logged into a subordinate unit of virtual cluster 1.</p> <p><code>vcluster 1</code> also lists the primary unit (master) and subordinate units (slave) in virtual cluster 1. The list includes the cluster index and serial number of each cluster unit in virtual cluster 1. The cluster unit that you have logged into is at the top of the list.</p> <p>If virtual domains are not enabled and you connect to the primary unit CLI, the HA state of the cluster unit in virtual cluster 1 is work. The display lists the cluster units starting with the primary unit.</p> <p>If virtual domains are not enabled and you connect to a subordinate unit CLI, the HA state of the cluster unit in virtual cluster 1 is standby. The display lists the cluster units starting with the subordinate unit that you have logged into.</p> <p>If virtual domains are enabled and you connect to the virtual cluster 1 primary unit CLI, the HA state of the cluster unit in virtual cluster 1 is work. The display lists the cluster units starting with the virtual cluster 1 primary unit.</p> <p>If virtual domains are enabled and you connect to the virtual cluster 1 subordinate unit CLI, the HA state of the cluster unit in virtual cluster 1 is standby. The display lists the cluster units starting with the subordinate unit that you are logged into.</p>
vcluster 2 Master Slave	<p><code>vcluster 2</code> only appears if virtual domains are enabled.</p> <p><code>vcluster 2</code> displays the HA state (hello, work, or standby) and HA heartbeat IP address of the cluster unit that you have logged into in virtual cluster 2. The HA heartbeat IP address is 169.254.0.2 if you are logged into the primary unit of virtual cluster 2 and 169.254.0.1 if you are logged into a subordinate unit of virtual cluster 2.</p> <p><code>vcluster 2</code> also lists the primary unit (master) and subordinate units (slave) in virtual cluster 2. The list includes the cluster index and serial number of each cluster unit in virtual cluster 2. The cluster unit that you have logged into is at the top of the list.</p> <p>If you connect to the virtual cluster 2 primary unit CLI, the HA state of the cluster unit in virtual cluster 2 is <code>work</code>. The display lists the cluster units starting with the virtual cluster 2 primary unit.</p> <p>If you connect to the virtual cluster 2 subordinate unit CLI, the HA state of the cluster unit in virtual cluster 2 is <code>standby</code>. The display lists the cluster units starting with the subordinate unit that you are logged into.</p>

## Examples

The following example shows `get system ha status` output for a cluster of two FortiGate-5001SX units operating in active-active mode. The cluster group ID, session pickup, load balance all, and the load balancing schedule are all set to the default values. The device

priority of the primary unit is also set to the default value. The device priority of the subordinate unit has been reduced to 100. The host name of the primary unit is 5001\_Slot\_4. The host name of the subordinate unit is 5001\_Slot\_3.

The command output was produced by connecting to the primary unit CLI (host name 5001\_Slot\_4).

```
Model: 5000
Mode: a-a
Group: 0
Debug: 0
ses_pickup: disable
load_balance: disable
schedule: round robin
Master:128 5001_Slot_4 FG50012204400045 1
Slave :100 5001_Slot_3 FG50012205400050 0
number of vcluster: 1
vcluster 1: work 169.254.0.2
Master:0 FG50012204400045
Slave :1 FG50012205400050
```

The following command output was produced by using `execute HA manage 0` to log into the subordinate unit CLI of the cluster shown in the previous example. The host name of the subordinate unit is 5001\_Slot\_3.

```
Model: 5000
Mode: a-a
Group: 0
Debug: 0
ses_pickup: disable
load_balance: disable
schedule: round robin
Slave :100 5001_Slot_3 FG50012205400050 0
Master:128 5001_Slot_4 FG50012204400045 1
number of vcluster: 1
vcluster 1: work 169.254.0.2
Slave :1 FG50012205400050
Master:0 FG50012204400045
```

The following example shows `get system ha status` output for a cluster of three FortiGate-5001 units operating in active-passive mode. The cluster group ID is set to 20 and session pickup is enabled. Load balance all and the load balancing schedule are set to the default value. The device priority of the primary unit is set to 200. The device priorities of the subordinate units are set to 128 and 100. The host name of the primary unit is 5001\_Slot\_5. The host names of the subordinate units are 5001\_Slot\_3 and 5001\_Slot\_4.

```
Model: 5000
Mode: a-p
Group: 20
Debug: 0
ses_pickup: enable
load_balance: disable
schedule: round robin
Master:200 5001_Slot_5 FG50012206400112 0
Slave :100 5001_Slot_3 FG50012205400050 1
Slave :128 5001_Slot_4 FG50012204400045 2
number of vcluster: 1
vcluster 1: work 169.254.0.1
Master:0 FG50012206400112
Slave :1 FG50012204400045
```



```
Slave :2 FG50012205400050
```

The following example shows `get system ha status` output for a cluster of two FortiGate-5001 units with virtual clustering enabled. This command output was produced by logging into the primary unit for virtual cluster 1 (hostname: 5001\_Slot\_4, serial number FG50012204400045).

The virtual clustering output shows that the cluster unit with host name 5001\_Slot\_4 and serial number FG50012204400045 is operating as the primary unit for virtual cluster 1 and the subordinate unit for virtual cluster 2.

For virtual cluster 1 the cluster unit that you have logged into is operating in the work state and the serial number of the primary unit for virtual cluster 1 is FG50012204400045. For virtual cluster 2 the cluster unit that you have logged into is operating in the standby state and the serial number of the primary unit for virtual cluster 2 is FG50012205400050.

```
Model: 5000
Mode: a-p
Group: 20
Debug: 0
ses_pickup: enable
load_balance: disable
schedule: round robin
Master:128 5001_Slot_4 FG50012204400045 1
Slave :100 5001_Slot_3 FG50012205400050 0
number of vcluster: 2
vcluster 1: work 169.254.0.2
Master:0 FG50012204400045
Slave :1 FG50012205400050
vcluster 2: standby 169.254.0.1
Slave :1 FG50012204400045
Master:0 FG50012205400050
```

The following example shows `get system ha status` output for the same cluster as shown in the previous example after using `execute ha manage 0` to log into the primary unit for virtual cluster 2 (hostname: 5001\_Slot\_3, serial number FG50012205400050).

```
Model: 5000
Mode: a-p
Group: 20
Debug: 0
ses_pickup: enable
load_balance: disable
schedule: round robin
Slave :100 5001_Slot_3 FG50012205400050 0
Master:128 5001_Slot_4 FG50012204400045 1
number of vcluster: 2
vcluster 1: standby 169.254.0.2
Slave :1 FG50012205400050
Master:0 FG50012204400045
vcluster 2: work 169.254.0.1
Master:0 FG50012205400050
Slave :1 FG50012204400045
```

The following example shows `get system ha status` output for a virtual cluster configuration where the cluster unit with hostname: 5001\_Slot\_4 and serial number FG50012204400045 is the primary unit for both virtual clusters. This command output is produced by logging into cluster unit with host name 5001\_Slot\_4 and serial number FG50012204400045.

```
Model: 5000
```

```
Mode: a-p
Group: 20
Debug: 0
ses_pickup: enable
load_balance: disable
schedule: round robin
Master:128 5001_Slot_4 FG50012204400045 1
Slave :100 5001_Slot_3 FG50012205400050 0
number of vcluster: 2
vcluster 1: work 169.254.0.2
Master:0 FG50012204400045
Slave :1 FG50012205400050
vcluster 2: work 169.254.0.2
Master:0 FG50012204400045
Slave :1 FG50012205400050
```

## About the HA cluster index and the execute ha manage command

When a cluster starts up, the FortiGate Cluster Protocol (FGCP) assigns a cluster index and a HA heartbeat IP address to each cluster unit based on the serial number of the cluster unit. The FGCP selects the cluster unit with the highest serial number to become the primary unit. The FGCP assigns a cluster index of 0 and an HA heartbeat IP address of 169.254.0.1 to this unit. The FGCP assigns a cluster index of 1 and an HA heartbeat IP address of 169.254.0.2 to the cluster unit with the second highest serial number. If the cluster contains more units, the cluster unit with the third highest serial number is assigned a cluster index of 2 and an HA heartbeat IP address of 169.254.0.3, and so on. You can display the cluster index assigned to each cluster unit using the `get system ha status` command. Also when you use the `execute ha manage` command you select a cluster unit to log into by entering its cluster index.

The cluster index and HA heartbeat IP address only change if a unit leaves the cluster or if a new unit joins the cluster. When one of these events happens, the FGCP resets the cluster index and HA heartbeat IP address of each cluster unit according to serial number in the same way as when the cluster first starts up.

Each cluster unit keeps its assigned cluster index and HA heartbeat IP address even as the units take on different roles in the cluster. After the initial cluster index and HA heartbeat IP addresses are set according to serial number, the FGCP checks other primary unit selection criteria such as device priority and monitored interfaces. Checking these criteria could result in selecting a cluster unit without the highest serial number to operate as the primary unit.

Even if the cluster unit without the highest serial number now becomes the primary unit, the cluster indexes and HA heartbeat IP addresses assigned to the individual cluster units do not change. Instead the FGCP assigns a second cluster index, which could be called the operating cluster index, to reflect this role change. The operating cluster index is 0 for the primary unit and 1 and higher for the other units in the cluster. By default both sets of cluster indexes are the same. But if primary unit selection selects the cluster unit that does not have the highest serial number to be the primary unit then this cluster unit is assigned an operating cluster index of 0. The operating cluster index is used by the FGCP only. You can display the operating cluster index assigned to each cluster unit using the `get system ha status` command. There are no CLI commands that reference the operating cluster index.



Even though there are two cluster indexes there is only one HA heartbeat IP address and the HA heartbeat address is not affected by a change in the operating cluster index.

---

## Using the execute ha manage command

When you use the CLI command `execute ha manage <index_integer>` to connect to the CLI of another cluster unit, the `<index_integer>` that you enter is the cluster index of the unit that you want to connect to.

## Using get system ha status to display cluster indexes

You can display the cluster index assigned to each cluster unit using the CLI command `get system ha status`. The following example shows the information displayed by the `get system ha status` command for a cluster consisting of two FortiGate-5001SX units operating in active-passive HA mode with virtual domains not enabled and without virtual clustering.

```
get system ha status
Model: 5000
Mode: a-p
Group: 0
Debug: 0
ses_pickup: disable
Master:128 5001_slot_7 FG50012205400050 0
Slave :128 5001_slot_11 FG50012204400045 1
number of vcluster: 1
vcluster 1: work 169.254.0.1
Master:0 FG50012205400050
Slave :1 FG50012204400045
```

In this example, the cluster unit with serial number FG50012205400050 has the highest serial number and so has a cluster index of 0 and the cluster unit with serial number FG50012204400045 has a cluster index of 1. From the CLI of the primary (or master) unit of this cluster you can connect to the CLI of the subordinate (or slave) unit using the following command:

```
execute ha manage 1
```

This works because the cluster unit with serial number FG50012204400045 has a cluster index of 1.

The `get system ha status` command output shows two similar lists of indexes and serial numbers. The listing on the sixth and seventh lines of the command output are the cluster indexes assigned according to cluster unit serial number. These are the cluster indexes that you enter when using the `execute ha manage` command. The cluster indexes shown in the last two lines of the command output are the operating cluster indexes that reflect how the cluster units are actually operating in the cluster. In this example both sets of cluster indexes are the same.

The last three lines of the command output display the status of vcluster 1. In a cluster consisting of two cluster units operating without virtual domains enabled all clustering actually takes place in virtual cluster 1. HA is designed to work this way to support virtual clustering. If this cluster was operating with virtual domains enabled, adding virtual cluster 2 is similar to adding a new copy of virtual cluster 1. Virtual cluster 2 is visible in the `get system ha status` command output when you add virtual domains to virtual cluster 2.

The HA heartbeat IP address displayed on line 8 is the HA heartbeat IP address of the cluster unit that is actually operating as the primary unit. For a default configuration this IP address will always be 169.254.0.1 because the cluster unit with the highest serial number will be the primary unit. This IP address changes if the operating primary unit is not the primary unit with the highest serial number.

## Example: actual and operating cluster indexes do not match

This example shows the output of the `get system ha status` command for the same cluster of two FortiGate-5001SX units. However, in this example the device priority of the cluster unit with the serial number FG50012204400045 is increased to 200. As a result the cluster unit with the lowest serial number becomes the primary unit. This means the actual and operating cluster indexes of the cluster units do not match.

```
get system ha status
Model: 5000
Mode: a-p
Group: 0
Debug: 0
ses_pickup: disable
Master:128 5001_slot_7 FG50012205400050 0
Slave :200 5001_slot_11 FG50012204400045 1
number of vcluster: 1
vcluster 1: work 169.254.0.2
Master:1 FG50012205400050
Slave :0 FG50012204400045
```

The actual cluster indexes have not changed but the operating cluster indexes have. Also, the HA heartbeat IP address displayed for vcluster 1 has changed to 169.254.0.2.

## Virtual clustering example output

The output of the `get system ha status` command is the same if a cluster is operating with virtual clustering turned on but with all virtual domains in virtual cluster 1. The following `get system ha status` command output example shows the same cluster operating as a virtual cluster with virtual domains in virtual cluster 1 and added to virtual cluster 2. In this example the cluster unit with serial number FG50012204400045 is the primary unit for virtual cluster 1 and the cluster unit with serial number FG50012205400050 is the primary unit for virtual cluster 2.

```
get system ha status
Model: 5000
Mode: a-p
Group: 0
Debug: 0
ses_pickup: disable
Master:128 5001_slot_7 FG50012205400050 0
Slave :200 5001_slot_11 FG50012204400045 1
number of vcluster: 2
vcluster 1: work 169.254.0.2
Master:1 FG50012205400050
Slave :0 FG50012204400045
vcluster 2: standby 169.254.0.1
Master:0 FG50012205400050
Slave :1 FG50012204400045
```

This example shows three sets of indexes. The indexes in lines six and seven are still used by the `execute ha manage` command. The indexes on lines ten and eleven are for the primary and subordinate units in virtual cluster 1 and the indexes on the last two lines are for virtual cluster 2.

## Managing individual cluster units

The following procedure describes how to use SSH to log into the primary unit CLI and from there to use the `execute ha manage` command to connect to the CLI of any other unit in the

cluster. The procedure is very similar if you use telnet, or the web-based manager dashboard CLI console.

You can use the `execute ha manage` command from the CLI of any cluster unit to log into the CLI of another the cluster unit. Usually you would use this command from the CLI of the primary unit to log into the CLI of a subordinate unit. However, if you have logged into a subordinate unit CLI, you can use this command to log into the primary unit CLI, or the CLI of another subordinate unit.

Using SSH or telnet or the web-based manager dashboard CLI console you can only log into the primary unit CLI. Using a direct console connection you can log into any cluster unit. In both cases you can use `execute ha manage` to connect to the CLI of other cluster units.



You log into the subordinate unit using the `FGT_ha_admin` administrator account. This built-in administrator account gives you read and write permission on the subordinate unit. Normally this built-in administrative account is not visible, however `FGT_ha_admin` does appear in event log messages.

1. Use SSH to connect to the cluster and log into the primary unit CLI.

Connect to any cluster interface configured for SSH administrative access to log into the cluster.

2. Enter the following command followed by a space and type a question mark (?):

```
execute ha manage
```

The CLI displays a list of all the subordinate units in the cluster. Each cluster unit is numbered, starting at 1. The information displayed for each cluster unit includes the unit serial number and the host name of the unit.

3. Complete the command with the number of the subordinate unit to log into. For example, to log into subordinate unit 1, enter the following command:

```
execute ha manage 1
```

Press Enter to connect to and log into the CLI of the selected subordinate unit. If this subordinate unit has a different host name, the CLI prompt changes to this host name.

You can use CLI commands to manage this subordinate unit. If you make changes to the configuration of any cluster unit (primary or subordinate unit) these changes are synchronized to all cluster units.

4. You can now use the `execute ha manage` command to connect to any other cluster unit (including the primary unit). You can also use the `exit` command to return to the primary unit CLI.

## Disconnecting a cluster unit from a cluster

Use the following procedures to disconnect a cluster unit from a functioning cluster without disrupting the operation of the cluster. You can disconnect a cluster unit if you need to use the disconnected FortiGate unit for another purpose, such as to act as a standalone firewall.

You can use the following procedures for a standard cluster and for a virtual clustering configuration. To use the following procedures from a virtual cluster you must be logged in as the admin administrator and you must have selected Global Configuration.

When you disconnect a cluster unit you must assign an IP address and netmask to one of the interfaces of the disconnected unit. You can disconnect any unit from the cluster even the primary unit. After the unit is disconnected, the cluster responds as if the disconnected unit has failed. The cluster may renegotiate and may select a new primary unit.

When the cluster unit is disconnected the HA mode is changed to standalone. In addition, all interface IP addresses of the disconnected unit are set to 0.0.0.0 except for the interface that you configure.

Otherwise the configuration of the disconnected unit is not changed. The HA configuration of the disconnected unit is not changed either (except to change the HA mode to Standalone).

### To disconnect a cluster unit from a cluster - web-based manager

1. Go to *System > Config > HA* to view the cluster members list.
2. Select the Disconnect from cluster icon for the cluster unit to disconnect from the cluster.
3. Select the interface that you want to configure. You also specify the IP address and netmask for this interface. When the FortiGate unit is disconnected, all management access options are enabled for this interface.
4. Specify an IP address and netmask for the interface. You can use this IP address to connect to the interface to configure the disconnected FortiGate unit.
5. Select OK.

The FortiGate unit is disconnected from the cluster and the cluster may renegotiate and select a new primary unit. The selected interface of the disconnected unit is configured with the specified IP address and netmask.

### To disconnect a cluster unit from a cluster - CLI

1. Enter the following command to disconnect a cluster unit with serial number FGT5002803033050. The internal interface of the disconnected unit is set to IP address 1.1.1.1 and netmask 255.255.255.0.

```
execute ha disconnect FGT5002803033050 internal 1.1.1.1 255.255.255.0
```

## Adding a disconnected FortiGate unit back to its cluster

If you disconnect a FortiGate unit from a cluster, you can re-connect the disconnected FortiGate unit to the cluster by setting the HA mode of the disconnected unit to match the HA mode of the cluster. Usually the disconnected unit rejoins the cluster as a subordinate unit and the cluster automatically synchronizes its configuration.



You do not have to change the HA password on the disconnected unit unless the HA password has been changed after the unit was disconnected. Disconnecting a unit from a cluster does not change the HA password.



You should make sure that the device priority of the disconnected unit is lower than the device priority of the current primary unit. You should also make sure that the HA *override* CLI option is not enabled on the disconnected unit. Otherwise, when the disconnected unit joins the cluster, the cluster will renegotiate and the disconnected unit may become the primary unit. If this happens, the configuration of the disconnected unit is synchronized to all other cluster units. This configuration change might disrupt the operation of the cluster.

The following procedure assumes that the disconnected FortiGate unit is correctly physically connected to your network and to the cluster but is not running in HA mode and not part of the cluster.

Before you start this procedure you should note the device priority of the primary unit.

### To add a disconnected FortiGate unit back to its cluster - web-based manager

1. Log into the disconnected FortiGate unit.  
If virtual domains are enabled, log in as the admin administrator and select Global Configuration.
2. Go to *System > Config > HA*.
3. Change Mode to match the mode of the cluster.
4. If required, change the group name and password to match the cluster.
5. Set the Device Priority lower than the device priority of the primary unit.
6. Select OK.

The disconnected FortiGate unit joins the cluster.

### To add a disconnected FortiGate unit back to its cluster - CLI

1. Log into the CLI of the FortiGate unit to be added back to the cluster.
2. Enter the following command to access the global configuration and add the FortiGate unit back to a cluster operating in active-passive mode and set the device priority to 50 (a low number) so that this unit will not become the primary unit:

```
config global
 config system ha
 set mode a-p
 set priority 50
 end
end
```

You may have to also change the group name, group id and password. However if you have not changed these for the cluster or the FortiGate unit after it was disconnected from the cluster you should not have to adjust them now.

## HA diagnose commands

You can use the following diagnose command to display a data about a cluster:

```
diagnose sys ha dump-by {all-xdb | all-vcluster| rcache | all-group |
memory | debug-zone | vdom | kernel | device | stat| sesync}
```

The example out put below is from a cluster of two FortiGate-5001Cs. In this cluster the base1 and base2 interfaces communicate the HA heartbeat and port monitoring has been added to poort1.

## all-xdb

This command displays information about the current configuration of the cluster and how its operating. You can use the out to determine the primary unit, the state of port monitoring as well as most cluster configuration details and status.

```
diagnose sys ha dump-by all-xdb
 HA information.
idx=1,nxentry=2,linkfails=7,flags=0,digest=7.72.e3.2e.8e.d1...
xentry FG-5KC3E13800046 nhbdev=2,nventry=0, hops=0.
 base1, 50, mac=0.9.f,bc.e.6c, neighbor=1.
 id=FG-5KC3E13800084, mac=0.9.f,bc.11.18.
 base2, 50, mac=0.9.f,bc.e.71, neighbor=1.
 id=FG-5KC3E13800084, mac=0.9.f,bc.11.1d.

xentry FG-5KC3E13800084 nhbdev=2,nventry=1, hops=1.
 base1, 50, mac=0.9.f,bc.11.18, neighbor=1.
 id=FG-5KC3E13800046, mac=0.9.f,bc.e.6c.
 base2, 50, mac=0.9.f,bc.11.1d, neighbor=1.
 id=FG-5KC3E13800046, mac=0.9.f,bc.e.71.
 npath=1,FG-5KC3E13800084
ventry
 idx=0,id=1,FG-5KC3E13800084,prio=128,0,claimed=0,override=0,flag=0,time=12974,mon=0
 mondev=port1,50

idx=0,nxentry=2,linkfails=7,flags=3,digest=7.95.b.9.a8.5d...
xentry FG-5KC3E13800084 nhbdev=2,nventry=1, hops=0.
 base1, 50, mac=0.9.f,bc.11.18, neighbor=1.
 id=FG-5KC3E13800046, mac=0.9.f,bc.e.6c.
 base2, 50, mac=0.9.f,bc.11.1d, neighbor=1.
 id=FG-5KC3E13800046, mac=0.9.f,bc.e.71.
ventry
 idx=0,id=1,FG-5KC3E13800084,prio=128,0,claimed=0,override=0,flag=0,time=12974,mon=0
 mondev=port1,50

xentry FG-5KC3E13800046 nhbdev=2,nventry=1, hops=1.
 base1, 50, mac=0.9.f,bc.e.6c, neighbor=1.
 id=FG-5KC3E13800084, mac=0.9.f,bc.11.18.
 base2, 50, mac=0.9.f,bc.e.71, neighbor=1.
 id=FG-5KC3E13800084, mac=0.9.f,bc.11.1d.
 npath=1,FG-5KC3E13800046
ventry
 idx=0,id=1,FG-5KC3E13800046,prio=128,0,claimed=0,override=0,flag=0,time=2,mon=0
 mondev=port1,50
```



## all-vcluster

This command displays the status and configuration of the individual cluster units. You can use the output of this command to determine the primary unit and the status of each cluster unit.

```
diagnose sys ha dump-by all-vcluster
 HA information.
vcluster id=1, nventry=2, state=work, digest=5.f8.d1.63.4d.d2...
ventry
 idx=0, id=1, FG-5KC3E13800046, prio=128, 0, claimed=0, override=0, flag=1, time=0, mon=0
 mondev=port1, 50
ventry
 idx=1, id=1, FG-5KC3E13800084, prio=128, 0, claimed=0, override=0, flag=0, time=12974, mon=0
```

## stat

This command displays some statistics about how well the cluster is functioning. Information includes packet counts, memory use, failed links and ping failures.

```
diagnose sys ha dump-by stat
 HA information.
packet count = 1, memory = 220.
check_linkfails = 0, linkfails = 0, check_pingsvrfails = 2822
bufcnt = -5, bufmem = 0
```

# HA and failover protection

In FortiGate active-passive HA, the FortiGate Clustering Protocol (FGCP) provides failover protection. This means that an active-passive cluster can provide FortiGate services even when one of the cluster units encounters a problem that would result in complete loss of connectivity for a stand-alone FortiGate unit. This failover protection provides a backup mechanism that can be used to reduce the risk of unexpected downtime, especially in a mission-critical environment.

The FGCP supports three kinds of failover protection. Device failover automatically replaces a failed device and restarts traffic flow with minimal impact on the network. Link failover maintains traffic flow if a link fails. Session failover resumes communication sessions with minimal loss of data if a device or link failover occurs.

This chapter describes how FGCP failover protection works and provides detailed NAT/Route and Transparent mode packet flow descriptions.

This chapter contains the following sections:

- [About active-passive failover](#)
- [About active-active failover](#)
- [Device failover](#)
- [HA heartbeat and communication between cluster units](#)
- [Cluster virtual MAC addresses](#)
- [Synchronizing the configuration](#)
- [Synchronizing kernel routing tables](#)
- [Synchronizing IPsec VPN SAs](#)
- [Link failover \(port monitoring or interface monitoring\)](#)
- [Subsecond failover](#)
- [Remote link failover](#)
- [Session failover \(session pick-up\)](#)
- [WAN optimization and HA](#)
- [Failover and attached network equipment](#)
- [Monitoring cluster units for failover](#)
- [NAT/Route mode active-passive cluster packet flow](#)
- [Transparent mode active-passive cluster packet flow](#)
- [Failover performance](#)

## About active-passive failover

To achieve failover protection in an active-passive cluster, one of the cluster units functions as the primary unit, while the rest of the cluster units are subordinate units, operating in an active stand-by mode. The cluster IP addresses and HA virtual MAC addresses are associated with the cluster interfaces of the primary unit. All traffic directed at the cluster is actually sent to and processed by the primary unit.

While the cluster is functioning, the primary unit functions as the FortiGate network security device for the networks that it is connected to. In addition, the primary unit and subordinate

units use the HA heartbeat to keep in constant communication. The subordinate units report their status to the cluster unit and receive and store connection and state table updates.

## Device failure

If the primary unit encounters a problem that is severe enough to cause it to fail, the remaining cluster units negotiate to select a new primary unit. This occurs because all of the subordinate units are constantly waiting to negotiate to become primary units. Only the heartbeat packets sent by the primary unit keep the subordinate units from becoming primary units. Each received heartbeat packet resets negotiation timers in the subordinate units. If this timer is allowed to run out because the subordinate units do not receive heartbeat packets from the primary unit, the subordinate units assume that the primary unit has failed, and negotiate to become primary units themselves.

Using the same FGCP negotiation process that occurs when the cluster starts up, after they determine that the primary unit has failed, the subordinate units negotiate amongst themselves to select a new primary unit. The subordinate unit that wins the negotiation becomes the new primary unit with the same MAC and IP addresses as the former primary unit. The new primary unit then sends gratuitous ARP packets out all of its interfaces to inform attached switches to send traffic to the new primary unit. Sessions then resume with the new primary unit.

## Link failure

If a primary unit interface fails or is disconnected while a cluster is operation, a link failure occurs. When a link failure occurs the cluster units negotiate to select a new primary unit. Since the primary unit has not stopped operating, it participates in the negotiation. The link failure means that a new primary unit must be selected and the cluster unit with the link failure joins the cluster as a subordinate unit.

Just as for a device failover, the new primary unit sends gratuitous arp packets out all of its interfaces to inform attached switches to send traffic to it. Sessions then resume with the new primary unit.

If a subordinate unit experiences a device failure its status in the cluster does not change. However, in future negotiations a cluster unit with a link failure is unlikely to become the primary unit.

## Session failover

If you enable session failover (also called session pickup) for the cluster, during cluster operation the primary unit informs the subordinate units of changes to the primary unit connection and state tables, keeping the subordinate units up-to-date with the traffic currently being processed by the cluster.

After a failover the new primary unit recognizes open sessions that were being handled by the cluster. The sessions continue to be processed by the new primary unit and are handled according to their last known state.

If you leave session pickup disabled, the cluster does not keep track of sessions and after a failover, active sessions have to be restarted or resumed.

## Primary unit recovery

If a primary unit recovers after a device or link failure, it will operate as a subordinate unit, unless the `override` CLI keyword is enabled and its device priority is set higher than the unit priority of other cluster units (see [“HA override” on page 1138](#)).

## About active-active failover

HA failover in a cluster running in active-active mode is similar to active-passive failover described above. Active-active subordinate units are constantly waiting to negotiate to become primary units and, if session failover is enabled, continuously receive connection state information from the primary unit. If the primary unit fails, or one of the primary unit interfaces fails, the cluster units use the same mechanisms to detect the failure and to negotiate to select a new primary unit. If session failover is enabled, the new primary unit also maintains communication sessions through the cluster using the shared connection state table.

Active-active HA load balances sessions among all cluster units. For session failover, the cluster must maintain all of these sessions. To load balance sessions, the functioning cluster uses a load balancing schedule to distribute sessions to all cluster units. The shared connection state table tracks the communication sessions being processed by all cluster units (not just the primary unit). After a failover, the new primary unit uses the load balancing schedule to re-distribute all of the communication sessions recorded in the shared connection state table among all of the remaining cluster units. The connections continue to be processed by the cluster, but possibly by a different cluster unit, and are handled according to their last known state.

## Device failover

The FGCP provides transparent device failover. Device failover is a basic requirement of any highly available system. Device failover means that if a device fails, a replacement device automatically takes the place of the failed device and continues operating in the same manner as the failed device.

In the case of FortiOS HA, the device is the primary unit. If the primary unit fails, device failover ensures that one of the subordinate units in the cluster automatically takes the place of the primary unit and can continue processing network traffic in the same way as the failed primary unit.



Device failover does not maintain communication sessions. After a device failover, communication sessions have to be restarted. To maintain communication sessions, you must enable session failover. See [“Session failover \(session pick-up\)” on page 1330](#).

---

FortiGate HA device failover is supported by the HA heartbeat, virtual MAC addresses, configuration synchronization, route synchronization and IPsec VPN SA synchronization.

The HA heartbeat makes sure that the subordinate units detect a primary unit failure. If the primary unit fails to respond on time to HA heartbeat packets the subordinate units assume that the primary unit has failed and negotiate to select a new primary unit.

The new primary unit takes the place of the failed primary unit and continues functioning in the same way as the failed primary unit. For the new primary unit to continue functioning like the failed primary unit, the new primary unit must be able to reconnect to network devices and the new primary unit must have the same configuration as the failed primary unit.

FortiGate HA uses virtual MAC addresses to reconnect the new primary unit to network devices. The FGCP causes the new primary unit interfaces to acquire the same virtual MAC addresses as the failed primary unit. As a result, the new primary unit has the same network identity as the failed primary unit.

The new primary unit interfaces have different physical connections than the failed primary unit. Both the failed and the new primary unit interfaces are connected to the same switches, but the new primary unit interfaces are connected to different ports on these switches. To make sure

that the switches send packets to the new primary unit, the new primary unit interfaces send gratuitous ARP packets to the connected switches. These gratuitous ARP packets notify the switches that the primary unit MAC and IP addresses are on different switch ports and cause the switches to send packets to the ports connected to the new primary unit. In this way, the new primary unit continues to receive packets that would otherwise have been sent to the failed primary unit.

Configuration synchronization means that the new primary unit always has the same configuration as the failed primary unit. As a result the new primary unit operates in exactly the same way as the failed primary unit. If configuration synchronization were not available the new primary unit may not process network traffic in the same way as the failed primary unit.

Route synchronization synchronizes the primary unit routing table to all subordinate units so that after a failover the new primary unit does not have to form a completely new routing table. IPsec VPN SA synchronization synchronizes IPsec VPN security associations (SAs) and other IPsec session data so that after a failover the new primary unit can resume IPsec tunnels without having to establish new SAs.

## HA heartbeat and communication between cluster units

The HA heartbeat keeps cluster units communicating with each other. The heartbeat consists of hello packets that are sent at regular intervals by the heartbeat interface of all cluster units. These hello packets describe the state of the cluster unit and are used by other cluster units to keep all cluster units synchronized.

HA heartbeat packets are non-TCP packets that use Ethertype values 0x8890, 0x8891, and 0x8892. The default time interval between HA heartbeats is 200 ms. The FGCP uses link-local IP4 addresses in the 169.254.0.x range for HA heartbeat interface IP addresses.

For best results, isolate the heartbeat devices from your user networks by connecting the heartbeat devices to a separate switch that is not connected to any network. If the cluster consists of two FortiGate units you can connect the heartbeat device interfaces directly using a crossover cable. Heartbeat packets contain sensitive information about the cluster configuration. Heartbeat packets may also use a considerable amount of network bandwidth. For these reasons, it is preferable to isolate heartbeat packets from your user networks.

On startup, a FortiGate unit configured for HA operation broadcasts HA heartbeat hello packets from its HA heartbeat interface to find other FortiGate units configured to operate in HA mode. If two or more FortiGate units operating in HA mode connect with each other, they compare HA configurations (HA mode, HA password, and HA group ID). If the HA configurations match, the units negotiate to form a cluster.

While the cluster is operating, the HA heartbeat confirms that all cluster units are functioning normally. The heartbeat also reports the state of all cluster units, including the communication sessions that they are processing.

### Heartbeat interfaces

A heartbeat interface is an Ethernet network interface in a cluster that is used by the FGCP for HA heartbeat communications between cluster units.

To change the HA heartbeat configuration go to *System > Config > HA* and select the *FortiGate interfaces to use as HA heartbeat interfaces*.



Do not use a switch port for the HA heartbeat traffic. This configuration is not supported.

---

From the CLI enter the following command to make port4 and port5 HA heartbeat interfaces and give both interfaces a heartbeat priority of 150:

```
config system ha
 set hbdev port4 150 port5 150
end
```

The following example shows how to change the default heartbeat interface configuration so that the port4 and port1 interfaces can be used for HA heartbeat communication and to give the port4 interface the highest heartbeat priority so that port4 is the preferred HA heartbeat interface.

```
config system ha
 set hbdev port4 100 port1 50
end
```

By default, for most FortiGate models two interfaces are configured to be heartbeat interfaces. You can change the heartbeat interface configuration as required. For example you can select additional or different heartbeat interfaces. You can also select only one heartbeat interface.

In addition to selecting the heartbeat interfaces, you also set the *Priority* for each heartbeat interface. In all cases, the heartbeat interface with the highest priority is used for all HA heartbeat communication. If the interface fails or becomes disconnected, the selected heartbeat interface that has the next highest priority handles all heartbeat communication.

If more than one heartbeat interface has the same priority, the heartbeat interface with the highest priority that is also highest in the heartbeat interface list is used for all HA heartbeat communication. If this interface fails or becomes disconnected, the selected heartbeat interface with the highest priority that is next highest in the list handles all heartbeat communication.

The default heartbeat interface configuration sets the priority of two heartbeat interfaces to 50. You can accept the default heartbeat interface configuration if one or both of the default heartbeat interfaces are connected. You can select different heartbeat interfaces, select more heartbeat interfaces and change heartbeat priorities according to your requirements.

For the HA cluster to function correctly, you must select at least one heartbeat interface and this interface of all of the cluster units must be connected together. If heartbeat communication is interrupted and cannot failover to a second heartbeat interface, the cluster units will not be able to communicate with each other and more than one cluster unit may become a primary unit. As a result the cluster stops functioning normally because multiple devices on the network may be operating as primary units with the same IP and MAC addresses creating a kind of split brain scenario.

The heartbeat interface priority range is 0 to 512. The default priority when you select a new heartbeat interface is 0. The higher the number the higher the priority.

In most cases you can maintain the default heartbeat interface configuration as long as you can connect the heartbeat interfaces together. Configuring HA heartbeat interfaces is the same for virtual clustering and for standard HA clustering.

You can enable heartbeat communications for physical interfaces, but not for VLAN subinterfaces, IPsec VPN interfaces, redundant interfaces, or for 802.3ad aggregate interfaces. You cannot select these types of interfaces in the heartbeat interface list.

Selecting more heartbeat interfaces increases reliability. If a heartbeat interface fails or is disconnected, the HA heartbeat fails over to the next heartbeat interface.

You can select up to 8 heartbeat interfaces. This limit only applies to FortiGate units with more than 8 physical interfaces.

HA heartbeat traffic can use a considerable amount of network bandwidth. If possible, enable HA heartbeat traffic on interfaces used only for HA heartbeat traffic or on interfaces connected to less busy networks.

## Connecting HA heartbeat interfaces

For most FortiGate models if you do not change the heartbeat interface configuration, you can isolate the default heartbeat interfaces of all of the cluster units by connecting them all to the same switch. Use one switch per heartbeat interface. If the cluster consists of two units you can connect the heartbeat interfaces together using crossover cables. For an example of how to connect heartbeat interfaces, see [“Connecting a FortiGate HA cluster” on page 1127](#).

HA heartbeat and data traffic are supported on the same cluster interface. In NAT/Route mode, if you decide to use heartbeat interfaces for processing network traffic or for a management connection, you can assign the interface any IP address. This IP address does not affect HA heartbeat traffic.

In Transparent mode, you can connect the heartbeat interface to your network and enable management access. You would then establish a management connection to the interface using the Transparent mode management IP address. This configuration does not affect HA heartbeat traffic.

## Heartbeat packets and heartbeat interface selection

HA heartbeat hello packets are constantly sent by all of the enabled heartbeat interfaces. Using these hello packets, each cluster unit confirms that the other cluster units are still operating. The FGCP selects one of the heartbeat interfaces to be used for communication between the cluster units. The FGCP selects the heartbeat interface for heartbeat communication based on the linkfail states of the heartbeat interfaces, on the priority of the heartbeat interfaces, and on the interface index.

The FGCP checks the linkfail state of all heartbeat interfaces to determine which ones are connected. The FGCP selects one of these connected heartbeat interfaces to be the one used for heartbeat communication. The FGCP selects the connected heartbeat interface with the highest priority for heartbeat communication.

If more than one connected heartbeat interface has the highest priority the FGCP selects the heartbeat interface with the lowest interface index. The web-based manager lists the FortiGate unit interfaces in alphabetical order. This order corresponds to the interface index order with lowest index at the top and highest at the bottom. If more than one heartbeat interface has the highest priority, the FGCP selects the interface that is highest in the heartbeat interface list (or first in alphabetical order) for heartbeat communication.

If the interface that is processing heartbeat traffic fails or becomes disconnected, the FGCP uses the same criteria to select another heartbeat interface for heartbeat communication. If the original heartbeat interface is fixed or reconnected, the FGCP again selects this interface for heartbeat communication.

The HA heartbeat communicates cluster session information, synchronizes the cluster configuration, synchronizes the cluster routing table, and reports individual cluster member status. The HA heartbeat constantly communicates HA status information to make sure that the cluster is operating properly.

## Interface index and display order

The web-based manager and CLI display interface names in alphanumeric order. For example, the sort order for a FortiGate unit with 10 interfaces (named port1 through port10) places port10 at the bottom of the list:

- port1
- port2 through 9
- port10

However, interfaces are indexed in hash map order, rather than purely by alphabetic order or purely by interface number value comparisons. As a result, the list is sorted primarily alphabetical by interface name (for example, base1 is before port1), then secondarily by index numbers:

- port1
- port10
- port2 through port9

## HA heartbeat interface IP addresses

The FGCP uses link-local IP4 addresses ([RFC 3927](#)) in the 169.254.0.x range for HA heartbeat interface IP addresses and for inter-VDOM link interface IP addresses. When a cluster initially starts up, the primary unit heartbeat interface IP address is 169.254.0.1. Subordinate units are assigned heartbeat interface IP addresses in the range 169.254.0.2 to 169.254.0.63. HA inter-VDOM link interfaces on the primary unit are assigned IP addresses 169.254.0.65 and 169.254.0.66.

The ninth line of the following CLI command output shows the HA heartbeat interface IP address of the primary unit.

```
get system ha status
 Model: 620
 Mode: a-p
 Group: 0
 Debug: 0
 ses_pickup: disable
 Master:150 head_office_upper FG600B3908600825 1
 Slave :150 head_office_lower FG600B3908600705 0
 number of vcluster: 1
 vcluster 1: work 169.254.0.1
 Master:0 FG600B3908600825
 Slave :1 FG600B3908600705
```

You can also use the `execute traceroute` command from the subordinate unit CLI to display HA heartbeat IP addresses and the HA inter-VDOM link IP addresses. For example, use `execute ha manage 1` to connect to the subordinate unit CLI and then enter the following command to trace the route to an IP address on your network:

```
execute traceroute 172.20.20.10
 traceroute to 172.20.20.10 (172.20.20.10), 32 hops max, 72 byte packets
 1 169.254.0.1 0 ms 0 ms 0 ms
 2 169.254.0.66 0 ms 0 ms 0 ms
 3 172.20.20.10 0 ms 0 ms 0 ms
```



Both HA heartbeat and data traffic are supported on the same FortiGate interface. All heartbeat communication takes place on a separate VDOM called `vsys_ha`. Heartbeat traffic uses a virtual interface called `port_ha` in the `vsys_ha` VDOM. Data and heartbeat traffic use the same physical interface, but they're logically separated into separate VDOMs.

## Heartbeat packet Ethertypes

Normal IP packets are 802.3 packets that have an Ethernet type (Ethertype) field value of 0x0800. Ether type values other than 0x0800 are understood as level2 frames rather than IP packets.

By default, HA heartbeat packets use the following Ethertypes:

- HA heartbeat packets for NAT/Route mode clusters use Ether type 0x8890. These packets are used by cluster units to find other cluster units and to verify the status of other cluster units while the cluster is operating. You can change the Ether type of these packets using the `ha-eth-type` option of the `config system ha` command.
- HA heartbeat packets for Transparent mode clusters use Ether type 0x8891. These packets are used by cluster units to find other cluster units and to verify the status of other cluster units while the cluster is operating. You can change the Ether type of these packets using the `hc-eth-type` option of the `config system ha` command.
- HA telnet sessions between cluster units over HA heartbeat links use Ether type 0x8893. The telnet sessions are used to synchronize the cluster configurations. Telnet sessions are also used when an administrator uses the `execute ha manage` command to connect from one cluster unit CLI to another. You can change the Ether type of these packets using the `l2ep-eth-type` option of the `config system ha` command.

Because heartbeat packets are recognized as level2 frames, the switches and routers on your heartbeat network that connect to heartbeat interfaces must be configured to allow them. If level2 frames are dropped by these network devices, heartbeat traffic will not be allowed between the cluster units.

Some third-party network equipment may use packets with these Ethertypes for other purposes. For example, Cisco N5K/Nexus switches use Ether type 0x8890 for some functions. When one of these switches receives Ether type 0x8890 packets from an attached cluster unit, the switch generates CRC errors and the packets are not forwarded. As a result, FortiGate units connected with these switches cannot form a cluster.

In some cases, if the heartbeat interfaces are connected and configured so regular traffic flows but heartbeat traffic is not forwarded, you can change the configuration of the switch that connects the HA heartbeat interfaces to allow level2 frames with Ethertypes 0x8890, 0x8893, and 0x8891 to pass.

Alternatively, you can use the following CLI options to change the Ethertypes of the HA heartbeat packets:

```
config system ha
 set ha-eth-type <ha_ethertype_4-digit_hex>
 set hc-eth-type <hc_ethertype_4-digit_hex>
 set l2ep-eth-type <l2ep_ethertype_4-digit_hex>
end
```

For example, use the following command to change the Ethertype of the HA heartbeat packets from 0x8890 to 0x8895 and to change the Ethertype of HA Telnet session packets from 0x8891 to 0x889f:

```
config system ha
 set ha-eth-type 8895
 set l2ep-eth-type 889f
end
```

## Modifying heartbeat timing

In an HA cluster, if a cluster unit CPU becomes very busy, the cluster unit may not be able to send heartbeat packets on time. If heartbeat packets are not sent on time other units in the cluster may think that the cluster unit has failed and the cluster will experience a failover.

A cluster unit CPU may become very busy if the cluster is subject to a syn flood attack, if network traffic is very heavy, or for other similar reasons. You can use the following CLI commands to configure how the cluster times HA heartbeat packets:

```
config system ha
 set hb-interval <interval_integer>
 set hb-lost-threshold <threshold_integer>
 set helo-holddown <holddown_integer>
end
```

## Changing the lost heartbeat threshold

The lost heartbeat threshold is the number of consecutive heartbeat packets that are not received from another cluster unit before assuming that the cluster unit has failed. The default value is 6, meaning that if the 6 heartbeat packets are not received from a cluster unit then that cluster unit is considered to have failed. The range is 1 to 60 packets.

If the primary unit does not receive a heartbeat packet from a subordinate unit before the heartbeat threshold expires, the primary unit assumes that the subordinate unit has failed.

If a subordinate unit does not receive a heartbeat packet from the primary unit before the heartbeat threshold expires, the subordinate unit assumes that the primary unit has failed. The subordinate unit then begins negotiating to become the new primary unit.

The lower the `hb-lost-threshold` the faster a cluster responds when a unit fails. However, sometimes heartbeat packets may not be sent because a cluster unit is very busy. This can lead to a false positive failure detection. To reduce these false positives you can increase the `hb-lost-threshold`.

Use the following CLI command to increase the lost heartbeat threshold to 12:

```
config system ha
 set hb-lost-threshold 12
end
```

## Changing the heartbeat interval

The heartbeat interval is the time between sending HA heartbeat packets. The heartbeat interval range is 1 to 20 (100\*ms). The heartbeat interval default is 2 (200 ms).

A heartbeat interval of 2 means the time between heartbeat packets is 200 ms. Changing the heartbeat interval to 5 changes the time between heartbeat packets to 500 ms (5 \* 100ms = 500ms).

The HA heartbeat packets consume more bandwidth if the heartbeat interval is short. But if the heartbeat interval is very long, the cluster is not as sensitive to topology and other network changes.

Use the following CLI command to increase the heartbeat interval to 10:

```
config system ha
 set hb-interval 10
end
```

The heartbeat interval combines with the lost heartbeat threshold to set how long a cluster unit waits before assuming that another cluster unit has failed and is no longer sending heartbeat packets. By default, if a cluster unit does not receive a heartbeat packet from a cluster unit for  $6 * 200 = 1200$  milliseconds or 1.2 seconds the cluster unit assumes that the other cluster unit has failed.

You can increase both the heartbeat interval and the lost heartbeat threshold to reduce false positives. For example, increasing the heartbeat interval to 20 and the lost heartbeat threshold to 30 means a failure will be assumed if no heartbeat packets are received after  $30 * 2000$  milliseconds = 60,000 milliseconds, or 60 seconds.

Use the following CLI command to increase the heartbeat interval to 20 and the lost heartbeat threshold to 30:

```
config system ha
 set hb-lost-threshold 20
 set hb-interval 30
end
```

## Changing the time to wait in the hello state

The hello state hold-down time is the number of seconds that a cluster unit waits before changing from hello state to work state. After a failure or when starting up, cluster units operate in the hello state to send and receive heartbeat packets so that all the cluster units can find each other and form a cluster. A cluster unit should change from the hello state to work state after it finds all of the other FortiGate units to form a cluster with. If for some reason all cluster units cannot find each other during the hello state then some cluster units may be joining the cluster after it has formed. This can cause disruptions to the cluster and affect how it operates.

One reason for a delay in all of the cluster units joining the cluster could be the cluster units are located at different sites or if for some other reason communication is delayed between the heartbeat interfaces.

If cluster units are joining your cluster after it has started up or if it takes a while for units to join the cluster you can increase the time that the cluster units wait in the hello state. The hello state hold-down time range is 5 to 300 seconds. The hello state hold-down time default is 20 seconds.

Use the following CLI command to increase the time to wait in the hello state to 1 minute (60 seconds):

```
config system ha
 set helo-holddown 60
end
```

## Enabling or disabling HA heartbeat encryption and authentication

You can enable HA heartbeat encryption and authentication to encrypt and authenticate HA heartbeat packets. HA heartbeat packets should be encrypted and authenticated if the cluster interfaces that send HA heartbeat packets are also connected to your networks.

If HA heartbeat packets are not encrypted the cluster password and changes to the cluster configuration could be exposed and an attacker may be able to sniff HA packets to get cluster information. Enabling HA heartbeat message authentication prevents an attacker from creating false HA heartbeat messages. False HA heartbeat messages could affect the stability of the cluster.

HA heartbeat encryption and authentication are disabled by default. Enabling HA encryption and authentication could reduce cluster performance. Use the following CLI command to enable HA heartbeat encryption and authentication.

```
config system ha
 set authentication enable
 set encryption enable
end
```

HA authentication and encryption uses AES-128 for encryption and SHA1 for authentication.

## Cluster virtual MAC addresses

When a cluster is operating, the FGCP assigns virtual MAC addresses to each primary unit interface. HA uses virtual MAC addresses so that if a failover occurs, the new primary unit interfaces will have the same virtual MAC addresses and IP addresses as the failed primary unit. As a result, most network equipment would identify the new primary unit as the exact same device as the failed primary unit.

If the MAC addresses changed after a failover, the network would take longer to recover because all attached network devices would have to learn the new MAC addresses before they could communicate with the cluster.

If a cluster is operating in NAT/Route mode, the FGCP assigns a different virtual MAC address to each primary unit interface. VLAN subinterfaces are assigned the same virtual MAC address as the physical interface that the VLAN subinterface is added to. Redundant interfaces or 802.3ad aggregate interfaces are assigned the virtual MAC address of the first interface in the redundant or aggregate list.

If a cluster is operating in Transparent mode, the FGCP assigns a virtual MAC address for the primary unit management IP address. Since you can connect to the management IP address from any interface, all of the FortiGate interfaces appear to have the same virtual MAC address.



A MAC address conflict can occur if two clusters are operating on the same network. See [“Diagnosing packet loss with two FortiGate HA clusters in the same broadcast domain” on page 1305](#) for more information.



Subordinate unit MAC addresses do not change. You can verify this by connecting to the subordinate unit CLI and using the `get hardware interface nic` command to display the MAC addresses of each FortiGate interface.

---

When the new primary unit is selected after a failover, the primary unit sends gratuitous ARP packets to update the devices connected to the cluster interfaces (usually layer-2 switches) with the virtual MAC address. Gratuitous ARP packets configure connected network devices to associate the cluster virtual MAC addresses and cluster IP address with primary unit physical interfaces and with the layer-2 switch physical interfaces. This is sometimes called using gratuitous ARP packets (sometimes called GARP packets) to train the network. The gratuitous

ARP packets sent from the primary unit are intended to make sure that the layer-2 switch forwarding databases (FDBs) are updated as quickly as possible.

Sending gratuitous ARP packets is not required for routers and hosts on the network because the new primary unit will have the same MAC and IP addresses as the failed primary unit. However, since the new primary unit interfaces are connected to different switch interfaces than the failed primary unit, many network switches will update their FDBs more quickly after a failover if the new primary unit sends gratuitous ARP packets.

## Changing how the primary unit sends gratuitous ARP packets after a failover

When a failover occurs it is important that the devices connected to the primary unit update their FDBs as quickly as possible to reestablish traffic forwarding.

Depending on your network configuration, you may be able to change the number of gratuitous ARP packets and the time interval between ARP packets to reduce the cluster failover time.

You cannot disable sending gratuitous ARP packets, but you can use the following command to change the number of packets that are sent. For example, enter the following command to send 20 gratuitous ARP packets:

```
config system ha
 set arps 20
end
```

You can use this command to configure the primary unit to send from 1 to 60 ARP packets. Usually you would not change the default setting of 5. In some cases, however, you might want to reduce the number of gratuitous ARP packets. For example, if your cluster has a large number of VLAN interfaces and virtual domains and because gratuitous ARP packets are broadcast, sending a higher number gratuitous ARP packets may generate a lot of network traffic. As long as the cluster still fails over successfully, you could reduce the number of gratuitous ARP packets that are sent to reduce the amount of traffic produced after a failover.

If failover is taking longer than expected, you may be able to reduce the failover time by increasing the number of gratuitous ARP packets sent.

You can also use the following command to change the time interval in seconds between gratuitous ARP packets. For example, enter the following command to change the time between ARP packets to 3 seconds:

```
config system ha
 set arps-interval 3
end
```

The time interval can be in the range of 1 to 20 seconds. The default is 8 seconds between gratuitous ARP packets. Normally you would not need to change the time interval. However, you could decrease the time to be able to send more packets in less time if your cluster takes a long time to failover.

There may also be a number of reasons to set the interval higher. For example, if your cluster has a large number of VLAN interfaces and virtual domains and because gratuitous ARP packets are broadcast, sending gratuitous ARP packets may generate a lot of network traffic. As long as the cluster still fails over successfully you could increase the interval to reduce the amount of traffic produced after a failover.

For more information about gratuitous ARP packets see [RFC 826](#) and [RFC 3927](#).

## Disabling gratuitous ARP packets after a failover

You can use the following command to turn off sending gratuitous ARP packets after a failover:

```
config system ha
 set gratuitous-arps disable
end
```

Sending gratuitous ARP packets is turned on by default.

In most cases you would want to send gratuitous ARP packets because its a reliable way for the cluster to notify the network to send traffic to the new primary unit. However, in some cases, sending gratuitous ARP packets may be less optimal. For example, if you have a cluster of FortiGate units in Transparent mode, after a failover the new primary unit will send gratuitous ARP packets to all of the addresses in its Forwarding Database (FDB). If the FDB has a large number of addresses it may take extra time to send all the packets and the sudden burst of traffic could disrupt the network.

If you choose to disable sending gratuitous ARP packets you must first enable the `link-failed-signal` setting. The cluster must have some way of informing attached network devices that a failover has occurred.

For more information about the `link-failed-signal` setting, see [“Updating MAC forwarding tables when a link failover occurs” on page 1323](#).

## How the virtual MAC address is determined

The virtual MAC address is determined based on following formula:

```
00-09-0f-09-<group-id_hex>-<vcluster_integer><idx>
```

where

`<group-id_hex>` is the HA Group ID for the cluster converted to hexadecimal. [Table 58](#) lists the virtual MAC address set for each group ID.

**Table 58:** HA group ID in integer and hexadecimal format

Integer Group ID	Hexadecimal Group ID
0	00
1	01
2	02
3	03
4	04
...	...
10	0a
11	0b
...	...
63	3f

**Table 58:** HA group ID in integer and hexadecimal format

...	...
255	ff

`<vcluster_integer>` is 0 for virtual cluster 1 and 2 for virtual cluster 2. If virtual domains are not enabled, HA sets the virtual cluster to 1 and by default all interfaces are in the root virtual domain. Including virtual cluster and virtual domain factors in the virtual MAC address formula means that the same formula can be used whether or not virtual domains and virtual clustering is enabled.

`<idx>` is the index number of the interface. Interfaces are numbered from 0 to x (where x is the number of interfaces). Interfaces are numbered according to their has map order. See [“Interface index and display order” on page 1296](#). The first interface has an index of 0. The second interface in the list has an index of 1 and so on.



Only the `<idx>` part of the virtual MAC address is different for each interface. The `<vcluster_integer>` would be different for different interfaces if multiple VDOMs have been added.



Between FortiOS releases interface indexing may change so the virtual MAC addresses assigned to individual FortiGate interfaces may also change.

### Example virtual MAC addresses

An HA cluster with HA group ID unchanged (default=0) and virtual domains not enabled would have the following virtual MAC addresses for interfaces port1 to port12:

- port1 virtual MAC: 00-09-0f-09-00-00
- port10 virtual MAC: 00-09-0f-09-00-01
- port2 virtual MAC: 00-09-0f-09-00-02
- port3 virtual MAC: 00-09-0f-09-00-03
- port4 virtual MAC: 00-09-0f-09-00-04
- port5 virtual MAC: 00-09-0f-09-00-05
- port6 virtual MAC: 00-09-0f-09-00-06
- port7 virtual MAC: 00-09-0f-09-00-07
- port8 virtual MAC: 00-09-0f-09-00-08
- port9 virtual MAC: 00-09-0f-09-00-09
- port11 virtual MAC: 00-09-0f-09-00-0a
- port12 virtual MAC: 00-09-0f-09-00-0b

If the group ID is changed to 34 these virtual MAC addresses change to:

- port1 virtual MAC: 00-09-0f-09-22-00
- port10 virtual MAC: 00-09-0f-09-22-01
- port2 virtual MAC: 00-09-0f-09-22-02
- port3 virtual MAC: 00-09-0f-09-22-03
- port4 virtual MAC: 00-09-0f-09-22-04
- port5 virtual MAC: 00-09-0f-09-22-05
- port6 virtual MAC: 00-09-0f-09-22-06
- port7 virtual MAC: 00-09-0f-09-22-07
- port8 virtual MAC: 00-09-0f-09-22-08
- port9 virtual MAC: 00-09-0f-09-22-09
- port11 virtual MAC: 00-09-0f-09-22-0a
- port12 virtual MAC: 00-09-0f-09-22-0b

A cluster with virtual domains enabled where the HA group ID has been changed to 23, port5 and port 6 are in the root virtual domain (which is in virtual cluster1), and port7 and port8 are in the vdom\_1 virtual domain (which is in virtual cluster 2) would have the following virtual MAC addresses:

port5 interface virtual MAC: 00-09-0f-09-23-05

port6 interface virtual MAC: 00-09-0f-09-23-06

port7 interface virtual MAC: 00-09-0f-09-23-27

port8 interface virtual MAC: 00-09-0f-09-23-28

## Displaying the virtual MAC address

Every FortiGate unit physical interface has two MAC addresses: the current hardware address and the permanent hardware address. The permanent hardware address cannot be changed, it is the actual MAC address of the interface hardware. The current hardware address can be changed. The current hardware address is the address seen by the network. For a FortiGate unit not operating in HA, you can use the following command to change the current hardware address of the port1 interface:

```
config system interface
 edit port1
 set macaddr <mac_address>
 end
end
```

For an operating cluster, the current hardware address of each cluster unit interface is changed to the HA virtual MAC address by the FGCP. The `macaddr` option is not available for a functioning cluster. You cannot change an interface MAC address and you cannot view MAC addresses from the `system interface` CLI command.

You can use the `get hardware nic <interface_name_str>` command to display both MAC addresses for any FortiGate interface. This command displays hardware information for the specified interface. Depending on their hardware configuration, this command may display different information for different interfaces. You can use this command to display the current hardware address as `Current_HWaddr` and the permanent hardware address as `Permanent_HWaddr`. For some interfaces the current hardware address is displayed as `MAC`.



The command displays a great deal of information about the interface so you may have to scroll the output to find the hardware addresses.



You can also use the `diagnose hardware deviceinfo nic <interface_str>` command to display both MAC addresses for any FortiGate interface.

Before HA configuration the current and permanent hardware addresses are the same. For example for one of the units in Cluster\_1:

```
FGT60B3907503171 # get hardware nic internal
.
.
.
MAC: 02:09:0f:78:18:c9
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

During HA operation the current hardware address becomes the HA virtual MAC address, for example for the units in Cluster\_1:

```
FGT60B3907503171 # get hardware nic internal
.
.
.
MAC: 00:09:0f:09:00:02
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

The following command output for Cluster\_2 shows the same current hardware address for port1 as for the internal interface of Cluster\_2, indicating a MAC address conflict.

```
FG300A2904500238 # get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:02
Permanent_HWaddr: 00:09:0F:85:40:FD
.
.
.
```

## Diagnosing packet loss with two FortiGate HA clusters in the same broadcast domain

A network may experience packet loss when two FortiGate HA clusters have been deployed in the same broadcast domain. Deploying two HA clusters in the same broadcast domain can result in packet loss because of MAC address conflicts. The packet loss can be diagnosed by pinging from one cluster to the other or by pinging both of the clusters from a device within the

broadcast domain. You can resolve the MAC address conflict by changing the HA Group ID configuration of the two clusters. The HA Group ID is sometimes also called the Cluster ID.

This section describes a topology that can result in packet loss, how to determine if packets are being lost, and how to correct the problem by changing the HA Group ID.



Packet loss on a network can also be caused by IP address conflicts. Finding and fixing IP address conflicts can be difficult. However, if you are experiencing packet loss and your network contains two FortiGate HA clusters you can use the information in this article to eliminate one possible source of packet loss.

## Changing the HA group ID to avoid MAC address conflicts

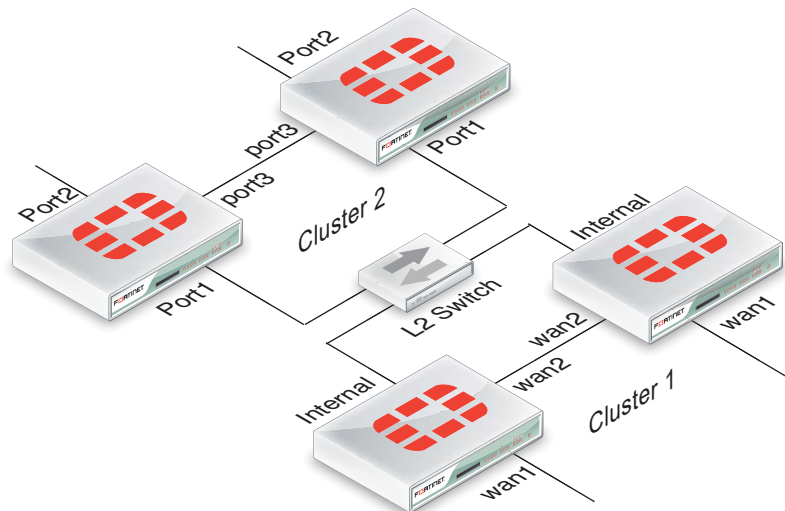
Change the Group ID to change the virtual MAC address of all cluster interfaces. You can change the Group ID from the FortiGate CLI using the following command:

```
config system ha
 set group-id <id_integer>
end
```

## Example topology

The topology below shows two clusters. The Cluster\_1 internal interfaces and the Cluster\_2 port 1 interfaces are both connected to the same broadcast domain. In this topology the broadcast domain could be an internal network. Both clusters could also be connected to the Internet or to different networks.

**Figure 209:**Example HA topology with possible MAC address conflicts



## Ping testing for packet loss

If the network is experiencing packet loss, it is possible that you will not notice a problem unless you are constantly pinging both HA clusters. During normal operation of the network you also might not notice packet loss because the loss rate may not be severe enough to timeout TCP sessions. Also many common types of TCP traffic, such as web browsing, may not be greatly affected by packet loss. However, packet loss can have a significant effect on real time protocols that deliver audio and video data.

To test for packet loss you can set up two constant ping sessions, one to each cluster. If packet loss is occurring the two ping sessions should show alternating replies and timeouts from each cluster.

Cluster_1	Cluster_2
reply	timeout
reply	timeout
reply	timeout
timeout	reply
timeout	reply
reply	timeout
reply	timeout
timeout	reply
timeout	reply
timeout	reply
timeout	reply

### Viewing MAC address conflicts on attached switches

If two HA clusters with the same virtual MAC address are connected to the same broadcast domain (L2 switch or hub), the MAC address will conflict and bounce between the two clusters. This example Cisco switch MAC address table shows the MAC address flapping between different interfaces (1/0/1 and 1/0/4).

1	0009.0f09.0002	DYNAMIC	Gi1/0/1
1	0009.0f09.0002	DYNAMIC	Gi1/0/4

## Synchronizing the configuration

The FGCP uses a combination of incremental and periodic synchronization to make sure that the configuration of all cluster units is synchronized to that of the primary unit.

### Configuration settings that are not synchronized

The following settings are not synchronized among cluster units:

- HA override.
- HA device priority.
- The virtual cluster priority.
- The FortiGate unit host name.
- The HA priority setting for a ping server (or dead gateway detection) configuration.
- The system interface settings of the HA reserved management interface.
- The HA default route for the reserved management interface, set using the `ha-mgt-interface-gateway` option of the `config system ha` command.

The primary unit synchronizes all other configuration settings, including the other HA configuration settings.

## Disabling automatic configuration synchronization

In some cases you may want to use the following command to disable automatic synchronization of the primary unit configuration to all cluster units.

```
config system ha
 set sync-config disable
end
```

When this option is disabled the cluster no longer synchronizes configuration changes. If a device failure occurs, the new primary unit may not have the same configuration as the failed primary unit. As a result, the new primary unit may process sessions differently or may not function on the network in the same way.

In most cases you should not disable automatic configuration synchronization. However, if you have disabled this feature you can use the `execute ha synchronize` command to manually synchronize a subordinate unit's configuration to that of the primary unit.

You must enter `execute ha synchronize` commands from the subordinate unit that you want to synchronize with the primary unit. Use the `execute ha manage` command to access a subordinate unit CLI. See [“Viewing cluster status from the CLI” on page 1277](#).

For example, to access the first subordinate unit and force a synchronization at any time, even if automatic synchronization is disabled enter:

```
execute ha manage 0
execute ha synchronize start
```

You can use the following command to stop a synchronization that is in progress.

```
execute ha synchronize stop
```

You can use the following command to a synchronization all parts of the configuration:

```
execute ha synchronize all
```

Individual options are also available to synchronize parts of the configuration. For example, enter the following command to synchronize CA certificates:

```
execute ha synchronize ca
```

## Incremental synchronization

When you log into the cluster web-based manager or CLI to make configuration changes, you are actually logging into the primary unit. All of your configuration changes are first made to the primary unit. Incremental synchronization then immediately synchronizes these changes to all of the subordinate units.

When you log into a subordinate unit CLI (for example using `execute ha manage`) all of the configuration changes that you make to the subordinate unit are also immediately synchronized to all cluster units, including the primary unit, using the same process.

Incremental synchronization also synchronizes other dynamic configuration information such as the DHCP server address lease database, routing table updates, IPsec SAs, MAC address tables, and so on. See [“FortiGate HA compatibility with PPPoE and DHCP” on page 1142](#) for more information about DHCP server address lease synchronization and [“Synchronizing kernel routing tables” on page 1315](#) for information about routing table updates.

Whenever a change is made to a cluster unit configuration, incremental synchronization sends the same configuration change to all other cluster units over the HA heartbeat link. An HA

synchronization process running on the each cluster unit receives the configuration change and applies it to the cluster unit. The HA synchronization process makes the configuration change by entering a CLI command that appears to be entered by the administrator who made the configuration change in the first place.

Synchronization takes place silently, and no log messages are recorded about the synchronization activity. However, log messages can be recorded by the cluster units when the synchronization process enters CLI commands. You can see these log messages on the subordinate units if you enable event logging and set the minimum severity level to *Information* and then check the event log messages written by the cluster units when you make a configuration change.

You can also see these log messages on the primary unit if you make configuration changes from a subordinate unit.

## Periodic synchronization

Incremental synchronization makes sure that as an administrator makes configuration changes, the configurations of all cluster units remain the same. However, a number of factors could cause one or more cluster units to go out of sync with the primary unit. For example, if you add a new unit to a functioning cluster, the configuration of this new unit will not match the configuration of the other cluster units. Its not practical to use incremental synchronization to change the configuration of the new unit.

Periodic synchronization is a mechanism that looks for synchronization problems and fixes them. Every minute the cluster compares the configuration file checksum of the primary unit with the configuration file checksums of each of the subordinate units. If all subordinate unit checksums are the same as the primary unit checksum, all cluster units are considered synchronized.

If one or more of the subordinate unit checksums is not the same as the primary unit checksum, the subordinate unit configuration is considered out of sync with the primary unit. The checksum of the out of sync subordinate unit is checked again every 15 seconds. This re-checking occurs in case the configurations are out of sync because an incremental configuration sequence has not completed. If the checksums do not match after 5 checks the subordinate unit that is out of sync retrieves the configuration from the primary unit. The subordinate unit then reloads its configuration and resumes operating as a subordinate unit with the same configuration as the primary unit.

The configuration of the subordinate unit is reset in this way because when a subordinate unit configuration gets out of sync with the primary unit configuration there is no efficient way to determine what the configuration differences are and to correct them. Resetting the subordinate unit configuration becomes the most efficient way to resynchronize the subordinate unit.

Synchronization requires that all cluster units run the same FortiOS firmware build. If some cluster units are running different firmware builds, then unstable cluster operation may occur and the cluster units may not be able to synchronize correctly.



Re-installing the firmware build running on the primary unit forces the primary unit to upgrade all cluster units to the same firmware build.

---

## Console messages when configuration synchronization succeeds

When a cluster first forms, or when a new unit is added to a cluster as a subordinate unit, the following messages appear on the CLI console to indicate that the unit joined the cluster and had its configuring synchronized with the primary unit.

```
slave's configuration is not in sync with master's, sequence:0
slave's configuration is not in sync with master's, sequence:1
slave's configuration is not in sync with master's, sequence:2
slave's configuration is not in sync with master's, sequence:3
slave's configuration is not in sync with master's, sequence:4
slave starts to sync with master
logout all admin users
slave succeeded to sync with master
```

## Console messages when configuration synchronization fails

If you connect to the console of a subordinate unit that is out of synchronization with the primary unit, messages similar to the following are displayed.

```
slave is not in sync with master, sequence:0. (type 0x3)
slave is not in sync with master, sequence:1. (type 0x3)
slave is not in sync with master, sequence:2. (type 0x3)
slave is not in sync with master, sequence:3. (type 0x3)
slave is not in sync with master, sequence:4. (type 0x3)
global compared not matched
```

If synchronization problems occur the console message sequence may be repeated over and over again. The messages all include a type value (in the example `type 0x3`). The type value can help Fortinet Support diagnose the synchronization problem.

**Table 59:** HA out of sync object messages and the configuration objects that they reference

Out of Sync Message	Configuration Object
HA_SYNC_SETTING_CONFIGURATION = 0x03	/data/config
HA_SYNC_SETTING_AV = 0x10	
HA_SYNC_SETTING_VIR_DB = 0x11	/etc/vir
HA_SYNC_SETTING_SHARED_LIB = 0x12	/data/lib/libav.so
HA_SYNC_SETTING_SCAN_UNIT = 0x13	/bin/scanunitd
HA_SYNC_SETTING_IMAP_PRXY = 0x14	/bin/imapd
HA_SYNC_SETTING_SMTP_PRXY = 0x15	/bin/smtp
HA_SYNC_SETTING_POP3_PRXY = 0x16	/bin/pop3
HA_SYNC_SETTING_HTTP_PRXY = 0x17	/bin/thttp
HA_SYNC_SETTING_FTP_PRXY = 0x18	/bin/ftpd
HA_SYNC_SETTING_FCNI = 0x19	/etc/fcni.dat

**Table 59:** HA out of sync object messages and the configuration objects that they reference

<b>Out of Sync Message</b>	<b>Configuration Object</b>
HA_SYNC_SETTING_FDNI = 0x1a	/etc/fdnservers.dat
HA_SYNC_SETTING_FSCI = 0x1b	/etc/sci.dat
HA_SYNC_SETTING_FSAE = 0x1c	/etc/fsae_adgrp.cache
HA_SYNC_SETTING_IDS = 0x20	/etc/ids.rules
HA_SYNC_SETTING_IDSUSER_RULES = 0x21	/etc/idsuser.rules
HA_SYNC_SETTING_IDSCUSTOM = 0x22	
HA_SYNC_SETTING_IDS_MONITOR = 0x23	/bin/ipsmonitor
HA_SYNC_SETTING_IDS_SENSOR = 0x24	/bin/ipsengine
HA_SYNC_SETTING_NIDS_LIB = 0x25	/data/lib/libips.so
HA_SYNC_SETTING_WEBLISTS = 0x30	
HA_SYNC_SETTING_CONTENTFILTER = 0x31	/data/cmdb/webfilter.bword
HA_SYNC_SETTING_URLFILTER = 0x32	/data/cmdb/webfilter.urlfilter
HA_SYNC_SETTING_FTGD_OVRD = 0x33	/data/cmdb/webfilter.ftgd-ovrd
HA_SYNC_SETTING_FTGD_LRATING = 0x34	/data/cmdb/webfilter.ftgd-ovrd
HA_SYNC_SETTING_EMAILLISTS = 0x40	
HA_SYNC_SETTING_EMAILCONTENT = 0x41	/data/cmdb/spamfilter.bword
HA_SYNC_SETTING_EMAILBWLIST = 0x42	/data/cmdb/spamfilter.emailbwl
HA_SYNC_SETTING_IPBWL = 0x43	/data/cmdb/spamfilter.ipbwl
HA_SYNC_SETTING_MHEADER = 0x44	/data/cmdb/spamfilter.mheader
HA_SYNC_SETTING_RBL = 0x45	/data/cmdb/spamfilter.rbl
HA_SYNC_SETTING_CERT_CONF = 0x50	/etc/cert/cert.conf
HA_SYNC_SETTING_CERT_CA = 0x51	/etc/cert/ca
HA_SYNC_SETTING_CERT_LOCAL = 0x52	/etc/cert/local
HA_SYNC_SETTING_CERT_CRL = 0x53	/etc/cert/crl

**Table 59:** HA out of sync object messages and the configuration objects that they reference

Out of Sync Message	Configuration Object
HA_SYNC_SETTING_DB_VER = 0x55	
HA_GET_DETAIL_CSUM = 0x71	
HA_SYNC_CC_SIG = 0x75	/etc/cc_sig.dat
HA_SYNC_CC_OP = 0x76	/etc/cc_op
HA_SYNC_CC_MAIN = 0x77	/etc/cc_main
HA_SYNC_FTGD_CAT_LIST = 0x7a	/migadmin/webfilter/ublock/ftgd/ data/

## Comparing checksums of cluster units

You can use the `diagnose sys ha showcsum` command to compare the configuration checksums of all cluster units. The output of this command shows checksums labelled `global` and `all` as well as checksums for each of the VDOMs including the `root` VDOM. The `get system ha-nonsync-csum` command can be used to display similar information; however, this command is intended to be used by FortiManager.

The primary unit and subordinate unit checksums should be the same. If they are not you can use the `execute ha synchronize` command to force a synchronization.

The following command output is for the primary unit of a cluster that does not have multiple VDOMs enabled:

```
diagnose sys ha showcsum
is_manage_master()=1, is_root_master()=1
debugzone
global: a0 7f a7 ff ac 00 d5 b6 82 37 cc 13 3e 0b 9b 77
root: 43 72 47 68 7b da 81 17 c8 f5 10 dd fd 6b e9 57
all: c5 90 ed 22 24 3e 96 06 44 35 b6 63 7c 84 88 d5

checksum
global: a0 7f a7 ff ac 00 d5 b6 82 37 cc 13 3e 0b 9b 77
root: 43 72 47 68 7b da 81 17 c8 f5 10 dd fd 6b e9 57
all: c5 90 ed 22 24 3e 96 06 44 35 b6 63 7c 84 88 d5
```

The following command output is for a subordinate unit of the same cluster:

```
diagnose sys ha showcsum
is_manage_master()=0, is_root_master()=0
debugzone
global: a0 7f a7 ff ac 00 d5 b6 82 37 cc 13 3e 0b 9b 77
root: 43 72 47 68 7b da 81 17 c8 f5 10 dd fd 6b e9 57
all: c5 90 ed 22 24 3e 96 06 44 35 b6 63 7c 84 88 d5

checksum
global: a0 7f a7 ff ac 00 d5 b6 82 37 cc 13 3e 0b 9b 77
root: 43 72 47 68 7b da 81 17 c8 f5 10 dd fd 6b e9 57
all: c5 90 ed 22 24 3e 96 06 44 35 b6 63 7c 84 88 d5
```



The following example shows using this command for the primary unit of a cluster with multiple VDOMs. Two VDOMs have been added named `test` and `Eng_vdm`.

From the primary unit:

```
config global
 sys ha showcsum
 is_manage_master()=1, is_root_master()=1
 debugzone
 global: 65 75 88 97 2d 58 1b bf 38 d3 3d 52 5b 0e 30 a9
 test: a5 16 34 8c 7a 46 d6 a4 1e 1f c8 64 ec 1b 53 fe
 root: 3c 12 45 98 69 f2 d8 08 24 cf 02 ea 71 57 a7 01
 Eng_vdm: 64 51 7c 58 97 79 b1 b3 b3 ed 5c ec cd 07 74 09
 all: 30 68 77 82 a1 5d 13 99 d1 42 a3 2f 9f b9 15 53

checksum
 global: 65 75 88 97 2d 58 1b bf 38 d3 3d 52 5b 0e 30 a9
 test: a5 16 34 8c 7a 46 d6 a4 1e 1f c8 64 ec 1b 53 fe
 root: 3c 12 45 98 69 f2 d8 08 24 cf 02 ea 71 57 a7 01
 Eng_vdm: 64 51 7c 58 97 79 b1 b3 b3 ed 5c ec cd 07 74 09
 all: 30 68 77 82 a1 5d 13 99 d1 42 a3 2f 9f b9 15 53
```

From the subordinate unit:

```
config global
 diagnose sys ha showcsum
 is_manage_master()=0, is_root_master()=0
 debugzone
 global: 65 75 88 97 2d 58 1b bf 38 d3 3d 52 5b 0e 30 a9
 test: a5 16 34 8c 7a 46 d6 a4 1e 1f c8 64 ec 1b 53 fe
 root: 3c 12 45 98 69 f2 d8 08 24 cf 02 ea 71 57 a7 01
 Eng_vdm: 64 51 7c 58 97 79 b1 b3 b3 ed 5c ec cd 07 74 09
 all: 30 68 77 82 a1 5d 13 99 d1 42 a3 2f 9f b9 15 53

checksum
 global: 65 75 88 97 2d 58 1b bf 38 d3 3d 52 5b 0e 30 a9
 test: a5 16 34 8c 7a 46 d6 a4 1e 1f c8 64 ec 1b 53 fe
 root: 3c 12 45 98 69 f2 d8 08 24 cf 02 ea 71 57 a7 01
 Eng_vdm: 64 51 7c 58 97 79 b1 b3 b3 ed 5c ec cd 07 74 09
 all: 30 68 77 82 a1 5d 13 99 d1 42 a3 2f 9f b9 15 53
```

## How to diagnose HA out of sync messages

This section describes how to use the commands `diagnose sys ha showcsum` and `diagnose debug` to diagnose the cause of HA out of sync messages.

If HA synchronization is not successful, use the following procedures on each cluster unit to find the cause.

### To determine why HA synchronization does not occur

1. Connect to each cluster unit CLI by connected to the console port.

2. Enter the following commands to enable debugging and display HA out of sync messages.

```
diagnose debug enable
diagnose debug console timestamp enable
diagnose debug application hataalk -1
diagnose debug application hasync -1
```

Collect the console output and compare the out of sync messages with the information in [Table 59 on page 1310](#).

3. Enter the following commands to turn off debugging.

```
diagnose debug disable
diagnose debug reset
```

### To determine what part of the configuration is causing the problem

If the previous procedure displays messages that include sync object 0x30 (for example, `HA_SYNC_SETTING_CONFIGURATION = 0x03`) there is a synchronization problem with the configuration. Use the following steps to determine the part of the configuration that is causing the problem.

If your cluster consists of two cluster units, use this procedure to capture the configuration checksums for each unit. If your cluster consists of more than two cluster units, repeat this procedure for all cluster units that returned messages that include 0x30 sync object messages.

1. Connect to each cluster unit CLI by connected to the console port.
2. Enter the following command to turn on terminal capture  
`diagnose debug enable`
3. Enter the following command to stop HA synchronization.  
`execute ha sync stop`
4. Enter the following command to display configuration checksums.  
`diagnose sys ha showcsum 1`
5. Copy the output to a text file.
6. Repeat for all affected units.
7. Compare the text file from the primary unit with the text file from each cluster unit to find the checksums that do not match.

You can use a diff function to compare text files.

8. Repeat steps 4 to 7 for each checksum level:

```
diagnose sys ha showcsum 2
diagnose sys ha showcsum 3
diagnose sys ha showcsum 4
diagnose sys ha showcsum 5
diagnose sys ha showcsum 6
diagnose sys ha showcsum 7
diagnose sys ha showcsum 8
```

9. When the non-matching checksum is found, attempt to drill down further. This is possible for objects that have sub-components.

For example you can enter the following commands:

```
diagnose sys ha showcsum system.global
diagnose sys ha showcsum system.interface
```

Generally it is the first non-matching checksum in one of the levels that is the cause of the synchronization problem.

10. Attempt to can remove/change the part of the configuration that is causing the problem. You can do this by making configuration changes from the primary unit or subordinate unit CLI.

11. Enter the following commands to start HA configuration and stop debugging:

```
execute ha sync start
diagnose debug disable
diagnose debug reset
```

## Recalculating the checksums to resolve out of sync messages

Sometimes an error can occur when checksums are being calculated by the cluster. As a result of this calculation error the CLI console could display out of sync error messages even though the cluster is otherwise operating normally. You can also sometimes see checksum calculation errors in `diagnose sys ha showcsum` command output when the checksums listed in the `debugzone` output don't match the checksums in the `checksum` part of the output.

One solution to this problem could be to re-calculate the checksums. The re-calculated checksums should match and the out of sync error messages should stop appearing.

You can use the following command to re-calculate HA checksums:

```
diagnose sys ha csum-recalculate [<vdom-name> | global]
```

Just entering the command without options recalculates all checksums. You can specify a VDOM name to just recalculate the checksums for that VDOM. You can also enter `global` to recalculate the global checksum.

## Synchronizing kernel routing tables

In a functioning cluster, the primary unit keeps all subordinate unit kernel routing tables (also called the forwarding information base FIB) up to date and synchronized with the primary unit. After a failover, because of these routing table updates the new primary unit does not have to populate its kernel routing table before being able to route traffic. This gives the new primary unit time to rebuild its regular routing table after a failover.

Use the following command to view the regular routing table. This table contains all of the configured routes and routes acquired from dynamic routing protocols and so on. This routing table is not synchronized. On subordinate units this command will not produce the same output as on the primary unit.

```
get router info routing-table
```

Use the following command to view the kernel routing table (FIB). This is the list of resolved routes actually being used by the FortiOS kernel. The output of this command should be the same on the primary unit and the subordinate units.

```
get router info kernel
```

This section describes how clusters handle dynamic routing failover and also describes how to use CLI commands to control the timing of routing table updates of the subordinate unit routing tables from the primary unit.

## Configuring graceful restart for dynamic routing failover

When an HA failover occurs, neighbor routers will detect that the cluster has failed and remove it from the network until the routing topology stabilizes. During the time the routers may stop sending IP packets to the cluster and communications sessions that would normally be processed by the cluster may time out or be dropped. Also the new primary unit will not receive routing updates and so will not be able to build and maintain its routing database.

You can configure graceful restart (also called nonstop forwarding (NSF)) as described in [RFC3623](#) (Graceful OSPF Restart) to solve the problem of dynamic routing failover. If graceful

restart is enabled on neighbor routers, they will keep sending packets to the cluster following the HA failover instead of removing it from the network. The neighboring routers assume that the cluster is experiencing a graceful restart.

After the failover, the new primary unit can continue to process communication sessions using the synchronized routing data received from the failed primary unit before the failover. This gives the new primary unit time to update its routing table after the failover.

You can use the following commands to enable graceful restart or NSF on Cisco routers:

```
router ospf 1
 log-adjacency-changes
 nsf ietf helper strict-lsa-checking
```

If the cluster is running BGP, use the following command to enable graceful restart for BGP:

```
config router bgp
 set graceful-restart enable
end
```

You can also add BGP neighbors and configure the cluster unit to notify these neighbors that it supports graceful restart.

```
config router bgp
 config neighbor
 edit <neighbor_address_Ipv4>
 set capability-graceful-restart enable
 end
end
```

If the cluster is running OSPF, use the following command to enable graceful restart for OSFP:

```
config router ospf
 set restart-mode graceful-restart
end
```

To make sure the new primary unit keeps its synchronized routing data long enough to acquire new routing data, you should also increase the HA route time to live, route wait, and route hold values to 60 using the following CLI command:

```
config system ha
 set route-ttl 60
 set route-wait 60
 set route-hold 60
end
```

## Controlling how the FGCP synchronizes kernel routing table updates

You can use the following commands to control some of the timing settings that the FGCP uses when synchronizing kernel routing table updates from the primary unit to subordinate units and maintaining routes on the primary unit after a failover.

```
config system ha
 set route-hold <hold_integer>
 set route-ttl <ttn_integer>
 set route-wait <wait_integer>
end
```

## Change how long routes stay in a cluster unit routing table

Change the `route-ttl` time to control how long routes remain in a cluster unit routing table. The time to live range is 0 to 3600 seconds. The default time to live is 10 seconds.

The time to live controls how long routes remain active in a cluster unit routing table after the cluster unit becomes a primary unit. To maintain communication sessions after a cluster unit becomes a primary unit, routes remain active in the routing table for the route time to live while the new primary unit acquires new routes.

If `route-ttl` is set to 0 the primary unit must acquire all new routes before it can continue processing traffic. By default, `route-ttl` is set to 10 which may mean that only a few routes will remain in the routing table after a failover. Normally keeping `route-ttl` to 10 or reducing the value to 0 is acceptable because acquiring new routes usually occurs very quickly, especially if graceful restart is enabled, so only a minor delay is caused by acquiring new routes.

If the primary unit needs to acquire a very large number of routes, or if for other reasons, there is a delay in acquiring all routes, the primary unit may not be able to maintain all communication sessions.

You can increase the route time to live if you find that communication sessions are lost after a failover so that the primary unit can use synchronized routes that are already in the routing table, instead of waiting to acquire new routes.

## Change the time between routing updates

Change the `route-hold` time to change the time that the primary unit waits between sending routing table updates to subordinate units. The route hold range is 0 to 3600 seconds. The default route hold time is 10 seconds.

To avoid flooding routing table updates to subordinate units, set `route-hold` to a relatively long time to prevent subsequent updates from occurring too quickly. Flooding routing table updates can affect cluster performance if a great deal of routing information is synchronized between cluster units. Increasing the time between updates means that this data exchange will not have to happen so often.

The `route-hold` time should be coordinated with the `route-wait` time.

## Change the time the primary unit waits after receiving a routing update

Change the `route-wait` time to change how long the primary unit waits after receiving routing updates before sending the updates to the subordinate units. For quick routing table updates to occur, set `route-wait` to a relatively short time so that the primary unit does not hold routing table changes for too long before updating the subordinate units.

The `route-wait` range is 0 to 3600 seconds. The default `route-wait` is 0 seconds.

Normally, because the `route-wait` time is 0 seconds the primary unit sends routing table updates to the subordinate units every time its routing table changes.

Once a routing table update is sent, the primary unit waits the `route-hold` time before sending the next update.

Usually routing table updates are periodic and sporadic. Subordinate units should receive these changes as soon as possible so `route-wait` is set to 0 seconds. `route-hold` can be set to a relatively long time because normally the next route update would not occur for a while.

In some cases, routing table updates can occur in bursts. A large burst of routing table updates can occur if a router or a link on a network fails or changes. When a burst of routing table updates occurs, there is a potential that the primary unit could flood the subordinate units with routing table updates. Flooding routing table updates can affect cluster performance if a great deal of routing information is synchronized between cluster units. Setting `route-wait` to a

longer time reduces the frequency of additional updates and prevents flooding of routing table updates from occurring.

## Synchronizing IPsec VPN SAs

The FGCP synchronizes IPsec security associations (SAs) between cluster members so that if a failover occurs, the cluster can resume IPsec sessions without having to establish new SAs. The result is improved failover performance because IPsec sessions are not interrupted to establish new SAs. Also, establishing a large number of SAs can reduce cluster performance.

The FGCP implements slightly different synchronization mechanisms for IKEv1 and IKEv2.

### Synchronizing SAs for IKEv1

When an SA is synchronized to the subordinate units, the sequence number is set to the maximum sequence number. After a failover, all inbound traffic that connects with the new primary unit and uses the SA will be accepted without needing to re-key. However, first outbound packet to use the SA causes the sequence number to overflow and so causes the new primary unit to re-key the SA.

Please note the following:

- The cluster synchronizes all IPsec SAs.
- IPsec SAs are not synchronized until the IKE process has finished synchronizing the ISAKMP SAs. This is required in for dialup tunnels since it is the synchronizing of the ISAKMP SA that creates the dialup tunnel.
- A dialup interface is created as soon as the phase1 is complete. This ensures that the when HA synchronizes phase1 information the dialup name is included.
- If the IKE process re-starts for any reason it deletes any dialup tunnels that exist. This forces the peer to re-key them.
- IPsec SA deletion happens immediately. Routes associated with a dialup tunnel that is being deleted are cleaned up synchronously as part of the delete, rather than waiting for the SA hard-expiry.
- The FGCP does not sync the IPsec tunnel MTU from the primary unit to the subordinate units. This means that after HA failover if the first packet received by the FortiGate unit arrives after the HA route has been deleted and before the new route is added and the packet is larger than the default MTU of 1024 then the FortiGate unit sends back an ICMP fragmentation required. However, as soon as routing is re-established then the MTU will be corrected and traffic will flow.

### Synchronizing SAs for IKEv2

Due to the way the IKEv2 protocol is designed the FGCP cannot use exactly the same solution that is used for synchronizing IKEv1 SAs, though it is similar.

For IKEv2, like IKEv1, the FGCP synchronizes IKE and ISAKMP SAs from the primary unit to the subordinate units. However, for IKEv2 the FGCP cannot actually use this IKE SA to send/receive IKE traffic because IKEv2 includes a sequence number in every IKE message and thus it would require synchronizing every message to the subordinate units to keep the sequence numbers on the subordinate units up to date.

After a failover when the new primary unit accepts incoming IKEv2 sessions, as in IKEv1, the primary unit uses the synchronized SA to decrypt the traffic before passing it through to its destination. For outgoing sessions, because the synchronized SA has an old sequence number, the primary unit negotiates a new SA. This is different from IKEv1 where the existing SA is re-keyed.

Normally for IKEv2 the new primary unit could just negotiate a CHILD\_SA using the synchronized SA. However, because the sequence numbers are not up-to-date, as noted above, the synchronized SA cannot be used and the primary unit must instead negotiate a whole new SA.

## Link failover (port monitoring or interface monitoring)

Link failover means that if a monitored interface fails, the cluster reorganizes to reestablish a link to the network that the monitored interface was connected to and to continue operating with minimal or no disruption of network traffic.

You configure monitored interfaces (also called interface monitoring or port monitoring) by selecting the interfaces to monitor as part of the cluster HA configuration.

You can monitor up to 64 interfaces.

The interfaces that you can monitor appear on the port monitor list. You can monitor all FortiGate interfaces including redundant interfaces and 802.3ad aggregate interfaces.

You cannot monitor the following types of interfaces (you cannot select the interfaces on the port monitor list):

- FortiGate interfaces that contain an internal switch.
- VLAN subinterfaces.
- IPsec VPN interfaces.
- Individual physical interfaces that have been added to a redundant or 802.3ad aggregate interface.
- FortiGate-5000 series backplane interfaces that have not been configured as network interfaces.

If you are configuring a virtual cluster you can create a different port monitor configuration for each virtual cluster. Usually for each virtual cluster you would monitor the interfaces that have been added to the virtual domains in each virtual cluster.



Wait until after the cluster is up and running to enable interface monitoring. You do not need to configure interface monitoring to get a cluster up and running and interface monitoring will cause failovers if for some reason during initial setup a monitored interface has become disconnected. You can always enable interface monitoring once you have verified that the cluster is connected and operating properly.



You should only monitor interfaces that are connected to networks, because a failover may occur if you monitor an unconnected interface.

---

### To enable interface monitoring - web-based manager

Use the following steps to monitor the port1 and port2 interfaces of a cluster.

1. Connect to the cluster web-based manager.
2. Go to *System > Config > HA* and edit the primary unit (*Role* is *MASTER*).
3. Select the *Port Monitor* check boxes for the *port1* and *port2* interfaces and select *OK*.

The configuration change is synchronized to all cluster units.

## To enable interface monitoring - CLI

Use the following steps to monitor the port1 and port2 interfaces of a cluster.

1. Connect to the cluster CLI.
2. Enter the following command to enable interface monitoring for port1 and port2.

```
configure system ha
 set monitor port1 port2
end
```

The following example shows how to enable monitoring for the external, internal, and DMZ interfaces.

```
config system ha
 set monitor external internal dmz
end
```

With interface monitoring enabled, during cluster operation, the cluster monitors each cluster unit to determine if the monitored interfaces are operating and connected. Each cluster unit can detect a failure of its network interface hardware. Cluster units can also detect if its network interfaces are disconnected from the switch they should be connected to.



Cluster units cannot determine if the switch that its interfaces are connected to is still connected to the network. However, you can use remote IP monitoring to make sure that the cluster unit can connect to downstream network devices. See [“Remote link failover” on page 1325](#).

---

Because the primary unit receives all traffic processed by the cluster, a cluster can only process traffic from a network if the primary unit can connect to it. So, if the link between a network and the primary unit fails, to maintain communication with this network, the cluster must select a different primary unit; one that is still connected to the network. Unless another link failure has occurred, the new primary unit will have an active link to the network and will be able to maintain communication with it.

To support link failover, each cluster unit stores link state information for all monitored cluster units in a link state database. All cluster units keep this link state database up to date by sharing link state information with the other cluster units. If one of the monitored interfaces on one of the cluster units becomes disconnected or fails, this information is immediately shared with all cluster units.

### If a monitored interface on the primary unit fails

If a monitored interface on the primary unit fails, the cluster renegotiates to select a new primary unit using the process described in [“Primary unit selection” on page 1131](#). Because the cluster unit with the failed monitored interface has the lowest monitor priority, a different cluster unit becomes the primary unit. The new primary unit should have fewer link failures.

After the failover, the cluster resumes and maintains communication sessions in the same way as for a device failure. See [“Device failover” on page 1292](#).

### If a monitored interface on a subordinate unit fails

If a monitored interface on a subordinate unit fails, this information is shared with all cluster units. The cluster does not renegotiate. The subordinate unit with the failed monitored interface continues to function in the cluster.

In an active-passive cluster after a subordinate unit link failover, the subordinate unit continues to function normally as a subordinate unit in the cluster.




In an active-active cluster after a subordinate unit link failure:

- The subordinate unit with the failed monitored interface can continue processing connections between functioning interfaces. However, the primary unit stops sending sessions to a subordinate unit that use any failed monitored interfaces on the subordinate unit.
- If session pickup is enabled, all sessions being processed by the subordinate unit failed interface that can be are failed over to other cluster units. Sessions that cannot be failed over are lost and have to be restarted.
- If session pickup is not enabled all sessions being processed by the subordinate unit failed interface are lost.

## How link failover maintains traffic flow

Monitoring an interface means that the interface is connected to a high priority network. As a high priority network, the cluster should maintain traffic flow to and from the network, even if a link failure occurs. Because the primary unit receives all traffic processed by the cluster, a cluster can only process traffic from a network if the primary unit can connect to it. So, if the link that the primary unit has to a high priority network fails, to maintain traffic flow to and from this network, the cluster must select a different primary unit. This new primary unit should have an active link to the high priority network.

**Figure 210:**A link failure causes a cluster to select a new primary unit



If a monitored interface on the primary unit fails, the cluster renegotiates and selects the cluster unit with the highest monitor priority to become the new primary unit. The cluster unit with the highest monitor priority is the cluster unit with the most monitored interfaces connected to networks.

After a link failover, the primary unit processes all traffic and all subordinate units, even the cluster unit with the link failure, share session and link status. In addition all configuration changes, routes, and IPsec SAs are synchronized to the cluster unit with the link failure.

In an active-active cluster, the primary unit load balances traffic to all the units in the cluster. The cluster unit with the link failure can process connections between its functioning interfaces (for example if the cluster has connections to an internal, external, and DMZ network, the cluster unit with the link failure can still process connections between the external and DMZ networks).

If a monitored interface on a subordinate unit fails, the subordinate unit shares this information with all cluster units. The cluster does not renegotiate. The subordinate unit with the failed monitored interface continues to function in the cluster. In an active-active cluster, the subordinate unit can continue processing connections between functioning interfaces. The primary unit re-distributes traffic that was being processed by the failed interface of the subordinate unit to other cluster units. If session pickup is enabled, similar to a failover, some of these sessions continue while others must restart. See [“Session failover \(session pick-up\)” on page 1330](#).

## Recovery after a link failover and controlling primary unit selection (controlling falling back to the prior primary unit)

If you find and correct the problem that caused a link failure (for example, re-connect a disconnected network cable) the cluster updates its link state database and re-negotiates to select a primary unit.

What happens next depends on how the cluster configuration affects primary unit selection:

- The former primary unit will once again become the primary unit (falling back to becoming the primary unit)
- The primary unit will not change.

As described in [“Displaying cluster unit age differences” on page 1134](#), when the link is restored, if no options are configured to control primary unit selection and the cluster age difference is less than 300 seconds the former primary unit will once again become the primary unit. If the age differences are greater than 300 seconds then a new primary unit is not selected. Since you have no control on the age difference the outcome can be unpredictable. This is not a problem in cases where its not important which unit becomes the primary unit.

## Preventing a primary unit change after a failed link is restored

Some organizations will not want the cluster to change primary units when the link is restored. Instead they would rather wait to restore the primary unit during a maintenance window. This functionality is not directly supported, but you can experiment with changing some primary unit selection settings. For example, in most cases it should work to enable override on all cluster units and make sure their priorities are the same. This should mean that the primary unit should not change after a failed link is restored.

Then, when you want to restore the original primary unit during a maintenance window you can just set its Device Priority higher. After it becomes the primary unit you can reset all device priorities to the same value. Alternatively during a maintenance window you could reboot the current primary unit and any subordinate units except the one that you want to become the primary unit.

If the `override` CLI keyword is enabled on one or more cluster units and the device priority of a cluster unit is set higher than the others, when the link failure is repaired and the cluster unit with the highest device priority will always become the primary unit.

## Testing link failover

You can test link failure by disconnecting the network cable from a monitored interface of a cluster unit. If you disconnect a cable from a primary unit monitored interface the cluster should renegotiate and select one of the other cluster units as the primary unit. You can also verify that

traffic received by the disconnected interface continues to be processed by the cluster after the failover.

If you disconnect a cable from a subordinate unit interface the cluster will not renegotiate.

## Updating MAC forwarding tables when a link failover occurs

When a FortiGate HA cluster is operating and a monitored interface fails on the primary unit, the primary unit usually becomes a subordinate unit and another cluster unit becomes the primary unit. After a link failover, the new primary unit sends gratuitous ARP packets to refresh the MAC forwarding tables (also called arp tables) of the switches connected to the cluster. This is normal link failover operation (for more information, see [“Link failover \(port monitoring or interface monitoring\)” on page 1319](#)).

Even when gratuitous ARP packets are sent, some switches may not be able to detect that the primary unit has become a subordinate unit and will keep sending packets to the former primary unit. This can occur if the switch does not detect the failure and does not clear its MAC forwarding table.

You have another option available to make sure the switch detects the failover and clears its MAC forwarding tables. You can use the following command to cause a cluster unit with a monitored interface link failure to briefly shut down all of its interfaces (except the heartbeat interfaces) after the failover occurs:

```
config system ha
 set link-failed-signal enable
end
```

Usually this means each interface of the former primary unit is shut down for about a second. When this happens the switch should be able to detect this failure and clear its MAC forwarding tables of the MAC addresses of the former primary unit and pickup the MAC addresses of the new primary unit. Each interface will shut down for a second but the entire process usually takes a few seconds. The more interfaces the FortiGate unit has, the longer it will take.

Normally, the new primary unit also sends gratuitous ARP packets that also help the switch update its MAC forwarding tables to connect to the new primary unit. If `link-failed-signal` is enabled, sending gratuitous ARP packets is optional and can be disabled if you don't need it or if its causing problems. See [“Disabling gratuitous ARP packets after a failover” on page 1302](#).

## Multiple link failures

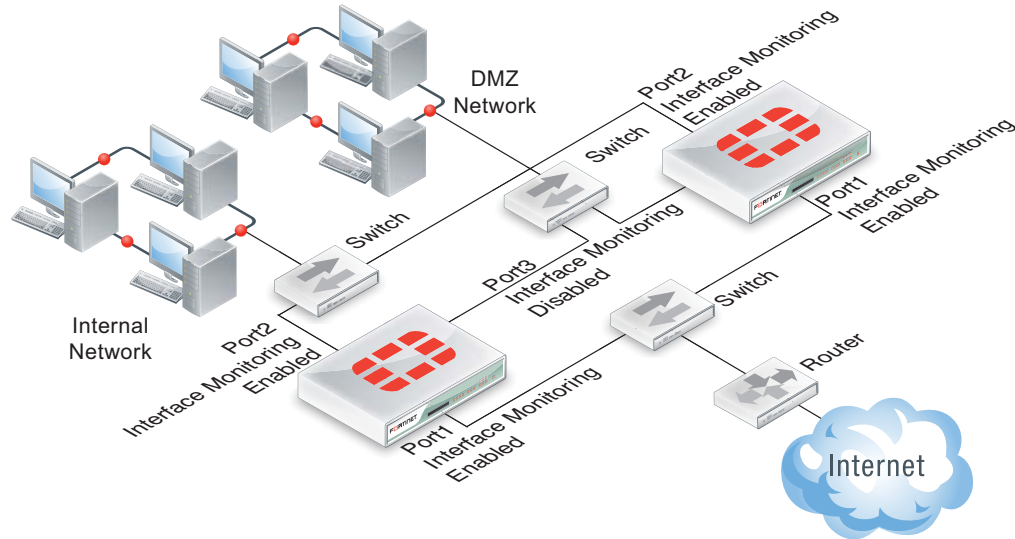
Every time a monitored interface fails, the cluster repeats the processes described above. If multiple monitored interfaces fail on more than one cluster unit, the cluster continues to negotiate to select a primary unit that can provide the most network connections.

## Example link failover scenarios

For the following examples, assume a cluster configuration consisting of two FortiGate units (FGT\_1 and FGT\_2) connected to three networks: internal using port2, external using port1, and DMZ using port3. In the HA configuration, the device priority of FGT\_1 is set higher than the unit priority of FGT\_2.

The cluster processes traffic flowing between the internal and external networks, between the internal and DMZ networks, and between the external and DMZ networks. If there are no link failures, FGT1 becomes the primary unit because it has the highest device priority.

**Figure 211:**Sample link failover scenario topology



### Example: the port1 link on FGT\_1 fails

If the port1 link on FGT\_1 fails, FGT\_2 becomes primary unit because it has fewer interfaces with a link failure. If the cluster is operating in active-active mode, the cluster load balances traffic between the internal network (port2) and the DMZ network (port3). Traffic between the Internet (port1) and the internal network (port2) and between the Internet (port1) and the DMZ network (port3) is processed by the primary unit only.

### Example: port2 on FGT\_1 and port1 on FGT\_2 fail

If port2 on FGT\_1 and port1 on FGT\_2 fail, then FGT\_1 becomes the primary unit. After both of these link failures, both cluster units have the same monitor priority. So the cluster unit with the highest device priority (FGT\_1) becomes the primary unit.

Only traffic between the Internet (port1) and DMZ (port3) networks can pass through the cluster and the traffic is handled by the primary unit only. No load balancing will occur if the cluster is operating in active-active mode.

## Subsecond failover

HA link failover supports subsecond failover (that is a failover time of less than one second). Subsecond failover is available for interfaces that can issue a link failure system call when the interface goes down. When an interface experiences a link failure and sends the link failure system call, the FGCP receives the system call and initiates a link failover.

For interfaces that do not support subsecond failover, port monitoring regularly polls the connection status of monitored interfaces. When a check finds that an interface has gone down, port monitoring causes a link failover. Subsecond failover results in a link failure being detected sooner because the system doesn't have to wait for the next poll to find out about the failure.

Subsecond failover requires interfaces that support sending the link failure system call. This functionality is available for:

- Interfaces with network processors (NPx)
- Interfaces with content processors (CP4, CP5, CP6, etc.)
- Interfaces in Fortinet Mezzanine Cards that include network and content processors (FMC-XD2, FMC-XG2, etc.)
- Accelerated interface modules (FortiGate-ASM-FB4, ADM-FB8, ADM-XB2, ADM-XD4, RTM-XD2 etc).
- Interfaces in security processor modules (FortiGate-ASM-CE4, ASM-XE2, etc)

Subsecond failover can accelerate HA failover to reduce the link failover time to less than one second under ideal conditions. Actual failover performance may vary depending on traffic patterns and network configuration. For example, some network devices may respond slowly to an HA failover.

No configuration changes are required to support subsecond failover. However, for best subsecond failover results, the recommended heartbeat interval is 100ms and the recommended lost heartbeat threshold is 5. (See [“Changing the heartbeat interval” on page 1298](#))

```
config system ha
 set hb-lost-threshold 5
 set hb-interval 1
end
```

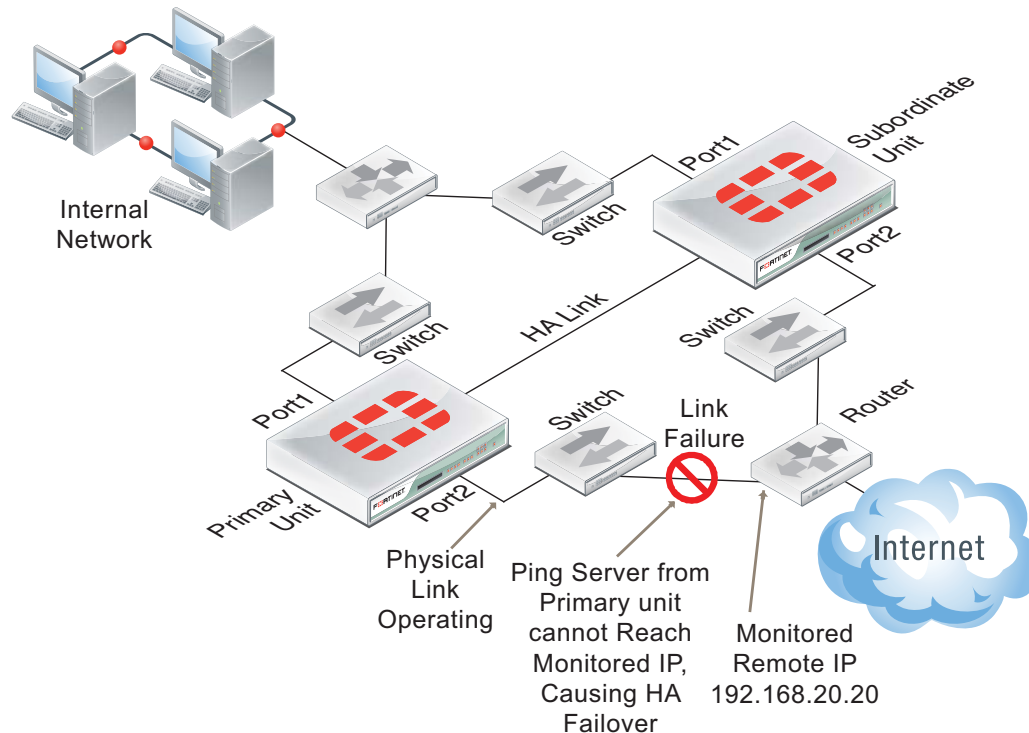
For information about how to reduce failover times, see [“Failover performance” on page 1340](#).

## Remote link failover

Remote link failover (also called remote IP monitoring) is similar to HA port monitoring and interface dead gateway detection. Port monitoring causes a cluster to failover if a monitored primary unit interface fails or is disconnected. Remote IP monitoring uses ping servers configured for FortiGate interfaces on the primary unit to test connectivity with IP addresses of network devices. Usually these would be IP addresses of network devices not directly connected to the cluster. For example, a downstream router. Remote IP monitoring causes a failover if one or more of these remote IP addresses does not respond to a ping server.

By being able to detect failures in network equipment not directly connected to the cluster, remote IP monitoring can be useful in a number of ways depending on your network configuration. For example, in a full mesh HA configuration, with remote IP monitoring, the cluster can detect failures in network equipment that is not directly connected to the cluster but that would interrupt traffic processed by the cluster if the equipment failed.

**Figure 212:**Example HA remote IP monitoring topology



In the simplified example topology shown in [Figure 212](#), the switch connected directly to the primary unit is operating normally but the link on the other side of the switches fails. As a result traffic can no longer flow between the primary unit and the Internet.

To detect this failure you can create a remote IP monitoring configuration consisting of a ping server dead gateway detection configuration for port2 of the cluster. The primary unit tests connectivity to 192.168.20.20. If the ping server cannot connect to 192.268.20.20 the cluster to fails over and the subordinate unit becomes the new primary unit. The remote HA monitoring ping server on the new primary unit can connect to 192.168.20.20 so the failover maintains connectivity between the internal network and the Internet through the cluster.

### To configure remote IP monitoring

1. Enter the following commands to configure HA remote monitoring for the example topology.

- Enter the `pingserver-monitor-interface` keyword to enable HA remote IP monitoring on port2.
- Leave the `pingserver-failover-threshold` set to the default value of 0. You can change this value if you do not want a failover to occur if only one ping server fails.
- Enter the `pingserver-flip-timeout` keyword to set the flip timeout to 120 minutes. After a failover, if HA remote IP monitoring on the new primary unit also causes a failover, the flip timeout prevents the failover from occurring until the timer runs out. Setting the `pingserver-flip-timeout` to 120 means that remote IP monitoring can only cause a failover every 120 minutes. This flip timeout is required to prevent repeating failovers if remote IP monitoring causes a failover from all cluster units because none of the cluster units can connect to the monitored IP addresses.

```
config system ha
 set pingserver-monitor-interface port2
 set pingserver-failover-threshold 0
 set pingserver-flip-timeout 120
end
```

2. Enter the following commands to add the ping server for the port2 interface and to set the HA remote IP monitoring priority for this ping server.
  - Enter the `detectserver` keyword to add the ping server and set the ping server IP address to 192.168.20.20.
  - Leave the `ha-priority` keyword set to the default value of 1. You only need to change this priority if you change the HA ping server failover threshold.



The `ha-priority` setting is not synchronized among cluster units. So if you want to change the `ha-priority` setting you must change it separately on each cluster unit. Otherwise it will remain set to the default value of 1.

- Use the `interval` keyword to set the time between ping server pings and use the `failtime` keyword to set the number of times that the ping can fail before a failure is detected (the failover threshold). The following example reduces the failover threshold to 2 but keeps the ping interval at the default value of 5.

```
config router gwdetect
 edit port2
 set server 192.168.20.20
 set ha-priority 1
 set interval 5
 set failtime 2
 end
```



You can also do this from the web-based manager by going to *Router > Static > Settings*, selecting *Create New* to add a new dead gateway detection configuration, setting *Ping Server* to 192.168.20.20, *HA Priority* to 1, *Ping Interval* to 5, and *Failover Threshold* to 2.

## Adding HA remote IP monitoring to multiple interfaces

You can enable HA remote IP monitoring on multiple interfaces by adding more interface names to the `pingserver-monitor-interface` keyword. If your FortiGate configuration includes VLAN interfaces, aggregate interfaces and other interface types, you can add the names of these interfaces to the `pingserver-monitor-interface` keyword to configure HA remote IP monitoring for these interfaces.

For example, enable remote IP monitoring for interfaces named port2, port20, and vlan\_234:

```
config system ha
 set pingserver-monitor-interface port2 port20 vlan_234
 set pingserver-failover-threshold 10
 set pingserver-flip-timeout 120
end
```

Then configure ping servers for each of these interfaces. In the following example, default values are accepted for all settings other than the server IP address.

```
config router gwdetect
 edit port2
 set server 192.168.20.20
 next
 edit port20
 set server 192.168.20.30
 next
 edit vlan_234
 set server 172.20.12.10
end
```

## Changing the ping server failover threshold

By default the ping server failover threshold is 0 and the HA priority is 1 so any HA remote IP monitoring ping server failure causes a failover. If you have multiple ping servers you may want a failover to occur only if more than one of them has failed.

For example, you may have 3 ping servers configured on three interfaces but only want a failover to occur if two of the ping servers fail. To do this you must set the HA priorities of the ping servers and the HA ping server failover threshold so that the priority of one ping server is less than the failover threshold but the added priorities of two ping servers is equal to or greater than the failover threshold. Failover occurs when the HA priority of all failed ping servers reaches or exceeds the threshold.

For example, set the failover threshold to 10 and monitor three interfaces:

```
config system ha
 set pingserver-monitor-interface port2 port20 vlan_234
 set pingserver-failover-threshold 10
 set pingserver-flip-timeout 120
end
```

Then set the HA priority of each ping server to 5.



The HA Priority (`ha-priority`) setting is not synchronized among cluster units. In the following example, you must set the HA priority to 5 by logging into each cluster unit.

---



```

config router gwdetect
 edit port2
 set server 192.168.20.20
 set ha-priority 5
 next
 edit port20
 set server 192.168.20.30
 set ha-priority 5
 next
 edit vlan_234
 set server 172.20.12.10
 set ha-priority 5
end

```

If only one of the ping servers fails, the total ping server HA priority will be 5, which is lower than the failover threshold so a failover will not occur. If a second ping server fails, the total ping server HA priority of 10 will equal the failover threshold, causing a failover.

By adding multiple ping servers to the remote HA monitoring configuration and setting the HA priorities for each, you can fine tune remote IP monitoring. For example, if it is more important to maintain connections to some networks you can set the HA priorities higher for these ping servers. And if it is less important to maintain connections to other networks you can set the HA priorities lower for these ping servers. You can also adjust the failover threshold so that if the cluster cannot connect to one or two high priority IP addresses a failover occurs. But a failover will not occur if the cluster cannot connect to one or two low priority IP addresses.

## Monitoring multiple IP addresses from one interface

You can add multiple IP addresses to a single ping server to use HA remote IP monitoring to monitor more than one IP address from a single interface. If you add multiple IP addresses, the ping will be sent to all of the addresses at the same time. The ping server only fails when no responses are received from any of the addresses.

```

config router gwdetect
 edit port2
 set server 192.168.20.20 192.168.20.30 172.20.12.10
 end

```

## Flip timeout

The HA remote IP monitoring configuration also involves setting a flip timeout. The flip timeout is required to reduce the frequency of failovers if, after a failover, HA remote IP monitoring on the new primary unit also causes a failover. This can happen if the new primary unit cannot connect to one or more of the monitored remote IP addresses. The result could be that until you fix the network problem that blocks connections to the remote IP addresses, the cluster will experience repeated failovers. You can control how often the failovers occur by setting the flip timeout. The flip timeout stops HA remote IP monitoring from causing a failover until the primary unit has been operating for the duration of the flip timeout.

If you set the flip timeout to a relatively high number of minutes you can find and repair the network problem that prevented the cluster from connecting to the remote IP address without the cluster experiencing very many failovers. Even if it takes a while to detect the problem, repeated failovers at relatively long time intervals do not usually disrupt network traffic.

## Detecting HA remote IP monitoring failovers

Just as with any HA failover, you can detect HA remote IP monitoring failovers by using SNMP to monitor for HA traps. You can also use alert email to receive notifications of HA status changes and monitor log messages for HA failover log messages. In addition, FortiGate units send the critical log message `Ping Server is down` when a ping server fails. The log message includes the name of the interface that the ping server has been added to.

## Session failover (session pick-up)

Session failover means that a cluster maintains active network TCP and IPsec VPN sessions (including NAT sessions) after a device or link failover. You can also configure session failover to maintain UDP and ICMP sessions. Session failover does not failover multicast, or SSL VPN sessions.

FortiGate HA does not support session failover by default. To enable session failover go to `System > Config > HA` and select *Enable Session Pick-up*.

From the CLI enter:

```
config system ha
 set session-pickup enable
end
```

To support session failover, when *Enable Session Pick-up* is selected, the FGCP maintains an HA session table for most TCP communication sessions being processed by the cluster and synchronizes this session table with all cluster units. If a cluster unit fails, the HA session table information is available to the remaining cluster units and these cluster units use this session table to resume most of the TCP sessions that were being processed by the failed cluster unit without interruption.

If session pickup is enabled, you can use the following command to also enable UDP and ICMP session failover:

```
config system ha
 set session-pickup-connectionless enable
end
```

You must enable session pickup for session failover protection. If you do not require session failover protection, leaving session pickup disabled may reduce CPU usage and reduce HA heartbeat network bandwidth usage.

## If session pickup is not selected

If *Enable Session Pick-up* is not selected, the FGCP does not maintain an HA session table and most TCP sessions do not resume after a failover. After a device or link failover all sessions are briefly interrupted and must be re-established at the application level after the cluster renegotiates.

Many protocols can successfully restart sessions with little, if any, loss of data. For example, after a failover, users browsing the web can just refresh their browsers to resume browsing. Since most HTTP sessions are very short, in most cases they will not even notice an interruption unless they are downloading large files. Users downloading a large file may have to restart their download after a failover.

Other protocols may experience data loss and some protocols may require sessions to be manually restarted. For example, a user downloading files with FTP may have to either restart downloads or restart their FTP client.

Some sessions may resume after a failover whether or not enable session pick-up is selected:

- [“UDP, ICMP, multicast and broadcast packet session failover” on page 1334](#)
- [“FortiOS Carrier GTP session failover” on page 1334](#)
- [“Active-active HA subordinate units sessions can resume after a failover” on page 1335.](#)

## Improving session synchronization performance

Two HA configuration options are available to reduce the performance impact of enabling session pickup. They include reducing the number of sessions that are synchronized by adding a session pickup delay and using more FortiGate interfaces for session synchronization.

### Reducing the number of sessions that are synchronized

Enable the `session-pickup-delay` CLI option to reduce the number of sessions that are synchronized by synchronizing sessions only if they remain active for more than 30 seconds. Enabling this option could greatly reduce the number of sessions that are synchronized if a cluster typically processes very many short duration sessions, which is typical of most HTTP traffic for example.

Use the following command to enable a 30 second session pickup delay:

```
config system ha
 set session-pickup-delay enable
end
```

Enabling session pickup delay means that if a failover occurs more sessions may not be resumed after a failover. In most cases short duration sessions can be restarted with only a minor traffic interruption. However, if you notice too many sessions not resuming after a failover you might want to disable this setting.

### Using multiple FortiGate interfaces for session synchronization

Using the `session-sync-dev` option you can select one or more FortiGate interfaces to use for synchronizing sessions as required for session pickup. Normally session synchronization occurs over the HA heartbeat link. Using this HA option means only the selected interfaces are used for session synchronization and not the HA heartbeat link. If you select more than one interface, session synchronization traffic is load balanced among the selected interfaces.

Moving sessions synchronization from the HA heartbeat interface reduces the bandwidth requirements of the HA heartbeat interface and may improve the efficiency and performance of the cluster, especially if the cluster is synchronizing a large number of sessions. Load balancing session synchronization among multiple interfaces can further improve performance and efficiency if the cluster is synchronizing a large number of sessions.

Use the following command to perform cluster session synchronization using the port10 and port12 interfaces.

```
config system ha
 set session-sync-dev port10 port12
end
```

Session synchronization packets use Ethertype 0x8892. The interfaces to use for session synchronization must be connected together either directly using the appropriate cable (possible if there are only two units in the cluster) or using switches. If one of the interfaces becomes disconnected the cluster uses the remaining interfaces for session synchronization. If all of the session synchronization interfaces become disconnected, session synchronization reverts back to using the HA heartbeat link. All session synchronization traffic is between the primary unit and each subordinate unit.

Since large amounts of session synchronization traffic can increase network congestion, it is recommended that you keep this traffic off of your network by using dedicated connections for it.

## Session failover not supported for all sessions

Most of the features applied to sessions by FortiGate UTM functionality require the FortiGate unit to maintain very large amounts of internal state information for each session. The FGCP does not synchronize internal state information for the following UTM features, so the following types of sessions will not resume after a failover:

- Virus scanning of HTTP, HTTPS, FTP, IMAP, IMAPS, POP3, POP3S, SMTP, SMTPS, IM, CIFS, and NNTP sessions,
- Web filtering and FortiGuard Web Filtering of HTTP and HTTPS sessions,
- Spam filtering of IMAP, IMAPS, POP3, POP3S, SMTP, and SMTPS sessions,
- DLP scanning of IMAP, IMAPS, POP3, POP3S, SMTP, SMTPS, SIP, SIMPLE, and SCCP sessions,
- DLP archiving of HTTP, HTTPS, FTP, IMAP, IMAPS, POP3, SMTP, SMTPS, IM, NNTP, AIM, ICQ, MSN, Yahoo! IM, SIP, SIMPLE, and SCCP signal control sessions,



Active-active clusters can resume some of these sessions after a failover. See [“Active-active HA subordinate units sessions can resume after a failover” on page 1335](#) for details.

---

If you use these features to protect most of the sessions that your cluster processes, enabling session failover may not actually provide significant session failover protection.

TCP sessions that are not being processed by these UTM features resume after a failover even if these sessions are accepted by security policies with UTM options configured. Only TCP sessions that are actually being processed by these UTM features do not resume after a failover. For example:

- TCP sessions that are not virus scanned, web filtered, spam filtered, content archived, or are not SIP, SIMPLE, or SCCP signal traffic resume after a failover, even if they are accepted by a security policy with UTM options enabled. For example, SNMP TCP sessions resume after a failover because FortiOS does not apply any UTM options to SNMP sessions.
- TCP sessions for a protocol for which UTM features have not been enabled resume after a failover even if they are accepted by a security policy with UTM features enabled. For example, if you have not enabled any antivirus or content archiving settings for FTP, FTP sessions resume after a failover.

The following UTM features do not affect TCP session failover:

- IPS does not affect session failover. Sessions being scanned by IPS resume after a failover. After a failover; however, IPS can only perform packet-based inspection of resumed sessions; reducing the number of vulnerabilities that IPS can detect. This limitation only applies to in-progress resumed sessions.
- Application control does not affect session failover. Sessions that are being monitored by application control resume after a failover.
- Logging enabled from UTM features does not affect session failover. UTM logging writes event log messages for UTM events; such as when a virus is found by antivirus scanning, when Web Filtering blocks a URL, and so on. Logging does not enable features that would prevent sessions from being failed over, logging just reports on the activities of enabled features.

If more than one UTM feature is applied to a TCP session, that session will not resume after a failover as long as one of the UTM features prevents session failover. For example:

- Sessions being scanned by IPS and also being virus scanned do not resume after a failover.
- Sessions that are being monitored by application control and that are being DLP archived or virus scanned will not resume after a failover.

## IPv6, NAT64, and NAT66 session failover

The FGCP supports IPv6, NAT64, and NAT66 session failover, if session pickup is enabled, these sessions are synchronized between cluster members and after an HA failover the sessions will resume with only minimal interruption.

## SIP session failover

The FGCP supports SIP session failover (also called stateful failover) for active-passive HA. To support SIP session failover, create a standard HA configuration and select *Enable Session Pick-up* option.

SIP session failover replicates SIP states to all cluster units. If an HA failover occurs, all in-progress SIP calls (setup complete) and their RTP flows are maintained and the calls will continue after the failover with minimal or no interruption.

SIP calls being set up at the time of a failover may lose signaling messages. In most cases the SIP clients and servers should use message retransmission to complete the call setup after the failover has completed. As a result, SIP users may experience a delay if their calls are being set up when an HA failover occurs. But in most cases the call setup should be able to continue after the failover.

## Explicit web proxy, WCCP, and WAN optimization session failover

Similar to UTM sessions, the explicit web proxy, WCCP and WAN optimization features all require the FortiGate unit to maintain very large amounts of internal state information for each session. This information is not maintained and these sessions do not resume after a failover.

## SSL offloading and HTTP multiplexing session failover

SSL offloading and HTTP multiplexing are both enabled from firewall virtual IPs and firewall load balancing. Similar to the features applied by UTM, SSL offloading and HTTP multiplexing requires the FortiGate unit to maintain very large amounts of internal state information for each session. Sessions accepted by security policies containing virtual IPs or virtual servers with SSL offloading or HTTP multiplexing enabled do not resume after a failover.

## IPsec VPN session failover

Session failover is supported for all IPsec VPN tunnels. To support IPsec VPN tunnel failover, when an IPsec VPN tunnel starts, the FGCP distributes the SA and related IPsec VPN tunnel data to all cluster units.

## SSL VPN session failover and SSL VPN authentication failover

Session failover is not supported for SSL VPN tunnels. However, authentication failover is supported for the communication between the SSL VPN client and the FortiGate unit. This means that after a failover, SSL VPN clients can re-establish the SSL VPN session between the SSL VPN client and the FortiGate unit without having to authenticate again.

However, all sessions inside the SSL VPN tunnel that were running before the failover are stopped and have to be restarted. For example, file transfers that were in progress would have to be restarted. As well, any communication sessions with resources behind the FortiGate unit that are started by an SSL VPN session have to be restarted.

To support SSL VPN cookie failover, when an SSL VPN session starts, the FGCP distributes the cookie created to identify the SSL VPN session to all cluster units.

## PPTP and L2TP VPN sessions

PPTP and L2TP VPNs are supported in HA mode. For a cluster you can configure PPTP and L2TP settings and you can also add security policies to allow PPTP and L2TP pass through. However, the FGCP does not provide session failover for PPTP or L2TP. After a failover, all active PPTP and L2TP sessions are lost and must be restarted.

## UDP, ICMP, multicast and broadcast packet session failover

By default, even with session pickup enabled, the FGCP does not maintain a session table for UDP, ICMP, multicast, or broadcast packets. So the cluster does not specifically support failover of these packets.

Some UDP traffic can continue to flow through the cluster after a failover. This can happen if, after the failover, a UDP packet that is part of an already established communication stream matches a security policy. Then a new session will be created and traffic will flow. So after a short interruption, UDP sessions can appear to have failed over. However, this may not be reliable for the following reasons:

- UDP packets in the direction of the security policy must be received before reply packets can be accepted. For example, if a port1 -> port2 policy accepts UDP packets, UDP packets received at port2 destined for the network connected to port1 will not be accepted until the policy accepts UDP packets at port1 that are destined for the network connected to port2. So, if a user connects from an internal network to the Internet and starts receiving UDP packets from the Internet (for example streaming media), after a failover the user will not receive any more UDP packets until the user re-connects to the Internet site.
- UDP sessions accepted by NAT policies will not resume after a failover because NAT will usually give the new session a different source port. So only traffic for UDP protocols that can handle the source port changing during a session will continue to flow.

You can however, enable session pickup for UDP and ICMP packets by enabling session pickup for TCP sessions and then enabling session pickup for connectionless sessions:

```
config system ha
 set session-pickup enable
 set session-pickup-connectionless enable
end
```

This configuration causes the cluster units to synchronize UDP and ICMP session tables and if a failover occurs UDP and ICMP sessions are maintained.

## FortiOS Carrier GTP session failover

FortiOS Carrier HA supports GTP session failover. The primary unit synchronizes the GTP tunnel state to all cluster units after the GTP tunnel setup is completed. After the tunnel setup is completed, GTP sessions use UDP and HA does not synchronize UDP sessions to all cluster units. However, similar to other UDP sessions, after a failover, since the new primary unit will have the GTP tunnel state information, GTP UDP sessions using the same tunnel can continue to flow with some limitations.

The limitation on packets continuing to flow is that there has to be a security policy to accept the packets. For example, if the FortiOS Carrier unit has an internal to external security policy, GTP UDP sessions using an established tunnel that are received by the internal interface are accepted by the security policy and can continue to flow. However, GTP UDP packets for an established tunnel that are received at the external interface cannot flow until packets from the same tunnel are received at the internal interface.

If you have bi-directional policies that accept GTP UDP sessions then traffic in either direction that uses an established tunnel can continue to flow after a failover without interruption.

## Active-active HA subordinate units sessions can resume after a failover

In an active-active cluster, subordinate units process sessions. After a failover, all cluster units that are still operating may be able to continue processing the sessions that they were processing before the failover. These sessions are maintained because after the failover the new primary unit uses the HA session table to continue to send session packets to the cluster units that were processing the sessions before the failover. Cluster units maintain their own information about the sessions that they are processing and this information is not affected by the failover. In this way, the cluster units that are still operating can continue processing their own sessions without loss of data.

The cluster keeps processing as many sessions as it can. But some sessions can be lost. Depending on what caused the failover, sessions can be lost in the following ways:

- A cluster unit fails (the primary unit or a subordinate unit). All sessions that were being processed by that cluster unit are lost.
- A link failure occurs. All sessions that were being processed through the network interface that failed are lost.

This mechanism for continuing sessions is not the same as session failover because:

- Only the sessions that can be maintained.
- The sessions are maintained on the same cluster units and not re-distributed.
- Sessions that cannot be maintained are lost.

## WAN optimization and HA

You can configure WAN optimization on a FortiGate HA cluster. The recommended HA configuration for WAN optimization is active-passive mode. Also, when the cluster is operating, all WAN optimization sessions are processed by the primary unit only. Even if the cluster is operating in active-active mode, HA does not load-balance WAN optimization sessions. HA also does not support WAN optimization session failover.

In a cluster, the primary unit only stores web cache and byte cache databases. These databases are not synchronized to the subordinate units. So, after a failover, the new primary unit must rebuild its web and byte caches. As well, the new primary unit cannot connect to a SAS partition that the failed primary unit used.

Rebuilding the byte caches can happen relatively quickly because the new primary unit gets byte cache data from the other FortiGate units that it is participating with in WAN optimization tunnels.

## Failover and attached network equipment

It normally takes a cluster approximately 6 seconds to complete a failover. However, the actual failover time experienced by your network users may depend on how quickly the switches

connected to the cluster interfaces accept the cluster MAC address update from the primary unit. If the switches do not recognize and accept the gratuitous ARP packets and update their MAC forwarding table, the failover time will increase.

Also, individual session failover depends on whether the cluster is operating in active-active or active-passive mode, and whether the content of the traffic is to be virus scanned. Depending on application behavior, it may take a TCP session a longer period of time (up to 30 seconds) to recover completely.

## Monitoring cluster units for failover

You can use logging and SNMP to monitor cluster units for failover. Both the primary and subordinate units can be configured to write log messages and send SNMP traps if a failover occurs. You can also log into the cluster web-based manager and CLI to determine if a failover has occurred. See [“Monitoring cluster units for failover” on page 1277](#).

## NAT/Route mode active-passive cluster packet flow

This section describes how packets are processed and how failover occurs in an active-passive HA cluster running in NAT/Route mode. In the example, the NAT/Route mode cluster acts as the internet firewall for a client computer’s internal network. The client computer’s default route points at the IP address of the cluster internal interface. The client connects to a web server on the Internet. Internet routing routes packets from the cluster external interface to the web server, and from the web server to the cluster external interface.

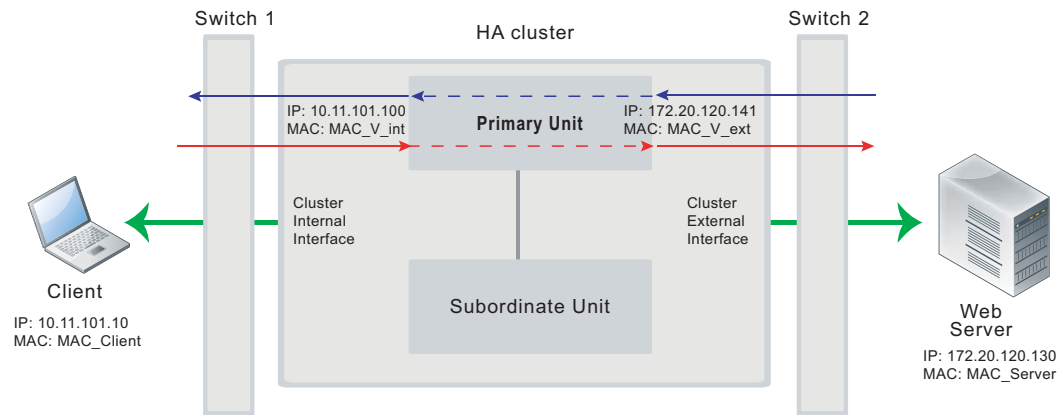
In an active-passive cluster operating in NAT/Route mode, four MAC addresses are involved in communication between the client and the web server when the primary unit processes the connection:

- Internal virtual MAC address (MAC\_V\_int) assigned to the primary unit internal interface,
- External virtual MAC address (MAC\_V\_ext) assigned to the primary unit external interface,
- Client MAC address (MAC\_Client),
- Server MAC address (MAC\_Server),

In NAT/Route mode, the HA cluster works as a gateway when it responds to ARP requests. Therefore, the client and server only know the gateway MAC addresses. The client only knows the cluster internal virtual MAC address (MAC\_V\_int) and the server only know the cluster external virtual MAC address (MAC\_V\_ext). Cluster virtual MAC addresses are described in [“Cluster virtual MAC addresses” on page 1300](#).



**Figure 213:**NAT/Route mode active-passive packet flow



### Packet flow from client to web server

1. The client computer requests a connection from 10.11.101.10 to 172.20.120.130.
2. The default route on the client computer recognizes 10.11.101.100 (the cluster IP address) as the gateway to the external network where the web server is located.
3. The client computer issues an ARP request to 10.11.101.100.
4. The primary unit intercepts the ARP request, and responds with the internal virtual MAC address (MAC\_V\_int) which corresponds to its IP address of 10.11.101.100.
5. The client's request packet reaches the primary unit internal interface.

	IP address	MAC address
<b>Source</b>	10.11.101.10	MAC_Client
<b>Destination</b>	172.20.120.130	MAC_V_int

6. The primary unit processes the packet.
7. The primary unit forwards the packet from its external interface to the web server.

	IP address	MAC address
<b>Source</b>	172.20.120.141	MAC_V_ext
<b>Destination</b>	172.20.120.130	MAC_Server

8. The primary unit continues to process packets in this way unless a failover occurs.

### Packet flow from web server to client

1. When the web server responds to the client's packet, the cluster external interface IP address (172.20.120.141) is recognized as the gateway to the internal network.
2. The web server issues an ARP request to 172.20.120.141.
3. The primary unit intercepts the ARP request, and responds with the external virtual MAC address (MAC\_V\_ext) which corresponds its IP address of 172.20.120.141.

- The web server then sends response packets to the primary unit external interface.

	IP address	MAC address
<b>Source</b>	172.20.120.130	MAC_Server
<b>Destination</b>	172.20.120.141	MAC_V_ext

- The primary unit processes the packet.
- The primary unit forwards the packet from its internal interface to the client.

	IP address	MAC address
<b>Source</b>	172.20.120.130	MAC_V_int
<b>Destination</b>	10.11.101.10	MAC_Client

- The primary unit continues to process packets in this way unless a failover occurs.

### When a failover occurs

The following steps are followed after a device or link failure of the primary unit causes a failover.

- If the primary unit fails the subordinate unit becomes the primary unit.
- The new primary unit changes the MAC addresses of all of its interfaces to the HA virtual MAC addresses.  
The new primary unit has the same IP addresses and MAC addresses as the failed primary unit.
- The new primary unit sends gratuitous ARP packets from the internal interface to the 10.11.101.0 network to associate its internal IP address with the internal virtual MAC address.
- The new primary unit sends gratuitous ARP packets to the 172.20.120.0 to associate its external IP address with the external virtual MAC address.
- Traffic sent to the cluster is now received and processed by the new primary unit.  
If there were more than two cluster units in the original cluster, these remaining units would become subordinate units.

## Transparent mode active-passive cluster packet flow

This section describes how packets are processed and how failover occurs in an active-passive HA cluster running in Transparent mode. The cluster is installed on an internal network in front of a mail server and the client connects to the mail server through the Transparent mode cluster.

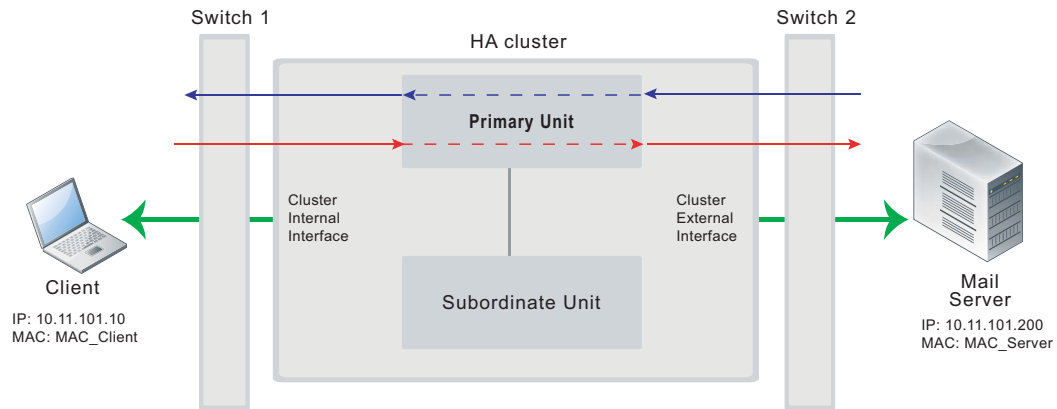
In an active-passive cluster operating in Transparent mode, two MAC addresses are involved in the communication between a client and a server when the primary unit processes a connection:

- Client MAC address (MAC\_Client)
- Server MAC address (MAC\_Server)

The HA virtual MAC addresses are not directly involved in communication between the client and the server. The client computer sends packets to the mail server and the mail server sends responses. In both cases the packets are intercepted and processed by the cluster.

The cluster's presence on the network is transparent to the client and server computers. The primary unit sends gratuitous ARP packets to Switch 1 that associate all MAC addresses on the network segment connected to the cluster external interface with the HA virtual MAC address. The primary unit also sends gratuitous ARP packets to Switch 2 that associate all MAC addresses on the network segment connected to the cluster internal interface with the HA virtual MAC address. In both cases, this results in the switches sending packets to the primary unit interfaces.

**Figure 214:**Transparent mode active-passive packet flow



### Packet flow from client to mail server

1. The client computer requests a connection from 10.11.101.10 to 10.11.101.200.
2. The client computer issues an ARP request to 10.11.101.200.
3. The primary unit forwards the ARP request to the mail server.
4. The mail server responds with its MAC address (MAC\_Server) which corresponds to its IP address of 10.11.101.200. The primary unit returns the ARP response to the client computer.
5. The client's request packet reaches the primary unit internal interface.

	IP address	MAC address
<b>Source</b>	10.11.101.10	MAC_Client
<b>Destination</b>	10.11.101.200	MAC_Server

6. The primary unit processes the packet.
7. The primary unit forwards the packet from its external interface to the mail server.

	IP address	MAC address
<b>Source</b>	10.11.101.10	MAC_Client
<b>Destination</b>	10.11.101.200	MAC_Server

8. The primary unit continues to process packets in this way unless a failover occurs.

### Packet flow from mail server to client

1. To respond to the client computer, the mail server issues an ARP request to 10.11.101.10.
2. The primary unit forwards the ARP request to the client computer.

3. The client computer responds with its MAC address (MAC\_Client) which corresponds to its IP address of 10.11.101.10. The primary unit returns the ARP response to the mail server.
4. The mail server's response packet reaches the primary unit external interface.

	IP address	MAC address
<b>Source</b>	10.11.101.200	MAC_Server
<b>Destination</b>	10.11.101.10	MAC_Client

5. The primary unit processes the packet.
6. The primary unit forwards the packet from its internal interface to the client.

	IP address	MAC address
<b>Source</b>	10.11.101.200	MAC_Server
<b>Destination</b>	10.11.101.10	MAC_Client

7. The primary unit continues to process packets in this way unless a failover occurs.

## When a failover occurs

The following steps are followed after a device or link failure of the primary unit causes a failover.

1. If the primary unit fails, the subordinate unit negotiates to become the primary unit.
2. The new primary unit changes the MAC addresses of all of its interfaces to the HA virtual MAC address.
3. The new primary unit sends gratuitous ARP packets to switch 1 to associate its MAC address with the MAC addresses on the network segment connected to the external interface.
4. The new primary unit sends gratuitous ARP packets to switch 2 to associate its MAC address with the MAC addresses on the network segment connected to the internal interface.
5. Traffic sent to the cluster is now received and processed by the new primary unit.  
If there were more than two cluster units in the original cluster, these remaining units would become subordinate units.

## Failover performance

This section describes the designed device and link failover times for a FortiGate cluster and also shows results of a failover performance test.

### Device failover performance

By design FGCP device failover time is 2 seconds for a two-member cluster with ideal network and traffic conditions. If subsecond failover is enabled the failover time can drop below 1 second.

All cluster units regularly receive HA heartbeat packets from all other cluster units over the HA heartbeat link. If any cluster unit does not receive a heartbeat packet from any other cluster unit for 2 seconds, the cluster unit that has not sent heartbeat packets is considered to have failed.

It may take another few seconds for the cluster to negotiate and re-distribute communication sessions. Typically if subsecond failover is not enabled you can expect a failover time of 9 to 15 seconds depending on the cluster and network configuration. The failover time can also be increased by more complex configurations and or configurations with network equipment that is slow to respond.

You can change the `hb-lost-threshold` to increase or decrease the device failover time. See [“Modifying heartbeat timing” on page 1298](#) for information about using `hb-lost-threshold`, and other heartbeat timing settings.

## Link failover performance

Link failover time is controlled by how long it takes for a cluster to synchronize the cluster link database. When a link failure occurs, the cluster unit that experienced the link failure uses HA heartbeat packets to broadcast the updated link database to all cluster units. When all cluster units have received the updated database the failover is complete.

It may take another few seconds for the cluster to negotiate and re-distribute communication sessions.

## Reducing failover times

You can do the following to help reduce failover times:

- Keep the network configuration as simple as possible with as few as possible network connections to the cluster.
- If possible operate the cluster in Transparent mode.
- Use high-performance switches to that the switches failover to interfaces connected to the new primary unit as quickly as possible.
- Use accelerated FortiGate interfaces. In some cases accelerated interfaces will reduce failover times.
- Make sure the FortiGate unit sends multiple gratuitous arp packets after a failover. In some cases, sending more gratuitous arp packets will cause connected network equipment to recognize the failover sooner. To send 10 gratuitous arp packets:

```
config system ha
 set arps 10
end
```

- Reduce the time between gratuitous arp packets. This may also caused connected network equipment to recognize the failover sooner. To send 50 gratuitous arp packets with 1 second between each packet:

```
config system ha
 set arp 50
 set arps-interval 1
end
```

- Reduce the number of lost heartbeat packets and reduce the heartbeat interval timers to be able to more quickly detect a device failure. To set the lost heartbeat threshold to 3 packets and the heartbeat interval to 100 milliseconds:

```
config system ha
 set hb-interval 3
 set hb-lost-threshold 1
end
```

- Reduce the hello state hold down time to reduce the amount of the time the cluster waits before transitioning from the hello to the work state. To set the hello state hold down time to 5 seconds:

```
config system ha
 set helo-holddown 5
end
```

- Enable sending a link failed signal after a link failover to make sure that attached network equipment responds a quickly as possible to a link failure. To enable the link failed signal:

```
config system ha
 set link-failed-signal enable
end
```

# HA and load balancing

FGCP active-active load balancing distributes network traffic among all of the units in a cluster. Load balancing can improve cluster performance because the processing load is shared among multiple cluster units.

This chapter describes how active-active load balancing works and provides detailed NAT/Route and Transparent mode packet flow descriptions.

This chapter contains the following sections:

- [Load balancing overview](#)
- [Configuring load balancing settings](#)
- [NAT/Route mode active-active cluster packet flow](#)
- [Transparent mode active-active cluster packet flow](#)

## Load balancing overview

In active-active HA, the FGCP uses a technique similar to unicast load balancing in which the primary unit is associated with the cluster HA virtual MAC addresses and cluster IP addresses. The primary unit is the only cluster unit to receive packets sent to the cluster.

An active-active HA cluster consists of a primary unit that processes communication sessions and one or more subordinate units that also process communication sessions. The primary unit receives all sessions and load balances sessions for security policies with UTM enabled to all cluster units. Communication between the cluster units uses the actual cluster unit MAC addresses.

Processing UTM sessions can be CPU and memory-intensive, load balancing UTM traffic may result in an active-active cluster having higher throughput than an active-passive cluster or a standalone FortiGate unit because resource-intensive UTM processing is distributed among all cluster units.

You can also enable the `load-balance-all` CLI keyword to have the primary unit load balance all TCP sessions. Load balancing TCP sessions is less likely to improve throughput because of extra overhead required for load balancing. So `load-balance-all` is disabled by default.

You can also enable the `load-balance-udp` CLI keyword to have the primary unit load balance all UDP sessions. Load balancing UDP sessions will also increase overhead so it is disabled by default.

During active-active HA load balancing operation, when the primary unit receives the first packet of a UTM session (or a TCP session if `load-balance-all` is enabled or a UDP session if `load-balance-udp` is enabled) the primary unit uses the configured load balancing schedule to determine the cluster unit that will process the session. The primary unit stores the load balancing information for each active load balanced session in the cluster load balancing session table. Using the information in this table, the primary unit can then forward all of the remaining packets in each session to the appropriate cluster unit. The load balancing session table is synchronized among all cluster units.

ICMP, multicast, and broadcast sessions are never load balanced and are always processed by the primary unit. VoIP, IM, P2P, IPsec VPN, HTTPS, SSL VPN, HTTP multiplexing, SSL offloading, WAN optimization, explicit web proxy, and WCCP sessions are also always processed only by the primary unit.

In addition to load balancing, active-active HA also provides device and link failover protection similar to active-passive HA. If the primary unit fails, a subordinate unit becomes the primary unit and resumes operating the cluster. See [“Device failover” on page 1292](#) and [“Link failover \(port monitoring or interface monitoring\)” on page 1319](#) for more information.

Active-active HA provides the same session failover protection as active-passive HA. See [“Session failover \(session pick-up\)” on page 1330](#) for more information about FortiGate session failover and its limitations.

Active-active HA also maintains as many UTM sessions as possible after a failover by continuing to process the UTM sessions that were being processed by the cluster units that are still operating. See [“Active-active HA subordinate units sessions can resume after a failover” on page 1335](#) for more information. Active-passive HA does not support maintaining UTM sessions after a failover.

## Load balancing schedules

The load balancing schedule controls how the primary unit distributes packets to all cluster units. You can select from the following load balancing schedules.

<b>None</b>	No load balancing. Select None when the cluster interfaces are connected to load balancing switches. If you select <i>None</i> , the Primary unit does not load balance traffic and the subordinate units process incoming traffic that does not come from the Primary unit. For all other load balancing schedules, all traffic is received first by the Primary unit, and then forwarded to the subordinate units. The subordinate units only receive and process packets sent from the primary unit.
<b>Hub</b>	Load balancing if the cluster interfaces are connected to a hub. Traffic is distributed to cluster units based on the source IP and destination IP of the packet.
<b>Least-Connection</b>	If the cluster units are connected using switches, select <i>Least Connection</i> to distribute network traffic to the cluster unit currently processing the fewest connections.
<b>Round-Robin</b>	If the cluster units are connected using switches, select Round-Robin to distribute network traffic to the next available cluster unit.
<b>Weighted Round-Robin</b>	Similar to round robin, but weighted values are assigned to each of the units in a cluster based on their capacity and on how many connections they are currently processing. For example, the primary unit should have a lower weighted value because it handles scheduling and forwards traffic. Weighted round robin distributes traffic more evenly because units that are not processing traffic will be more likely to receive new connections than units that are very busy.
<b>Random</b>	If the cluster units are connected using switches, select Random to randomly distribute traffic to cluster units.
<b>IP</b>	Load balancing according to IP address. If the cluster units are connected using switches, select <i>IP</i> to distribute traffic to units in a cluster based on the source IP and destination IP of the packet.
<b>IP Port</b>	Load balancing according to IP address and port. If the cluster units are connected using switches, select IP Port to distribute traffic to units in a cluster based on the source IP, source port, destination IP, and destination port of the packet.



Once a packet has been propagated to a subordinate unit, all packets are part of that same communication session and are also propagated to that same subordinate unit. Traffic is distributed according to communication session, not just according to individual packet.

Any subordinate unit that receives a forwarded packet processes it, without applying load balancing. Note that subordinate units are still considered to be active, because they perform routing, virus scanning, and other FortiGate unit tasks on their share of the traffic. Active subordinate units also share their session and link status information with all cluster units. The only things that active members do not do is make load balancing decisions.

Even though the primary unit is responsible for the load balancing process, the primary unit still acts like a FortiGate unit in that it processes packets, performing, routing, firewall, virus scanning, and other FortiGate unit tasks on its share of the traffic. Depending on the load balancing schedule used, the primary unit may assign itself a smaller share of the total load.

## Selecting which packets are load balanced

The primary unit processes all ICMP traffic. By default, the primary unit also processes all TCP and UDP traffic and load balances virus scanning traffic among all cluster units. You can change the default configuration so that the cluster load balances TCP, UDP traffic, and virus scanning traffic among all cluster units.

Load balancing increases network bandwidth usage and also increases the load on the primary unit CPU. Because of this, in some network environments, load balancing TCP and UDP traffic may not result in an overall cluster performance increase. However, in other network environments, TCP and UDP load balancing may improve cluster performance.

If the cluster is configured to load balance virus scanning sessions, the primary unit uses the load balancing schedule to distribute HTTP, FTP, SMTP, POP3, and IMAP packets to be virus scanned, among the primary unit and the subordinate units. Load balancing virus scanning traffic is much more likely to increase cluster performance. Virus scanning is processor intensive for the cluster unit that is performing the virus scanning. Distributing virus scanning over the cluster units significantly reduces the processing load on the primary unit. As a result overall cluster performance should improve. See [“Load balancing UTM sessions, TCP sessions, and UDP sessions” on page 1347](#).

## More about active-active failover

If a subordinate unit fails, the primary unit re-distributes the connections that the subordinate unit was processing among the remaining active cluster members. If the primary unit fails, the subordinate units negotiate to select a new primary unit. The new primary unit continues to distribute packets among the remaining active cluster units.

Failover works in a similar way if the cluster consists of only two units. If the primary unit fails the subordinate unit negotiates and becomes the new primary unit. If the subordinate unit fails, the primary unit processes all traffic. In both cases, the single remaining unit continues to function as a primary unit, maintaining the HA virtual MAC address for all of its interfaces.

## HTTPS sessions, active-active load balancing, and proxy servers

To prevent HTTPS web filtering problems active-active HA does not load balance HTTPS sessions. The FortiGate unit identifies HTTPS sessions as all sessions received on the HTTPS TCP port. The default HTTPS port is 443. You can use the CLI command `config antivirus service` to configure the FortiGate unit to use a custom port for HTTPS sessions. If you change the HTTPS port using this CLI command, the FGCP stops load balancing all sessions that use the custom HTTPS port.

Normally you would not change the HTTPS port. However, if your network uses a proxy server for HTTPS traffic you may have to use the `config antivirus service` command to configure your cluster to use a custom HTTPS port. If your network uses a proxy server you might also use the same port for both HTTP and HTTPS traffic. In this case you would use `config antivirus service` to configure the FortiGate unit to use custom ports for both HTTP and HTTPS traffic.

Using the same port for HTTP and HTTPS traffic can cause problems with active-active clusters because active-active clusters always load balance HTTP traffic. If both HTTP and HTTPS use the same port, the active-active cluster cannot tell the difference between HTTP and HTTPS traffic and will load balance both HTTP and HTTPS traffic.

As mentioned above, load balancing HTTPS traffic may cause problems with HTTPS web filtering. To avoid this problem, you should configure your proxy server to use different ports for HTTP and HTTPS traffic. Then use the `config antivirus service` command to configure your cluster to also use different ports for HTTP and HTTPS.

## Using FortiGate network processor interfaces to accelerate active-active HA performance

Many FortiGate models and FortiGate AMC modules include network processors that can provide hardware acceleration for active-active HA load balancing by offloading load balancing from the primary unit CPU. HA load balancing can be accelerated by interfaces accelerated by NP network processors.

In some cases, performance of the primary unit can be reduced by active-active HA load balancing. Primary unit CPU cycles and bus bandwidth are required to receive, calculate load balancing schedules, and send balanced packets to the subordinate units. In very busy active-active clusters the primary unit may not be able to keep up with the processing load. This can result in lost traffic and can also cause the primary unit to delay sending heartbeat packets possibly reducing the stability and reliability of the active-active HA cluster.

Adding network processors to busy cluster unit interfaces increases load balancing performance by offloading load balancing to the network processors. The first packet of every new session is received by the primary unit and the primary unit uses its load balancing schedule to select the cluster unit that will process the new session. This information is passed back to the network processor and all subsequent packets of the same sessions are received by the primary unit interface network processor which sends the packet directly to a subordinate unit without using the primary unit CPU. Load balancing is effectively offloaded from the primary unit to the network processor resulting in a faster and more stable active-active cluster.

Using network processors to accelerate load balancing is especially useful if the `load-balance-all` and `load-balance-udp` options are enabled and the cluster is load balancing all TCP and UDP sessions because this could mean that the cluster is load balancing an excessive number of sessions.

To take advantage of network processor load balancing acceleration, connect the cluster unit interfaces with network processors to the busiest networks. Connect non-accelerated interfaces to less busy networks. No special FortiOS or HA configuration is required. Network processor acceleration of active-active HA load balancing is supported for any active-active HA configuration or active-active HA load balancing schedule.

## Configuring load balancing settings

This section describes how to configure the following load balancing settings:

- [Selecting a load balancing schedule](#)
- [Load balancing UTM sessions, TCP sessions, and UDP sessions](#)
- [Configuring weighted-round-robin weights](#)
- [Dynamically optimizing weighted load balancing according to how busy cluster units are](#)

### Selecting a load balancing schedule

You can select the load balancing schedule when initially configuring the cluster and you can change the load balancing schedule at any time while the cluster is operating without affecting cluster operation.

You can select a load balancing schedule from the CLI. Use the following command to select a load balancing schedule:

```
config system ha
 set schedule {hub | ip | ipport | leastconnection | none | random
 | round-robin | weight-round-robin}
end
```

### Load balancing UTM sessions, TCP sessions, and UDP sessions

By default a FortiGate active-active cluster load balances UTM sessions among all cluster units. UTM processing applies protocol recognition, virus scanning, IPS, web filtering, email filtering, data leak prevention (DLP), application control, and VoIP content scanning and protection to HTTP, HTTPS, FTP, IMAP, IMAPS, POP3, POP3S, SMTP, SMTPS, IM, NNTP, SIP, SIMPLE, and SCCP sessions accepted by security policies. By load balancing this resource-intensive UTM processing among all cluster units, an active-active HA cluster may provide better UTM performance than a standalone FortiGate unit. Other features enabled in security policies such as Endpoint security, traffic shaping and authentication (identity-based policies) have no effect active-active load balancing.

All other sessions are processed by the primary unit. Using the CLI, you can configure the cluster to load balance TCP sessions among all cluster units in addition to UTM sessions. All UDP, ICMP, multicast, and broadcast sessions are not load balanced, but are processed by the primary unit.

Use the following command to enable load balancing UTM and TCP sessions.

```
config system ha
 set load-balance-all enable
end
```

Enabling `load-balance-all` to load balance TCP sessions may not improve throughput because the cluster requires additional overhead to load balance sessions. The primary unit receives all sessions and load balances some TCP sessions to the subordinate units. Load balancing UTM sessions can improve performance because UTM session performance is limited by CPU performance. However, load balancing a non-UTM session usually requires about as much overhead as just processing it.

If your active-active cluster is processing TCP sessions and not performing UTM, you can enable `load-balance-all` and monitor network performance to see if it improves. If performance is not improved, you should change the HA mode to active-passive since active-active HA is not providing any benefit.

Using the CLI, you can also configure the cluster to load balance UDP sessions among all cluster units in addition to UTM sessions (and optionally TCP sessions).

Use the following command to enable load balancing UTM and UDP sessions.

```
config system ha
 set load-balance-udp enable
end
```

Enabling `load-balance-udp` to load balance UDP sessions may not improve throughput because the cluster requires additional overhead to load balance sessions. The primary unit receives all sessions and load balances some UDP sessions to the subordinate units. Load balancing UTM sessions can improve performance because UTM session performance is limited by CPU performance. However, load balancing a non-UTM session usually requires about as much overhead as just processing it.

If your active-active cluster is processing UDP sessions and not performing UTM, you can enable `load-balance-udp` and monitor network performance to see if it improves. If performance is not improved, you should change the HA mode to active-passive since active-active HA is not providing any benefit.

## Configuring weighted-round-robin weights

You can configure weighted round-robin load balancing for a cluster and configure the static weights for each of the cluster units according to their priority in the cluster. When you set `schedule` to `weight-round-robin` you can use the `weight` option to set the static weight of each cluster unit. The static weight is set according to the priority of each unit in the cluster. A FortiGate HA cluster can contain up to four FortiGate units so you can set up to 4 static weights.

The priority of a cluster unit is determined by its device priority, the number of monitored interfaces that are functioning, its age in the cluster and its serial number. Priorities are used to select a primary unit and to set an order of all of the subordinate units. Thus the priority order of a cluster unit can change depending on configuration settings, link failures and so on. Since weights are also set using this priority order the weights are independent of specific cluster units but do depend on the role of the each unit in the cluster.

You can use the following command to display the priority order of units in a cluster. The following example displays the priority order for a cluster of 5 FortiGate-620B units:

```
get system ha status
 Model: 620
 Mode: a-p
 Group: 0
 Debug: 0
 ses_pickup: disable
 Master:150 head_office_cla FG600B3908600825 0
 Slave :150 head_office_clb FG600B3908600705 1
 Slave :150 head_office_clc FG600B3908600702 2
 Slave :150 head_office_cld FG600B3908600605 3
 Slave :150 head_office_cle FG600B3908600309 4
 number of vcluster: 1
 vcluster 1: work 169.254.0.1
 Master:0 FG600B3908600825
 Slave :1 FG600B3908600705
 Slave :2 FG600B3908600702
 Slave :3 FG600B3908600605
 Slave :4 FG600B3908600309
```

The cluster units are listed in priority order starting at the 6th output line. The primary unit always has the highest priority and is listed first followed by the subordinate units in priority order. The last 5 output lines list the cluster units in vcluster 1 and are not always in priority order. For more information about the `get system ha status` command, see [“Viewing cluster status from the CLI” on page 1277](#).

The default static weight for each cluster unit is 40. This means that sessions are distributed evenly among all cluster units. You can use the `weight` option to change the static weights of cluster units to distribute sessions depending on each unit’s priority in the cluster. The weight can be between 0 and 255. Increase the weight to increase the number of connections processed by the cluster unit with that priority.

You set the weight for each unit separately. For the example cluster of 5 FortiGate-620B units you can set the weight for each unit as follows:

```
config system ha
 set mode a-a
 set schedule weight-round-robin
 set weight 0 5
 set weight 1 10
 set weight 2 15
 set weight 3 20
 set weight 4 30
end
```

If you enter the `get` command to view the HA configuration the output for `weight` would be:

```
weight 5 10 15 20 30 40 40 40 40 40 40 40 40 40 40
```

This configuration has the following results if the output of the `get system ha status` command is that shown above:

- The first five connections are processed by the primary unit (host name `head_office_cla`, priority 0, weight 5). From the output of the
- The next 10 connections are processed by the first subordinate unit (host name `head_office_clb`, priority 1, weight 10)
- The next 15 connections are processed by the second subordinate unit (host name `head_office_clc`, priority 2, weight 15)
- The next 20 connections are processed by the third subordinate unit (host name `head_office_cld`, priority 3, weight 20)
- The next 30 connections are processed by the fourth subordinate unit (host name `head_office_cle`, priority 4, weight 30)

## Dynamically optimizing weighted load balancing according to how busy cluster units are

In conjunction with using static weights to load balance sessions among cluster units you can configure a cluster to dynamically load balance sessions according to individual cluster unit CPU usage, memory usage, and number of HTTP, FTP, IMAP, POP3, SMTP, or NNTP proxy sessions. If any of these system loading indicators increases above configured thresholds, weighted load balancing dynamically sends fewer new sessions to the busy unit until it recovers.

High CPU or memory usage indicates that a unit is under increased load and may not be able to process more sessions. HTTP, FTP, IMAP, POP3, SMTP, or NNTP proxy use are also good indicators of how busy a cluster unit is, since processing high numbers of these proxy sessions can quickly reduce overall cluster unit performance.

For example, you can set a CPU usage high watermark threshold. When a cluster unit reaches this high watermark threshold fewer sessions are sent to it. With fewer sessions to process the cluster unit's CPU usage should fall back to the low watermark threshold. When the low watermark threshold is reached the cluster resumes normal load balancing of sessions to the cluster unit.

You can set individual high and low watermark thresholds and weights for CPU usage, memory usage, and for the number of HTTP, FTP, IMAP, POP3, SMTP, or NNTP proxy sessions.

The CPU usage, memory usage, and UTM proxy weights determine how the cluster load balances sessions when a high watermark threshold is reached and also affect how the cluster load balances sessions when multiple cluster units reach different high watermark thresholds at the same time. For example, you might be less concerned about a cluster unit reaching the memory usage high watermark threshold than reaching the CPU usage high watermark threshold. If this is the case you can set the weight lower for memory usage. Then, if one cluster unit reaches the CPU usage high watermark threshold and a second cluster unit reaches the memory usage high watermark threshold the cluster will load balance more sessions to the unit with high memory usage and fewer sessions to the cluster unit with high CPU usage. As a result, reaching the CPU usage high watermark will have a greater affect on how sessions are redistributed than reaching the memory usage high watermark.

When a high watermark threshold is reached, the corresponding weight is subtracted from the static weight of the cluster unit. The lower the weight the fewer the number of sessions that are load balanced to that unit. Subsequently when the low watermark threshold is reached, the static weight of the cluster returns to its configured value. For the weights to all be effective the weights assigned to the load indicators should usually be lower than or equal to the static weights assigned to the cluster units.

Use the following command to set thresholds and weights for CPU and memory usage and HTTP, FTP, IMAP, POP3, SMTP, or NNTP proxy sessions:

```
config system ha
 set mode a-a
 set schedule weight-round-robin
 set cpu-threshold <weight> <low> <high>
 set memory-threshold <weight> <low> <high>
 set http-proxy-threshold <weight> <low> <high>
 set ftp-proxy-threshold <weight> <low> <high>
 set imap-proxy-threshold <weight> <low> <high>
 set nntp-proxy-threshold <weight> <low> <high>
 set pop3-proxy-threshold <weight> <low> <high>
 set smtp-proxy-threshold <weight> <low> <high>
end
```

For each option, the weight range is 0 to 255 and the default weight is 5. The low and high watermarks are a percent (0 to 100). The default low and high watermarks are 0 which means they are disabled. The default configuration when weighted load balancing is enabled looks like the following:

```
config system ha
 set mode a-a
 set schedule weight-round-robin
 set cpu-threshold 5 0 0
 set memory-threshold 5 0 0
 set http-proxy-threshold 5 0 0
 set ftp-proxy-threshold 5 0 0
 set imap-proxy-threshold 5 0 0
 set nntp-proxy-threshold 5 0 0
 set pop3-proxy-threshold 5 0 0
 set smtp-proxy-threshold 5 0 0
end
```



When you first enable HA weighted load balancing, the weighted load balancing configuration is synchronized to all cluster units and each cluster unit has the default configuration shown above. Changes to the CPU, memory, HTTP, FTP, IMAP, NNTP, POP3, and SMTP thresholds and low and high watermarks must be made for each cluster unit and are not synchronized to the other cluster units.

---

When you configure them, the high watermarks must be greater than their corresponding low watermarks.

For CPU and memory usage the low and high watermarks are compared with the percentage CPU and memory use of the cluster unit. For each of the UTM proxies the high and low watermarks are compared to a number that represents percent of the max number of proxy sessions being used by a proxy. This number is calculated using the formula:

$$\text{proxy usage} = (\text{current sessions} * 100) / \text{max sessions}$$

where:

`current sessions` is the number of active sessions for the proxy type.

`max sessions` is the session limit for the proxy type. The session limit depends on the FortiGate unit and its configuration.

You can use the following command to display the maximum and current number of sessions for a UTM proxy:

```
get test { ftpd | http | imap | nntp | pop3 | smtp } 4
```

You can use the following command to display the maximum number of sessions and the and current number of sessions for all of the proxies:

```
get test proxyworker 4
```

The command output includes lines similar to the following:

```
get test http 4
HTTP Common
Current Connections 5000/8032
```


In the example, 5000 is the current number of proxy connections being used by HTTP and 8032 is the maximum number of proxy sessions allowed. For this example the proxy usage would be:

```
proxy usage = (5000 * 100) / 8032
proxy usage = 62%
```

### Example weighted load balancing configuration

Consider a cluster of three FortiGate-620B units with host names 620\_ha\_1, 620\_ha\_2, and 620\_ha\_3 as shown in [Figure 215](#). This example describes how to configure weighted load balancing settings for CPU and memory usage for the cluster and then to configure UTM proxy weights to send most HTTP and POP3 proxy sessions to different cluster units.

**Figure 215:**Example HA weighted load balancing configuration



Connect to the cluster CLI and use the following command to set the CPU usage threshold weight to 30, low watermark to 60, and high watermark to 80. This command also sets the memory usage threshold weight to 10, low watermark to 60, and high watermark to 90.

```
config system ha
 set mode a-a
 set schedule weight-round-robin
 set cpu-threshold 30 60 80
 set memory-threshold 10 60 90
end
```



The static weights for the cluster units remain at the default values of 40. Since this command changes the mode to `a-a` and the schedule to `weight-round-robin` for the first time, the weight settings are synchronized to all cluster units.

As a result of this configuration, if the CPU usage of any cluster unit (for example, `620_ha_1`) reaches 80% the static weight for that cluster unit is reduced from 40 to 10 and only 10 of every 120 new sessions are load balanced to this cluster unit. If the memory usage of `620_ha_1` also reaches 90% the static weight further reduces to 0 and no new sessions are load balanced to `620_ha_1`. Also, if the memory usage of `620_ha_2` reaches 90% the static weight of `620_ha_2` reduces to 30 and 30 of every 120 new sessions are load balanced to `620_ha_2`.

Now that you have established the weight load balancing configuration for the entire cluster you can monitor the cluster to verify that processing gets distributed evenly to all cluster units. From the web-based manager you can go do *System > Config > HA > View HA Statistics* and see the CPU usage, active sessions, memory usage and other statistics for all of the cluster units. If you notice that one cluster unit is more or less busy than others you can adjust the dynamic weights separately for each cluster unit.

For example, in some active-active clusters the primary unit may tend to be busier than other cluster units because in addition to processing sessions the primary unit also receives all packets sent to the cluster and performs load balancing to distribute the sessions to other cluster units. To reduce the load on the primary unit you could reduce the CPU and memory usage high watermark thresholds for the primary unit so that fewer sessions are distributed to the primary unit. You could also reduce the primary unit's high watermark setting for the proxies to distribute more proxy sessions to other cluster units.



Note that this would only be useful if you are using device priorities and override settings to make sure the same unit always becomes the primary unit. See [“Controlling primary unit selection using device priority and override” on page 1140](#).

---

If the example cluster is configured for `620_ha_2` to be the primary unit, connect to the `620_ha_2`'s CLI and enter the following command to set CPU usage, memory usage, and proxy usage high watermark thresholds lower.

```
config system ha
 set cpu-threshold 30 60 70
 set memory-threshold 30 60 70
 set http-proxy-threshold 30 60 70
 set ftp-proxy-threshold 30 60 70
 set imap-proxy-threshold 30 60 70
 set nntp-proxy-threshold 30 60 70
 set pop3-proxy-threshold 30 60 70
 set smtp-proxy-threshold 30 60 70
end
```

As a result, when any of these factors reaches 70% on the primary unit, fewer sessions will be processed by the primary unit, preventing the number of sessions being processed from rising.

## NAT/Route mode active-active cluster packet flow

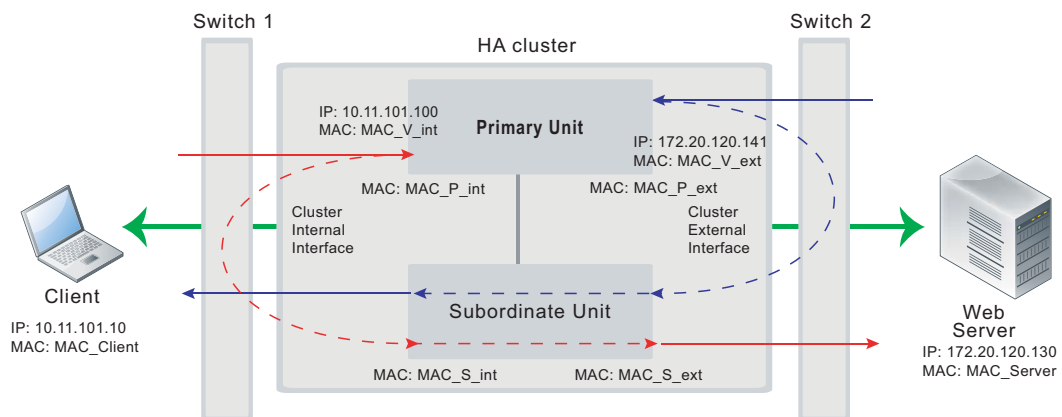
This section describes an example of how packets are load balanced and how failover occurs in an active-active HA cluster running in NAT/Route mode. In the example, the NAT/Route mode cluster acts as the internet firewall for a client computer's internal network. The client computer's default route points at the IP address of the cluster internal interface. The client connects to a web server on the Internet. Internet routing routes packets from the cluster external interface to the web server, and from the web server to the cluster external interface.

In NAT/Route mode, eight MAC addresses are involved in active-active communication between the client and the web server when the primary unit load balances packets to the subordinate unit:

- Internal virtual MAC address (MAC\_V\_int) assigned to the primary unit internal interface,
- External virtual MAC address (MAC\_V\_ext) assigned to the primary unit external interface,
- Client MAC address (MAC\_Client),
- Server MAC address (MAC\_Server),
- Primary unit original internal MAC address (MAC\_P\_int),
- Primary unit original external MAC address (MAC\_P\_ext),
- Subordinate unit internal MAC address (MAC\_S\_int),
- Subordinate unit external MAC address (MAC\_S\_ext).

In NAT/Route mode, the HA cluster works as a gateway when it responds to ARP requests. Therefore, the client and server only know the gateway MAC addresses. The client only knows the cluster internal virtual MAC address (MAC\_V\_int) and the server only knows the cluster external virtual MAC address (MAC\_V\_ext). The cluster virtual MAC address is described in “Cluster virtual MAC addresses” on page 1300.

**Figure 216:**NAT/Route mode active-active packet flow



### Packet flow from client to web server

1. The client computer requests a connection from 10.11.101.10 to 172.20.120.130.
2. The default route on the client computer recognizes 10.11.101.100 (the cluster IP address) as the gateway to the external network where the web server is located.
3. The client computer issues an ARP request to 10.11.101.100.
4. The primary unit intercepts the ARP request, and responds with the internal virtual MAC address (MAC\_V\_int) which corresponds to its IP address of 10.11.101.100.

- The client's request packet reaches the primary unit internal interface.

	<b>IP address</b>	<b>MAC address</b>
<b>Source</b>	10.11.101.10	MAC_Client
<b>Destination</b>	172.20.120.130	MAC_V_int

- The primary unit decides that the subordinate unit should handle this packet, and forwards it to the subordinate unit internal interface. The source MAC address of the forwarded packet is changed to the actual MAC address of the primary unit internal interface.

	<b>IP address</b>	<b>MAC address</b>
<b>Source</b>	10.11.101.10	MAC_P_int
<b>Destination</b>	172.20.120.130	MAC_S_int

- The subordinate unit recognizes that the packet has been forwarded from the primary unit and processes it.
- The subordinate unit forwards the packet from its external interface to the web server.

	<b>IP address</b>	<b>MAC address</b>
<b>Source</b>	172.20.120.141	MAC_S_ext
<b>Destination</b>	172.20.120.130	MAC_Server

- The primary unit forwards further packets in the same session to the subordinate unit.
10. Packets for other sessions are load balanced by the primary unit and either sent to the subordinate unit or processed by the primary unit.

### Packet flow from web server to client

- When the web server responds to the client's packet, the cluster external interface IP address (172.20.120.141) is recognized as the gateway to the internal network.
- The web server issues an ARP request to 172.20.120.141.
- The primary unit intercepts the ARP request, and responds with the external virtual MAC address (MAC\_V\_ext) which corresponds its IP address of 172.20.120.141.
- The web server then sends response packets to the primary unit external interface.

	<b>IP address</b>	<b>MAC address</b>
<b>Source</b>	172.20.120.130	MAC_Server
<b>Destination</b>	172.20.120.141	MAC_V_ext

- The primary unit decides that the subordinate unit should handle this packet, and forwards it to the subordinate unit external interface. The source MAC address of the forwarded packet is changed to the actual MAC address of the primary unit external interface.

	IP address	MAC address
<b>Source</b>	172.20.120.130	MAC_P_ext
<b>Destination</b>	172.20.120.141	MAC_S_ext

- The subordinate unit recognizes that packet has been forwarded from the primary unit and processes it.
- The subordinate unit forwards the packet from its internal interface to the client.

	IP address	MAC address
<b>Source</b>	172.20.120.130	MAC_S_int
<b>Destination</b>	10.11.101.10	MAC_Client

- The primary unit forwards further packets in the same session to the subordinate unit.
- Packets for other sessions are load balanced by the primary unit and either sent to the subordinate unit or processed by the primary unit.

## When a failover occurs

The following steps are followed after a device or link failure of the primary unit causes a failover.

- If the primary unit fails, the subordinate unit negotiates to become the primary unit.
- The new primary unit changes the MAC addresses of all of its interfaces to the HA virtual MAC addresses.  
The new primary unit has the same IP addresses and MAC addresses as the failed primary unit.
- The new primary unit sends gratuitous ARP packets to the 10.10.101.0 network to associate its internal IP address with the internal virtual MAC address.
- The new primary unit sends gratuitous ARP packets to the 172.20.120.0 network to associate its external IP address with the external virtual MAC address.
- Traffic sent to the cluster is now received and processed by the new primary unit.  
If there were more than two cluster units in the original cluster, the new primary unit would load balance packets to the remaining cluster members.

## Transparent mode active-active cluster packet flow

This section describes an example of how packets are load balanced and how failover occurs in an active-active HA cluster running in Transparent mode. The cluster is installed on an internal network in front of a mail server and the client connects to the mail server through the Transparent mode cluster.

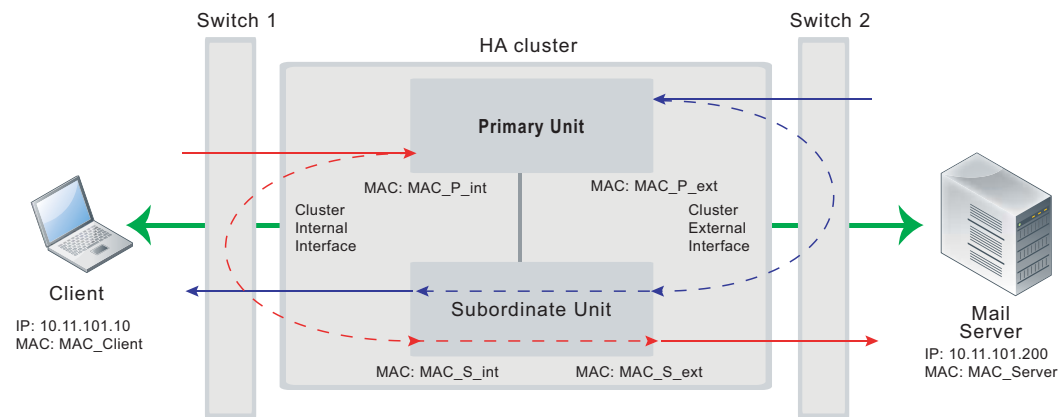
In Transparent mode, six MAC addresses are involved in active-active communication between a client and a server when the primary unit load balances packets to the subordinate unit:

- Client MAC address (MAC\_Client),
- Server MAC address (MAC\_Server),
- Primary unit original internal MAC address (MAC\_P\_int),
- Primary unit original external MAC address (MAC\_P\_ext),
- Subordinate unit internal MAC address (MAC\_S\_int),
- Subordinate unit external MAC address (MAC\_S\_ext).

The HA virtual MAC addresses are not directly involved in communicate between the client and the server. The client computer sends packets to the mail server and the mail server sends responses. In both cases the packets are intercepted and load balanced among cluster members.

The cluster’s presence on the network and its load balancing are transparent to the client and server computers. The primary unit sends gratuitous ARP packets to Switch 1 that associate all MAC addresses on the network segment connected to the cluster external interface with the external virtual MAC address. The primary unit also sends gratuitous ARP packets to Switch 2 that associate all MAC addresses on the network segment connected to the cluster internal interface with the internal virtual MAC address. In both cases, this results in the switches sending packets to the primary unit interfaces.

**Figure 217:**Transparent mode active-active packet flow



### Packet flow from client to mail server

1. The client computer requests a connection from 10.11.101.10 to 10.11.101.200.
2. The client computer issues an ARP request to 10.11.101.200.
3. The primary unit forwards the ARP request to the mail server.
4. The mail server responds with its MAC address (MAC\_Server) which corresponds to its IP address of 10.11.101.200. The primary unit returns the ARP response to the client computer.
5. The client’s request packet reaches the primary unit internal interface.

	IP address	MAC address
<b>Source</b>	10.11.101.10	MAC_Client
<b>Destination</b>	10.11.101.200	MAC_Server

- The primary unit decides that the subordinate unit should handle this packet, and forwards it to the subordinate unit internal interface. The source MAC address of the forwarded packet is changed to the actual MAC address of the primary unit internal interface.

	<b>IP address</b>	<b>MAC address</b>
<b>Source</b>	10.11.101.10	MAC_P_int
<b>Destination</b>	10.11.101.200	MAC_S_int

- The subordinate unit recognizes that packet has been forwarded from the primary unit and processes it.
- The subordinate unit forwards the packet from its external interface to the mail server.

	<b>IP address</b>	<b>MAC address</b>
<b>Source</b>	10.11.101.10	MAC_S_ext
<b>Destination</b>	10.11.101.200	MAC_Server

- The primary unit forwards further packets in the same session to the subordinate unit.
- Packets for other sessions are load balanced by the primary unit and either sent to the subordinate unit or processed by the primary unit.

### Packet flow from mail server to client

- To respond to the client computer, the mail server issues an ARP request to 10.11.101.10.
- The primary unit forwards the ARP request to the client computer.
- The client computer responds with its MAC address (MAC\_Client) which corresponds to its IP address of 10.11.101.10. The primary unit returns the ARP response to the mail server.
- The mail server's response packet reaches the primary unit external interface.

	<b>IP address</b>	<b>MAC address</b>
<b>Source</b>	10.11.101.200	MAC_Server
<b>Destination</b>	10.11.101.10	MAC_Client

- The primary unit decides that the subordinate unit should handle this packet, and forwards it to the subordinate unit external interface. The source MAC address of the forwarded packet is changed to the actual MAC address of the primary unit external interface.

	<b>IP address</b>	<b>MAC address</b>
<b>Source</b>	10.11.101.200	MAC_P_ext
<b>Destination</b>	10.11.101.10	MAC_S_ext

- The subordinate unit recognizes that packet has been forwarded from the primary unit and processes it.

7. The subordinate unit forwards the packet from its internal interface to the client.

	<b>IP address</b>	<b>MAC address</b>
<b>Source</b>	10.11.101.200	MAC_S_int
<b>Destination</b>	10.11.101.10	MAC_Client

8. The primary unit forwards further packets in the same session to the subordinate unit.
9. Packets for other sessions are load balanced by the primary unit and either sent to the subordinate unit or processed by the primary unit.

### When a failover occurs

The following steps are followed after a device or link failure of the primary unit causes a failover.

1. If the primary unit fails the subordinate unit negotiates to become the primary unit.
2. The new primary unit changes the MAC addresses of all of its interfaces to the HA virtual MAC address.
3. The new primary unit sends gratuitous ARP requests to switch 1 to associate its MAC address with the MAC addresses on the network segment connected to the external interface.
4. The new primary unit sends gratuitous ARP requests to switch 2 to associate its MAC address with the MAC addresses on the network segment connected to the internal interface.
5. Traffic sent to the cluster is now received and processed by the new primary unit.

If there were more than two cluster units in the original cluster, the new primary unit would load balance packets to the remaining cluster members.

# HA with FortiGate-VM and third-party products

This chapter provides information about operating FortiOS VM cluster and operating FortiGate clusters with third party products such as layer-2 and layer-3 switches. This chapter describes:

- [FortiGate-VM for VMware HA configuration](#)
- [FortiGate VM for Hyper-V HA configuration](#)
- [Failover issues with layer-3 switches](#)
- [Changing spanning tree protocol settings for some switches](#)
- [Failover and attached network equipment](#)
- [Ethertype conflicts with third-party switches](#)
- [LACP, 802.3ad aggregation and third-party switches](#)

## FortiGate-VM for VMware HA configuration

If you want to combine two or more FortiGate-VM instances into a FortiGate Clustering Protocol (FGSP) High Availability (HA) cluster the VMware server's virtual switches used to connect the heartbeat interfaces must operate in promiscuous mode. This permits HA heartbeat communication between the heartbeat interfaces. HA heartbeat packets are non-TCP packets that use Ethertype values 0x8890, 0x8891, and 0x8890. The FGCP uses link-local IP4 addresses in the 169.254.0.x range for HA heartbeat interface IP addresses.

To enable promiscuous mode in VMware:

1. In the vSphere client, select your VMware server in the left pane and then select the Configuration tab in the right pane.
2. In Hardware, select Networking.
3. Select Properties of a virtual switch used to connect heartbeat interfaces.
4. In the Properties window left pane, select vSwitch and then select Edit.
5. Select the Security tab, set Promiscuous Mode to Accept, then select OK.
6. Select Close.

You must also set the virtual switches connected to other FortiGate interfaces to allow MAC address changes and to accept forged transmits. This is required because the FGCP sets virtual MAC addresses for all FortiGate interfaces and the same interfaces on the different VM instances in the cluster will have the same virtual MAC addresses.

To make the required changes in VMware:

1. In the vSphere client, select your VMware server in the left pane and then select the Configuration tab in the right pane.
2. In Hardware, select Networking.
3. Select Properties of a virtual switch used to connect FortiGate VM interfaces.
4. Set MAC Address Change to Accept.
5. Set Forged Transmits to Accept.



## FortiGate VM for Hyper-V HA configuration

Promiscuous mode and support for MAC address spoofing is required for FortiGate-VM for Hyper-V to support FortiGate Clustering Protocol (FGCP) high availability (HA). By default the FortiGate-VM for Hyper-V has promiscuous mode enabled in the XML configuration file in the FortiGate-VM Hyper-V image. If you have problems with HA mode, confirm that this is still enabled.

In addition, because the FGCP applies virtual MAC addresses to FortiGate data interfaces and because these virtual MAC addresses mean that matching interfaces of different FortiGate-VM instances will have the same virtual MAC addresses you have to configure Hyper-V to allow MAC spoofing. But you should only enable MAC spoofing for FortiGate-VM data interfaces. You should not enable MAC spoofing for FortiGate HA heartbeat interfaces.

With promiscuous mode enabled and the correct MAC spoofing settings you should be able to configure HA between two or more FortiGate-VM for Hyper-V instances.

## Troubleshooting layer-2 switches

Issues may occur because of the way an HA cluster assigns MAC addresses to the primary unit. Two clusters with the same group ID cannot connect to the same switch and cannot be installed on the same network unless they are separated by a router.

### Forwarding delay on layer 2 switches

You must ensure that if there is a switch between the FortiGate HA cluster and the network it is protecting and the switch has a forwarding delay (even if spanning tree is disabled) when one of its interfaces is activated then the forwarding delay should be set as low as possible. For example, some versions of Cisco IOS have a forwarding delay of 15 seconds even when spanning tree is disabled. If left at this default value then TCP session pickup can fail because traffic is not forwarded through the switch on HA failover.

## Failover issues with layer-3 switches

After a failover, the new primary unit sends gratuitous ARP packets to refresh the MAC forwarding tables of the switches connected to the cluster. If the cluster is connected using layer-2 switches, the MAC forwarding tables (also called arp tables) are refreshed by the gratuitous ARP packets and the switches start directing packets to the new primary unit.

In some configurations that use layer-3 switches, after a failover, the layer-3 switches may not successfully re-direct traffic to the new primary unit. The possible reason for this is that the layer-3 switch might keep a table of IP addresses and interfaces and may not update this table for a relatively long time after the failover (the table is not updated by the gratuitous ARP packets). Until the table is updated, the layer-3 switch keeps forwarding packets to the now failed cluster unit. As a result, traffic stops and the cluster does not function.

As of the release date of this document, Fortinet has not developed a workaround for this problem. One possible solution would be to clear the forwarding table on the layer-3 switch.

The `config system ha link-failed-signal` command described in [“Updating MAC forwarding tables when a link failover occurs” on page 1323](#) can be used to resolve link failover issues similar to those described here.

## Changing spanning tree protocol settings for some switches

Configuration changes may be required when you are running an active-active HA cluster that is connected to a switch that operates using the spanning tree protocol. For example, the following spanning tree parameters may need to be changed:

---

<b>Maximum Age</b>	The time that a bridge stores the spanning tree bridge control data unit (BPDU) before discarding it. A maximum age of 20 seconds means it may take 20 seconds before the switch changes a port to the listening state.
<b>Forward Delay</b>	The time that a connected port stays in listening and learning state. A forward delay of 15 seconds assumes a maximum network size of seven bridge hops, a maximum of three lost BPDUs and a hello-interval of 2 seconds.

---

For an active-active HA cluster to be compatible with the spanning tree algorithm, the FGCP requires that the sum of maximum age and forward delay should be less than 20 seconds. The maximum age and forward delay settings are designed to prevent layer 2 loops. If there is no possibility of layer 2 loops in the network, you could reduce the forward delay to the minimum value.

For some Dell 3348 switches the default maximum age is 20 seconds and the default forward delay is 15 seconds. In this configuration the switch cannot work with a FortiGate HA cluster. However, the switch and cluster are compatible if the maximum age is reduced to 10 seconds and the forward delay is reduced to 5 seconds.

### Spanning Tree protocol (STP)

Spanning tree protocol is an IEEE 802.1 standard link management protocol that for media access control bridges. STP uses the spanning tree algorithm to provide path redundancy while preventing undesirable loops in a network that are created by multiple active paths between stations. Loops can be created if there are more than route between two hosts. To control path redundancy, STP creates a tree that spans all of the switches in an extended network. Using the information in the tree, the STP can force redundant paths into a standby, or blocked, state. The result is that only one active path is available at a time between any two network devices (preventing looping). Redundant links are used as backups if the initial link should fail. Without spanning tree in place, it is possible that two connections may be simultaneously live, which could result in an endless loop of traffic on the network.

### Bridge Protocol Data Unit (BPDU)

BPDUs are spanning tree data messages exchanged across switches within an extended network. BPDU packets contain information on ports, addresses, priorities and costs and ensure that the data ends up where it was intended to go. BPDU messages are exchanged across bridges to detect loops in a network topology. The loops are then removed by shutting down selected bridge interfaces and placing redundant switch ports in a backup, or blocked, state.

## Failover and attached network equipment

It normally takes a cluster approximately 6 seconds to complete a failover. However, the actual failover time may depend on how quickly the switches connected to the cluster interfaces accept the cluster MAC address update from the primary unit. If the switches do not recognize

and accept the gratuitous ARP packets and update their MAC forwarding table, the failover time will increase.

Also, individual session failover depends on whether the cluster is operating in active-active or active-passive mode, and whether the content of the traffic is to be virus scanned. Depending on application behavior, it may take a TCP session a longer period of time (up to 30 seconds) to recover completely.

## Ethertype conflicts with third-party switches

Some third-party network equipment may use packets with Ethertypes that are the same as the ethertypes used for HA heartbeat packets. For example, Cisco N5K/Nexus switches use Ethertype 0x8890 for some functions. When one of these switches receives Ethertype 0x8890 heartbeat packets from an attached cluster unit, the switch generates CRC errors and the packets are not forwarded. As a result, FortiGate units connected with these switches cannot form a cluster.

In some cases, if the heartbeat interfaces are connected and configured so regular traffic flows but heartbeat traffic is not forwarded, you can change the configuration of the switch that connects the HA heartbeat interfaces to allow level2 frames with Ethertypes 0x8890, 0x8893, and 0x8891 to pass.

You can also use the following CLI commands to change the Ethertypes of the HA heartbeat packets:

```
config system ha
 set ha-eth-type <ha_ethertype_4-digit_hex>
 set hc-eth-type <hc_ethertype_4-digit_hex>
 set l2ep-eth-type <l2ep_ethertype_4-digit_hex>
end
```

For more information, see [“Heartbeat packet Ethertypes” on page 1297](#).

## LACP, 802.3ad aggregation and third-party switches

If a cluster contains 802.3ad aggregated interfaces you should connect the cluster to switches that support configuring multiple Link Aggregation (LAG) groups.

The primary and subordinate unit interfaces have the same MAC address, so if you cannot configure multiple LAG groups a switch may place all interfaces with the same MAC address into the same LAG group; disrupting the operation of the cluster.

You can change the FortiGate configuration to prevent subordinate units from participating in LACP negotiation. For example, use the following command to do this for an aggregate interface named Port1\_Port2:

```
config system interface
 edit Port1_Port2
 set lacp-ha-slave disable
 end
```


This configuration prevents the subordinate unit interfaces from sending or receiving packets. Resulting in the cluster not being able to operate in active-active mode. As well, failover may be slower because after a failover the new primary unit has to perform LACP negotiation before being able to process network traffic.

For more information, see [“Example: HA and 802.3ad aggregated interfaces” on page 1188](#).

# VRRP

A Virtual Router Redundancy Protocol (VRRP) configuration can be used as a high availability solution to make sure that a network maintains connectivity with the Internet (or with other networks) even if the default router for the network fails. Using VRRP, if a router or a FortiGate unit fails all traffic to this router transparently fails over to another router or FortiGate unit that takes over the role of the router or FortiGate unit that failed. If the failed router or FortiGate unit is restored, it will once again take over processing traffic for the network. VRRP is described by [RFC 3768](#).

**Figure 218:**Example VRRP configuration



To configure VRRP you create a VRRP group that contains two or more routers. Some or all of these routers can be FortiGate units. You can include different FortiGate models in the same VRRP group. The group members are configured to be the master router and one or more backup routers of the VRRP group. The network directs all traffic to the master's IP address and MAC address. If the master fails, VRRP dynamically shifts packet forwarding to a backup router. VRRP provides this redundancy without user intervention or additional configuration to any of the devices on the network.

The VRRP redundancy scheme means that devices on the network keep a single IP address for the default gateway and this IP address maps to a well-known virtual MAC address. If the VRRP master fails, one of the backup units becomes the new master and acquires virtual IP and MAC addresses that match the addresses of the master. The network then automatically directs all traffic to the backup unit. VRRP uses the broadcast capabilities of Ethernet networks. As long as one of the routers in a VRRP group is running, ARP requests for the default gateway IP address always receive replies. Additionally, hosts can send packets outside their subnet without interruption.

FortiGate units support VRRP and can be quickly and easily integrated into a network that has already deployed a group of routers using VRRP. You can also create a new VRRP configuration consisting of a FortiGate unit acting as a VRRP master with one or more VRRP-compatible routers acting as backup routers. Some or all of those backup routers can be FortiGate units.

During normal operation the VRRP master unit sends VRRP advertisement messages to the backup units. A backup unit will not attempt to become a master unit while it is receiving these messages. When a FortiGate unit operating as a VRRP master fails, a backup unit takes its place and continues processing network traffic. The backup unit assumes the master unit has

failed if it stops receiving the advertisement messages from the master unit. The backup unit with the highest priority becomes the new master unit after a short delay. During this delay the new master unit sends gratuitous ARPs to the network to map the virtual router IP address to its MAC address. As a result, all packets sent to the default route IP address are sent to the new master unit. If the backup unit is a FortiGate unit, the network continues to benefit from FortiOS security features. If the backup unit is a router, after a failure traffic will continue to flow, but FortiOS security features will be unavailable until the FortiGate unit is back on line.

During a VRRP failover, as the backup unit starts to forward traffic it will not have session information for all of the failed over in-progress sessions. If the backup unit is operating as a normal FortiGate unit it will not be able to forward this traffic because of the lack of session information. To resolve this problem, immediately after a failover and for a short time as it is taking over traffic processing, the backup unit operates with asymmetric routing enabled. This allows the backup unit to re-create all of the in-progress sessions and add them to the session table. While operating with asymmetric routing enabled, the backup unit cannot apply security functions. When the start-time ends the backup unit disables asymmetric routing and returns to normal operation including applying security functions.

## Adding a VRRP virtual router to a FortiGate interface

Use the following command to add a VRRP virtual router to the port10 interface of a FortiGate unit. This VRRP virtual router has a virtual router ID of 200, uses IP address 10.31.101.200 and has a priority of 255. Since this is the highest priority this interface is configured to be the master of the VRRP group with ID number 200.

```
config system interface
 edit port10
 config vrrp
 edit 200
 set vrip 10.31.101.200
 set priority 255
 end
 end
 end
```

## VRRP virtual MAC address

The VRRP virtual MAC address (or virtual router MAC address) is a shared MAC address adopted by the VRRP master. If the master fails the same virtual MAC master fails over to the new master. As a result, all packets for VRRP routers can continue to use the same virtual MAC address. You must enable the VRRP virtual MAC address feature on all members of a VRRP group.

Each VRRP router is associated with its own virtual MAC address. The last part of the virtual MAC depends on the VRRP virtual router ID using the following format:

```
00-00-5E-00-01-<VRID_hex>
```

Where <VRID\_hex> is the VRRP virtual router ID in hexadecimal format in internet standard bit-order. For more information about the format of the virtual MAC see RFC 3768.

Some examples:

- If the VRRP virtual router ID is 10 the virtual MAC would be 00-00-5E-00-01-0a.
- If the VRRP virtual router ID is 200 the virtual MAC would be 00-00-5E-00-01-c8.

The VRRP virtual MAC address feature is disabled by default. When you enable the feature on a FortiGate interface, all of the VRRP routers added to that interface use their own VRRP virtual MAC address. Each virtual MAC address will be different because each virtual router has its own ID.

Use the following command to enable the VRRP virtual MAC address on the port2 interface:

```
config system interface
 edit port2
 set vrrp-virtual-mac enable
 end
end
```

The port2 interface will now accept packets sent to the MAC addresses of the VRRP virtual routers added to this interface.

Using the VRRP virtual MAC address can improve network efficiency especially on large and complex LANs because when a failover occurs devices on the LAN do not have to learn a new MAC address for the new VRRP router.

If the VRRP virtual MAC address feature is disabled, the VRRP group uses the MAC address of the master. In the case of a FortiGate VRRP virtual router this is the MAC address of the FortiGate interface that the VRRP virtual routers are added to. If a master fails, when the new master takes over it sends gratuitous ARPs to associate the VRRP virtual router IP address with the MAC address of the new master (or the interface of the FortiGate unit that has become the new master). If the VRRP virtual MAC address is enabled the new master uses the same MAC address as the old master.

## Configuring VRRP

To configure VRRP you must configure two or more FortiGate interfaces or routers with the same virtual router ID and IP address. Then these FortiGate units or routers can automatically join the same VRRP group. You must also assign priorities to each of the FortiGate units or routers in the VRRP group. One of the FortiGate units or routers must have the highest priority to become the master. The other FortiGate units or routers in the group are assigned lower priorities and become backup units. All of the units in the VRRP group should have different priorities. If the master unit fails, VRRP automatically fails over to the remaining unit in the group with the highest priority.

You configure VRRP from the FortiGate CLI by adding a VRRP virtual router to a FortiGate interface. You can add VRRP virtual routers to multiple FortiGate interfaces and you can add more than one virtual router to the same interface.

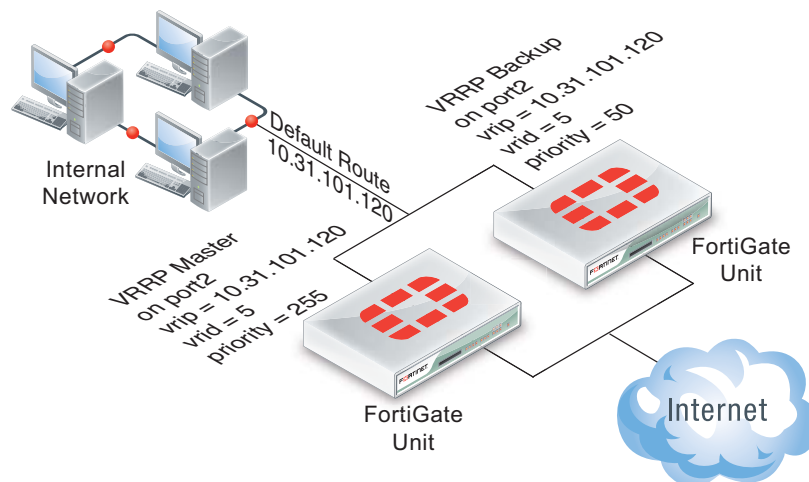
### Example VRRP configuration: two FortiGate units in a VRRP group

This example includes a VRRP group consisting of two FortiGate units that connect an internal network to the Internet. As shown in [Figure 219](#), the internal network's default route is 10.31.101.120.

The FortiGate port2 interfaces connect to the internal network. A VRRP virtual router is added to each FortiGate unit's port2 interface. The virtual router IP address is 10.31.101.120 (the internal network's default route) and the virtual router's ID is 5. The VRRP priority of the master unit is set to 255 and the VRRP priority of the backup unit is 50. The port2 interface of each FortiGate unit should have an IP address that is different from the virtual router IP address and the port2 interface IP addresses should be different from each other.

This example also includes enabling the VRRP virtual MAC address on both FortiGate unit port2 interfaces so that the VRRP group uses the VRRP virtual MAC address.

**Figure 219:**Example VRRP configuration with two FortiGate units



### To configure the FortiGate units for VRRP

1. Select one of the FortiGate units to be the VRRP master and the other to be the backup unit.
2. From the master unit's CLI, enter the following command to enable the VRRP virtual MAC address on the port2 interface and add the VRRP virtual router to the port2 interface:

```
config system interface
 edit port2
 set vrrp-virtual-mac enable
 config vrrp
 edit 5
 set vrip 10.31.101.120
 set priority 255
 end
 end
```

3. From the backup unit's CLI, enter the following command to enable the VRRP virtual MAC address on the port2 interface and add the VRRP virtual router to the port2 interface:

```
config system interface
 edit port2
 set vrrp-virtual-mac enable
 config vrrp
 edit 5
 set vrip 10.31.101.120
 set priority 50
 end
 end
```

### Example VRRP configuration: VRRP load balancing two FortiGate units and two VRRP groups


In this configuration two VRRP groups are involved. Each FortiGate unit participates in both of them. One FortiGate unit is the master of one group and the other FortiGate unit is the master of the other group. The network distributes traffic between two different default routes (10.31.101.120 and 10.31.101.130). One VRRP group is configured with one of the default route IP addresses and the other VRRP group get the other default route IP address. So during

normal operation both FortiGate units are processing traffic and the VRRP groups are used to load balance the traffic between the two FortiGate units.

If one of the FortiGate units fails, the remaining FortiGate unit becomes the master of both VRRP groups. The network sends all traffic for both default routes to this FortiGate unit. The result is a configuration that under normal operation load balances traffic between two FortiGate units, but if one of the FortiGate units fails, all traffic fails over to the unit that is still operating.

This example also includes enabling the VRRP virtual MAC address on both FortiGate unit port2 interfaces so that the VRRP groups use their VRRP virtual MAC addresses.

**Figure 220:**Example VRRP configuration with two FortiGate units and two VRRP groups



### To configure the FortiGate units

1. Log into the CLI of FortiGate unit A.
2. Enter the following command to enable the VRRP virtual MAC address feature and add the VRRP groups to the port2 interface of FortiGate unit A:

```
config system interface
 edit port2
 set vrrp-virtual-mac enable
 config vrrp
 edit 50 (32)
 set vrip 10.31.101.120
 set priority 255
 next
 edit 100 (64)
 set vrip 10.31.101.130
 set priority 50
 end
 end
end
```

3. Log into the CLI of FortiGate unit B.



4. Enter the following command to enable the VRRP virtual MAC address feature and add the VRRP groups to the port2 interface of FortiGate unit B:

```
config system interface
 edit port2
 set vrrp-virtual-mac enable
 config vrrp
 edit 50
 set vrip 10.31.101.120
 set priority 50
 next
 edit 100
 set vrip 10.31.101.130
 set priority 255
 end
 end
end
```

## Optional VRRP configuration settings

In addition to the basic configuration settings, you can change to the VRRP configuration to:

- Adjust the virtual router advertisement message interval between 1 and 255 seconds using the `adv-interval` option.
- Adjust the startup time using the `start-time` option. The default start time is 3 seconds and the range is 1 to 255 seconds. The start time is the maximum time that the backup unit waits between receiving advertisement messages from the master unit. If the backup unit does not receive an advertisement message during this time it assumes the master has failed and becomes the new master unit. In some cases the advertisement messages may be delayed. For example, some switches with spanning tree enabled may delay some of the advertisement message packets. If you find that backup units are attempting to become master units without the master unit failing, you can extend the start time to make sure the backup units wait long enough for the advertisement messages.
- Enable or disable individual virtual router configurations using the `status` option. Normally virtual router configurations are enabled but you can temporarily disable one if its not required.
- Enable or disable preempt mode using the `preempt` option. In preempt mode a higher priority backup unit can preempt a lower priority master unit. This can happen if a master has failed, a backup unit has become the master unit, and the failed master is restarted. Since the restarted unit will have a higher priority, if preempt mode is enabled the restarted unit will replace the current master unit. Preempt mode is enabled by default.
- Monitor the route to a destination IP address using the `vrdst` option.

# FortiGate Session Life Support Protocol (FGSP)

You can use the `config system session-sync` command to configure the FortiGate Session Life Support Protocol (FGSP) (previously called TCP session synchronization or standalone session synchronization) between two FortiGate units. The two FortiGate units must be the same model. The FGSP synchronizes both IPv4 and IPv6 TCP, UDP, ICMP, expectation, and NAT sessions. You can use this feature with external routers or load balancers configured to distribute or load balance sessions between two peer FortiGate units. If one of the peers fails, session failover occurs and active sessions fail over to the peer that is still operating. This failover occurs without any loss of data. As well, the external routers or load balancers will detect the failover and re-distribute all sessions to the peer that is still operating.



In previous versions of FortiOS the FGSP was called TCP session synchronization or standalone session synchronization. However, the FGSP has been expanded to include configuration synchronization and session synchronization of connectionless sessions, expectation sessions, and NAT sessions.



You cannot configure FGSP HA when FGCP HA is enabled. However FCSP HA is compatible with VRRP.

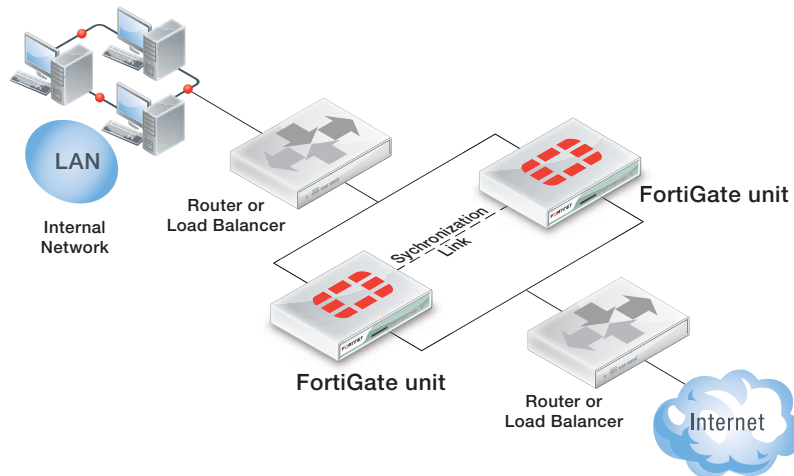


FCSP or standalone session synchronization is not supported if the FortiGate units are running different firmware versions.

---

The FGSP can be used instead of FGCP HA to provide **session synchronization** between two peer FortiGate units. If the external load balancers direct all sessions to one peer the affect is similar to active-passive FGCP HA. If external load balancers or routers load balance traffic to both peers, the effect is similar to active-active FGCP HA. The load balancers should be configured so that all of the packets for any given session are processed by the same peer. This includes return packets.

**Figure 221:FGSP HA**



By default, FGSP synchronizes all IPv4 and IPv6 TCP sessions and also synchronizes the configuration of the FortiGate units.

You can optionally enable session pickup to synchronize connectionless (UDP and ICMP) sessions, expectation sessions, and NAT sessions. If you do not enable session pickup, the FGSP does not share session tables for the particular session type and sessions do not resume after a failover. All sessions that are interrupted by the failover and must be re-established at the application level. Many protocols can successfully restart sessions with little, or no, loss of data. Others may not recover easily. Enable session pickup for sessions that may be difficult to reestablish. Since session pickup requires FortiGate resources, only enable this feature for sessions that you need to have synchronized.

You can also optionally add filters to control which sessions are synchronized. You can add filters to only synchronize packets from specified source and destination addresses, specified source and destination interfaces, and specified services.

Load balancing and session failover is done by external routers or load balancers instead of by the FGSP. The FortiGate units just perform session synchronization to support session failover.

## Synchronizing the configuration

The FGSP also includes configuration synchronization, allowing you to make configuration changes once for both FortiGate units instead of requiring you to make duplicate configuration changes on each FortiGate unit. Settings that identify the FortiGate unit to the network, for example, interface IP addresses and BGP neighbor settings, are not synchronized so each FortiGate unit maintains its identity on the network.

By default configuration synchronization is disabled. You can use the following command to enable it.

```
config system ha
 set standalone-config-sync enable
end
```

## Synchronizing UDP and ICMP (connectionless) sessions

In many configurations, due to their non-stateful nature, UDP and ICMP sessions don't need to be synchronized to naturally failover. However, if its required you can configure the FGSP to synchronize UDP and ICMP sessions by entering the following command:

```
config system ha
 set session-pickup enable
 set session-pickup-connectionless enable
end
```

## Synchronizing NAT sessions

By default, NAT session are not synchronized. However, the FGSP can synchronize NAT session if you enter the following command:

```
config system ha
 set session-pickup enable
 set session-pickup-nat enable
end
```

However, if you want NAT sessions to resume after a failover you should not configure NAT to use the destination interface IP address since the FGSP FortiGate units have different IP addresses. With this configuration, after a failover all sessions that include the IP addresses of interfaces on the failed FortiGate unit will have nowhere to go since the IP addresses of the failed FortiGate unit will no longer be on the network.

Instead, in an FGSP configuration, if you want NAT sessions to failover you should use IP pools with the type set to overload (which is the default IP pool type). For example:

```
config firewall ippool
 edit FGSP-pool
 set type overload
 set startip 172.20.120.10
 set endip 172.20.120.20
 end
```

Then when you configure NAT firewall policies, turn on NAT and select to use dynamic IP pool and select the IP Pool that you added. Add the same IP pools and firewall policies to both FortiGate units.

## Synchronizing expectation (asymmetric) sessions

By default, expectation sessions (or asymmetric sessions) are not synchronized. Normally, session synchronization cannot be asymmetric because it is stateful. So all of the packets of a given session must be processed on the same peer. This includes return packets.

However, if you have an asymmetric routing configuration, you can enter the following command to synchronize asymmetric sessions by dynamically detecting asymmetric sessions and disabling anti-reply for these sessions.

```
config system ha
 set session-pickup enable
 set session-pickup-expectation enable
end
```

The FGSP enforces firewall policies for asymmetric traffic, including cases where the TCP 3-way handshake is split between two FortiGates. For example, FGT-A receives the TCP-SYN, FGT-B receives the TCP-SYN-ACK, and FGT-A receives the TCP-ACK. Under normal conditions a firewall will drop this connection since the 3-way handshake was not seen by the same firewall. However two FortiGates with FGSP configured will be able to properly pass this traffic since the firewall sessions are synchronized.

If traffic will be highly asymmetric, as described above, the following command must be enabled on both FortiGates.

```
config system ha
 set session-pickup enable
 set session-pickup-expectation enable
end
```

This asymmetric function can also work with connectionless UDP and ICMP traffic. The following command needs to be enabled on both FortiGates.

```
config system ha
 set session-pickup enable
 set session-pickup-connectionless enable
end
```

Synchronizing asymmetric traffic can be very useful in situations where multiple Internet connections from different ISPs are spread across two FortiGates. Since it is typically not possible to guarantee Internet bound traffic leaving via an ISP will return using the exact same ISP, the FGSP provides critical firewall functions in this situation.

The FGSP also has applications in virtualized computing environments where virtualized hosts move between data centers. The firewall session synchronization features of FGSP allow for more flexibility than in traditional firewalling functions.

## UTM Flow-based Inspection and Asymmetric Traffic

UTM inspection (flow or proxy based) for a session is not expected to work properly if the traffic in the session is balanced across more than one FortiGate in either direction. Flow-based UTM should be used in FGSP deployments.

For an environment where traffic is symmetric, UTM can be used with the following limitations:

- No session synchronization for the sessions inspected using proxy-based UTM. Sessions will drop and need to be reestablished after data path failover.
- Sessions with flow-based UTM will failover; however, inspection of failed over sessions after the failover may not work.

A single FortiGate must see both the request and reply traffic for UTM inspection to function correctly. For environments where asymmetric traffic is expected, UTM inspection should not be used.

## Notes and limitations

FGSP HA has the following limitations:

- The FGSP is a global configuration option. As a result you can only add one service to a filter configuration. You cannot add custom services or service groups even if virtual domains are not enabled.

- You can only add one filter configuration to a given FGSP configuration. However, you can add multiple filters by adding multiple identical FGSP configurations, each one with a different filter configuration.
- Sessions accepted by security policies with UTM options configured are not synchronized.
- FGSP HA is configured from the CLI.
- FGSP HA is available for FortiGate units or virtual domains operating in NAT/Route or Transparent mode. NAT sessions are not synchronized in either mode (unless NAT synchronization is enabled as described in [“Synchronizing NAT sessions” on page 1372](#)). In NAT/Route mode, only sessions for route mode security policies are synchronized. In Transparent mode, only sessions for normal Transparent mode policies are synchronized.
- FGSP HA is supported for traffic on physical interfaces, VLAN interfaces, zones, aggregate interfaces, and NPx (NP4, NP6 etc.) accelerated interfaces. The FGSP has not been tested for inter-vdom links, between HA clusters, and for redundant interfaces.
- The names of the matching interfaces, including VLAN interfaces, aggregate interfaces and so on, must be the same on both peers.

## Configuring FGSP HA

You configure FGSP HA separately for each virtual domain to be synchronized. If virtual domain configuration is not enabled, you configure FGSP HA for the root virtual domain. When virtual domain configuration is enabled and you have added virtual domains you configure FGSP HA for each virtual domain to be synchronized. You don't have to synchronize all virtual domains.

You must configure FGSP HA and network settings on both peers. Once you establish the initial configuration, the configurations of both FortiGate units are synchronized so when you change the configuration of one, the changes are synchronized to the other.

On each FortiGate unit, configuring FGSP HA consists of selecting the virtual domains to be synchronized using the `syncvd` field, selecting the virtual domain on the other peer that receives the synchronization packets using the `peervd` field, and setting the IP address of the interface in the peer unit that receives the synchronization packets using the `peerip` field. The interface with the `peerip` must be in the `peervd` virtual domain.

The `syncvd` and `peervd` settings must be the same on both peers. However, the `peerip` settings will be different because the `peerip` setting on the first peer includes the IP address of an interface on the second peer. And the `peerip` setting on the second peer includes the IP address of an interface on the first peer.

For FGSP HA to work properly all synchronized virtual domains must be added to both peers. The names of the matching interfaces in each virtual domain must also be the same; this includes the names of matching VLAN interfaces. Note that the index numbers of the matching interfaces and VLAN interfaces can be different. Also the VLAN IDs of the matching VLAN interfaces can be different.

For a configuration example, see [“Basic example configuration” on page 1375](#).

## Configuring the session synchronization link

When FGSP HA is operating, the peers share session information over an Ethernet link between the peers similar to an HA heartbeat link. Usually you would use the same interface on each peer for session synchronization. You should connect the session synchronization interfaces directly without using a switch or other networking equipment. For FortiGate-5000 systems you can use a backplane interface as the session synchronization link.

You can use different interfaces on each peer for session synchronization links. Also, if you have multiple sessions synchronization configurations, you can have multiple links between the peers. In fact if you are synchronizing a lot of sessions, you may want to configure and connect multiple session synchronization links to distribute session synchronization traffic to these multiple links.

You cannot configure backup session synchronization links. Each configuration only includes one session synchronization link.

The session synchronization link should always be maintained. If session synchronization communication is interrupted and a failure occurs, sessions will not failover and data could be lost.

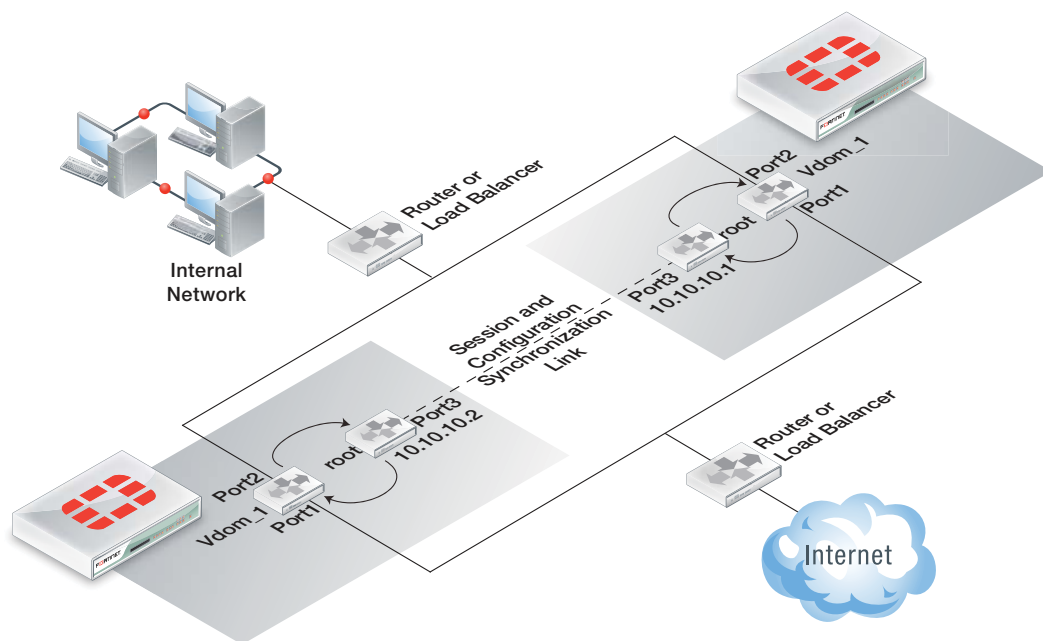
Session synchronization traffic can use a considerable amount of network bandwidth. If possible, session synchronization link interfaces should only be used for session synchronization traffic and not for data traffic.

## Basic example configuration

The following configuration example shows how to configure basic FGSP HA for the two peer FortiGate units shown in [Figure 222 on page 1375](#). The host names of peers are peer\_1 and peer\_2. Both peers are configured with two virtual domains: root and vdom\_1. All sessions processed by vdom\_1 are synchronized. The synchronization link interface is port3 which is in the root virtual domain. The IP address of port3 on peer\_1 is 10.10.10.1. The IP address of port3 on peer\_2 is 10.10.10.2.

Also on both peers, port1 and port2 are added to vdom\_1. On peer\_1 the IP address of port1 is set to 192.168.20.1 and the IP address of port2 is set to 172.110.20.1. On peer\_2 the IP address of port1 is set to 192.168.20.2 and the IP address of port2 is set to 172.110.20.2.

**Figure 222:**Example FGSP HA network configuration



### To configure FGSP HA

1. Configure the load balancer or router to send all sessions to peer\_1.
2. Configure the load balancer or router to send all traffic to peer\_2 if peer\_1 fails.

3. Use normal FortiGate configuration steps on peer\_1:
  - Enable virtual domain configuration.
  - Add the vdom\_1 virtual domain.
  - Add port1 and port2 to the vdom\_1 virtual domain and configure these interfaces.
  - Set the IP address of port1 to 192.168.20.1.
  - Set the IP address of port2 to 172.110.20.1.
  - Set the IP address of port3 to 10.10.10.1.
  - Add route mode security policies between port1 and port2 to vdom\_1.

4. Enter the following commands to configure session synchronization for peer\_1

```
config system session-sync
 edit 1
 set peerip 10.10.10.2
 set peervd root
 set syncvd vdom_1
 end
```

5. Use normal FortiGate configuration steps on peer\_2:

- Enable virtual domain configuration.
- Add the vdom\_1 virtual domain.
- Add port1 and port2 to the vdom\_1 virtual domain and configure these interfaces.
- Set the IP address of port1 to 192.168.20.2.
- Set the IP address of port2 to 172.110.20.2.
- Set the IP address of port3 to 10.10.10.1.
- Add route mode security policies between port1 and port2 to vdom\_1.

6. Enter the following command to configure session synchronization for peer\_1

```
config system session-sync
 edit 1
 set peerip 10.10.10.1
 set peervd root
 set syncvd vdom_1
 end
```

Now that the FortiGate units are connected and configured their configurations are synchronized, so when you make a configuration change on one FortiGate unit it is synchronized to the other one.

### To add filters

You can add a filter to this basic configuration if you only want to synchronize some TCP sessions. For example you can enter the following command to add a filter so that only HTTP sessions are synchronized:

```
config system session-sync
 edit 1
 config filter
 set service HTTP
 end
 end
```



You can also add a filter to control the source and destination addresses of the IPv4 packets that are synchronized. For example you can enter the following command to add a filter so that only sessions with source addresses in the range 10.10.10.100 to 10.10.10.200 are synchronized.

```
config system session-sync
 edit 1
 config filter
 set srcaddr 10.10.10.100 10.10.10.200
 end
 end
```

You can also add a filter to control the source and destination addresses of the IPv6 packets that are synchronized. For example you can enter the following command to add a filter so that only sessions with destination addresses in the range 2001:db8:0:2::/64 are synchronized.

```
config system session-sync
 edit 1
 config filter
 set dstaddr6 2001:db8:0:2::/64
 end
 end
```

### **To synchronize UDP and ICMP sessions**

You enter the following command to add synchronization of UDP and ICMP sessions to this configuration:

```
config system ha
 set session-pickup enable
 set session-pickup-connectionless enable
end
```

### **To synchronize the configuration**

Enter the following command to enable configuration synchronization.

```
config system ha
 set standalone-config-sync enable
end
```

## **Verifying FGSP configuration and synchronization**

You can use the following diagnose commands to verify that the FGSP and its synchronization functions are operating correctly.

## FGSP configuration summary and status

Enter the following command to display a summary of the FGSP configuration and synchronization status:

```
diagnose sys session sync
sync_ctx: sync_started=1, sync_tcp=1, sync_others=1,
sync_expectation=1, sync_redir=0, sync_nat=1.
sync: create=12:0, update=0, delete=0:0, query=14
recv: create=14:0, update=0, delete=0:0, query=12
ses pkts: send=0, alloc_fail=0, recv=0, recv_err=0 sz_err=0
nCfgr_sess_sync_num=5, mtu=16000
sync_filter:
1: vd=0, szone=0, dzone=0, saddr=0.0.0.0:0.0.0.0,
 daddr=0.0.0.0:0.0.0.0,
```

`sync_started=1` shows that synchronization is working. If this is set to 0 then something is not correct with session synchronization and synchronization has not been able to start because of it.

`sync_tcp=1`, `sync_others=1`, `sync_expectation=1`, and `sync_nat=1` show that the FGSP has been configured to synchronize TCP, connectionless, asymmetric, and NAT sessions.

`sync: create=12:0` and `recv: create=14:0` show that this FortiGate has synchronized 12 sessions to its peer and has received 14 sessions from its peer.

`sync_filter` shows the configured FGSP filter. In this case no filter has been created so all sessions are synchronized.

`vd=0` indicates that root VDOM sessions are synchronized.

## Verifying that sessions are synchronized

Enter the command `diagnose sys session list` to display information about the sessions being processed by the FortiGate. In the command output look for sessions that should be synchronized and make sure they contain output lines that include `synced` (for example, `state=log may_dirty ndr synced`) to confirm that they are being synchronized by the FGSP.

```
diagnose sys session list
session info: proto=6 proto_state=05 duration=469 expire=0 timeout=3600
flags=00000000 sockflag=00000000 sockport=21 av_idx=0 use=4
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 policy_dir=0 tunnel=/
state=log may_dirty ndr synced
statistic(bytes/packets/allow_err): org=544/9/1 reply=621/7/0 tuples=2
origin->sink: org pre->post, reply pre->post
 dev=46->45/45->46
gwy=10.2.2.1/10.1.1.1
hook=pre dir=org act=noop
 192.168.1.50:45327->172.16.1.100:21(0.0.0.0:0)
hook=post dir=reply act=noop
 172.16.1.100:21->192.168.1.50:45327(0.0.0.0:0)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=1 id_policy_id=0 auth_info=0 chk_client_info=0 vd=0
serial=00002deb tos=ff/ff ips_view=1 app_list=2000 app=16427
dd_type=0 dd_mode=0
per_ip_bandwidth meter: addr=192.168.1.50, bps=633
```

# Configuring FRUP

The FortiGate Redundant UTM Protocol (FRUP) provides similar redundancy to FGCP full mesh HA in a single unified design that includes redundant switching and routing. FRUP is available on the FortiGate-100D and will be expanded to other models in future releases.

A FRUP cluster consists of 2 (and only 2) identical FortiGate-100D units that have dual redundant links to all connected devices and networks and can include redundant FortiAP units. Connections to the Internet normally use the wan1 and wan2 interfaces for redundant connections. Connections to internal networks and servers use redundant connections to FortiGate-100D switch ports. FRUP uses the FortiGate-100D switch ports for full mesh HA instead of external redundant switches.

Each device or network has a default active connection to one of the FortiGate units and a default backup connection to the other. Ideally, the default active and backup connections should balance traffic between the FortiGate units in the cluster so that both FortiGate units are processing the same amount of traffic.

FRUP uses virtual IPs and virtual MACs so that when a failover occurs, network devices do not have to learn new IP or MAC addresses. FRUP also synchronizes the configuration between the units in the cluster.

Use the following CLI command on both FortiGate-100D units to configure FRUP.

```
config system ha
 set hbdev "ha1" 50 "ha2" 100
 set override disable
 set priority 128
 set frup enable
 config frup-settings
 set active-interface "wan2"
 set active-switch-port 14
 set backup-interface "wan1"
 end
end
```


Both units must have the same heartbeat device configuration and have FRUP enabled to form a FRUP cluster. Active interface and switch ports must be complementary according to your configuration (see the following example).

## FRUP configuration example

This example includes the following:

- Two FortiGate-100D units (FGT-A and FGT-B)
- HA1 and HA2 for redundant HA heartbeat connectivity between FGT-1 and FGT-2
- Dual gateways (router1 and router2):
  - FGT-A has an active connection from WAN1 to router-1 and a backup connection from WAN2 to router-2
  - FGT-B has an active connection from WAN2 to router-2 and a backup connection from WAN1 to router-1
- Dual connections to the internal network and an internal server:
  - FGT-A has an active connection to an internal server using switch port 16 and a backup connection to the internal network using switch port 14
  - FGT-B has a backup connection to an internal server using switch port 16 and an active connection to the internal network using switch port 14
- FortiGate interfaces use virtual IP addresses and pseudo-MAC physical addresses, all devices continue to send to the same IP/Mac and don't need to re-learn after a failover
- Both FortiGate units will handle and process traffic
- Backup links are normally administratively down
- Sessions and the FortiGate configuration are synchronized between the cluster units

**Figure 223:**Example FRUP configuration



### Configuring FGT-A

Change the host name to FGT-A. Set up FGT-A with dual redundant internet links using WAN1 and WAN2. Set WAN1 and WAN2 interfaces to use static addressing and set both static routes to same priority and distance.

From CLI enter the following command:

```
config system ha
 set hbdev "ha1" 50 "ha2" 100
 set override disable
 set priority 255
 set frup enable
 config frup-settings
 set active-interface "wan1"
 set active-switch-port 16
 set backup-interface "wan2"
 end
end
```

## Configuring FGT-B

Use the same firmware version as FGT-A and set the FortiGate unit to factory defaults. Change the host name to FGT-B.

From CLI enter the following command:

```
config system ha
 set hbdev "ha1" 50 "ha2" 100
 set override disable
 set priority 128
 set frup enable
 config frup-settings
 set active-interface "wan2"
 set active-switch-port 14
 set backup-interface "wan1"
 end
end
```

## Connecting, testing and operating the FRUP cluster

Connect to the FortiGate-100D units to the network as shown in the diagram and power them on. The FortiGate-100D units should find each other and form a cluster.

Traffic from the server (IP: 192.168.1.200) connected to port16 should pass through FGT-A (since port16 is the active-switch-port in FGT-A) and to the Internet using WAN1 (since WAN1 is the active-interface in FGT-A).

Run a sniffer on both FortiGate units using the following command:

```
diagnose sniffer packet any 'icmp' 4
interfaces=[any]
filters=[icmp]
```

Then run a ping to 4.2.2.2 from the server. The sniffer should show results similar to the following:

```

FGT-A # diagnose sniffer packet any 'icmp' 4
interfaces=[any]
filters=[icmp]
0.231160 wan1 in 4.2.2.2 -> 172.20.120.123: icmp: echo reply
0.231202 internal out 4.2.2.2 -> 192.168.1.200: icmp: echo reply
0.231209 eth1 out 4.2.2.2 -> 192.168.1.200: icmp: echo reply
1.198520 internal in 192.168.1.200 -> 4.2.2.2: icmp: echo request
1.198555 wan1 out 172.20.120.123 -> 4.2.2.2: icmp: echo request
1.222569 wan1 in 4.2.2.2 -> 172.20.120.123: icmp: echo reply
1.222589 internal out 4.2.2.2 -> 192.168.1.200: icmp: echo reply
1.222595 eth1 out 4.2.2.2 -> 192.168.1.200: icmp: echo reply
2.199916 internal in 192.168.1.200 -> 4.2.2.2: icmp: echo request
2.199952 wan1 out 172.20.120.123 -> 4.2.2.2: icmp: echo request
2.232998 wan1 in 4.2.2.2 -> 172.20.120.123: icmp: echo reply
2.233017 internal out 4.2.2.2 -> 192.168.1.200: icmp: echo reply
2.233023 eth1 out 4.2.2.2 -> 192.168.1.200: icmp: echo reply
3.201347 internal in 192.168.1.200 -> 4.2.2.2: icmp: echo request
3.201385 wan1 out 172.20.120.123 -> 4.2.2.2: icmp: echo request
3.235406 wan1 in 4.2.2.2 -> 172.20.120.123: icmp: echo reply
3.235425 internal out 4.2.2.2 -> 192.168.1.200: icmp: echo reply
3.235430 eth1 out 4.2.2.2 -> 192.168.1.200: icmp: echo reply

18 packets received by filter
0 packets dropped by kernel

FGT-A # █

```

Traffic initiating from a host on the internal network (for example, IP: 192.168.1.114) connected to port14 should flow through FGT-B (since port14 is the active-switch-port in FGT-B) to the Internet using WAN2 (since WAN2 is FGT-B's active interface).

Run the same sniffer on both FortiGate units and then run a ping to 74.125.226.1 from the internal network. The sniffer should show results similar to the following:

```

FGT-B $ diagnose sniffer packet any 'icmp' 4
interfaces=[any]
filters=[icmp]
1.887458 internal in 192.168.1.114 -> 74.125.226.1: icmp: echo request
1.887488 wan2 out 172.20.120.130 -> 74.125.226.1: icmp: echo request
1.887492 eth0 out 172.20.120.130 -> 74.125.226.1: icmp: echo request
1.898137 wan2 in 74.125.226.1 -> 172.20.120.130: icmp: echo reply
1.898153 internal out 74.125.226.1 -> 192.168.1.114: icmp: echo reply
1.898159 eth1 out 74.125.226.1 -> 192.168.1.114: icmp: echo reply
2.885644 internal in 192.168.1.114 -> 74.125.226.1: icmp: echo request
2.885682 wan2 out 172.20.120.130 -> 74.125.226.1: icmp: echo request
2.885687 eth0 out 172.20.120.130 -> 74.125.226.1: icmp: echo request
2.896175 wan2 in 74.125.226.1 -> 172.20.120.130: icmp: echo reply
2.896194 internal out 74.125.226.1 -> 192.168.1.114: icmp: echo reply
2.896201 eth1 out 74.125.226.1 -> 192.168.1.114: icmp: echo reply
3.884046 internal in 192.168.1.114 -> 74.125.226.1: icmp: echo request
3.884091 wan2 out 172.20.120.130 -> 74.125.226.1: icmp: echo request
3.884096 eth0 out 172.20.120.130 -> 74.125.226.1: icmp: echo request
3.894192 wan2 in 74.125.226.1 -> 172.20.120.130: icmp: echo reply
3.894213 internal out 74.125.226.1 -> 192.168.1.114: icmp: echo reply
3.894220 eth1 out 74.125.226.1 -> 192.168.1.114: icmp: echo reply

18 packets received by filter
0 packets dropped by kernel

```

Shutdown FGT-A.

Traffic from the internal network should be handled by FGT-B using WAN1 and WAN2 interfaces:

```
FGT-B # diagnose sniffer packet any 'icmp' 4
interfaces=[any]
filters=[icmp]
0.954086 internal in 192.168.1.200 -> 4.2.2.2: icmp: echo request
0.954226 wan1 out 172.20.120.123 -> 4.2.2.2: icmp: echo request
0.968696 wan1 in 4.2.2.2 -> 172.20.120.123: icmp: echo reply
0.968780 internal out 4.2.2.2 -> 192.168.1.200: icmp: echo reply
0.968796 eth1 out 4.2.2.2 -> 192.168.1.200: icmp: echo reply
1.166934 internal in 192.168.1.114 -> 74.125.226.1: icmp: echo request
1.166960 wan2 out 172.20.120.130 -> 74.125.226.1: icmp: echo request
1.166966 eth0 out 172.20.120.130 -> 74.125.226.1: icmp: echo request
1.177525 wan2 in 74.125.226.1 -> 172.20.120.130: icmp: echo reply
1.177541 internal out 74.125.226.1 -> 192.168.1.114: icmp: echo reply
1.177547 eth1 out 74.125.226.1 -> 192.168.1.114: icmp: echo reply
1.955117 internal in 192.168.1.200 -> 4.2.2.2: icmp: echo request
1.955259 wan1 out 172.20.120.123 -> 4.2.2.2: icmp: echo request
1.987992 wan1 in 4.2.2.2 -> 172.20.120.123: icmp: echo reply
1.988084 internal out 4.2.2.2 -> 192.168.1.200: icmp: echo reply
1.988101 eth1 out 4.2.2.2 -> 192.168.1.200: icmp: echo reply
2.165320 internal in 192.168.1.114 -> 74.125.226.1: icmp: echo request
2.165346 wan2 out 172.20.120.130 -> 74.125.226.1: icmp: echo request
2.165352 eth0 out 172.20.120.130 -> 74.125.226.1: icmp: echo request
2.175081 wan2 in 74.125.226.1 -> 172.20.120.130: icmp: echo reply
2.175098 internal out 74.125.226.1 -> 192.168.1.114: icmp: echo reply
2.175105 eth1 out 74.125.226.1 -> 192.168.1.114: icmp: echo reply
2.956439 internal in 192.168.1.200 -> 4.2.2.2: icmp: echo request
2.956583 wan1 out 172.20.120.123 -> 4.2.2.2: icmp: echo request
2.973142 wan1 in 4.2.2.2 -> 172.20.120.123: icmp: echo reply
2.973237 internal out 4.2.2.2 -> 192.168.1.200: icmp: echo reply
2.973255 eth1 out 4.2.2.2 -> 192.168.1.200: icmp: echo reply
3.163683 internal in 192.168.1.114 -> 74.125.226.1: icmp: echo request
3.163709 wan2 out 172.20.120.130 -> 74.125.226.1: icmp: echo request
3.163714 eth0 out 172.20.120.130 -> 74.125.226.1: icmp: echo request
3.174329 wan2 in 74.125.226.1 -> 172.20.120.130: icmp: echo reply
3.174362 internal out 74.125.226.1 -> 192.168.1.114: icmp: echo reply
3.174369 eth1 out 74.125.226.1 -> 192.168.1.114: icmp: echo reply
3.957570 internal in 192.168.1.200 -> 4.2.2.2: icmp: echo request
3.957711 wan1 out 172.20.120.123 -> 4.2.2.2: icmp: echo request
3.979899 wan1 in 4.2.2.2 -> 172.20.120.123: icmp: echo reply
3.979990 internal out 4.2.2.2 -> 192.168.1.200: icmp: echo reply
3.980012 eth1 out 4.2.2.2 -> 192.168.1.200: icmp: echo reply

38 packets received by filter
0 packets dropped by kernel

FGT-B #
```

To re-establish the cluster after a failover you need to restart both FortiGate-100D units. Just re-starting FGT-A may not bring the cluster back online.



# Chapter 10 Install and System Administration for FortiOS 5.0

This guide contains the following sections:

[Differences between Models and Firmware](#) highlights key differences that exist between FortiGate models and firmware versions.

[Using the web-based manager](#) provides an overview of the web-based manager interface for FortiOS. If you are new to the FortiOS web-based manager, this chapter provides a high level overview of how to use this method of administration.

[Using the CLI](#) provides an overview of the command line interface (CLI) for FortiOS. If you are new to the FortiOS CLI, this chapter provides a high level overview of how to use this method of administration.

[Basic Administration](#) describes the simple setup requirements an administrator should do to get the FortiGate unit on the network and enabling the flow of traffic.

[Best practices](#) discusses methods to make the various components of FortiOS more efficient and offers suggestions on ways to configure the FortiGate unit.

[FortiGuard](#) discusses the FortiGuard network services and configuration examples.

[FortiCloud](#) discusses the FortiCloud hosted security management and log retention service.

[Interfaces](#) describes the FortiGate interface options and configuration choices.

[Central management](#) describes how to configure the FortiGate unit to use FortiManager as a method of maintaining the device and other features that FortiManager has to facilitate the administration of multiple devices.

[Monitoring](#) describes various methods of collecting log data and tracking traffic flows and trends.

[VLANs](#) discusses the implementation of virtual local area networks (VLANs) in FortiOS and how to configure and use them.

[PPTP and L2TP](#) describes these virtual private network (VPN) types and how to configure them.

[Session helpers](#) describes what session helpers are and how to view and configure them.

[Advanced concepts](#) describes more involved administrative topics to enhance network security and traffic efficiency.

# Differences between Models and Firmware

This section examines some of the key differences that exist between different FortiGate models and different versions of the FortiOS 5.0 firmware. It is important to keep these differences in mind when reading the FortiOS 5.0 Handbook, in order to understand how the documentation relates to your specific FortiGate unit.

## Differences between Models

- There are certain features that are not available on all models. For example, the Switch Controller, which allows a FortiGate unit to manage a FortiSwitch unit, is only available on FortiGate models 100D, 140D, 200D, 240D, 600C, 800C, and 1000C.

Other features may be available only through the CLI on models, while other models have options in the web-based manager. For example, SSL content inspection is a CLI-only feature on FortiGate models 20C, 30C, and 40C, while models 60C+ have options in the web-based manager.

For more information about some of the features that vary by model, please see the [Feature/Platform Matrix](#).

- Naming conventions may vary between FortiGate models. For example, on some models the interface used for the local area network is called *lan*, while on other units it is called *internal*.
- Menus may vary by model. For example, on some FortiGate units, the menu option *Router* is not available. Instead, routing is configured by going to *System > Network > Routing*.

## Differences between Firmware Versions

- Many changes are introduced in new patches to the FortiOS 5.0 firmware. For more information about these changes, please see [What's New for FortiOS 5.0](#) and the FortiOS [Release Notes](#).
- In FortiOS 5.0 Patch 3, *Feature Select* was added, which controls which menus are visible in the web-based manager. If a feature you wish to use does not appear in the web-based manager, go to *System > Config > Features* to ensure that the feature has not been turned off.
- Menu names may change between firmware versions. For example, In FortiOS 5.0 Patch 3, the features formerly known *UTM Security Profiles* was renamed *Security Profiles*.
- Menus may move between firmware versions. For example, in FortiOS Patch 5, Endpoint Control moved from being part of the *User* menu to having a menu option of its own, found at *User & Device > Endpoint Protection*.
- Options may also be removed in a firmware patch. For example, in FortiOS Patch 5, the *Client Reputation Monitor* was removed. Client Reputation results can now be found in the *Threat History* widget.

# Using the web-based manager

This section describes the features of the web-based manager administrative interface (sometimes referred to as a graphical user interface, or GUI) of your unit. This section also explains common web-based manager tasks that an administrator does on a regular basis.

The following topics are included in this section:

- [Web-based manager overview](#)
- [Web-based manager menus and pages](#)
- [Entering text strings](#)
- [Dashboard](#)
- [Basic configurations](#)

## Web-based manager overview

The web-based manager is a user-friendly interface for configuring settings and managing the FortiGate unit. Accessing the web-based manager is easy and can be done by using either HTTP or a secure HTTPS connection from any management computer, using a web browser.

The recommended minimum screen resolution for properly displaying the web-based manager is 1280 by 1024. Some web browsers do not correctly display the windows within the web-based manager interface. Verify that you have a supported web browser by reviewing the Knowledge Base articles: [Microsoft Windows web browsers supported by Fortinet products web-based manager \(GUI\) web browsers](#) and [Mac OS browsers for use with Fortinet hardware web-based manager \(GUI\)](#).

The web-based manager also provides the CLI Console widget, which enables you to connect to the command line interface (CLI) without exiting out of the web-based manager.

## Web-based manager menus and pages

The web-based manager provides access to configuration options for most of the FortiOS features from the main menus. The web-based manager contains the following main menus:

<b>System</b>	Configure system settings, such as network interfaces, virtual domains, DHCP and DNS services, administrators, certificates, High Availability (HA), system time, set system options and set display options on the web-based manager.
<b>Router</b>	Configure static, dynamic and multicast routing and view the router monitor.
<b>Policy</b>	Configure firewall policies, protocol options and Central NAT Table.
<b>Firewall Objects</b>	Configure supporting content for firewall policies including scheduling, services, traffic shapers, addresses, virtual IP and load balancing.
<b>Security Profiles</b>	Configure antivirus and email filtering, web filtering, intrusion protection, data leak prevention, application control, VOIP, ICAP and Client Reputation.

<b>VPN</b>	Configure IPsec and SSL virtual private networking.
<b>User &amp; Device</b>	Configure user accounts and user authentication including external authentication servers. This menu also includes endpoint security features, such as FortiClient configuration and application detection patterns.
<b>WAN Opt. &amp; Cache</b>	Configure WAN optimization and web caching to improve performance and security of traffic passing between locations on your wide area network (WAN) or from the Internet to your web servers.
<b>WiFi &amp; Switch Controller</b>	Configure the unit to act as a wireless network controller, managing the wireless Access Point (AP) functionality of FortiWiFi and FortiAP units.
<b>Log &amp; Report</b>	Configure logging and alert email as well as reports. View log messages and reports.
<b>Current VDOM</b>	Appears only when VDOMs are enabled on the unit to switch between VDOMs.

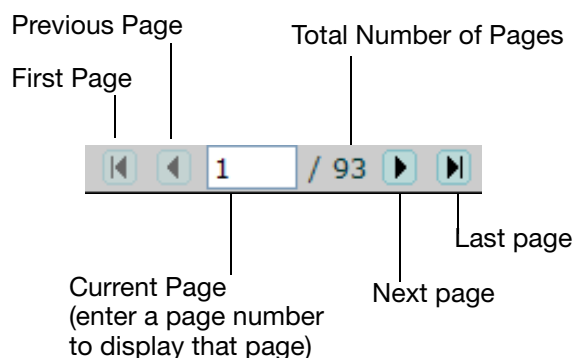
## Using information tables

Many of the web-based manager pages contain tables of information that you can filter to display specific information. Administrators with read and write access can define the filters.

## Using page navigation

Some pages contain information and lists that span multiple pages. At the bottom of the page is the page navigation controls that enables you to move between pages.

**Figure 224:**Page controls



## Adding filters to web-based manager lists

Filters are used to locate a specific set of information or content within multiple pages. These are especially useful in locating specific log entries. The specific filtering options vary, depending on the type of information in the log.

To create a filter, select *Filter Settings* or the filter icon in a column heading. When a filter is applied to a column, the filter icon becomes green. Filter settings are stored in the unit's configuration and will be maintained the next time that you access any list for which you have added filters.

Filtering variables can include: a numeric range (such as 25-50), an IP address or part of an address or any text string combination, including special characters.

Note that the filtering ignores characters following a “<” unless the followed by a space. For example, the filtering ignores `<string` but not `< string`. Filtering also ignores matched opening and closing (`<` and `>`) characters and any characters between them. For example, filtering will ignore `<string>`.

For columns that contain only specific content, such as log message severity, a list of terms is provided from which options can be selected.

## Using column settings

Column settings are used to select the types of information which are displayed on a certain page. Some pages have a large amounts of information is available and not all content can be displayed on a single screen. Also, some pages may contain content that is not of use to you. Using column settings, you can display only that content which is important to your requirements.

To configure column settings, right-click the header of a column and select *Column Settings*. Any changes that you make to the column settings of a list are stored in the unit’s configuration and will display the next time that you access the list.

To return a page’s columns to their default state, select *Reset All Columns*, located at the bottom of the *Column Settings* menu.

## Entering text strings

The configuration of a FortiGate unit is stored in the FortiOS configuration database. To change the configuration, you can use the web-based manager or CLI to add, delete, or change configuration settings. These changes are stored in the database as you make them.

Individual settings in the configuration database can be text strings, numeric values, selections from a list of allowed options, or on/off (enable/disable) settings.

### Entering text strings (names)

Text strings are used to name entities in the configuration. For example, the name of a firewall address, administrative user, and so on. You can enter any character in a FortiGate configuration text string except, to prevent Cross-Site Scripting (XSS) vulnerabilities, the following characters:

“ (double quote), & (ampersand), ' (single quote), < (less than) and > (greater than)

Most web-based manager text string fields make it easy to add an acceptable number of characters and prevent you from adding the XSS vulnerability characters.



There is a different character limitation for VDOM names and hostnames. For both, the only legal characters are numbers (0-9), letters (a-z, A-Z), and special characters - and \_.

---

From the CLI, you can also use the `tree` command to view the number of characters that are allowed in a name field. For example, firewall address names can contain up to 64 characters. When you add a firewall address to the web-based manager, you are limited to entering 64

characters in the firewall address name field. From the CLI you can enter the following `tree` command to confirm that the firewall address `name` field allows 64 characters.

```
config firewall address
 tree
 -- [address] --*name (64)
 |- subnet
 |- type
 |- start-ip
 |- end-ip
 |- fqdn (256)
 |- cache-ttl (0,86400)
 |- wildcard
 |- comment (64 xss)
 |- associated-interface (16)
 +- color (0,32)
```

The `tree` command output also shows the number of characters allowed for other firewall address name settings. For example, the fully-qualified domain name (`fqdn`) field can contain up to 256 characters.

## Entering numeric values

Numeric values set various sizes, rates, numeric addresses, and other numeric values. For example, a static routing priority of 10, a port number of 8080, or an IP address of 10.10.10.1. Numeric values can be entered as a series of digits without spaces or commas (for example, 10 or 64400), in dotted decimal format (for example the IP address 10.10.10.1) or, as in the case of MAC or IPv6 addresses, separated by colons (for example, the MAC address 00:09:0F:B7:37:00). Most numeric values are standard base-10 numbers, but some fields (again, such as MAC addresses) require hexadecimal numbers.

Most web-based manager numeric value fields make it easy to add the acceptable number of digits within the allowed range. CLI help includes information about allowed numeric value ranges. Both the web-based manager and the CLI prevent you from entering invalid numbers.

## Enabling or disabling options

If a configuration option can only be on or off (enabled or disabled), the web-based manager presents a check box or other control that can only be enabled or disabled. From the CLI you can set the option to `enable` or `disable`.

## Dashboard

The various dashboard menus provides a way to access information about network activity and events, as well as configure basic system settings. You can easily add more dashboards and edit existing ones to easily view the content you need.

Each information “chunk” is found within a widget. Widgets provide an easy and quick way to view a variety of information, such as statistical information or network activity. There are a selection of widgets to choose from by selecting the *Widgets* option.

Administrators must have read and write privileges for adding and configuring dashboards and widgets.

## Adding dashboards and widgets

Dashboards that you create are automatically added under the default status and usage dashboards. You can add, remove, or rename a dashboard, regardless of whether it is default. You can also reset the Dashboard menu to its default settings by selecting *Reset Dashboards*.

If VDOMs are enabled, only the dashboards within Global are available for configuration.

### To add a dashboard

1. Go to *System > Dashboard > Status*.
2. Select *Dashboard*, located at the top left of the page.
3. Select *Add Dashboard*.

To add a widget to a dashboard, select *Widget* located at the top left of the dashboard page.

## System Information widget

The *System Information* widget shows status information on the FortiGate unit and provides the access point to update the firmware and backup the configurations.

<b>Host Name</b>	The name of the FortiGate unit. For details on changing the name, see <a href="#">Changing the FortiGate unit's host name</a> .  If the FortiGate unit is in HA mode, this information is not displayed.
<b>Serial Number</b>	The serial number of the FortiGate unit. The serial number is specific to that FortiGate unit and does not change with firmware upgrades.
<b>Operation Mode</b>	The current operating mode of the FortiGate unit. A FortiGate unit can operate in NAT mode or Transparent mode. Select <i>Change</i> to switch between NAT and transparent mode. For more information, see <a href="#">Changing the operation mode</a> .  If virtual domains are enabled, this field shows the operating mode of the current virtual domain. The Global System Status dashboard does not include this information.
<b>HA Status</b>	The status of High Availability (HA) within the cluster. Standalone indicates the FortiGate unit is not operating in HA mode. Active-Passive or Active-Active indicate the FortiGate unit is operating in HA mode. Select <i>Configure</i> , to change the HA configuration.
<b>Cluster Name</b>	The name of the HA cluster for this FortiGate unit. The FortiGate unit must be operating in HA mode to display this field.
<b>Cluster Members</b>	The FortiGate units in the HA cluster. Information displayed about each member includes host name, serial number, and whether the FortiGate unit is a primary (master) or subordinate (slave) FortiGate unit in the cluster.  The FortiGate unit must be operating in HA mode with virtual domains disabled to display this information.
<b>Virtual Cluster 1</b>	The role of each FortiGate unit in virtual cluster 1 and virtual cluster 2.
<b>Virtual Cluster 2</b>	The FortiGate unit must be operating in HA mode with virtual domains enabled to display this information.

<b>System Time</b>	The current date and time. Select <i>Change</i> , to configure the system time. For more information, see <a href="#">Configuring system time</a> .
<b>Firmware Version</b>	The version of the current firmware installed on the FortiGate unit. Select <i>Update</i> to upload a different firmware version. For more information, see <a href="#">Changing the firmware</a> .
<b>System Configuration</b>	The time period of when the configuration file was backed up. Select <i>Backup</i> to back up the current configuration. For more information, see <a href="#">Backing up the configuration</a> .  To restore a configuration file, select <i>Restore</i> . For more information, see <a href="#">Restoring your firmware configuration</a> .
<b>Current Administrator</b>	The number of administrators currently logged into the FortiGate unit.  Select <i>Details</i> to view more information about each administrator that is currently logged in  If you want to changed the current administrator's password, see <a href="#">Changing the currently logged in administrator's password</a> .
<b>Uptime</b>	The time in days, hours, and minutes since the FortiGate unit was started or rebooted.
<b>Virtual Domain</b>	Status of virtual domains on your FortiGate unit. Select <i>Enable</i> or <i>Disable</i> to change the status of virtual domains feature.  If you enable or disable virtual domains, your session will be terminated and you will need to log in again.
<b>Explicit Proxy Load Balance</b>	The status of each feature. Select <i>Enable</i> or <i>Disable</i> to change the status of the feature. When enabled, the menu option appears.

## Changing the FortiGate unit's host name

The host name appears in the *Host Name* row, in the *System Information* widget. The host name also appears at the CLI prompt when you are logged in to the CLI and as the SNMP system name.

To change the host name on the FortiGate unit, in the *System Information* widget, select *Change* in the *Host Name* row. The only administrators that can change a FortiGate unit's host name are administrators whose admin profiles permit system configuration write access. If the FortiGate unit is part of an HA cluster, you should use a unique host name to distinguish the FortiGate unit from others in the cluster.

## Changing the operation mode

FortiGate units and individual VDOMs can operate in NAT or Transparent mode. From the *System Information* dashboard widget, you can change the operating mode for your FortiGate unit or for a VDOM and perform sufficient network configuration to ensure that you can connect to the web-based manager in the new mode.

### NAT mode

In NAT mode, the FortiGate unit is visible to the network that it is connected to and all of its interfaces are on different subnets. Each interface that is connected to a network must be configured with an IP address that is valid for that subnet.

You would typically use NAT mode when the FortiGate unit is deployed as a gateway between private and public networks (or between any networks). In its default NAT mode configuration,



the FortiGate unit functions as a router, routing traffic between its interfaces. Security policies control communications through the FortiGate unit to both the Internet and between internal networks. In NAT mode, the FortiGate unit performs network address translation before IP packets are sent to the destination network.

For example, a company has a FortiGate unit as their interface to the Internet. The FortiGate unit also acts as a router to multiple subnets within the company. In this situation, the FortiGate unit is set to NAT mode and has a designated port for the Internet, wan1, with an address of 172.20.120.129, which is the public IP address. The internal network segments are behind the FortiGate unit and invisible to the public access, for example port 2 has an address of 10.10.10.1. The FortiGate unit translates IP addresses passing through it to route the traffic to the correct subnet or to the Internet.

### **Transparent Mode**

In transparent mode, the FortiGate unit is invisible to the network. All of its interfaces are on the same subnet and share the same IP address. To connect the FortiGate unit to your network, all you have to do is configure a management IP address and a default route.

You would typically use the FortiGate unit in transparent mode on a private network behind an existing firewall or behind a router. In transparent mode, the FortiGate unit also functions as a firewall. Security policies control communications through the FortiGate unit to the Internet and internal network. No traffic can pass through the FortiGate unit until you add security policies.

For example, the company has a router or other firewall in place. The network is simple enough that all users are on the same internal network. They need the FortiGate unit to perform application control, antivirus, intrusion protection, and similar traffic scanning. In this situation, the FortiGate unit is set to transparent mode. The traffic passing through the FortiGate unit does not change the addressing from the router to the internal network. Security policies and security profiles define the type of scanning the FortiGate unit performs on traffic entering the network.

#### **To switch from NAT to transparent mode**

1. From the *System Information* dashboard widget select *Change* beside *Operation Mode*.
2. From the *Operation Mode* list, select *Transparent*.
3. Enter the *Management IP* address and *Netmask*. This is the IP address to connect to when configuring and maintaining the device.
4. Enter the *Default Gateway*.
5. Select *OK*.

#### **To change the transparent mode management IP address**

1. From the *System Information* dashboard widget select *Change* beside *Operation Mode*.
2. Enter a new IP address and netmask in the *Management IP/Network* field as required and select *OK*.

Your web browser is disconnected from the web-based manager. To reconnect to the web-based manager browse to the new management IP address.

#### **To switch from transparent to NAT mode**

1. From the *System Information* dashboard widget select *Change* beside *Operation Mode*.
2. From the *Operation Mode* list, select *NAT*.
3. Enter valid IP address and netmask for the network from which you want to manage the FortiGate unit.
4. Select the interface to which the *Interface IP/Netmask* settings apply
5. Enter the IP address default gateway required to reach other networks from the FortiGate unit.

6. After the FortiGate unit switches to NAT mode, you may need to go *Router > Static Route* and edit this default route.  
For low-end FortiGate units, go to *System > Network > Routing*.
7. Select *OK*.

## Configuring system time

The FortiGate unit's system time can be changed using the *System Information* widget by selecting *Change* in the *System Time* row.

<b>System Time</b>	The current system date and time on the FortiGate unit.
<b>Refresh</b>	Update the display of the FortiGate unit's current system date and time.
<b>Time Zone</b>	Select the current system time zone for the FortiGate unit.
<b>Set Time</b>	Select to set the system date and time to the values.
<b>Synchronize with NTP Server</b>	Select to use a Network Time Protocol (NTP) server to automatically set the system date and time. You must specify the server and synchronization interval.  FortiGate units use NTP Version 4. For more information about NTP see <a href="http://www.ntp.org">http://www.ntp.org</a> .
<b>Server</b>	Enter the IP address or domain name of an NTP server. To find an NTP server that you can use, see <a href="http://www.ntp.org">http://www.ntp.org</a> .
<b>Sync Interval</b>	Specify how often the FortiGate unit should synchronize its time with the NTP server.
<b>Enable NTP Server</b>	Select to make the FortiGate unit an NTP server that client computers can ping for time synchronization. When selected, the Listen on Interfaces option appears. Add the interfaces the FortiGate unit will listen for time requests.

Daylight savings time is enabled by default. You can disable daylight savings time using the CLI commands:

```
config system global
 set dst disable
end
```

## Changing the firmware



To avoid losing configuration settings you should always back up your configuration before changing the firmware image.

Also, when updating firmware, you should first refer to [Supported Upgrade Paths for FortiOS Firmware](#) to help ensure a successful upgrade.

Administrators whose admin profiles permit maintenance read and write access can change the FortiGate unit's firmware. Firmware images can be installed from a number of sources including a local hard disk, a local USB disk, or the FortiGuard Network.

To change the firmware, go to *System > Dashboard > Status > System Information* widget and select the *Update* link on the *Firmware Version* row.

---

<b>Upgrade From</b>	Select the firmware source from the drop down list of available sources.
<b>Firmware Version</b>	<p>This appears only when selecting <i>FortiGuard Network</i> is selected from the <i>Upgrade From</i> drop-down list. Select a firmware version from the drop-down list.</p> <p>If downgrading the firmware on the FortiGate unit, select the check box beside <i>Allow Firmware Downgrade</i>.</p>
<b>Upgrade File</b>	<p>Browse to the location of the firmware image on your local hard disk.</p> <p>This field is available for local hard disk and USB only.</p>
<b>Allow Firmware Downgrade</b>	<p>Select to confirm the installation of an older firmware image (downgrade).</p> <p>This appears only when selecting <i>FortiGuard Network</i> is selected from the <i>Upgrade From</i> drop-down list.</p>
<b>Upgrade Partition</b>	<p>The number of the partition being updated.</p> <p>This field is available only if your FortiGate unit has more than one firmware partition.</p>
<b>Boot the New Firmware</b>	<p>By default, this is enabled. Select to disable the FortiGate unit's reboot process when installing a firmware image to a partition.</p> <p>This option enables you to install a firmware image to a partition without the FortiGate unit rebooting itself and making the firmware image the default firmware that is currently running.</p>

---



You need to register your FortiGate unit with Customer Support to access firmware updates for your model. For more information, go to <http://support.fortinet.com> or contact Customer Support.

---

## Backing up the configuration

Administrators can back up the FortiGate unit's configuration file from the *System Information* widget. Select *Backup* in the *System Configuration* row, to back up the firmware configuration file to a local computer, USB disk or to a FortiManager unit.

You should always back up your configuration whenever you make any modifications to the device configuration or performing any firmware updates or changes.

---

<b>Local PC</b>	Select to back up the configuration file to a local management computer.
<b>FortiManager</b>	<p>Select to back up the configuration file to a FortiManager unit. The Central Management settings must be enabled and a FortiManager unit connected with the FortiGate unit so that the FortiGate unit can send the configuration file to the FortiManager unit.</p> <p>To enable central management, go to <i>System &gt; Admin &gt; Settings</i>.</p>

---

<b>USB Disk</b>	Select to back up the configuration file to a USB key that is connected to the FortiGate unit.
<b>Full Config</b>	Select to backup the full VDOM configuration. This appears only when the FortiGate unit has VDOM configuration enabled.
<b>VDOM Config</b>	Select to backup the only the VDOM configuration file. This option backs up only the configuration file within that VDOM. Select the VDOM from the drop-down list, and select <i>Backup</i> .
<b>Encrypt configuration file</b>	Select to enable a password to the configuration file for added security.
<b>Password</b>	Enter the password that will be used to restore the configuration file.
<b>Confirm</b>	Re-enter the password.

## Formatting USB

The FortiGate unit enables you to back up the configuration of the device to a USB flash drive. The USB flash drive must be formatted as a FAT16 disk.

To format the USB flash drive, either use the CLI command `exe usb-disk format.` or within Windows at a command prompt, enter the command

```
format <drive_letter>: /FS:FAT /V:<drive_label>
```

where *<drive\_letter>* is the letter of the connected USB flash drive and *<drive\_label>* is the name to give the USB drive.

## Remote FortiManager backup and restore options

After successfully connecting to the FortiManager unit from your FortiGate unit, you can back up and restore your configuration to and from the FortiManager unit.

A list of revisions is displayed when restoring the configuration from a remote location. The list allows you to choose the configuration to restore. To use the FortiManager unit as a method of backup and restore of configuration files, you must first configure a connection between the two devices. For more information, see [Central management](#).

## Remote FortiGuard backup and restore options

Your FortiGate unit can be remotely managed by a central management server that is available when you register for the FortiGuard Analysis and Management Service. FortiGuard Analysis and Management Service is a subscription-based service and is purchased by contacting support.

After registering, you can back up or restore your configuration. FortiGuard Analysis and Management Service is useful when administering multiple FortiGate units without having a FortiManager unit. Using this service, you can also upgrade the firmware. Upgrading the firmware is available in the *Firmware Upgrade* section of the backup and restore menu.

When restoring the configuration from a remote location, a list of revisions is displayed so that you can choose the configuration file to restore.



The FortiGuard-FortiManager protocol is used when connecting to the FortiGuard Analysis and Management Service. This protocol runs over SSL using IPv4/TCP port 541 and includes the following functions:

- detects FortiGate unit dead or alive status
- detects management service dead or alive status
- notifies the FortiGate units about configuration changes, AV/IPS database update and firewall changes.

## Restoring your firmware configuration

Administrators can restore a configuration file that was backed up using the *System Information* widget. If the configuration file was encrypted, you will need the password to restore the configuration file.

<b>Local PC</b>	Select to back up the configuration file to a local management computer.
<b>FortiManager</b>	Select to back up the configuration file to a FortiManager unit. The Central Management settings must be enabled and a FortiManager unit connected with the FortiGate unit so that the FortiGate unit can send the configuration file to the FortiManager unit.  To enable central management, go to <i>System &gt; Admin &gt; Settings</i> .
<b>USB Disk</b>	Select to back up the configuration file to a USB key that is connected to the FortiGate unit.
<b>Filename</b>	Select Browse to locate the configuration file.
<b>Password</b>	If a password was set when saving the configuration file, enter the password.

## Viewing online administrators

The *System Information* widget enables you to view information about the administrators logged into the FortiGate unit. To view logged in administrators, in the *System Information* widget, select *Details*.

## Changing the currently logged in administrator's password

Use the *System Information* widget, to change your password. To do this, select the *Change Password* option in the *Current Administrator* row.

## License Information widget

The *License Information* widget displays the status of your technical support contract and FortiGuard subscriptions. The FortiGate unit updates the license information status indicators automatically when attempting to connect to the FortiGuard Distribution Network (FDN). FortiGuard Subscriptions status indicators are green if the FDN was reachable and the license was valid during the last connection attempt, grey if the FortiGate unit cannot connect to the FDN, and orange if the FDN is reachable but the license has expired.

When a new FortiGate unit is powered on, it automatically searches for FortiGuard services. If the FortiGate unit is configured for central management, it will look for FortiGuard services on the configured FortiManager system. The FortiGate unit sends its serial number to the FortiGuard service provider, which then determines whether the FortiGate unit is registered and has valid contracts for FortiGuard subscriptions and FortiCare support services. If the FortiGate unit is registered and has a valid contract, the License Information is updated.

If the FortiGate unit is not registered, any administrator with the `super_admin` profile sees a reminder message that provides access to a registration form.

When a contract is due to expire within 30 days, any administrator with the `super_admin` profile sees a notification message that provides access to an Add Contract form. Simply enter the new contract number and select *Add*. Fortinet Support also sends contract expiry reminders.

You can optionally disable notification for registration or contract inquiry using the `config system global` command in the CLI. Selecting any of the *Configure* options will take you to the Maintenance page.

---

<b>Support Contract</b>	Displays details about your current Fortinet Support contract. <ul style="list-style-type: none"><li>• If <i>Not Registered</i> appears, select <i>Register</i> to register the FortiGate unit.</li><li>• If <i>Expired</i> appears, select <i>Renew</i> for information on renewing your technical support contract. Contact your local reseller.</li><li>• If <i>Registered</i> appears, the name of the support that registered this FortiGate unit is also displayed. The various types of contracts that you currently have and the expiry date for each type.</li><li>• You can select <i>Login Now</i> to log into the Fortinet Support account that registered this FortiGate unit.</li></ul>
<b>FortiGuard Services</b>	Displays your current licenses for services from FortiGuard. Select <i>Renew</i> to update any of the licenses.
<b>FortiCloud</b>	Displays details about your current FortiCloud subscription. If the green <i>Activate</i> button appears, select it to either create a new account or add the FortiGate unit to an existing account. <p>If you have already activated FortiCloud, the name of the <i>Account</i> will be listed. Select <i>Launch Portal</i> to view your FortiCloud account in a web browser.</p> <p>Information on the current <i>Type</i> and <i>Storage</i> is also listed. You can select <i>Upgrade</i> to change the type of your FortiCloud account.</p>
<b>FortiClient Software</b>	Displays FortiClient license details and the number of <i>Register</i> and <i>Allowed</i> FortiClient users. You can select <i>Details</i> for more information about the current FortiClient users.
<b>FortiToken Mobile</b>	Displays the number of <i>Assigned</i> and <i>Allowed</i> FortiTokens.
<b>SMS</b>	Displays the number of <i>Sent</i> and <i>Allowed</i> SMS messages. You can select <i>Add Messages</i> to configure a new SMS message.
<b>Virtual Domain</b>	Displays the maximum number of virtual domains the FortiGate unit supports with the current license. <p>For high-end models, you can select the <i>Purchase More</i> link to purchase a license key through Fortinet technical support to increase the maximum number of VDOMs.</p>

---

## FortiGate unit Operation widget

The *Unit Operation* widget is an illustrated version of the FortiGate unit's front panel that shows the status of the FortiGate unit's network interfaces. Interfaces appear green when connected. Hover the mouse pointer over an interface to view further details.

Icons around the front panel indicate when the FortiGate unit is connected to a FortiAnalyzer or FortiManager device, or FortiClient installations. Select the icon in the widget to jump to the configuration page for each device. When connected to one of these devices, a green check mark icon appears next to the icon. If the device communication is configured but the device is unreachable, a red X appears.

## System Resources widget

The *System Resources* widget displays basic FortiGate unit resource usage. This widget displays the information for CPU and memory in either real-time or historical data. For FortiGate units with multiple CPUs, you can view the CPU usage as an average of all CPUs or each one individually.

This widget also is where you reboot or shutdown the FortiGate unit.



The options to reboot or shutdown the FortiGate unit are not available for an admin using the *prof\_admin* profile.

---

Use the *Refresh* icon when you want to view current system resource information, regardless of whether you are viewing real-time or historical type format.

To change the resource view from real-time to historical, or change the CPU view (for multiple CPU FortiGate units), select the *Edit* icon (visible when you hover the mouse over the widget).

When viewing CPU and memory usage in the web-based manager, only the information for core processes displays. CPU for management processes, is excluded. For example, HTTPS connections to the web-based manager.

## Alert Message Console widget

The *Alert Messages Console* widget helps you monitor system events on your FortiGate unit such as firmware changes, network security events, or virus detection events. Each message shows the date and time that the event occurred.

You can configure the alert message console settings to control what types of messages are displayed on the console.

### To configure the Alert Message Console

1. Locate the *Alert Message Console* widget within the Dashboard menu.
2. Select the *Edit* icon in the *Alert Message Console* title bar.
3. Select the types of alerts that you do not want to be displayed in the widget.
4. Select *OK*.

## CLI Console widget

The *CLI Console* widget enables you to access the CLI without exiting from the web-based manager.

The two controls located on the CLI Console widget title bar are *Customize*, and *Detach*.

- *Detach* moves the CLI Console widget into a pop-up window that you can resize and reposition. Select *Attach*. to move the widget back to the dashboard's page.
- *Customize* enables you to change the appearance of the console by selecting fonts and colors for the text and background.

## Session History widget

The *Session History* widget displays the total session activity on the device. Activity displays on a per second basis. Select the *Edit* icon in the title bar (which appears when you hover the mouse over the widget) to change the time period for the widget.

## Top Sessions widget

The *Top Sessions* widget polls the FortiGate unit for session information for IPv4 or IPv6 addresses, or both. Rebooting the FortiGate unit will reset the Top Session statistics to zero.

When you select *Details* to view the current sessions list, a list of all sessions currently processed by the FortiGate unit.

Detailed information is available in *System > Monitor > Sessions*. Use the following table to modify the default settings of the Top Sessions widget.

## USB Modem widget

The *USB modem* widget enables you to monitor the status of your USB modem, and configure it as needed.

## Advanced Threat Protection Statistics widget

The *Advanced Threat Protection Statistics* widget displays a count of detected malware and files scanned for these types of intrusions. It also displays statics on the number of files sent to FortiSandbox and the results from sandboxing.

## Features widget

The *Features* widget displays a number of *Basic Features* and *Security Features* and whether or not each feature is currently enabled or disabled. Options for features that are disabled will not appear in the web-based manager.

For *Security Features*, several *Preset* options are available, which can be selected from the dropdown menu in the widget:

- *UTM* enabled all security features and should be chosen for networks that require full protection from FortiOS. *UTM* is the default setting.
- *WF* enables web filtering features.
- *ATP* enables protection against viruses and other external threats.
- *NGFW* enables application control and protection from external attacks.
- *NGFW + ATP* enables features that protect against external threats and attacks.



## RAID monitor widget

The *RAID Monitor* widget displays the current state of the RAID array and each RAID disk. This widget does not display unless the FortiGate unit has more than one disk installed and is not available for FortiOS Carrier.

---

<b>Array status icon</b>	<p>Displays the status of the RAID array.</p> <ul style="list-style-type: none"><li>• Green with a check mark shows a healthy RAID array.</li><li>• Yellow triangle shows the array is in a degraded state but it is still functioning. A degraded array is slower than a healthy array. Rebuild the array to fix the degraded state.</li><li>• A wrench shows the array is being rebuilt.</li></ul> <p>Positioning the mouse over the array status icon displays a text message of the status of the array.</p>
<b>Disk status icon</b>	<p>There is one icon for each disk in the array.</p> <ul style="list-style-type: none"><li>• Green with a check mark shows a healthy disk.</li><li>• Red with an X shows the disk has failed and needs attention.</li></ul> <p>Positioning the mouse over the disk status icon displays the status of the disk, and the storage capacity of the disk.</p>
<b>RAID Level</b>	<p>The RAID level of this RAID array. The RAID level is set as part of configuring the RAID array.</p>
<b>Status bar</b>	<p>The bar shows the percentage of the RAID array that is currently in use.</p>
<b>Used/Free/Total</b>	<p>Displays the amount of RAID array storage that is being used, the amount of storage that is free, and the total storage in the RAID array. The values are in gigabytes.</p>

---

## RAID disk configuration

The RAID disk is configured from the Disk Configuration page.

---

<b>RAID level</b>	<p>Select the level of RAID. Options include:</p> <ul style="list-style-type: none"><li>• <b>RAID-0</b> — (striping) better performance, no redundancy</li><li>• <b>RAID-1</b> — (mirroring) half the storage capacity, with redundancy</li><li>• <b>RAID-5</b> — striping with parity checking, and redundancy</li></ul> <p>Available RAID level options depend on the available number of hard disks. Two or more disks are required for RAID 0 or RAID 1. Three or more disks are required for RAID 5.</p> <p>Changing the RAID level will erase any stored log information on the array, and reboot the FortiGate unit. The FortiGate unit will remain offline while it reconfigures the RAID array. When it reboots, the array will need to synchronize before being fully operational.</p>
-------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

---

<b>Status</b>	<p>The status, or health, of RAID array. This status can be one of:</p> <ul style="list-style-type: none"> <li>• <b>OK</b> — standard status, everything is normal</li> <li>• <b>OK (Background-Synchronizing) (%)</b> — synchronizing the disks after changing RAID level, Synchronizing progress bar shows percent complete</li> <li>• <b>Degraded</b> — One or more of the disks in the array has failed, been removed, or is not working properly. A warning is displayed about the lack of redundancy in this state. Also, a degraded array is slower than a healthy array. Select <i>Rebuild RAID</i> to fix the array.</li> <li>• <b>Degraded (Background-Rebuilding) (%)</b> — The same as degraded, but the RAID array is being rebuilt in the background. The array continues to be in a fragile state until the rebuilding is completed.</li> </ul>
<b>Size</b>	<p>The size of the RAID array in gigabytes (GB). The size of the array depends on the RAID level selected, and the number of disks in the array.</p>
<b>Rebuild RAID</b>	<p>Select to rebuild the array after a new disk has been added to the array, or after a disk has been swapped in for a failed disk.</p> <p>If you try to rebuild a RAID array with too few disks you will get a rebuild error. After inserting a functioning disk, the rebuild will start.</p> <p>This button is only available when the RAID array is in a degraded state and has enough disks to be rebuilt.</p> <p>You cannot restart a rebuild once a rebuild is already in progress.</p> <p><b>Note:</b> If a disk has failed, the number of working disks may not be enough for the RAID level to function. In this case, replace the failed disk with a working disk to rebuild the RAID array.</p>
<b>Disk#</b>	<p>The disk's position in the array. This corresponds to the physical slot of the disk.</p> <p>If a disk is removed from the FortiGate unit, the disk is marked as not a member of the array and its position is retained until a new disk is inserted in that drive bay.</p>
<b>Status</b>	<p>The status of this disk. Options include OK, and unavailable.</p> <p>A disk is unavailable if it is removed or has failed.</p>

## Basic configurations

Before going ahead and configuring security policies, users, and security profiles, you should perform some basic configurations to set up your FortiGate unit.

### Changing your administrator password

By default, you can log in to the web-based manager by using the admin administrator account and no password. It is highly recommended that you add a password to the admin administrator account. For improved security, you should regularly change the admin administrator account password and the passwords for any other administrator accounts that you add.

To change an administrator's password, go to *System > Admin > Administrators*, edit the administrator account, and then change the password.

For details on selecting a password, and password best practices, see “[Passwords](#)” on [page 1429](#).

For information about resetting a lost administrator's password, see [docs.fortinet.com/sysadmin.html](https://docs.fortinet.com/sysadmin.html).

## Changing the web-based manager language

The default language of the web-based manager is English. To change the language, go to *System > Admin > Settings*. In the *Display Settings* section, select the language you want from the *Language* drop-down list.

For best results, you should select the language that the management computer operating system uses.

## Changing administrative access

Through administrative access, an administrator can connect to the FortiGate unit. Access is available through a number of services, including HTTPS and SSH. The default configuration allows administrative access to one or more of the unit's interfaces as described in the [QuickStart Guide](#).

### To change administrative access

1. Go to *System > Network > Interface*.
2. Select the interface.
3. Select the administrative access type or types for that interface.
4. Select *OK*.

## Changing the web-based manager idle timeout

By default, the web-based manager disconnects administrative sessions if no activity occurs for five minutes. This prevents someone from using the web-based manager if the management PC is left unattended.

### To change the idle timeout

1. Go to *System > Admin > Settings*.
2. In the *Administration Settings* section, enter the time in minutes in the *Idle Timeout* field.
3. Select *Apply*.

## Switching VDOMs

When VDOMs are enabled, a menu appears in the left column called *Current VDOM*. This menu displays a drop-down list that lists the configured VDOMs.

To switch to a VDOM using the *Current VDOM* menu, select the VDOM that you want to switch to from the drop-down list. You are automatically redirected to that VDOM.

VDOMs are enabled on the *System Information* Dashboard Widget.

## Connecting to the CLI from the web-based manager

You can use the CLI to configure all configuration options available from the web-based manager. Some configuration options are available only from the CLI.

To connect to the CLI console, go to *System > Dashboard > Status* and select inside the window of the *CLI Console* widget to automatically connect. For more information on using the CLI, see [“Using the CLI” on page 1405](#).

## Logging out

Select the Logout icon to quit your administrative session. If you only close the browser or leave the web-based manager to surf to another web site, you remain logged in until the idle timeout (default 5 minutes) expires. To change the timeout, see [“Changing the web-based manager idle timeout” on page 1403](#).

# Using the CLI

The command line interface (CLI) is an alternative configuration tool to the web-based manager. While the configuration of the web-based manager uses a point-and-click method, the CLI requires typing commands or uploading batches of commands from a text file, like a configuration script.

This section also explains common CLI tasks that an administrator does on a regular basis and includes the topics:

- [Connecting to the CLI](#)
- [Command syntax](#)
- [Sub-commands](#)
- [Permissions](#)
- [Tips](#)

## Connecting to the CLI

You can access the CLI in two ways:

- **Locally** — Connect your computer directly to the FortiGate unit's console port. Local access is required in some cases:
  - If you are installing your FortiGate unit for the first time and it is not yet configured to connect to your network, unless you reconfigure your computer's network settings for a peer connection, you may only be able to connect to the CLI using a local serial console connection. For more information, see [“Connecting to the CLI” on page 1427](#).
  - Restoring the firmware utilizes a boot interrupt. Network access to the CLI is not available until after the boot process has completed, making local CLI access the only viable option.
- **Through the network** — Connect your computer through any network attached to one of the FortiGate unit's network ports. The network interface must have enabled Telnet or SSH administrative access if you will connect using an SSH/Telnet client, or HTTP/HTTPS administrative access if you will connect using the *CLI Console* widget in the web-based manager.

### Connecting to the CLI using a local console

Local console connections to the CLI are formed by directly connecting your management computer or console to the FortiGate unit, using its DB-9 or RJ-45 console port. To connect to the local console you need:

- a computer with an available serial communications (COM) port
- the RJ-45-to-DB-9 or null modem cable included in your FortiGate package
- terminal emulation software such as HyperTerminal for Microsoft Windows

The following procedure describes connection using Microsoft HyperTerminal software; steps may vary with other terminal emulators.

### To connect to the CLI using a local serial console connection

1. Using the null modem or RJ-45-to-DB-9 cable, connect the FortiGate unit's console port to the serial communications (COM) port on your management computer.
2. On your management computer, start HyperTerminal.
3. For the *Connection Description*, enter a *Name* for the connection, and select *OK*.
4. On the *Connect using* drop-down list box, select the communications (COM) port on your management computer you are using to connect to the FortiGate unit.
5. Select *OK*.
6. Select the following *Port* settings and select *OK*.

<b>Bits per second</b>	9600
<b>Data bits</b>	8
<b>Parity</b>	None
<b>Stop bits</b>	1
<b>Flow control</b>	None

7. Press Enter or Return on your keyboard to connect to the CLI.
8. Type a valid administrator account name (such as `admin`) and press Enter.
9. Type the password for that administrator account and press Enter. (In its default state, there is no password for the `admin` account.)

The CLI displays the following text:

```
Welcome!
```

```
Type ? to list available commands.
```

You can now enter CLI commands, including configuring access to the CLI through SSH or Telnet. For details, see [“Enabling access to the CLI through the network \(SSH or Telnet\)” on page 1406](#).

## Enabling access to the CLI through the network (SSH or Telnet)

SSH or Telnet access to the CLI is accomplished by connecting your computer to the FortiGate unit using one of its RJ-45 network ports. You can either connect directly, using a peer connection between the two, or through any intermediary network.



If you do not want to use an SSH/Telnet client and you have access to the web-based manager, you can alternatively access the CLI through the network using the *CLI Console* widget in the web-based manager.

You must enable SSH and/or Telnet on the network interface associated with that physical network port. If your computer is not connected directly or through a switch, you must also configure the FortiGate unit with a static route to a router that can forward packets from the FortiGate unit to your computer. You can do this using either a local console connection or the web-based manager.

## Requirements

- a computer with an available serial communications (COM) port and RJ-45 port
- terminal emulation software such as HyperTerminal for Microsoft Windows
- the RJ-45-to-DB-9 or null modem cable included in your FortiGate package
- a network cable
- prior configuration of the operating mode, network interface, and static route (for details, see)

## To enable SSH or Telnet access to the CLI using a local console connection

1. Using the network cable, connect the FortiGate unit's network port either directly to your computer's network port, or to a network through which your computer can reach the FortiGate unit.
2. Note the number of the physical network port.
3. Using a local console connection, connect and log into the CLI. For details, see [“Connecting to the CLI using a local console” on page 1405](#).

4. Enter the following command:

```
config system interface
 edit <interface_str>
 set allowaccess <protocols_list>
 next
end
```

where:

- <interface\_str> is the name of the network interface associated with the physical network port and containing its number, such as port1
- <protocols\_list> is the complete, space-delimited list of permitted administrative access protocols, such as https ssh telnet

For example, to exclude HTTP, HTTPS, SNMP, and PING, and allow only SSH and Telnet administrative access on port1:

```
set system interface port1 config allowaccess ssh telnet
```

5. To confirm the configuration, enter the command to display the network interface's settings.

```
get system interface <interface_str>
```

The CLI displays the settings, including the allowed administrative access protocols, for the network interfaces.

To connect to the CLI through the network interface, see [“Connecting to the CLI using SSH” on page 1407](#) or [“Connecting to the CLI using Telnet” on page 1408](#).

## Connecting to the CLI using SSH

Once the FortiGate unit is configured to accept SSH connections, you can use an SSH client on your management computer to connect to the CLI.

Secure Shell (SSH) provides both secure authentication and secure communications to the CLI. FortiGate units support 3DES and Blowfish encryption algorithms for SSH.

Before you can connect to the CLI using SSH, you must first configure a network interface to accept SSH connections. For details, see [“Enabling access to the CLI through the network \(SSH or Telnet\)” on page 1406](#). The following procedure uses PuTTY. Steps may vary with other SSH clients.

### To connect to the CLI using SSH

1. On your management computer, start an SSH client.
2. In *Host Name (or IP Address)*, enter the IP address of a network interface on which you have enabled SSH administrative access.
3. In *Port*, enter 22.
4. For the *Connection type*, select *SSH*.
5. Select *Open*.

The SSH client connects to the FortiGate unit.

The SSH client may display a warning if this is the first time you are connecting to the FortiGate unit and its SSH key is not yet recognized by your SSH client, or if you have previously connected to the FortiGate unit but used a different IP address or SSH key. This is normal if your management computer is directly connected to the FortiGate unit with no network hosts between them.

6. Click *Yes* to verify the fingerprint and accept the FortiGate unit's SSH key. You will not be able to log in until you have accepted the key.
7. The CLI displays a login prompt.
8. Type a valid administrator account name (such as `admin`) and press *Enter*.
9. Type the password for this administrator account and press *Enter*.

The FortiGate unit displays a command prompt (its host name followed by a #) . You can now enter CLI commands.



If three incorrect login or password attempts occur in a row, you will be disconnected. If this occurs, wait one minute, then reconnect to attempt the login again.

---

### Connecting to the CLI using Telnet

Once the FortiGate unit is configured to accept Telnet connections, you can use a Telnet client on your management computer to connect to the CLI.



Telnet is not a secure access method. SSH should be used to access the CLI from the Internet or any other untrusted network.

---

Before you can connect to the CLI using Telnet, you must first configure a network interface to accept SSH connections. For details, see [“Enabling access to the CLI through the network \(SSH or Telnet\)” on page 1406](#).

### To connect to the CLI using Telnet

1. On your management computer, start a Telnet client.
2. Connect to a FortiGate network interface on which you have enabled Telnet.
3. Type a valid administrator account name (such as `admin`) and press *Enter*.



4. Type the password for this administrator account and press Enter.

The FortiGate unit displays a command prompt (its host name followed by a #) . You can now enter CLI commands.



If three incorrect login or password attempts occur in a row, you will be disconnected. If this occurs, wait one minute, then reconnect to attempt the login again.

## Command syntax

When entering a command, the command line interface (CLI) requires that you use valid syntax and conform to expected input constraints. It will reject invalid commands.

Fortinet documentation uses the following conventions to describe valid command syntax


### Terminology

Each command line consists of a command word that is usually followed by words for the configuration data or other specific item that the command uses or affects:

```
get system admin
```

To describe the function of each word in the command line, especially if that nature has changed between firmware versions, Fortinet uses terms with the following definitions.

**Figure 225:**Command syntax terminology



- **command** — A word that begins the command line and indicates an action that the FortiGate unit should perform on a part of the configuration or host on the network, such as `config` or `execute`. Together with other words, such as fields or values, that end when you press the Enter key, it forms a command line. Exceptions include multiline command lines, which can be entered using an escape sequence. (See [“Shortcuts and key commands” on page 1416.](#))  
Valid command lines must be unambiguous if abbreviated. (See [“Command abbreviation” on page 1417.](#)) Optional words or other command line permutations are indicated by syntax notation. (See [“Notation” on page 1410.](#))
- **sub-command** — A kind of command that is available only when nested within the scope of another command. After entering a command, its applicable sub-commands are available to you until you exit the scope of the command, or until you descend an additional level into

another sub-command. Indentation is used to indicate levels of nested commands. (See “Indentation” on page 1410.)

Not all top-level commands have sub-commands. Available sub-commands vary by their containing scope. (See “Sub-commands” on page 1412.)

- **object** — A part of the configuration that contains tables and/or fields. Valid command lines must be specific enough to indicate an individual object.
- **table** — A set of fields that is one of possibly multiple similar sets which each have a name or number, such as an administrator account, policy, or network interface. These named or numbered sets are sometimes referenced by other parts of the configuration that use them. (See “Notation” on page 1410.)
- **field** — The name of a setting, such as `ip` or `hostname`. Fields in some tables must be configured with values. Failure to configure a required field will result in an invalid object configuration error message, and the FortiGate unit will discard the invalid table.
- **value** — A number, letter, IP address, or other type of input that is usually your configuration setting held by a field. Some commands, however, require multiple input values which may not be named but are simply entered in sequential order in the same command line. Valid input types are indicated by constraint notation. (See “Notation” on page 1410.)
- **option** — A kind of value that must be one or more words from of a fixed set of options. (See “Notation” on page 1410.)

## Indentation

Indentation indicates levels of nested commands, which indicate what other subcommittees are available from within the scope. For example, the `edit` sub-command is available only within a command that affects tables, and the `next` sub-command is available only from within the `edit` sub-command:

```
config system interface
 edit port1
 set status up
 next
end
```

For information about available sub-commands, see “Sub-commands” on page 1412.

## Notation

Brackets, braces, and pipes are used to denote valid permutations of the syntax. Constraint notations, such as `<address_ipv4>`, indicate which data types or string patterns are acceptable value input.

**Table 60:** Command syntax notation

Convention	Description
<b>Square brackets</b> [ ]	A non-required word or series of words. For example: <code>[verbose {1   2   3}]</code> indicates that you may either omit or type both the <code>verbose</code> word and its accompanying option, such as <code>verbose 3</code> .

**Table 60:** Command syntax notation

<p><b>Angle brackets</b> &lt; &gt;</p>	<p>A word constrained by data type. The angled brackets contain a descriptive name followed by an underscore ( _ ) and suffix that indicates the valid data type. For example, &lt;retries_int&gt;, indicates that you should enter a number of retries, such as 5.</p> <p>Data types include:</p> <ul style="list-style-type: none"> <li>• &lt;xxx_name&gt;: A name referring to another part of the configuration, such as policy_A.</li> <li>• &lt;xxx_index&gt;: An index number referring to another part of the configuration, such as 0 for the first static route.</li> <li>• &lt;xxx_pattern&gt;: A regular expression or word with wild cards that matches possible variations, such as *@example.com to match all email addresses ending in @example.com.</li> <li>• &lt;xxx_fqdn&gt;: A fully qualified domain name (FQDN), such as mail.example.com.</li> <li>• &lt;xxx_email&gt;: An email address, such as admin@example.com.</li> <li>• &lt;xxx_ipv4&gt;: An IPv4 address, such as 192.168.1.99.</li> <li>• &lt;xxx_v4mask&gt;: A dotted decimal IPv4 netmask, such as 255.255.255.0.</li> <li>• &lt;xxx_ipv4mask&gt;: A dotted decimal IPv4 address and netmask separated by a space, such as 192.168.1.99 255.255.255.0.</li> <li>• &lt;xxx_ipv4/mask&gt;: A dotted decimal IPv4 address and CIDR-notation netmask separated by a slash, such as 192.168.1.1/24.</li> <li>• &lt;xxx_ipv4range&gt;: A hyphen ( - )-delimited inclusive range of IPv4 addresses, such as 192.168.1.1-192.168.1.255.</li> <li>• &lt;xxx_ipv6&gt;: A colon ( : )-delimited hexadecimal IPv6 address, such as 3f2e:6a8b:78a3:0d82:1725:6a2f:0370:6234.</li> <li>• &lt;xxx_v6mask&gt;: An IPv6 netmask, such as /96.</li> <li>• &lt;xxx_ipv6mask&gt;: A dotted decimal IPv6 address and netmask separated by a space.</li> <li>• &lt;xxx_str&gt;: A string of characters that is not another data type, such as P@ssw0rd. Strings containing spaces or special characters must be surrounded in quotes or use escape sequences. See <a href="#">“Special characters” on page 1418</a>.</li> <li>• &lt;xxx_int&gt;: An integer number that is <b>not</b> another data type, such as 15 for the number of minutes.</li> </ul>
<p><b>Curly braces</b> { }</p>	<p>A word or series of words that is constrained to a set of options delimited by either vertical bars or spaces. You must enter at least one of the options, unless the set of options is surrounded by square brackets [ ].</p>

**Table 60:** Command syntax notation

<b>Options delimited by vertical bars  </b>	Mutually exclusive options. For example: <code>{enable   disable}</code> indicates that you must enter either <code>enable</code> or <code>disable</code> , but must not enter both.
<b>Options delimited by spaces</b>	Non-mutually exclusive options. For example: <code>{http https ping snmp ssh telnet}</code> indicates that you may enter all or a subset of those options, in any order, in a space-delimited list, such as: <code>ping https ssh</code>

## Sub-commands

Each command line consists of a command word that is usually followed by words for the configuration data or other specific item that the command uses or affects:

```
get system admin
```

Sub-commands are available from within the scope of some commands. When you enter a sub-command level, the command prompt changes to indicate the name of the current command scope. For example, after entering:

```
config system admin
```

the command prompt becomes:

```
(admin) #
```

Applicable sub-commands are available to you until you exit the scope of the command, or until you descend an additional level into another sub-command.

For example, the `edit` sub-command is available only within a command that affects tables; the `next` sub-command is available only from within the `edit` sub-command:

```
config system interface
 edit port1
 set status up
 next
end
```

Sub-command scope is indicated by indentation. See [“Indentation” on page 1410](#).

Available sub-commands vary by command. From a command prompt within `config`, two types of sub-commands might become available:

- commands affecting fields
- commands affecting tables

**Table 61:** Commands for tables

<b>clone &lt;table&gt;</b>	<p>Clone (or make a copy of) a table from the current object.</p> <p>For example, in <code>config firewall policy</code>, you could enter the following command to clone security policy 27 to create security policy 30:</p> <pre>clone 27 to 39</pre> <p>In <code>config antivirus profile</code>, you could enter the following command to clone an antivirus profile named <code>av_pro_1</code> to create a new antivirus profile named <code>av_pro_2</code>:</p> <pre>clone av_pro_1 to av_pro_2</pre> <p><code>clone</code> may not be available for all tables.</p>
<b>delete &lt;table&gt;</b>	<p>Remove a table from the current object.</p> <p>For example, in <code>config system admin</code>, you could delete an administrator account named <code>newadmin</code> by typing <code>delete newadmin</code> and pressing Enter. This deletes <code>newadmin</code> and all its fields, such as <code>newadmin</code>'s first-name and email-address.</p> <p><code>delete</code> is only available within objects containing tables.</p>
<b>edit &lt;table&gt;</b>	<p>Create or edit a table in the current object.</p> <p>For example, in <code>config system admin</code>:</p> <ul style="list-style-type: none"> <li>edit the settings for the default <code>admin</code> administrator account by typing <code>edit admin</code>.</li> <li>add a new administrator account with the name <code>newadmin</code> and edit <code>newadmin</code>'s settings by typing <code>edit newadmin</code>.</li> </ul> <p><code>edit</code> is an interactive sub-command: further sub-commands are available from within <code>edit</code>.</p> <p><code>edit</code> changes the prompt to reflect the table you are currently editing.</p> <p><code>edit</code> is only available within objects containing tables.</p> <p>In objects such as security policies, <code>&lt;table&gt;</code> is a sequence number. To create a new entry without the risk of overwriting an existing one, enter <code>edit 0</code>. The CLI initially confirms the creation of entry 0, but assigns the next unused number after you finish editing and enter <code>end</code>.</p>
<b>end</b>	<p>Save the changes to the current object and exit the <code>config</code> command. This returns you to the top-level command prompt.</p>
<b>get</b>	<p>List the configuration of the current object or table.</p> <ul style="list-style-type: none"> <li>In objects, <code>get</code> lists the table names (if present), or fields and their values.</li> <li>In a table, <code>get</code> lists the fields and their values.</li> </ul> <p>For more information on <code>get</code> commands, see the <a href="#">CLI Reference</a>.</p>

**Table 61:** Commands for tables

<b><i>purge</i></b>	<p>Remove all tables in the current object.</p> <p>For example, in <code>config forensic user</code>, you could type <code>get</code> to see the list of user names, then type <code>purge</code> and then <code>y</code> to confirm that you want to delete all users.</p> <p><code>purge</code> is only available for objects containing tables.</p> <p><b>Caution:</b> Back up the FortiGate unit before performing a <code>purge</code>. <code>purge</code> cannot be undone. To restore purged tables, the configuration must be restored from a backup.</p> <p><b>Caution:</b> Do not <code>purge system interface</code> or <code>system admin</code> tables. <code>purge</code> does not provide default tables. This can result in being unable to connect or log in, requiring the FortiGate unit to be formatted and restored.</p>
<b><i>rename &lt;table&gt; to &lt;table&gt;</i></b>	<p>Rename a table.</p> <p>For example, in <code>config system admin</code>, you could rename <code>admin3</code> to <code>fwadmin</code> by typing <code>rename admin3 to fwadmin</code>.</p> <p><code>rename</code> is only available within objects containing tables.</p>
<b><i>show</i></b>	<p>Display changes to the default configuration. Changes are listed in the form of configuration commands.</p>

## Example of table commands

From within the `system admin` object, you might enter:

```
edit admin_1
```

The CLI acknowledges the new table, and changes the command prompt to show that you are now within the `admin_1` table:

```
new entry 'admin_1' added
(admin_1)#
```

**Table 62:** Commands for fields

<b><i>abort</i></b>	Exit both the <code>edit</code> and/or <code>config</code> commands without saving the fields.
<b><i>append</i></b>	Add an option to an existing list.
<b><i>end</i></b>	Save the changes made to the current table or object fields, and exit the <code>config</code> command. (To exit without saving, use <code>abort</code> instead.)
<b><i>get</i></b>	<p>List the configuration of the current object or table.</p> <ul style="list-style-type: none"> <li>• In objects, <code>get</code> lists the table names (if present), or fields and their values.</li> <li>• In a table, <code>get</code> lists the fields and their values.</li> </ul>
<b><i>move</i></b>	Move an object within a list, when list order is important. For example, rearranging security policies within the policy list.

**Table 62:** Commands for fields

<b><i>next</i></b>	<p>Save the changes you have made in the current table's fields, and exit the <code>edit</code> command to the object prompt. (To save and exit completely to the root prompt, use <code>end</code> instead.)</p> <p><code>next</code> is useful when you want to create or edit several tables in the same object, without leaving and re-entering the <code>config</code> command each time.</p> <p><code>next</code> is only available from a table prompt; it is not available from an object prompt.</p>
<b><i>select</i></b>	<p>Clear all options except for those specified.</p> <p>For example, if a group contains members A, B, C, and D and you remove all users except for B, use the command <code>select member B</code>.</p>
<b><i>set &lt;field&gt;</i></b> <b><i>&lt;value&gt;</i></b>	<p>Set a field's value.</p> <p>For example, in <code>config system admin</code>, after typing <code>edit admin</code>, you could type <code>set password newpass</code> to change the password of the admin administrator to <code>newpass</code>.</p> <p><b>Note:</b> When using <code>set</code> to change a field containing a space-delimited list, type the whole new list. For example, <code>set &lt;field&gt; &lt;new-value&gt;</code> will replace the list with the <code>&lt;new-value&gt;</code> rather than appending <code>&lt;new-value&gt;</code> to the list.</p>
<b><i>show</i></b>	<p>Display changes to the default configuration. Changes are listed in the form of configuration commands.</p>
<b><i>unselect</i></b>	<p>Remove an option from an existing list.</p>
<b><i>unset &lt;field&gt;</i></b>	<p>Reset the table or object's fields to default values.</p> <p>For example, in <code>config system admin</code>, after typing <code>edit admin</code>, typing <code>unset password</code> resets the password of the admin administrator account to the default (in this case, no password).</p>

Example of field commands

From within the `admin_1` table, you might enter:

```
set password my1stExamplePassword
```

to assign the value `my1stExamplePassword` to the `password` field. You might then enter the `next` command to save the changes and edit the next administrator's table.

## Permissions

Depending on the account that you use to log in to the FortiGate unit, you may not have complete access to all CLI commands. Access profiles control which CLI commands an administrator account can access. Access profiles assign either read, write, or no access to each area of the FortiGate software. To view configurations, you must have read access. To make changes, you must have write access.

Unlike other administrator accounts, the administrator account named `admin` exists by default and cannot be deleted. The `admin` administrator account is similar to a root administrator account. This administrator account always has full permission to view and change all FortiGate configuration options, including viewing and changing **all** other administrator accounts. Its name and permissions cannot be changed. It is the only administrator account that can reset

another administrator's password without being required to enter that administrator's existing password.



Set a strong password for the admin administrator account, and change the password regularly. By default, this administrator account has no password. Failure to maintain the password of the admin administrator account could compromise the security of your FortiGate unit.

For complete access to all commands, you must log in with the administrator account named `admin`.

## Tips

Basic features and characteristics of the CLI environment provide support and ease of use for many CLI tasks.

## Help

To display brief help during command entry, press the question mark (?) key.

- Press the question mark (?) key at the command prompt to display a list of the commands available and a description of each command.
- Type a word or part of a word, then press the question mark (?) key to display a list of valid word completions or subsequent words, and to display a description of each.

## Shortcuts and key commands

**Table 63:** Shortcuts and key commands

Action	Keys
List valid word completions or subsequent words. If multiple words could complete your entry, display all possible completions with helpful descriptions of each.	?
Complete the word with the next available match. Press the key multiple times to cycle through available matches.	Tab
Recall the previous command. Command memory is limited to the current session.	Up arrow, or Ctrl + P
Recall the next command.	Down arrow, or Ctrl + N
Move the cursor left or right within the command line.	Left or Right arrow
Move the cursor to the beginning of the command line.	Ctrl + A
Move the cursor to the end of the command line.	Ctrl + E



**Table 63:** Shortcuts and key commands

Move the cursor backwards one word.	Ctrl + B
Move the cursor forwards one word.	Ctrl + F
Delete the current character.	Ctrl + D
Abort current interactive commands, such as when entering multiple lines. If you are not currently within an interactive command such as <code>config</code> or <code>edit</code> , this closes the CLI connection.	Ctrl + C
Continue typing a command on the next line for a multi-line command. For each line that you want to continue, terminate it with a backslash (\). To complete the command line, terminate it by pressing the spacebar and then the Enter key, without an immediately preceding backslash.	\ then Enter

## Command abbreviation

You can abbreviate words in the command line to their smallest number of non-ambiguous characters.

For example, the command `get system status` could be abbreviated to `g sy st`.

## Adding and removing options from lists

When adding options to a list, such as a user group, using the `set` command will remove the previous configuration. For example, if you wish to add user D to a user group that already contains members A, B, and C, the command would need to be `set member A B C D`. If only `set member D` was used, then all former members would be removed from the group.

However, there are additional commands which can be used instead of `set` for changing options in a list.

**Table 64:** Additional commands for lists

<b><i>append</i></b>	Add an option to an existing list. For example, <code>append member</code> would add user D to a user group while all previous group members are retained
<b><i>select</i></b>	Clear all options except for those specified. For example, if a group contains members A, B, C, and D and you remove all users except for B, use the command <code>select member B</code> .
<b><i>unselect</i></b>	Remove an option from an existing list. For example, <code>unselect member A</code> would remove member A from a group will all previous group members are retained.

## Environment variables

The CLI supports the following environment variables. Variable names are case-sensitive.

**Table 65:** Environment variables

<b>\$USERFROM</b>	The management access type ( <code>ssh</code> , <code>telnet</code> , <code>jsconsole</code> for the <i>CLI Console</i> widget in the web-based manager, and so on) and the IP address of the administrator that configured the item.
<b>\$USERNAME</b>	The account name of the administrator that configured the item.
<b>\$SerialNum</b>	The serial number of the FortiGate unit.

For example, the FortiGate unit's host name can be set to its serial number.

```
config system global
 set hostname $SerialNum
end
```

As another example, you could log in as `admin1`, then configure a restricted secondary administrator account for yourself named `admin2`, whose `first-name` is `admin1` to indicate that it is another of your accounts:

```
config system admin
 edit admin2
 set first-name $USERNAME
```

## Special characters

The characters `<`, `>`, `(`, `)`, `#`, `'`, and `"` are not permitted in most CLI fields. These characters are special characters, also known as reserved characters.

You may be able to enter special character as part of a string's value by using a special command, enclosing it in quotes, or preceding it with an escape sequence — in this case, a backslash (`\`) character.

In other cases, different keystrokes are required to input a special character. If you need to enter `?` as part of config, you first need to input CTRL-V. If you enter the question mark (`?`) without first using CTRL-V, the question mark has a different meaning in CLI: it will show available command options in that section.

For example, if you enter `?` without CTRL-V:

```
edit "*.xe
token line: Unmatched double quote.
```

If you enter `?` with CTRL-V:

```
edit "*.xe?"
new entry '*.xe?' added
```

**Table 66:** Entering special characters

Character	Keys
<code>?</code>	Ctrl + V then <code>?</code>
Tab	Ctrl + V then Tab

**Table 66:** Entering special characters

Space (to be interpreted as part of a string value, not to end the string)	Enclose the string in quotation marks: "Security Administrator".  Enclose the string in single quotes: 'Security Administrator'.  Precede the space with a backslash: Security\ Administrator.
' (to be interpreted as part of a string value, not to end the string)	\'
" (to be interpreted as part of a string value, not to end the string)	\"
\	\\

## Using grep to filter get and show command output

In many cases, the `get` and `show` (and `diagnose`) commands may produce a large amount of output. If you are looking for specific information in a large `get` or `show` command output, you can use the `grep` command to filter the output to only display what you are looking for. The `grep` command is based on the standard UNIX `grep`, used for searching text output based on regular expressions.

Use the following command to display the MAC address of the FortiGate unit internal interface:

```
get hardware nic internal | grep Current_HWaddr
Current_HWaddr 00:09:0f:cb:c2:75
```

Use the following command to display all TCP sessions in the session list and include the session list line number in the output

```
get system session list | grep -n tcp
```

Use the following command to display all lines in HTTP replacement message commands that contain URL (upper or lower case):

```
show system replacemsg http | grep -i url
```

There are three additional options that can be applied to `grep`:

```
-A <num> After
-B <num> Before
-C <num> Context
```

The option `-f` is also available to support Fortinet contextual output, in order to show the complete configuration. The following example shows the difference in output when `-f` option is used versus when it is not.

### Using -f:

```
show | grep -f ldap-group1
 config user group
 edit "ldap-group1"
 set member "pc40-LDAP"
 next
 end
 config firewall policy
 edit 2
 set srcintf "port31"
 set dstintf "port32"
 set srcaddr "all"
 set action accept
 set identity-based enable
 set nat enable
 config identity-based-policy
 edit 1
 set schedule "always"
 set groups "ldap-group1"
 set dstaddr "all"
 set service "ALL"
 next
 end
 next
 end
end
```

### Without using -f:

```
show | grep ldap-group1
 edit "ldap-group1"
 set groups "ldap-group1"
```

## Language support and regular expressions

Characters such as ñ, é, symbols, and ideographs are sometimes acceptable input. Support varies by the nature of the item being configured. CLI commands, objects, field names, and options must use their exact ASCII characters, but some items with arbitrary names or values may be input using your language of choice. To use other languages in those cases, you must use the correct encoding.

Input is stored using Unicode UTF-8 encoding but is not normalized from other encodings into UTF-8 before it is stored. If your input method encodes some characters differently than in UTF-8, your configured items may not display or operate as expected.

Regular expressions are especially impacted. Matching uses the UTF-8 character values. If you enter a regular expression using another encoding, or if an HTTP client sends a request in an encoding other than UTF-8, matches may not be what you expect.

For example, with Shift-JIS, backslashes ( \ ) could be inadvertently interpreted as the symbol for the Japanese yen ( ¥ ) and vice versa. A regular expression intended to match HTTP requests containing money values with a yen symbol therefore may not work if the symbol is entered using the wrong encoding.

For best results, you should:

- use UTF-8 encoding, or
- use only the characters whose numerically encoded values are the same in UTF-8, such as the US-ASCII characters that are also encoded using the same values in ISO 8859-1, Windows code page 1252, Shift-JIS and other encodings, or
- for regular expressions that must match HTTP requests, use the same encoding as your HTTP clients.



HTTP clients may send requests in encodings other than UTF-8. Encodings usually vary by the client's operating system or input language. If you cannot predict the client's encoding, you may only be able to match any parts of the request that are in English, because regardless of the encoding, the values for English characters tend to be encoded identically. For example, English words may be legible regardless of interpreting a web page as either ISO 8859-1 or as GB2312, whereas simplified Chinese characters might only be legible if the page is interpreted as GB2312.

---

If you configure your FortiGate unit using other encodings, you may need to switch language settings on your management computer, including for your web browser or Telnet/SSH client. For instructions on how to configure your management computer's operating system language, locale, or input method, see its documentation.

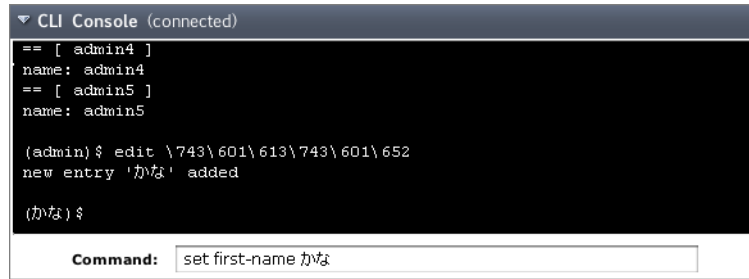
If you choose to configure parts of the FortiGate unit using non-ASCII characters, verify that all systems interacting with the FortiGate unit also support the same encodings. You should also use the same encoding throughout the configuration if possible in order to avoid needing to switch the language settings of the web-based manager and your web browser or Telnet/SSH client while you work.

Similarly to input, your web browser or CLI client should usually interpret display output as encoded using UTF-8. If it does not, your configured items may not display correctly in the web-based manager or CLI. Exceptions include items such as regular expressions that you may have configured using other encodings in order to match the encoding of HTTP requests that the FortiGate unit receives.

#### **To enter non-ASCII characters in the CLI Console widget**

1. On your management computer, start your web browser and go to the URL for the FortiGate unit's web-based manager.
2. Configure your web browser to interpret the page as UTF-8 encoded.
3. Log in to the FortiGate unit.
4. Go to *System > Dashboard > Status*.
5. In title bar of the *CLI Console* widget, click *Edit* (the pencil icon).
6. Enable *Use external command input box*.
7. Select *OK*.
8. The *Command* field appears below the usual input and display area of the *CLI Console* widget.
9. In *Command*, type a command.

**Figure 226:**Entering encoded characters (*CLI Console* widget)



**10. Press Enter.**

In the display area, the *CLI Console* widget displays your previous command interpreted into its character code equivalent, such as:


```
edit \743\601\613\743\601\652
```

and the command's output.

**To enter non-ASCII characters in a Telnet/SSH client**

1. On your management computer, start your Telnet or SSH client.
2. Configure your Telnet or SSH client to send and receive characters using UTF-8 encoding. Support for sending and receiving international characters varies by each Telnet/SSH client. Consult the documentation for your Telnet/SSH client.
3. Log in to the FortiGate unit.
4. At the command prompt, type your command and press Enter.

**Figure 227:**Entering encoded characters (PuTTY)



You may need to surround words that use encoded characters with single quotes ( ' ).

Depending on your Telnet/SSH client's support for your language's input methods and for sending international characters, you may need to interpret them into character codes before pressing Enter.

For example, you might need to enter:

```
edit '\743\601\613\743\601\652'
```

5. The CLI displays your previous command and its output.

## Screen paging

You can configure the CLI to pause after displaying each page's worth of text when displaying multiple pages of output. When the display pauses, the last line displays `--More--`. You can then either:

- press the spacebar to display the next page.
- type `Q` to truncate the output and return to the command prompt.

This may be useful when displaying lengthy output, such as the list of possible matching commands for command completion, or a long list of settings. Rather than scrolling through or possibly exceeding the buffer of your terminal emulator, you can simply display one page at a time.

To configure the CLI display to pause when the screen is full:

```
config system console
 set output more
end
```

## Baud rate

You can change the default baud rate of the local console connection.

To change the baud rate enter the following commands:

```
config system console
 set baudrate {115200 | 19200 | 38400 | 57600 | 9600}
end
```

## Editing the configuration file on an external host

You can edit the FortiGate configuration on an external host by first backing up the configuration file to a TFTP server. Then edit the configuration file and restore it to the FortiGate unit.

Editing the configuration on an external host can be timesaving if you have many changes to make, especially if your plain text editor provides advanced features such as batch changes.

### To edit the configuration on your computer

1. Use `execute backup` to download the configuration file to a TFTP server, such as your management computer.
2. Edit the configuration file using a plain text editor that supports Unix-style line endings.



Do not edit the first line. The first line(s) of the configuration file (preceded by a # character) contains information about the firmware version and FortiGate model. If you change the model number, the FortiGate unit will reject the configuration file when you attempt to restore it.

3. Use `execute restore` to upload the modified configuration file back to the FortiGate unit. The FortiGate unit downloads the configuration file and checks that the model information is correct. If it is, the FortiGate unit loads the configuration file and checks each command for errors. If a command is invalid, the FortiGate unit ignores the command. If the configuration file is valid, the FortiGate unit restarts and loads the new configuration.

## Using Perl regular expressions

Some FortiGate features, such as spam filtering and web content filtering can use either wildcards or Perl regular expressions.

See <http://perldoc.perl.org/perlretut.html> for detailed information about using Perl regular expressions. For more information on using Perl expressions see the [Security Profiles](#) chapter of *The Handbook*.

## Differences between regular expression and wildcard pattern matching

In Perl regular expressions, the period (‘.’) character refers to any single character. It is similar to the question mark (‘?’) character in wildcard pattern matching. As a result:

- `example.com` not only matches `example.com` but also matches `examplecom`, `examplebcom`, `exampleccom` and so on.

To match a special character such as the period (‘.’) and the asterisk (‘\*’), regular expressions use the slash (‘\’) escape character. For example:

- To match `example.com`, the regular expression should be `example\.com`.

In Perl regular expressions, the asterisk (‘\*’) means match 0 or more times of the character before it, not 0 or more times of any character. For example:

- `exam*\ .com` matches `exammm.com` but does not match `example.com`.

To match any character 0 or more times, use ‘.\*’ where ‘.’ means any character and the ‘\*’ means 0 or more times. For example:

- the wildcard match pattern `exam* .com` is equivalent to the regular expression `exam.*\ .com`.

## Word boundary

In Perl regular expressions, the pattern does not have an implicit word boundary. For example, the regular expression “test” not only matches the word “test” but also matches any word that contains the word “test” such as “atest”, “mytest”, “testimony”, “atestb”. The notation “\b” specifies the word boundary. To match exactly the word “test”, the expression should be `\btest\b`.

## Case sensitivity

Regular expression pattern matching is case sensitive in the Web and Spam filters. To make a word or phrase case insensitive, use the regular expression `/i`. For example, `/bad language/i` will block all instances of “bad language” regardless of case.

**Table 67:** Perl regular expression examples

Expression	Matches
<code>abc</code>	abc (that exact character sequence, but anywhere in the string)
<code>^abc</code>	abc at the beginning of the string
<code>abc\$</code>	abc at the end of the string
<code>a b</code>	either of a and b
<code>^abc abc\$</code>	the string abc at the beginning or at the end of the string
<code>ab{2,4}c</code>	an a followed by two, three or four b's followed by a c
<code>ab{2,}c</code>	an a followed by at least two b's followed by a c
<code>ab*c</code>	an a followed by any number (zero or more) of b's followed by a c
<code>ab+c</code>	an a followed by one or more b's followed by a c
<code>ab?c</code>	an a followed by an optional b followed by a c; that is, either abc or ac
<code>a.c</code>	an a followed by any single character (not newline) followed by a c



**Table 67:** Perl regular expression examples

<b>a\.c</b>	a.c exactly
<b>[abc]</b>	any one of a, b and c
<b>[Aa]bc</b>	either of Abc and abc
<b>[abc]+</b>	any (nonempty) string of a's, b's and c's (such as a, abba, acbabcaaaa)
<b>[^abc]+</b>	any (nonempty) string which does not contain any of a, b and c (such as defg)
<b>\d\d</b>	any two decimal digits, such as 42; same as <code>\d{2}</code>
<b>/i</b>	makes the pattern case insensitive. For example, <code>/bad language/i</code> blocks any instance of "bad language" regardless of case.
<b>\w+</b>	a "word": a nonempty sequence of alphanumeric characters and low lines (underscores), such as foo and 12bar8 and foo_1
<b>100\s*mk</b>	the strings 100 and mk optionally separated by any amount of white space (spaces, tabs, newlines)
<b>abc\b</b>	abc when followed by a word boundary (e.g. in abc! but not in abcd)
<b>perl\b</b>	perl when not followed by a word boundary (e.g. in perlert but not in perl stuff)
<b>\x</b>	tells the regular expression parser to ignore white space that is neither backslashed nor within a character class. You can use this to break up your regular expression into (slightly) more readable parts.

# Basic Administration

The FortiGate unit requires some basic configuration to add it to your network. These basic steps include assigning IP addresses, adding routing, and configuring security policies. Until the administrator completes these steps, inter-network and Internet traffic will not flow through the unit.

There are two methods of configuring the FortiGate unit: the web-based manager or the command line interface (CLI). This chapter will step through both methods to complete the basic configurations to put the device on your network. Use whichever you are most comfortable with.

This chapter also provides guidelines for password and administrator best practices as well as how to upgrade the firmware.

This section includes the topics:

- [Connecting to the FortiGate unit](#)
- [System configuration](#)
- [Passwords](#)
- [Administrators](#)
- [Configuration backups](#)
- [Firmware](#)
- [Controlled upgrade](#)

For setup and configuration of your specific FortiGate models, see the [QuickStart Guide](#) for that model.

## Connecting to the FortiGate unit

To configure, maintain and administer the FortiGate unit, you need to connect to it from a management computer. There are two ways to do this:

- using the web-based manager: a GUI interface that you connect to using a current web browser, such as Firefox or Internet Explorer.
- using the command line interface (CLI): a command line interface similar to DOS or UNIX commands that you connect to using SSH or a Telnet terminal.

### Connecting to the web-based manager

To connect to the web-based manager, you require:

- a computer with an Ethernet connection
- Microsoft Internet Explorer version 6.0 or higher or any recent version of a common web browser
- an Ethernet cable.

### To connect to the web-based manager

1. Set the IP address of the management computer to the static IP address 192.168.1.2 with a netmask of 255.255.255.0.
2. Using the Ethernet cable, connect the internal or port 1 interface of the FortiGate unit to the computer Ethernet connection.
3. Start your browser and enter the address `https://192.168.1.99`. (remember to include the “s” in `https://`).

To support a secure HTTPS authentication method, the FortiGate unit ships with a self-signed security certificate that is offered to remote clients whenever they initiate a HTTPS connection to the FortiGate unit. When you connect, the FortiGate unit displays two security warnings in a browser.

The first warning prompts you to accept and optionally install the FortiGate unit’s self-signed security certificate. If you do not accept the certificate, the FortiGate unit refuses the connection. If you accept the certificate, the FortiGate login page appears. The credentials entered are encrypted before they are sent to the FortiGate unit. If you choose to accept the certificate permanently, the warning is not displayed again.

Just before the FortiGate login page is displayed, a second warning informs you that the FortiGate certificate distinguished name differs from the original request. This warning occurs because the FortiGate unit redirects the connection. This is an informational message. Select *OK* to continue logging in.

4. Type `admin` in the Name field and select *Login*.

## Connecting to the CLI

The command line interface (CLI) is an alternative method of configuring the FortiGate unit. The CLI compliments the web-based manager in that it not only has the same configuration options, but also contains additional settings not available through the web-based manager.

If you are new to FortiOS or a command line interface configuration tool, see [“Using the CLI” on page 1405](#) for an overview of the CLI, how to connect to it, and how to use it.

## System configuration

Once the FortiGate unit is connected and traffic can pass through, several more configuration options are available. While not mandatory, they will help to ensure better control with the firewall.

### Setting the time and date

For effective scheduling and logging, the FortiGate system date and time should be accurate. You can either manually set the system date and time or configure the FortiGate unit to automatically keep its time correct by synchronizing with a Network Time Protocol (NTP) server.

#### To set the date and time - web-based manager

1. Go to *System > Dashboard > Status*.
2. Under *System Information > System Time*, select *Change*.
3. Select your *Time Zone*.
4. Select *Set Time* and set the FortiGate system date and time.
5. Select *OK*.

## Set the time and date - CLI

```
config system global
 set timezone <zone_value>
end
execute date [<date_str>]
execute time [<time_str>]
```



By default, FortiOS has the daylight savings time configuration enabled. The system time must be manually adjusted after daylight saving time ends. To disable DST, enter the following command in the CLI:

```
config system global
 set dst disable
end
```

## Using the NTP Server

The Network Time Protocol enables you to keep the FortiGate time in sync with other network systems. By enabling NTP on the FortiGate unit, FortiOS will check with the NTP server you select at the configured intervals. This will also ensure that logs and other time-sensitive settings on the FortiGate unit are correct.

The FortiGate unit maintains its internal clock using a built-in battery. At startup, the time reported by the FortiGate unit will indicate the hardware clock time, which may not be accurate. When using NTP, the system time might change after the FortiGate has successfully obtained the time from a configured NTP server.

For the NTP server, you can identify a specific port/IP address for this self-originating traffic. The configuration is performed in the CLI with the command `set source-ip`. For example, to set the source IP of NTP to be on the DMZ1 port with an IP of 192.168.4.5, the commands are:

```
config system ntp
 set ntpsyn enable
 set syncinterval 5
 set source-ip 192.168.4.5
end
```

## Configuring FortiGuard

FortiGuard is Fortinet's threat research and response team. With more than 200 security engineers and forensic analysts around the globe providing 24 hours a day, 365 days a year analysis of current threats on the Internet, the FortiGuard team's sole purpose is to protect customers.

The FortiGuard Distribution Network (FDN) is a world-wide network of FortiGuard Distribution Servers (FDS). When the FortiGate unit connects to the FDN, it connects to the nearest FDS. To do this, all FortiGate units are programmed with a list of FDS addresses sorted by nearest time zone according to the time zone configured for the FortiGate unit.

Before you can begin receiving updates, you must register your FortiGate unit from the Fortinet webpage. After registering, you need to configure the FortiGate unit to connect to the FDN to update antivirus, antispam, and IPS attack definitions.

## Updating antivirus and IPS definitions

After you have registered your FortiGate unit, you can update the definitions for antivirus and IPS. The FortiGuard Center enables you to receive push updates, allow push update to a specific IP address, and schedule updates for daily, weekly, or hourly intervals.

### To update antivirus definitions and IPS signatures

1. Go to *System > Config > FortiGuard*.
2. Select the expand arrow for *AV and IPS Options*.
3. Select *Update Now* to update the antivirus definitions.

If the connection to the FDN is successful, the web-based manager displays a message similar to the following:

Your update request has been sent. Your database will be updated in a few minutes. Please check your update page for the status of the update.

After a few minutes, if an update is available, the FortiGuard Center Services information on the Dashboard lists new version information for antivirus definitions. The System Status page also displays new dates and version numbers for the antivirus definitions. Messages are recorded to the event log indicating whether or not the update was successful or not.



Updating antivirus definitions can cause a very short disruption in traffic currently being scanned while the FortiGate unit applies the new signature database. Schedule updates when traffic is light, for example overnight, to minimize any disruption.

## Passwords

The FortiGate unit ships with a default admin account that has no password. You will want to apply a password to prevent anyone from logging into the FortiGate unit and changing configuration options.

### To change the administrator password - web-based manager

1. Go to *System > Admin > Administrators*.
2. Select the admin account and select *Change Password*.
3. Enter a new password and select *OK*.

### Set the admin password - CLI

```
config system admin
 edit admin
 set password <admin_password>
 end
```

## Password considerations

When changing the password, consider the following to ensure better security.

- Do not make passwords that are obvious, such as the company name, administrator names, or other obvious word or phrase.
- Use numbers in place of letters, for example, `passw0rd`. Alternatively, spell words with extra letters, for example, `password`.
- Administrative passwords can be up to 64 characters.

- Include a mixture of letters, numbers, and upper and lower case.
- Use multiple words together, or possibly even a sentence, for example keytothehighway.
- Use a password generator.
- Change the password regularly and always make the new password unique and not a variation of the existing password, such as changing from `password` to `password1`.
- Write the password down and store it in a safe place away from the management computer, in case you forget it or ensure that at least two people know the password in the event that one person becomes ill, is away on vacation or leaves the company. Alternatively, have two different admin logins.

## Password policy

The FortiGate unit includes the ability to enforce a password policy for administrator login. With this policy, you can enforce regular changes and specific criteria for a password including:

- minimum length between 8 and 64 characters.
- if the password must contain uppercase (A, B, C) and/or lowercase (a, b, c) characters.
- if the password must contain numbers (1, 2, 3).
- if the password must contain non-alphanumeric characters (!, @, #, \$, %, ^, &, \*, (),).
- where the password applies (admin or IPsec or both).
- the duration of the password before a new one must be specified.

### To apply a password policy - web-based manager

1. Go to *System > Admin > Settings*.
2. Select *Enable Password Policy* and configure the settings as required.

### To apply a password policy - CLI

```
config system password-policy
 set status enable
```

Configure the other settings as required.

If you add a password policy or change the requirements on an existing policy, the next time that administrator logs into the FortiGate unit, they are prompted to update their password to meet the new requirements before proceeding to log in.

**Figure 228:** Password policy dialog box

Your password does not conform to the password policy, please input a new password.

New Password:  ?

Confirm Password:

---

**Password Policy**

Minimum Length:	15 Characters
Must Contain:	1 Upper Case Letters

## Lost Passwords

If an administrator password has been lost, refer to the SysAdmin's Notebook article "Resetting a lost admin password," found at [docs.fortinet.com/p/sysadmin-s-notebook-and-tech-notes](https://docs.fortinet.com/p/sysadmin-s-notebook-and-tech-notes).

## Administrators

By default, the FortiGate unit has a super administrator called "admin". This user login cannot be deleted and always has ultimate access over the FortiGate unit. Additional administrators can be added for various functions, each with a unique username, password, and set of access privileges.

There are two levels of administrator accounts; regular administrators and system administrators. Regular administrators are administrators with any admin profile other than the default super\_admin. System administrators are administrators that are assigned the super\_admin profile, which has the highest level of access.

### Adding administrators



The name of the administrator should not contain the characters <> () # " ' . Using these characters in the administrator account name can result in a cross site scripting (XSS) vulnerability.

Only the default "admin" account or an administrator with read-write access control to add new administrator accounts and control their permission levels can create a new administrator account. If you log in with an administrator account that does not have the super\_admin admin profile, the administrators list will show only the administrators for the current virtual domain.

When adding administrators, you are setting up the administrator's user account. An administrator account comprises of an administrator's basic settings as well as their access profile. The access profile is a definition of what the administrator is capable of viewing and editing.

#### To add an administrator - web-based manager

1. Go to *System > Admin > Administrators*.
2. Select *Create New*.
3. Enter the administrator name.
4. Select the type of account it will be. If you select *Remote*, the FortiGate unit can reference a RADIUS, LDAP or TACAS+ server.
5. When selecting *Remote* or *PKI* accounts, select the User Group the account will access.  
For information on logging in using remote authentication servers, see the [User Authentication Guide](#). For an example of setting up a user with LDAP, see "[LDAP Admin Access and Authorization](#)" on page 1432
6. Enter the password for the user.  
This may be a temporary password that the administrator can change later. Passwords can be up to 256 characters in length. For more information on passwords, see "[Passwords](#)" on page 1429.
7. Select *OK*.

### To add an administrator - CLI

```
config system admin
 edit <admin_name>
 set password <password>
 set accprofile <profile_name>
 end
```

## LDAP Admin Access and Authorization

You can use the LDAP server as a means to add administrative users, saving the time to add users to the FortiGate unit administrator list. After configuring, any user within the selected LDAP group server can automatically log into the FortiGate unit as an administrator. Ensure that the admin profile is the correct level of access, or the users within the LDAP group are the only ones authorized to configure or modify the configuration of the FortiGate unit.

To do this, requires three steps:

- configure the LDAP server
- add the LDAP server to a user group
- configure the administrator account

### Configure the LDAP server

First set up the LDAP server as you normally would, and include a group to bind to.

#### To configure the LDAP server - web-based manager

1. Go to *User & Device > Remote > LDAP* and select *Create New*.
2. Enter a *Name* for the server.
3. Enter the *Server IP* address or name.
4. Enter the *Common Name Identifier* and *Distinguished Name*.
5. Set the *Bind Type* to *Regular* and enter the *User DN* and *Password*.
6. Select *OK*.

#### To configure the LDAP server - CLI

```
config user ldap
 edit <ldap_server_name>
 set server <server_ip>
 set cnid cn
 set dn DC=XYZ,DC=COM
 set type regular
 set username CN=Administrator,CN=Users,DC=XYZ,DC=COM
 set password <password>
 set member-attr <group_binding>
 end
```

### Add the LDAP server to a user group

Next, create a user group that will include the LDAP server that was created above.

#### To create a user group - web-based manager

1. Go to *User & Device > User Group > User Group* and select *Create New*.
2. Enter a *Name* for the group.



3. In the section labelled *Remote authentication servers*, select *Add*.
4. Select the *Remote Server* from the drop-down list.
5. Select *OK*.

#### To create a user group - CLI

```
config user group
 edit <group_name>
 config match
 edit 1
 set server-name <LDAP_server>
 set group-name <group_name>
 end
 end
 end
```

### Configure the administrator account

Now you can create a new administrator, where rather than entering a password, you will use the new user group and the wildcard option for authentication.

#### To create an administrator - web-based manager

1. Go to *System > Admin > Administrators* and select *Create New*.
2. In the *Administrator* field, enter the name for the administrator.
3. For *Type*, select *Remote*.
4. Select the *User Group* created above from the drop-down list.
5. Select *Wildcard*.
6. The *Wildcard* option allows for LDAP users to connect as this administrator.
7. Select an *Admin Profile*.
8. Select *OK*.

#### To create an administrator - CLI

```
config system admin
 edit <admin_name>
 set remote-auth enable
 set accprofile super_admin
 set wildcard enable
 set remote-group ldap
 end
```

### Monitoring administrators

You can view the administrators logged in using the *System Information* widget on the Dashboard. On the widget is the *Current Administrator* row that shows the administrator logged in and the total logged in. Selecting *Details* displays the administrators), where they are logging in from and how (CLI, web-based manager) and when they logged in.

You are also able to monitor the activities the administrators perform on the FortiGate unit using the logging of events. Event logs include a number of options to track configuration changes.

### To set logging - web-based manager

1. Go to *Log & Report > Log Config > Log Settings*.
- 2 Under *Event Logging*, ensure *System activity event* is selected.
- 3 Select *Apply*.

### To set logging - CLI

```
config log eventfilter
 set event enable
 set system enable
end
```

To view the logs go to *Log & Report > Event Log*.

## Administrator profiles

Administer profiles define what the administrator user can do when logged into the FortiGate unit. When you set up an administrator user account, you also assign an administrator profile, which dictates what the administrator user will see. Depending on the nature of the administrator's work, access level or seniority, you can allow them to view and configure as much, or as little, as required.

### super\_admin profile

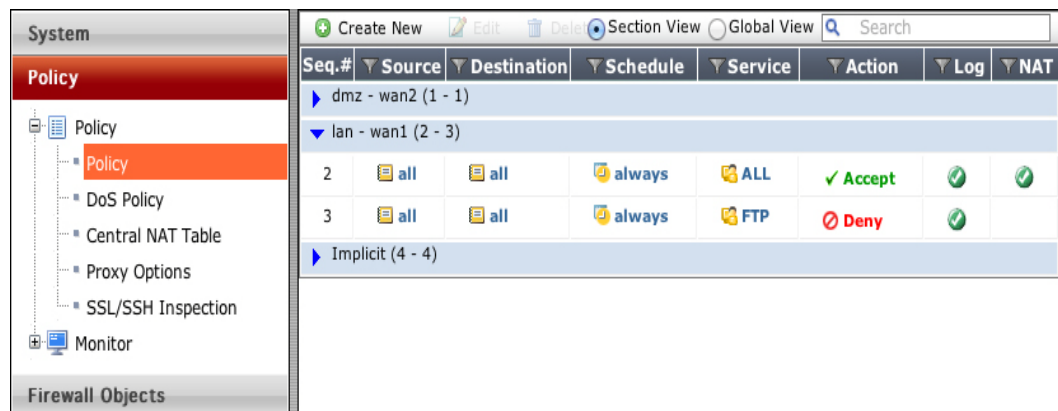
The super\_admin administrator is the administrative account that the primary administrator should have to log into the FortiGate unit. The profile can not be deleted or modified to ensure there is always a method to administer the FortiGate unit. This user profile has access to all components of FortiOS, including the ability to add and remove other system administrators. For some administrative functions, such as backing up and restoring the configuration using SCP, super\_admin access is required.

### Creating profiles

To configure administrator profiles go to *System > Admin > Admin Profiles*. You can only assign one profile to an administrator user.

On the *New Admin Profile* page, you define the components of FortiOS that will be available to view and/or edit. For example, if you configure a profile so that the administrator can only access the firewall components, when an administrator with that profile logs into the FortiGate unit, they will only be able to view and edit any firewall components including policies, addresses, schedules and any other settings that directly affect security policies.

**Figure 229:**The view of an administrator with firewall-only access



## Global and vdom profiles

By default, when you add a new administrative profile, it is set to have a vdom scope. That is, only the `super_admin` has a global profile that enables configuration of the entire FortiGate unit.

There may be instances where additional global administrative profiles may be required. To add more global profiles, use the following CLI command to set or change an administrative profile to be global.

```
config system accprofile
 set scope global
 ...
end
```

Once the scope is set, you can enable the read and read/write settings.

## Regular (password) authentication for administrators

You can use a password stored on the local FortiGate unit to authenticate an administrator. When you select *Regular* for *Type*, you will see *Local* as the entry in the *Type* column when you view the list of administrators.

## Management access

Management access defines how administrators are able to log on to the FortiGate unit to perform management tasks such as configuration and maintenance. Methods of access can include local access through the console connection or remote access over a network or modem interface using various protocols including Telnet and HTTPS.

You can configure management access on any interface in your VDOM. In NAT mode, the interface IP address is used for management access. In transparent mode, you configure a single management IP address that applies to all interfaces in your VDOM that permit management access. The FortiGate unit also uses this IP address to connect to the FDN for virus and attack updates.

The system administrator (`admin`) can access all VDOMs, and create regular administrator accounts. A regular administrator account can access only the VDOM to which it belongs and the management computer must connect to an interface in that VDOM. In both cases, the management computer must connect to an interface that permits management access and its IP address must be on the same network. Management access can be via HTTP, HTTPS, Telnet, or SSH sessions, if those services are enabled on the interface. HTTPS and SSH are preferred as they are more secure.

You can allow remote administration of the FortiGate unit. However, allowing remote administration from the Internet could compromise the security of the FortiGate unit. You should avoid this unless it is required for your configuration. The following precautions can be taken to improve the security of a FortiGate unit that allows remote administration from the Internet:

- Use secure administrative user passwords.
- Change these passwords regularly.
- Enable two-factor authentication for administrators.
- Enable secure administrative access to this interface using only HTTPS or SSH.
- Use Trusted Hosts to limit where the remote access can originate from.
- Do not change the system idle timeout from the default value of 5 minutes.

## Security Precautions

One potential point of a security breach is at the management computer. Administrators who leave their workstations for a prolonged amount of time while staying logged into the web-based manager or CLI leave the firewall open to malicious intent.

### Change the admin username and password

The default super administrator user name, `admin`, is a very standard, making it easy for someone with malicious intent to determine or guess. Having the correct user name is one half of the key to the FortiGate unit being compromised and so the default name should be changed.

To do this, you need to create another super user with full access and log in as that user. Then go to *System > Admin > Administrator*, select the *admin* account, and select *Edit* to change the user name.

If it has not been done already, the password should also be changed at this time. For tips about changing the password, see [“Password considerations” on page 1429](#).

### Preventing unwanted login attempts

Setting trusted hosts for an administrator limits what computers an administrator can log in from, causing the FortiGate unit to only accept the administrator's login from the configured IP address. Any attempt to log in with the same credentials from any other IP address will be dropped.

Trusted hosts are configured when adding a new administrator by going to *System > Admin > Administrators* in the web-based manager or `config system admin` in the CLI.

To ensure the administrator has access from different locations, you can enter up to ten IP addresses, though ideally this should be kept to a minimum. For higher security, use an IP address with a net mask of `255.255.255.255`, and enter an IP address (non-zero) in each of the three default trusted host fields. Also ensure all entries contain actual IP addresses, not the default `0.0.0.0`.

The trusted hosts apply to the web-based manager, ping, snmp and the CLI when accessed through Telnet or SSH. CLI access through the console port is not affected.

### Prevent multiple admin sessions

Multiple admin sessions can occur when multiple users access the FortiGate using the same admin account. By default, the FortiGate unit enables multiple logins of administrators using the same login credentials from different locations. To control admin log ins, and minimize the potential of configuration collisions, you can disable concurrent admin sessions. When disabled, only one user can use the admin account at a time. When a second admin attempts to connect, connection is denied with a message that the login attempt failed.

**To disable concurrent admin sessions, enter the following command in the CLI:**

```
config system global
 set admin-concurrent disable
end
```

On 2U FortiGate units, this option is also available in the Web-Based Manager by going to *System > Admin > Settings* and select *Allow each admin to log in with multiple sessions*.

## Segregated administrative roles

To minimize the effect of an administrator causing errors to the FortiGate configuration and possibly jeopardizing the network, create individual administrative roles where none of the administrators have super-admin permissions. For example, one admin account is used solely to create security policies, another for users and groups, another for VPN, and so on.

## Disable admin services

On untrusted networks, turn off the weak administrative services such as Telnet and HTTP. With these services, passwords are passed in the clear, not encrypted. These services can be disabled by going to *System > Network > Interface* and unselecting the required check boxes.

## SSH login time out

When logging into the console using SSH, the default time of inactivity is 120 seconds (2 minutes) to successfully log into the FortiGate unit. To enhance security, you can configure the time to be shorter. Using the CLI, you can change the length of time the command prompt remains idle before the FortiGate unit will log the administrator out. The range can be between 10 and 3600 seconds.

**To set the logout time enter the following commands:**

```
config system global
 set admin-ssh-grace-time <number_of_seconds>
end
```

## Administrator lockout

By default, the FortiGate unit includes set number of three password retries, allowing the administrator a maximum of three attempts to log into their account before they are locked out for a set amount of time (by default, 60 seconds).

The number of attempts can be set to an alternate value, as well as the default wait time before the administrator can try to enter a password again. You can also change this to further sway would-be hackers. Both settings are must be configured with the CLI

**To configure the lockout options:**

```
config system global
 set admin-lockout-threshold <failed_attempts>
 set admin-lockout-duration <seconds>
end
```

For example, to set the lockout threshold to one attempt and a five minute duration before the administrator can try again to log in enter the commands”

```
config system global
 set admin-lockout-threshold 1
 set admin-lockout-duration 300
end
```

## Idle time-out

To avoid the possibility of an administrator walking away from the management computer and leaving it exposed to unauthorized personnel, you can add an idle time-out that will automatically log the user out if the web-based manager is not used for a specified amount of time. This will cause the administrator to log in to the device again in order to continue their work.

The time-out can be set as high as 480 minutes, or eight hours, although this is not recommend.

#### To set the idle time out - web-based manager

1. Go to *System > Admin > Settings*.
2. In the *Administration Settings*, enter the amount of time the Administrator login can remain idle in the *Idle Timeout* field.
3. Select *Apply*.

#### To set the idle time out - CLI

```
config system global
 set admintimeout <minutes>
end
```

### Administrative ports

You can set the web-based manager access to use HTTP, HTTPS, SSH, and Telnet. In these cases, the default ports for these protocols are 80, 443, 22, and 23 respectively. You can change the ports used for network administration to a different, unused port to further limit potential hackers.



Ensure the port you select is not a port you will be using for other applications. For a list of assigned port numbers, see <http://www.iana.org/assignments/port-numbers>.

---

#### To change the administrative ports - web-based manager

1. Go to *System > Admin > Settings*.
2. In the *Web Administration Ports* section, change the port numbers.
3. Select *Apply*.

#### To change the administrative ports - CLI

```
config system global
 set admin-port <http_port_number>
 set admin-sport <https_port_number>
 set admin-ssh-port <ssh_port_number>
 set admin-telnet-port <telnet_port_number>
end
```

When logging into the FortiGate unit, by default FortiOS will automatically use the default ports. That is, when logging into the FortiGate IP address, you only need to enter the address, for example:

```
https://192.168.1.1
```

When you change the administrative port number, the port number must be added to the url. For example, if the port number for HTTPS access is 2112, the administrator must enter the following address:

```
https://192.168.1.1:2112
```

### HTTPS redirect

When selecting port numbers for various protocols, you can also enable or disable the Redirect to HTTPS option. When enabled, if you select the Administrative Access for an interface to be

only HTTP, HTTPS will automatically be enabled, allowing the administrator access with either HTTP or HTTPS. The administrator can then log in using HTTPS for better security.

Note that if an SSL VPN is also configured for the same port, the SSL connection is over the HTTPS protocol. In these situations, the FortiGate unit will not redirect an HTTP address to the SSL VPN HTTPS address. Ideally, the administrator should not have the management address and an SSL VPN portal on the same interface.

### Log in/out warning message

For administrators logging in and out of the FortiGate unit, you can include a login disclaimer. This disclaimer provides a statement that must be accepted or declined where corporations are governed by strict usage policies for forensics and legal reasons.

The disclaimer is enabled through the CLI.

#### To disable an interface:

```
config system global
 set pre-login-banner enable
 set post-login-banner enable
end
```

When set, once the administrator enters their user name and password, the disclaimer appears. They must select either Accept or Decline to proceed. When the post login is enabled, once the administrator logs out they are presented with the same message.

The banner is a default message that you can customize by going to *System > Config > Replacement Messages*. Select *Extended View* to see the *Admin* category and messages.

### Disable the console interface

To prevent any unwanted login attempts using the COM communication port, you can disable login connections to the FortiGate unit.

This command is specifically for the COM port. You can still use FortiExplorer to connect and configure the FortiGate unit if required.

#### To disable an interface:

```
config system console
 set login disable
end
```

### Disable interfaces

If any of the interfaces on the FortiGate unit are not being used, disable traffic on that interface. This avoids someone plugging in network cables and potentially causing network bypass or loop issues.

#### To disable an interface - web-based manager

1. Go to *System > Network > Interface*.
2. Select the interface from the list and select *Edit*.
3. For *Administrative Access*, select *Down*.
4. Select *OK*.

## To disable an interface - CLI

```
config system interface
 edit <interface_name>
 set status down
 end
```

## RADIUS authentication for administrators

Remote Authentication and Dial-in User Service (RADIUS) servers provide authentication, authorization, and accounting functions. FortiGate units use the authentication and authorization functions of the RADIUS server. To use the RADIUS server for authentication, you must configure the server before configuring the FortiGate users or user groups that will need it.

If you have configured RADIUS support and a user is required to authenticate using a RADIUS server, the FortiGate unit sends the user's credentials to the RADIUS server for authentication. If the RADIUS server can authenticate the user, the user is successfully authenticated with the FortiGate unit. If the RADIUS server cannot authenticate the user, the FortiGate unit refuses the connection.

If you want to use a RADIUS server to authenticate administrators in your VDOM, you must configure the authentication before you create the administrator accounts. To do this you need to:

- configure the FortiGate unit to access the RADIUS server
- create the RADIUS user group
- configure an administrator to authenticate with a RADIUS server.

## Configuring LDAP authentication for administrators

Lightweight Directory Access Protocol (LDAP) is an Internet protocol used to maintain authentication data that may include departments, people, groups of people, passwords, email addresses, printers, etc.

If you have configured LDAP support and an administrator is required to authenticate using an LDAP server, the FortiGate unit contacts the LDAP server for authentication. If the LDAP server cannot authenticate the administrator, the FortiGate unit refuses the connection.

If you want to use an LDAP server to authenticate administrators in your VDOM, you must configure the authentication before you create the administrator accounts. To do this you need to:

- configure an LDAP server
- create an LDAP user group
- configure an administrator to authenticate with an LDAP server.

To view the LDAP server list, go to *User & Device > Remote > LDAP*.

For more information, see [“LDAP Admin Access and Authorization”](#) on page 1432.

## TACACS+ authentication for administrators

Terminal Access Controller Access-Control System (TACACS+) is a remote authentication protocol that provides access control for routers, network access servers, and other network computing devices via one or more centralized servers.

If you have configured TACACS+ support and an administrator is required to authenticate using a TACACS+ server, the FortiGate unit contacts the TACACS+ server for authentication. If the TACACS+ server cannot authenticate the administrator, the connection is refused by the FortiGate unit.



If you want to use an TACACS+ server to authenticate administrators in your VDOM, you must configure the authentication before you create the administrator accounts. To do this you need to:

- configure the FortiGate unit to access the TACACS+ server
- create a TACACS+ user group
- configure an administrator to authenticate with a TACACS+ server.

### PKI certificate authentication for administrators

Public Key Infrastructure (PKI) authentication uses a certificate authentication library that takes a list of peers, peer groups, and user groups and returns authentication successful or denied notifications. Users only need a valid certificate for successful authentication; no username or password is necessary.

To use PKI authentication for an administrator, you must configure the authentication before you create the administrator accounts. To do this you need to:

- configure a PKI user
- create a PKI user group
- configure an administrator to authenticate with a PKI certificate.

## General Settings

Go to *System > Admin > Settings* to configure basic settings for administrative access, password policies and displaying additional options in the web-based manager.

### Administrative port settings

The Administrative Settings enable you to change the default port configurations for administrative connections to the FortiGate unit for added security. When connecting to the FortiGate unit when the port has changed, the port must be included, such as `https://<ip_address>:<port>`. For example, if you are connecting to the FortiGate unit using port 99, the url would be `https://192.168.1.99:99`.

If you make a change to the default port number for HTTP, HTTPS, Telnet, or SSH, ensure that the port number is unique.

### Password policies

Password policies, available by going to *System > Admin > Settings*, enable you to create a password policy that any administrator or user who updates their password, must follow. Using the available options you can define the required length of the password, what it must contain (numbers, upper and lower case, and so on) and an expiry time frame.

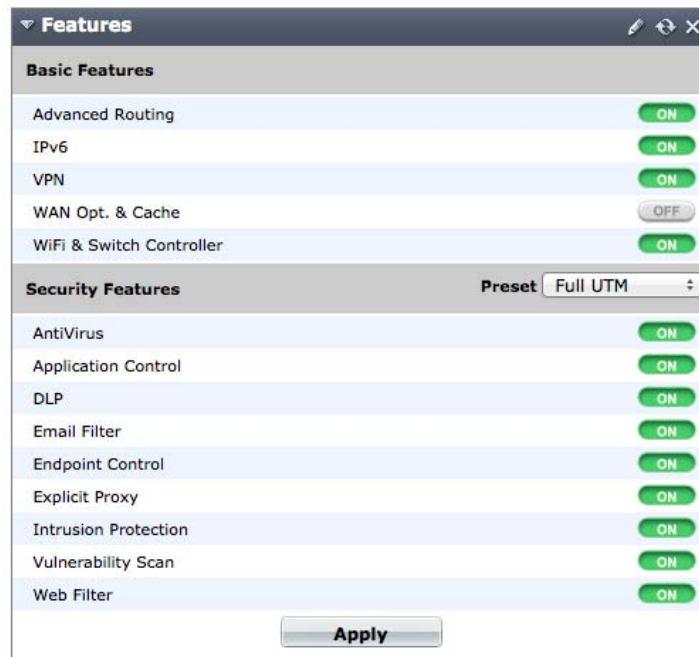
The FortiGate unit will warn of any password that is added and does not meet the criteria.

### Feature Select

Feature Select is used to disable features which are not required for network administration. Disabling features also removes all related configuration options from the web-based manager.

Feature Select can be managed using the *Features* widget on the *Status* page. They can also be found at *System > Config > Features*, where additional features are also available by selecting *Show More*.

**Figure 230:**The Features widget



If a feature, such as IPv6, has been configured before being removed from the web-based manager, this configuration will still exist as part of the network, even though it is no longer visible using the web-based manager.

## Configuration backups

Once you configure the FortiGate unit and it is working correctly, it is extremely important that you backup the configuration. In some cases, you may need to reset the FortiGate unit to factory defaults or perform a TFTP upload of the firmware, which will erase the existing configuration. In these instances, the configuration on the device will have to be recreated, unless a backup can be used to restore it.

It is also recommended that once *any* further changes are made that you backup the configuration immediately, to ensure you have the most current configuration available. Also, ensure you backup the configuration before upgrading the FortiGate unit's firmware. Should anything happen during the upgrade that changes the configuration, you can easily restore the saved configuration.

Always backup the configuration and store it on the management computer or off-site. You have the option to save the configuration file to various locations including the local PC, USB key, FTP and TFTP site. The latter two are configurable through the CLI only.

If you have VDOMs, you can back up the configuration of the entire FortiGate unit or only a specific VDOM. Note that if you are using FortiManager or FortiCloud, full backups are performed and the option to backup individual VDOMs will not appear.

### To back up the FortiGate configuration - web-based manager

1. Go to *System > Dashboard > Status*.
2. On the *System Information* widget, select *Backup* for the *System Configuration*.

3. Select to backup to your *Local PC* or to a *USB key*.  
The *USB Disk* option will be grayed out if no USB drive is inserted in the USB port. You can also backup to the FortiManager using the CLI.
4. If VDOMs are enabled, select to backup the entire FortiGate configuration (*Full Config*) or only a specific VDOM configuration (*VDOM Config*).
5. If backing up a VDOM configuration, select the VDOM name from the list.
6. Select *Encrypt configuration file*.  
Encryption must be enabled on the backup file to back up VPN certificates.
7. Enter a password and enter it again to confirm it. You will need this password to restore the file.
8. Select *Backup*.
9. The web browser will prompt you for a location to save the configuration file. The configuration file will have a .conf extension.

### To back up the FortiGate configuration - CLI

```
execute backup config management-station <comment>
```

... or ...

```
execute backup config usb <backup_filename> [<backup_password>]
```

... or for FTP, note that port number, username are optional depending on the FTP site...

```
execute backup config ftp <backup_filename> <ftp_server> [<port>]
 [<user_name>] [<password>]
```

... or for TFTP ...

```
execute backup config tftp <backup_filename> <tftp_servers>
 <password>
```

Use the same commands to backup a VDOM configuration by first entering the commands:

```
config vdom
 edit <vdom_name>
```

## Backup and restore a configuration file using SCP

You can use secure copy protocol (SCP) to download the configuration file from the FortiGate unit as an alternative method of backing up the configuration file or an individual VDOM configuration file. This is done by enabling SCP for an administrator account and enabling SSH on a port used by the SCP client application to connect to the FortiGate unit. SCP is enabled using the CLI commands:

```
config system global
 set admin-scp enable
end
```

Use the same commands to backup a VDOM configuration by first entering the commands:

```
config global
 set admin-scp enable
end
config vdom
 edit <vdom_name>
```

## Enable SSH access on the interface

SCP uses the SSH protocol to provide secure file transfer. The interface you use for administration must allow SSH access.

### To enable SSH - web-based manager:

1. Go to *System > Network > Interface*.
2. Select the interface you use for administrative access and select *Edit*.
3. In the *Administrative Access* section, select *SSH*.
4. Select *OK*.

### To enable SSH - CLI:

```
config system interface
 edit <interface_name>
 set allowaccess ping https ssh
 end
```



When adding to, or removing a protocol, you must type the entire list again. For example, if you have an access list of HTTPS and SSH, and you want to add PING, typing:

```
set allowaccess ping
```

...only PING will be set. In this case, you must type...

```
set allowaccess https ssh ping
```

---

## Using the SCP client

The FortiGate unit downloads the configuration file as `sys_conf`. Use the following syntax to download the file:

### Linux

```
scp admin@<FortiGate_IP>:fgt-config <location>
```

### Windows

```
pscp admin@<FortiGate_IP>:fgt-config <location>
```

The following examples show how to download the configuration file from a FortiGate-100D, at IP address 172.20.120.171, using Linux and Windows SCP clients.

### Linux client example

To download the configuration file to a local directory called `~/config`, enter the following command:

```
scp admin@172.20.120.171:fgt-config ~/config
```

Enter the admin password when prompted.

### Windows client example

To download the configuration file to a local directory called `c:\config`, enter the following command in a Command Prompt window:

```
pscp admin@172.20.120.171:fgt-config c:\config
```

Enter the admin password when prompted.

## SCP public-private key authentication

SCP authenticates itself to the FortiGate unit in the same way as an administrator using SSH accesses the CLI. Instead of using a password, you can configure the SCP client and the FortiGate unit with a public-private key pair.

### To configure public-private key authentication

1. Create a public-private key pair using a key generator compatible with your SCP client.
2. Save the private key to the location on your computer where your SSH keys are stored.  
This step depends on your SCP client. The Secure Shell key generator automatically stores the private key.

3. Copy the public key to the FortiGate unit using the CLI commands:

```
config system admin
 edit admin
 set ssh-public-key1 "<key-type> <key-value>"
 end
```

<key-type> must be the ssh-dss for a DSA key or ssh-rsa for an RSA key. For the <key-value>, copy the public key data and paste it into the CLI command.

If you are copying the key data from Windows Notepad, copy one line at a time and ensure that you paste each line of key data at the end of the previously pasted data. As well:

- Do not copy the end-of-line characters that appear as small rectangles in Notepad.
- Do not copy the ---- BEGIN SSH2 PUBLIC KEY ---- or Comment: "[2048-bit dsa, ...]" lines.
- Do not copy the ---- END SSH2 PUBLIC KEY ---- line.

4. Type the closing quotation mark and press Enter.

Your SCP client can now authenticate to the FortiGate unit based on SSH keys rather than the administrator password.

## Restoring a configuration using SCP

To restore the configuration using SCP, use the commands:

```
scp <local_file> <admin_user>@<FGT_IP>:fgt_restore_config
```

To use this command/method of restoring the FortiGate configuration, you need to log in as the "admin" administrator.

## Restoring a configuration

Should you need to restore a configuration file, use the following steps:

### To restore the FortiGate configuration - web-based manager

1. Go to *System > Dashboard > Status*.
2. On the *System Information* widget, select *Restore* for the *System Configuration*.
3. Select to upload the configuration file to be restored from your *Local PC* or a *USB key*.  
The *USB Disk* option will be grayed out if no USB drive is inserted in the USB port. You can restore from the FortiManager using the CLI.
4. Enter the path and file name of the configuration file, or select *Browse* to locate the file.
5. Enter a password if required.
6. Select *Restore*.

### To back up the FortiGate configuration - CLI

```
execute restore config management-station normal 0
```

... or ...

```
execute restore config usb <filename> [<password>]
```

... or for FTP, note that port number, username are optional depending on the FTP site...

```
execute backup config ftp <backup_filename> <ftp_server> [<port>]
 [<user_name>] [<password>]
```

... or for TFTP ...

```
execute backup config tftp <backup_filename> <tftp_server> <password>
```

The FortiGate unit will load the configuration file and restart. Once the restart has completed, verify that the configuration has been restored.

## Configuration revisions

The *Revisions* options on the *System Information* widget enables you to manage multiple versions of configuration files. Revision control requires either a configured central management server, or FortiGate units with 512 MB or more of memory. The central management server can either be a FortiManager unit or the FortiCloud.

When revision control is enabled on your unit, and configurations have been backed up, a list of saved revisions of those backed-up configurations appears.

## Restore factory defaults

There may be a point where need to reset the FortiGate unit to its original defaults; for example, to begin with a fresh configuration. There are two options when restoring factory defaults. The first resets the entire device to the original out-of-the-box configuration:

### To reset the FortiGate unit to its factory default settings - web-based manager

1. Go to *System > Dashboard > Status*.
2. In the *System Information* widget, select *Restore* for the *System Configuration*.
3. Select *Restore Factory Defaults* at the top of the page.

### You can reset using the CLI by entering the command:

```
execute factoryreset
```

When prompted, type *y* to confirm the reset.

Alternatively, in the CLI you can reset the factory defaults but retain the interface and VDOM configuration.

Use the command:

```
execute factoryreset2
```

## Firmware

Fortinet periodically updates the FortiGate firmware to include new features and resolve important issues. After you have registered your FortiGate unit, you can download firmware updates from the support web site, <http://support.fortinet.com>.

The FortiGate unit includes a number of firmware installation options that enables you to test new firmware without disrupting the existing installation, and load it from different locations as required.

Before you install any new firmware, be sure to follow the steps below:

- review the [Release Notes](#) for a new firmware release.
- review the [Supported Upgrade Paths](#) document to make sure the upgrade from your current image to the desired new image is supported.
- backup the current configuration.
- download the new firmware image.
- test the patch release until you are satisfied that it applies to your configuration.

Installing a patch release without reviewing release notes or testing the firmware may result in changes to settings or unexpected issues.

Only FortiGate admin user and administrators whose access profiles contain system read and write privileges can change the FortiGate firmware.

## Downloading firmware

Firmware images for all FortiGate units is available on the Fortinet Customer Support website. You must register your FortiGate unit to access firmware images. Register the FortiGate unit by visiting <http://support.fortinet.com> and select Product Registration.

### To download firmware

1. Log into the site using your user name and password.
2. Go to *Firmware Images > FortiGate*.
3. Select the most recent FortiOS version.
4. Locate the firmware for your FortiGate unit, right-click the link and select the Download option for your browser.

## Testing new firmware before installing

FortiOS enables you to test a new firmware image by installing the firmware image from a system reboot and saving it to system memory. After completing this procedure, the FortiGate unit operates using the new firmware image with the current configuration. This new firmware image is not permanently installed. The next time the FortiGate unit restarts, it operates with the originally installed firmware image using the current configuration. If the new firmware image operates successfully, you can install it permanently using the procedure “[Upgrading the firmware - web-based manager](#)” on page 1449.

To use this procedure, you must connect to the CLI using the FortiGate console port and a RJ-45 to DB-9 or null modem cable. This procedure temporarily installs a new firmware image using your current configuration.

For this procedure you install a TFTP server that you can connect to from the FortiGate internal interface. The TFTP server should be on the same subnet as the internal interface.

### To test the new firmware image

1. Connect to the CLI using a RJ-45 to DB-9 or null modem cable.
2. Make sure the TFTP server is running.
3. Copy the new firmware image file to the root directory of the TFTP server.
4. Make sure the FortiGate unit can connect to the TFTP server using the `execute ping` command.

5. Enter the following command to restart the FortiGate unit:  
`execute reboot`
6. As the FortiGate unit reboots, press any key to interrupt the system startup. As the FortiGate unit starts, a series of system startup messages appears.  
When the following messages appears:  
`Press any key to display configuration menu...`
7. Immediately press any key to interrupt the system startup.



You have only 3 seconds to press any key. If you do not press a key soon enough, the FortiGate unit reboots and you must login and repeat the `execute reboot` command.

---

If you successfully interrupt the startup process, the following messages appears:

```
[G]: Get firmware image from TFTP server.
[F]: Format boot device.
[B]: Boot with backup firmware and set as default
[C]: Configuration and information
[Q]: Quit menu and continue to boot with default
firmware.
[H]: Display this list of options.
```

Enter G, F, Q, or H:

8. Type G to get the new firmware image from the TFTP server.  
The following message appears:  
`Enter TFTP server address [192.168.1.168]:`
9. Type the address of the TFTP server and press Enter:  
The following message appears:  
`Enter Local Address [192.168.1.188]:`
10. Type an IP address of the FortiGate unit to connect to the TFTP server.  
The IP address must be on the same network as the TFTP server.



Make sure you do not enter the IP address of another device on this network.

---

The following message appears:

```
Enter File Name [image.out]:
```

11. Enter the firmware image file name and press Enter.

The TFTP server uploads the firmware image file to the FortiGate unit and the following appears.

```
Save as Default firmware/Backup firmware/Run image without saving:
[D/B/R]
```

12. Type R.

The FortiGate image is installed to system memory and the FortiGate unit starts running the new firmware image, but with its current configuration.



You can test the new firmware image as required. When done testing, you can reboot the FortiGate unit, and the FortiGate unit will resume using the firmware that was running before you installed the test firmware.

## Upgrading the firmware - web-based manager

Installing firmware replaces your current antivirus and attack definitions, along with the definitions included with the firmware release you are installing. After you install new firmware, make sure that antivirus and attack definitions are up to date.



Always remember to back up your configuration before making any changes to the firmware.

---

### To upgrade the firmware

1. Log into the web-based manager as the admin administrative user.
2. Go to *System > Dashboard > Status*.
3. Under *System Information > Firmware Version*, select *Update*.
4. Type the path and filename of the firmware image file, or select *Browse* and locate the file.
5. Select *OK*.

The FortiGate unit uploads the firmware image file, upgrades to the new firmware version, restarts, and displays the FortiGate login. This process takes a few minutes.

## Upgrading the firmware - CLI

Installing firmware replaces your current antivirus and attack definitions, along with the definitions included with the firmware release you are installing. After you install new firmware, make sure that antivirus and attack definitions are up to date. You can also use the CLI command `execute update-now` to update the antivirus and attack definitions. For more information, see the [FortiGate Administration Guide](#).

Before you begin, ensure you have a TFTP server running and accessible to the FortiGate unit.



Always remember to back up your configuration before making any changes to the firmware.

---

### To upgrade the firmware using the CLI

1. Make sure the TFTP server is running.
2. Copy the new firmware image file to the root directory of the TFTP server.
3. Log into the CLI.
4. Make sure the FortiGate unit can connect to the TFTP server.

You can use the following command to ping the computer running the TFTP server. For example, if the IP address of the TFTP server is 192.168.1.168:

```
execute ping 192.168.1.168
```

5. Enter the following command to copy the firmware image from the TFTP server to the FortiGate unit:

```
execute restore image tftp <filename> <tftp_ipv4>
```

Where <name\_str> is the name of the firmware image file and <tftp\_ip4> is the IP address of the TFTP server. For example, if the firmware image file name is `image.out` and the IP address of the TFTP server is `192.168.1.168`, enter:

```
execute restore image tftp image.out 192.168.1.168
```

The FortiGate unit responds with the message:

```
This operation will replace the current firmware version!
```

```
Do you want to continue? (y/n)
```

6. Type `y`.
7. The FortiGate unit uploads the firmware image file, upgrades to the new firmware version, and restarts. This process takes a few minutes.
8. Reconnect to the CLI.
9. Update antivirus and attack definitions, by entering:

```
execute update-now
```

## Installing firmware from a system reboot using the CLI

There is a possibility that the firmware upgrade does not load properly and the FortiGate unit will not boot, or continuously reboots. If this occurs, it is best to perform a fresh install of the firmware from a reboot using the CLI.

This procedure installs a firmware image and resets the FortiGate unit to default settings. You can use this procedure to upgrade to a new firmware version, revert to an older firmware version, or re-install the current firmware.

To use this procedure, you must connect to the CLI using the FortiGate console port and a RJ-45 to DB-9, or null modem cable. This procedure reverts the FortiGate unit to its factory default configuration.

For this procedure you install a TFTP server that you can connect to from the FortiGate internal interface. The TFTP server should be on the same subnet as the internal interface.

Before beginning this procedure, ensure you back up the FortiGate unit configuration.

If you are reverting to a previous FortiOS version, you might not be able to restore the previous configuration from the backup configuration file.

Installing firmware replaces your current antivirus and attack definitions, along with the definitions included with the firmware release you are installing. After you install new firmware, make sure that antivirus and attack definitions are up to date.

### To install firmware from a system reboot

1. Connect to the CLI using the RJ-45 to DB-9 or null modem cable.
2. Make sure the TFTP server is running.
3. Copy the new firmware image file to the root directory of the TFTP server.
4. Make sure the internal interface is connected to the same network as the TFTP server.
5. To confirm the FortiGate unit can connect to the TFTP server, use the following command to ping the computer running the TFTP server. For example, if the IP address of the TFTP server is `192.168.1.168`:

```
execute ping 192.168.1.168
```

6. Enter the following command to restart the FortiGate unit.

```
execute reboot
```

The FortiGate unit responds with the following message:

```
This operation will reboot the system!
Do you want to continue? (y/n)
```

7. Type `y`.

As the FortiGate unit starts, a series of system startup messages appears. When the following messages appears:

```
Press any key to display configuration menu.....
```

Immediately press any key to interrupt the system startup.



You have only 3 seconds to press any key. If you do not press a key soon enough, the FortiGate unit reboots and you must log in and repeat the `execute reboot` command.

---

If you successfully interrupt the startup process, the following messages appears:

```
[G]: Get firmware image from TFTP server.
[F]: Format boot device.
[B]: Boot with backup firmware and set as default
[C]: Configuration and information
[Q]: Quit menu and continue to boot with default
firmware.
[H]: Display this list of options.
```

```
Enter G, F, Q, or H:
```

8. Type `G` to get to the new firmware image form the TFTP server.

The following message appears:

```
Enter TFTP server address [192.168.1.168]:
```

9. Type the address of the TFTP server and press `Enter`:

The following message appears:

```
Enter Local Address [192.168.1.188]:
```

10. Type an IP address the FortiGate unit can use to connect to the TFTP server. The IP address can be any IP address that is valid for the network the interface is connected to.



Make sure you do not enter the IP address of another device on this network.

---

The following message appears:

```
Enter File Name [image.out]:
```

11. Enter the firmware image filename and press `Enter`.

The TFTP server uploads the firmware image file to the FortiGate unit and a message similar to the following appears:

```
Save as Default firmware/Backup firmware/Run image without saving:
[D/B/R]
```

12 Type D.

The FortiGate unit installs the new firmware image and restarts. The installation might take a few minutes to complete.

## Reverting to a previous firmware version - web-based manager

The following procedures revert the FortiGate unit to its factory default configuration and deletes any configuration settings.

If you are reverting to a previous FortiOS version, you might not be able to restore the previous configuration from the backup configuration file.



Always remember to back up your configuration before making any changes to the firmware.

### To revert to a previous firmware version

1. Copy the firmware image file to the management computer.
2. Log into the FortiGate web-based manager.
3. Go to *System > Dashboard > Status*.
4. Under *System Information > Firmware Version*, select *Update*.
5. Type the path and filename of the firmware image file, or select *Browse* and locate the file.
6. Select *OK*.

The FortiGate unit uploads the firmware image file, reverts to the old firmware version, resets the configuration, restarts, and displays the FortiGate login. This process takes a few minutes.

7. Log into the web-based manager.
8. Restore your configuration.

For information about restoring your configuration see [“Restoring a configuration” on page 1445](#).

## Reverting to a previous firmware version - CLI

This procedure reverts the FortiGate unit to its factory default configuration and deletes IPS custom signatures, web content lists, email filtering lists, and changes to replacement messages.

Before beginning this procedure, it is recommended that you:

- back up the FortiGate unit system configuration using the command `execute backup config`
- back up the IPS custom signatures using the command `execute backup ipsuserdefsig`
- back up web content and email filtering lists

To use the following procedure, you must have a TFTP server the FortiGate unit can connect to.

### To revert to a previous firmware version using the CLI

1. Make sure the TFTP server is running
2. Copy the firmware image file to the root directory of the TFTP server.
3. Log into the FortiGate CLI.

4. Make sure the FortiGate unit can connect to the TFTP server execute by using the `execute ping` command.
5. Enter the following command to copy the firmware image from the TFTP server to the FortiGate unit:
 

```
execute restore image tftp <name_str> <tftp_ipv4>
```

 Where `<name_str>` is the name of the firmware image file and `<tftp_ipv4>` is the IP address of the TFTP server. For example, if the firmware image file name is `imagev28.out` and the IP address of the TFTP server is `192.168.1.168`, enter:
 

```
execute restore image tftp image28.out 192.168.1.168
```

 The FortiGate unit responds with this message:
 

```
This operation will replace the current firmware version!
Do you want to continue? (y/n)
```
6. Type `y`.
 The FortiGate unit uploads the firmware image file. After the file uploads, a message similar to the following appears:
 

```
Get image from tftp server OK.
Check image OK.
This operation will downgrade the current firmware version!
Do you want to continue? (y/n)
```
7. Type `y`.
8. The FortiGate unit reverts to the old firmware version, resets the configuration to factory defaults, and restarts. This process takes a few minutes.
9. Reconnect to the CLI.
10. To restore your previous configuration, if needed, use the command:
 

```
execute restore config <name_str> <tftp_ip4>
```
11. Update antivirus and attack definitions using the command:
 

```
execute update-now.
```

## Configuration Revision

The *Configuration Revisions* menu enables you to manage multiple versions of configuration files on models that have a 512 flash memory and higher. Revision control requires either a configured central management server or the local hard drive. The central management server can either be a FortiManager unit or FortiCloud.

If central management is not configured on your FortiGate unit, a message appears to tell you to do one of the following:

- enable central management (see [Central management](#))
- obtain a valid license.

When revision control is enabled on your FortiGate unit, and configurations backups have been made, a list of saved revisions of those backed-up configurations appears.

Configuration revisions are viewed in the *System Information* widget on the Dashboard.

## Backup and Restore from a USB key

Use a USB key to either backup a configuration file or restore a configuration file. You should always make sure a USB key is properly install before proceeding since the FortiGate unit must recognize that the key is installed in its USB port.

You can only save VPN certificates if you encrypt the file. Make sure the configuration encryption is enabled so you can save the VPN certificates with the configuration file. An encrypted file is ineffective if selected for the USB Auto-Install feature.

### To backup configuration using the CLI

1. Log into the CLI.
2. Enter the following command to backup the configuration files:  

```
exec backup config usb <filename>
```
3. Enter the following command to check the configuration files are on the key:  

```
exec usb-disk list
```

### To restore configuration using the CLI

1. Log into the CLI.
2. Enter the following command to restore the configuration files:  

```
exec restore image usb <filename>
```

The FortiGate unit responds with the following message:  
This operation will replace the current firmware version!  
Do you want to continue? (y/n)
3. Type `y`.
- 4.

## Backup and Restore an encrypted config file from a USB key

You can save and boot an encrypted configuration file from a USB key. The configuration will only load when rebooting the FortiGate unit with the USB key inserted.

The administrator must back up the configuration to the USB key using the command:

```
execute backup config usb-mode <password>
```

administrator backup the configuration, to the USB key ("exec backup config usb-mode")

Insert the USB key into any FortiGate unit running the same image/patch release as the FortiGate unit that created the configuration file

The Administrator runs the CLI command below to reboot the FortiGate unit and load the configuration file from the USB key:

```
execute restore config usb-mode <password>
```

The FortiGate unit saves the password into the flash memory. When system boots, the FortiGate unit loads the configurations from the USB key using the saved password in the flash. This configuration is read-only. That is, no configuration changes can be made while running with the configuration. The administrator is not permitted to make any configuration changes while configurations are loaded from USB (read-only)

If the USB key is removed while the FortiGate unit is running, the FortiGate unit deletes the password from the flash memory and reboots.

## Controlled upgrade

Using a controlled upgrade, you can upload a new version of the FortiOS firmware to a separate partition in the FortiGate memory for later upgrade. The FortiGate unit can also be configured so that when it is rebooted, it will automatically load the new firmware (CLI only). Using this option, you can stage a number of FortiGate units to do an upgrade simultaneously to all devices using FortiManager or script.

### To load the firmware for later installation - web-based manager

1. Go to *System > Dashboard > Status*.
2. Under *System Information > Firmware Version*, select *Update*.
3. Type the path and filename of the firmware image file, or select *Browse* and locate the file.
4. Deselect the *Boot the New Firmware* option
5. Select *OK*.

### To load the firmware for later installation - CLI

```
execute restore secondary-image {ftp | tftp | usb}
```

To set the FortiGate unit so that when it reboots, the new firmware is loaded, use the CLI command...

```
execute set-next-reboot {primary | secondary}
```

... where {primary | secondary} is the partition with the preloaded firmware.

### To trigger the upgrade using the web-based manager

1. Go to *System > Dashboard > Status*.
2. Under *System Information > Firmware Version*, select *Details*.
3. Select the check box for the new firmware version.  
The *Comments* column indicates which firmware version is the current active version.
4. Select *Upgrade* icon.

# Best practices

The FortiGate unit is installed, and traffic is moving. With your network sufficiently protected, you can now fine-tune the firewall for the best performance and efficiency. This chapter describes configuration options that can ensure your FortiGate unit is running at its best performance.

This section includes the topics on:

- [Hardware](#)
- [Shutting down](#)
- [Performance](#)
- [Firewall](#)
- [Intrusion protection](#)
- [Antivirus](#)
- [Web filtering](#)
- [Antispam](#)

## Hardware

### Environmental specifications

Keep the following environmental specifications in mind when installing and setting up your FortiGate unit.

- Operating temperature: 32 to 104°F (0 to 40°C) (temperatures may vary, depending on the FortiGate model)
- If you install the FortiGate unit in a closed or multi-unit rack assembly, the operating ambient temperature of the rack environment may be greater than room ambient temperature. Therefore, make sure to install the equipment in an environment compatible with the manufacturer's maximum rated ambient temperature.
- Storage temperature: -13 to 158°F (-25 to 70°C) (temperatures may vary, depending on the FortiGate model)
- Humidity: 5 to 90% non-condensing
- Air flow - For rack installation, make sure that the amount of air flow required for safe operation of the equipment is not compromised.
- For free-standing installation, make sure that the appliance has at least 1.5 in. (3.75 cm) of clearance on each side to allow for adequate air flow and cooling.

This device complies with part FCC Class A, Part 15, UL/CUL, C Tick, CE and VCCI. Operation is subject to the following two conditions:

- This device may not cause harmful interference, and
- This device must accept any interference received, including interference that may cause undesired operation.

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with



the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio/TV technician for help.

The equipment compliance with FCC radiation exposure limit set forth for uncontrolled Environment.



Risk of Explosion if battery is replaced by an incorrect type. Dispose of used batteries according to the instructions.

To reduce the risk of fire, use only No. 26 AWG or larger UL Listed or CSA Certified Telecommunication Line Cord

---

## Grounding

- Ensure the FortiGate unit is connected and properly grounded to a lightning and surge protector. WAN or LAN connections that enter the premises from outside the building should be connected to an Ethernet CAT5 (10/100 Mb/s) surge protector.
- Shielded Twisted Pair (STP) Ethernet cables should be used whenever possible rather than Unshielded Twisted Pair (UTP).
- Do not connect or disconnect cables during lightning activity to avoid damage to the FortiGate unit or personal injury.

## Rack mount instructions

**Elevated Operating Ambient** - If installed in a closed or multi-unit rack assembly, the operating ambient temperature of the rack environment may be greater than room ambient. Therefore, consideration should be given to installing the equipment in an environment compatible with the maximum ambient temperature (T<sub>ma</sub>) specified by the manufacturer.

**Reduced Air Flow** - Installation of the equipment in a rack should be such that the amount of air flow required for safe operation of the equipment is not compromised.

**Mechanical Loading** - Mounting of the equipment in the rack should be such that a hazardous condition is not achieved due to uneven mechanical loading.

**Circuit Overloading** - Consideration should be given to the connection of the equipment to the supply circuit and the effect that overloading of the circuits might have on overcurrent protection and supply wiring. Appropriate consideration of equipment nameplate ratings should be used when addressing this concern.

**Reliable Earthing** - Reliable earthing of rack-mounted equipment should be maintained.

Particular attention should be given to supply connections other than direct connections to the branch circuit (e.g. use of power strips).

## Shutting down

Always shut down the FortiGate operating system properly before turning off the power switch to avoid potential hardware problems.

### To power off the FortiGate unit - web-based manager

1. Go to *System > Status*.
2. In the *System Resources* widget, select *Shutdown*.

### To power off the FortiGate unit - CLI

```
execute shutdown
```

Once this has been done, you can safely turn off the power switch or disconnect the power cables from the power supply.

## Performance

- Disable any management features you do not need. If you don't need SSH or SNMP, disable them. SSH also provides another possibility for would-be hackers to infiltrate your FortiGate unit.
- Put the most used firewall rules to the top of the interface list.
- Log only necessary traffic. The writing of logs, especially if to an internal hard disk, slows down performance.
- Enable only the required application inspections.
- Keep alert systems to a minimum. If you send logs to a syslog server, you may not need SNMP or email alerts, making for redundant processing.
- Establish scheduled FortiGuard updates at a reasonable rate. Daily updates occurring every 4-5 hours are sufficient for most situations. In more heavy-traffic situations, schedule updates for the evening when more bandwidth can be available.
- Keep security profiles to a minimum. If you do not need a profile on a firewall rule, do not include it.
- Keep VDOMs to a minimum. On low-end FortiGate units, avoid using them if possible.
- Avoid traffic shaping if you need maximum performance. Traffic shaping, by definition, slows down traffic.

## Firewall

- Avoid using the *All* selection for the source and destination addresses. Use addresses or address groups.
- Avoid using *Any* for the services.
- Use logging on a policy only when necessary and be aware of the performance impact. For example, you may want to log all dropped connections but can choose to use this sparingly

by sampling traffic data rather than have it continually storing log information you may not use.

- Use the comment field to input management data, for example: who requested the rule, who authorized it, etc.
- Avoid FQDN addresses if possible, unless they are internal. It can cause a performance impact on DNS queries and security impact from DNS spoofing.
- If possible, avoid port ranges on services for security reasons.
- Use groups whenever possible.
- To ensure that all AV push updates occur, ensure you have an AV profile enabled in a security policy.

## Intrusion protection

- Create and use security profiles with specific signatures and anomalies you need per-interface and per-rule.
- Do not use predefined or generic profiles. While these profiles are convenient to supply immediate protection, you should create profiles to suit your network environment.
- If you do use the default profiles, reduce the IPS signatures/anomalies enabled in the profile to conserve processing time and memory.
- If you are going to enable anomalies, make sure you tune thresholds according to your environment.
- If you need protection, but not audit information, disable the logging option.
- Tune the IP-protocol parameter accordingly.

## Antivirus

- Enable only the protocols you need to scan. If you have antivirus scans occurring on the SMTP server, or use FortiMail, it is redundant to have scanning occur on the FortiGate unit as well.
- Reduce the maximum file size to be scanned. Viruses usually travel in small files of around 1 to 2 megabytes.
- Antivirus scanning within an HA cluster can impact performance.
- Enable grayware scanning on security profiles tied to Internet browsing.
- Do not quarantine files unless you regularly monitor and review them. This is otherwise a waste of space and impacts performance.
- Use file patterns to avoid scanning where it is not required.
- Enable heuristics from the CLI if high security is required using the command `config antivirus heuristic`.

## Web filtering

- Web filtering within an HA cluster impacts performance.
- Always review the DNS settings to ensure the servers are fast.
- Content blocking may cause performance overhead.
- Local URL filters are faster than FortiGuard web filters, because the filter list is local and the FortiGate unit does not need to go out to the Internet to get the information from a FortiGuard web server.

## Antispam

- If possible use, a FortiMail unit. The antispam engines are more robust.
- Use fast DNS servers.
- Use specific security profiles for the rule that will use antispam.
- DNS checks may cause false positive with HELO DNS lookup.
- Content analysis (banned words) may impose performance overhead.

## Security

- Use NTP to synchronize time on the FortiGate and the core network systems, such as email servers, web servers, and logging services.
- Enable log rules to match corporate policy. For example, log administration authentication events and access to systems from untrusted interfaces.
- Minimize adhoc changes to live systems, if possible, to minimize interruptions to the network. When not possible, create backup configurations and implement sound audit systems using FortiAnalyzer and FortiManager.
- If you only need to allow access to a system on a specific port, limit the access by creating the strictest rule possible.

# FortiGuard

The FortiGuard Distribution Network (FDN) of servers provides updates to antivirus, antispam and IPS definitions to your FortiGate unit. Worldwide coverage of FortiGuard services is provided by FortiGuard service points. FortiGuard Subscription Services provide comprehensive Unified Threat Management (UTM) security solutions to enable protection against content and network level threats.

The FortiGuard team can be found around the globe, monitoring virus, spyware and vulnerability activities. As vulnerabilities are found, signatures are created and pushed to the subscribed FortiGate units. The Global Threat Research Team enables Fortinet to deliver a combination of multi-layered security intelligence and provide true zero-day protection from new and emerging threats. The FortiGuard Network has data centers around the world located in secure, high availability locations that automatically deliver updates to the Fortinet security platforms to and protect the network with the most up-to-date information.

To ensure optimal response and updates, the FortiGate unit will contact a FortiGuard service point closest to the FortiGate installation, using the configured time zone information. FortiGuard services are continuously updated year-round, 24x7.

Every FortiGate unit includes a free 30-day FortiGuard trial license. FortiGuard license management is performed by Fortinet servers. The FortiGate unit automatically contacts a FortiGuard service point when enabling FortiGuard services. Contact Fortinet Technical Support to renew a FortiGuard license after the free trial.

This section includes the topics:

- [FortiGuard Services](#)
- [Antivirus and IPS](#)
- [Web filtering](#)
- [Email filtering](#)
- [Security tools](#)
- [Troubleshooting](#)

## FortiGuard Services

The FortiGuard services provide a number of services to monitor world-wide activity and provide the best possible security.

### Next Generation Firewall

The Next Generation Firewall (NGFW) offers integrated, high-performance protection against today's wide range of advanced threats targeting your applications, data, and users.

NGFW services include:

- **Intrusion Prevention System (IPS)**- The FortiGuard Intrusion Prevention System (IPS) uses a customizable database of more than 4000 known threats to stop attacks that evade conventional firewall defenses. It also provides behavior-based heuristics, enabling the

system to recognize threats when no signature has yet been developed. It also provides more than 1000 application identity signatures for complete application control.

- **Application Control** - Application Control allows you to identify and control applications on networks and endpoints regardless of port, protocol, and IP address used. It gives you unmatched visibility and control over application traffic, even traffic from unknown applications and sources.

## Advanced Threat Protection

Advanced Threat Protect (ATP) provides protection against a variety of threats, including Advanced Targeted Attacks (ATA), also known as Advanced Persistent Threats (APT).

ATP services include:

- **Antivirus** -The FortiGuard Antivirus Service provides fully automated updates to ensure protection against the latest content level threats. It employs advanced virus, spyware, and heuristic detection engines to prevent both new and evolving threats from gaining access to your network and protects against vulnerabilities.
- **Web Filtering** - Web Filtering provides Web URL filtering to block access to harmful, inappropriate, and dangerous web sites that may contain phishing/pharming attacks, malware such as spyware, or objectionable content that can expose your organization to legal liability. Based on automatic research tools and targeted research analysis, real-time updates enable you to apply highly-granular policies that filter web access based on 78 web content categories, over 45 million rated web sites, and more than two billion web pages - all continuously updated.

## Other Services

FortiGuard provides a number of additional services, including:

- **Vulnerability Scanning** - FortiGuard Services provide comprehensive and continuous updates for vulnerabilities, remediation, patch scan, and configuration benchmarks.
- **Email Filtering** - The FortiGuard Antispam Service uses both a sender IP reputation database and a spam signature database, along with sophisticated spam filtering tools on Fortinet appliances and agents, to detect and block a wide range of spam messages. Updates to the IP reputation and spam signature databases are provided continuously via the FDN.
- **Messaging Services** - Messaging Services allow a secure email server to be automatically enabled on your FortiGate unit to send alert email or send email authentication tokens. With the SMS gateway, you can enter phone numbers where the FortiGate unit will send the SMS messages.  
Note that depending on your carrier, there may be a slight time delay on receiving messages.
- **DNS and DDNS** - The FortiGuard DNS and DDNS services provide an efficient method of DNS lookups once subscribed to the FortiGuard network. This is the default option. The FortiGate unit connects automatically to the FortiGuard DNS server. If you do not register, you need to configure an alternate DNS server.

Configure the DDNS server settings using the CLI commands:

```
config system fortiguard
 set ddns-server-ip
 set ddns-server-port
end
```

## Support Contract and FortiGuard Subscription Services

The *Support Contract* and *FortiGuard Subscription Services* sections are displayed in abbreviated form within the *License Information* widget. A detailed version is available by going to *System > Config > FortiGuard*.

The Support Contract area displays the availability or status of your FortiGate unit's support contract. The status displays can be either *Unreachable*, *Not Registered*, or *Valid Contract*.

The FortiGuard Subscription Services area displays detailed information about your FortiGate unit's support contract and FortiGuard subscription services. On this page, you can also manually update the antivirus and IPS engines.

The status icons for each section indicates the state of the subscription service. The icon corresponds to the availability description.

- **Gray (Unreachable)** – the FortiGate unit is not able to connect to service.
- **Orange (Not Registered)** – the FortiGate unit can connect, but not subscribed.
- **Yellow (Expired)** – the FortiGate unit had a valid license that has expired.
- **Green (Valid license)** – the FortiGate unit can connect to FDN and has a registered support contract. If the Status icon is green, the expiry date also appears.

## FortiCloud

FortiCloud is a hosted security management and log retention service for FortiGate products. It gives you a centralized reporting, traffic analysis, configuration and log retention without the need for additional hardware and software.

A subscription to FortiCloud also includes Cloud Sandbox, a service in which suspicious files can be inspected in isolation from your network.

For more information about FortiCloud, see [“FortiCloud” on page 1474](#).

## Antivirus and IPS

The FortiGuard network is an always updating service, including grayware and signatures for application control. There are two methods of updating the virus and IPS signatures on your FortiGate unit: manually or through push updates.

### Detection during update

During an update, the FortiGate unit will continue to detect to scan network traffic. Sessions occurring right before an update will be scanned using the current signatures. Sessions that occur during the update, when the signature database is reloading, will be on hold until the signatures load, at which point the new signatures are used to scan these sessions. Sessions occur i ng right after the update will also use the new signatures.

## Antivirus and IPS Options

Go to *System > Config > FortiGuard*, and expand the *AV and IPS Options* section to configure the antivirus and IPS options for connecting and downloading definition files.

---

<b>Use override server address</b>	Select to configure an override server if you cannot connect to the FDN or if your organization provides updates using their own FortiGuard server.
<b>Allow Push Update</b>	Select to allow updates sent automatically to your FortiGate unit when they are available
<b>Allow Push Update status icon</b>	The status of the FortiGate unit for receiving push updates: <ul style="list-style-type: none"><li>• <b>Gray (Unreachable)</b> - the FortiGate unit is not able to connect to push update service</li><li>• <b>Yellow (Not Available)</b> - the push update service is not available with your current support license</li><li>• <b>Green (Available)</b> - the push update service is allowed.</li></ul>
<b>Use override push IP and Port</b>	Available only if both <i>Use override server address</i> and <i>Allow Push Update</i> are enabled.  Enter the IP address and port of the NAT device in front of your FortiGate unit. FDS will connect to this device when attempting to reach the FortiGate unit.  The NAT device must be configured to forward the FDS traffic to the FortiGate unit on UDP port 9443.
<b>Schedule Updates</b>	Select this check box to enable updates to be sent to your FortiGate unit at a specific time. For example, to minimize traffic lag times, you can schedule the update to occur on weekends or after work hours.  Note that a schedule of once a week means any urgent updates will not be pushed until the scheduled time. However, if there is an urgent update required, select the <i>Update Now</i> button.
<b>Update Now</b>	Select to manually initiate an FDN update.
<b>Submit attack characteristics... (recommended)</b>	Select to help Fortinet maintain and improve IPS signatures. The information sent to the FortiGuard servers when an attack occurs and can be used to keep the database current as variants of attacks evolve.

---

## Manual updates

To manually update the signature definitions file, you need to first go to the Support web site at <https://support.fortinet.com>. Once logged in, select FortiGuard Service Updates from the Download area of the web page. The browser will present you the most current antivirus and IPS signature definitions which you can download.

Once downloaded to your computer, log into the FortiGate unit to load the definition file.



### To load the definition file onto the FortiGate unit

1. Go to *System > Config > FortiGuard*.
2. Select the *Update* link for either *AV Definitions* or *IPS Definitions*.
3. Locate the downloaded file and select *OK*.

The upload may take a few minutes to complete.

## Automatic updates

The FortiGate unit can be configured to request updates from the FortiGuard Distribution Network. You can configure this to be on a scheduled basis, or with push notifications.

### Scheduling updates

Scheduling updates ensures that the virus and IPS definitions are downloaded to your FortiGate unit on a regular basis, ensuring that you do not forget to check for the definition files yourself. As well, by scheduling updates during off-peak hours, such as evenings or weekends, when network usage is minimal, ensures that the network activity will not suffer from the added traffic of downloading the definition files.

If you require the most up-to-date definitions as viruses and intrusions are found in the wild, the FortiGuard Distribution Network can push updates to the FortiGate units as they are developed. This ensures that your network will be protected from any breakouts of a virus within the shortest amount of time, minimizing any damaging effect that can occur. Push updates require that you have registered your FortiGate unit.

Once push updates are enabled, the next time new antivirus or IPS attack definitions are released, the FDN notifies all the FortiGate unit that a new update is available. Within 60 seconds of receiving a push notification, the unit automatically requests the update from the FortiGuard servers.

### To enable scheduled updates - web-based manager

1. Go to *System > Config > FortiGuard*.
2. Click the *Expand Arrow* for *AV and IPS Options*.
3. Select the *Scheduled Update* check box.
4. Select the frequency of the updates and when within that frequency.
5. Select *Apply*.

### To enable scheduled updates - CLI

```
config system autoupdate schedule
 set status enable
 set frequency {every | daily | weekly}
 set time <hh:mm>
 set day <day_of_week>
end
```

## Push updates

Push updates enable you to get immediate updates when new virus or intrusions have been discovered and new signatures are created. This ensures that when the latest signature is available it will be sent to the FortiGate.

When a push notification occurs, the FortiGuard server sends a notice to the FortiGate unit that there is a new signature definition file available. The FortiGate unit then initiates a download of the definition file, similar to the scheduled update.

To ensure maximum security for your network, you should have a scheduled update as well as enable the push update, in case an urgent signature is created, and your cycle of the updates only occurs weekly.

#### To enable push updates - web-based manager

1. Got to *System > Config > FortiGuard*.
2. Click the Expand Arrow for *AV and IPS Options*.
3. Select *Allow Push Update*.
4. Select *Apply*.

#### To enable push updates - CLI

```
config system autoupdate push-update
 set status enable
end
```

### Push IP override


If the FortiGate unit is behind another NAT device (or another FortiGate unit), to ensure it receives the push update notifications, you need to use an override IP address for the notifications. To do this, you create a virtual IP to map to the external port of the NAT device.

Generally speaking, if there are two FortiGate devices as in the diagram below, the following steps need to be completed on the FortiGate NAT device to ensure the FortiGate unit on the internal network receives the updates:

- Add a port forwarding virtual IP to the FortiGate NAT device that connects to the Internet by going to *Firewall Objects > Virtual IP*.
- Add a security policy to the FortiGate NAT device that connects to the Internet that includes the port forwarding virtual IP.
- Configure the FortiGate unit on the internal network with an override push IP and port.

On the FortiGate internal device, the virtual IP is entered as the *Use push override IP* address.

**Figure 231:**Using a virtual IP for a FortiGate unit behind a NAT device



#### To enable push update override- web-based manager

1. Got to *System > Config > FortiGuard*.
2. Click the Expand Arrow for *AV and IPS Options*.
3. Select *Allow Push Update*.
4. Select *Use push override IP*.
5. Enter the virtual IP address configured on the NAT device.
6. Select *Apply*.

#### To enable push updates - CLI

```
config system autoupdate push-update
 set status enable
 set override enable
 set address <vip_address>
end
```

## Web filtering

The multiple FortiGuard data centers around the world hold the entire categorized URL database and receive rating requests from customer-owned FortiGate units, typically triggered by browser-based URL requests. When these rating requests are responded to with the categories stored for specific URLs, the requesting FortiGate unit will then use its own local profile configuration to determine what action to take, for example blocking, monitoring, or permitting the URL request.

Rating responses can also be cached locally on the FortiGate unit, providing a quicker response time while easing load on the FortiGuard servers and aiding in a quicker response time for less common URL requests. This is a very effective method for common sites such as search

engines and other frequently visited sites. Other sites that are less frequently visited can also be cached locally for a determined amount of time.

By default, the web filtering cache is enabled. The cache includes a time-to-live value, which is the amount of time a URL will stay in the cache before expiring. You can change this value to shorten or extend the time between 300 and 86400 seconds. For a site such as Google, the frequency of its access can keep it in the cache, while other sites can remain in the cache up to 24 hours, or less depending on the configuration.

## Web Filtering and Email Filtering Options

Go to *System > Config > FortiGuard*, and expand arrow to view *Web Filtering and Email Filtering Options* for setting the size of the caches and ports used.

---

<b>Web Filter cache TTL</b>	Set the Time To Live value. This is the number of seconds the FortiGate unit will store a blocked IP or URL locally, saving time and network access traffic, checking the FortiGuard server. Once the TTL has expired, the FortiGate unit will contact an FDN server to verify a web address. The TTL must be between 300 and 86400 seconds.
<b>Antispam cache TTL</b>	Set the Time To Live value. This is the number of seconds the FortiGate unit will store a blocked IP or URL locally, saving time and network access traffic, checking the FortiGuard server. Once the TTL has expired, the FortiGate unit will contact an FDN server to verify a web address. The TTL must be between 300 and 86400 seconds.
<b>Port Section</b>	Select the port assignments for contacting the FortiGuard servers. Select the <i>Test Availability</i> button to verify the connection using the selected port.
<b>To have a URL's category rating re-evaluated, please click here</b>	Select to re-evaluate a URL's category rating on the FortiGuard Web Filter service.

---

## URL verification

If you discover a URL - yours or one you require access to has been incorrectly flagged as an inappropriate site - you can ask the FortiGuard team to re-evaluate the site. To do this, go to *System > Config > FortiGuard*, select the blue arrow for *Web Filtering and Email Filtering Options* and select the link for re-evaluation.

### To modify the web filter cache size - web-based manager

1. Got to *System > Config > FortiGuard*.
2. Click the Expand Arrow for *Web Filtering and Email Filtering Options*.
3. Enter the TTL value for the *Web filter cache*.
4. Select *Apply*.

### To modify the web filter cache size - CLI

```
config system fortiguard
 set webfilter-cache-ttl <integer>
end
```

Further web filtering options can be configured to block specific URLs, and allow others through. These configurations are available through the *Security Profiles > Web Filter* menu. For more information, see [Security](#) chapter of The Handbook.

## Email filtering

The FortiGuard data centers monitor and update email databases of known spam sources. With FortiGuard Antispam enabled, the FortiGate unit verifies incoming email sender address and IPs against the database, and take the necessary action as defined within the antivirus profiles.

Spam source IP addresses can also be cached locally on the FortiGate unit, providing a quicker response time, while easing load on the FortiGuard servers, aiding in a quicker response time for less common email address requests.

By default, the antispam cache is enabled. The cache includes a time-to-live value, which is the amount of time an email address will stay in the cache before expiring. You can change this value to shorten or extend the time between 300 and 86400 seconds.

### To modify the antispam filter cache size - web-based manager

1. Got to *System > Config > FortiGuard*.
2. Click the Expand Arrow for *Web Filtering and Email Filtering Options*.
3. Enter the TTL value for the *antispam cache*.
4. Select *Apply*.

### To modify the web filter cache size - CLI

```
config system fortiguard
 set antispam-cache-ttl <integer>
end
```

Further antispam filtering options can be configured to block, allow or quarantine, specific email addresses. These configurations are available through the *Security Profiles > Antispam* menu. For more information, see [Security Profiles](#) chapter of The Handbook.

## Security tools

The FortiGuard online center provides a number of online security tools that enable you to verify or check ratings of web sites, email addresses as well as check file for viruses. These features are available at <http://www.fortiguard.com>.

### URL lookup

By entering a web site address, you can see if it has been rated and what category and classification it is filed as. If you find your web site or a site you commonly go to has been wrongly categorized, you can use this page to request that the site be re-evaluated.

<http://www.fortiguard.com/webfiltering/webfiltering.html>

### IP and signature lookup

The IP and signature lookup enables you to check whether an IP address is blacklisted in the FortiGuard IP reputation database or whether a URL or email address is in the signature database.

<http://www.fortiguard.com/antispam/antispam.html>

## Online virus scanner

If you discover a suspicious file on your machine, or suspect that a program you downloaded from the Internet might be malicious you can scan it using the FortiGuard online scanner. The questionable file can be uploaded from your computer to a dedicated server where it will be scanned using FortiClient Antivirus. Only one file of up to 1 MB can be checked at any one time. All files will be forwarded to our research labs for analysis.

[http://www.fortiguard.com/antivirus/virus\\_scanner.html](http://www.fortiguard.com/antivirus/virus_scanner.html)

## Malware removal tools

Tools have been developed by FortiGuard Labs to disable and remove the specific malware and related variants. Some tools have been developed to remove specific malware, often tough to remove. A universal cleaning tool, FortiCleanup, is also available for download.

The FortiCleanup is a tool developed to identify and cleanse systems of malicious rootkit files and their associated malware. Rootkits consist of code installed on a system with kernel level privileges, often used to hide malicious files, keylog and thwart detection / security techniques. The aim of this tool is to reduce the effectiveness of such malware by finding and eliminating rootkits. The tool offers a quick memory scan as well as a full system scan. FortiCleanup will not only remove malicious files, but also can cleanse registry entries, kernel module patches, and other tricks commonly used by rootkits - such as SSDT hooks and process enumeration hiding.

A license to use these applications is provided free of charge, courtesy of Fortinet.

[http://www.fortiguard.com/antivirus/malware\\_removal.html](http://www.fortiguard.com/antivirus/malware_removal.html)

## FortiSandbox

A FortiSandbox unit can be used for automated sample tracking, or sandboxing, for files from a FortiGate unit. This allows suspicious files to be sent to be inspected without risking network security. If the file exhibits risky behavior, or is found to contain a virus, a new virus signature is created and added to the FortiGuard antivirus signature database.

Cloud Sandbox can also be used for sandboxing if you have an active FortiCloud subscription. For more information, see [“FortiCloud” on page 1474](#).

## Troubleshooting

If you are not getting FortiGuard web filtering or antispam services, there are a few things to verify communication to the FortiGuard Distribution Network (FDN) is working. Before any troubleshooting, ensure that the FortiGate unit has been registered and you or your company, has subscribed to the FortiGuard services.

## Web-based manager verification

The simplest method to check that the FortiGate unit is communicating with the FDN, is to check the *License Information* dashboard widget. Any subscribed services should have a green check mark beside them indicating that connections are successful. Any other icon indicates a problem with the connection, or you are not subscribed to the FortiGuard services.

**Figure 232:**License Information widget showing FortiGuard availability

License Information		
<b>Support Contract</b>		
Registration	Registered (Login: [redacted]) [Login Now]	✓
Hardware	8 x 5 support (Expired: 2012-11-24) [Renew]	✗
Firmware	8 x 5 support (Expired: 2012-11-24) [Renew]	✗
Enhanced Support	24 x 7 support (Expired: 2012-11-24) [Renew]	✗
Comprehensive Support	24 x 7 support (Expired: 2012-11-24) [Renew]	✗
<b>FortiGuard Services</b>		
AntiVirus	Expired [Renew]	✗
IPS	Expired [Renew]	✗
Vulnerability Scan	Expired [Renew]	✗
Web Filtering	Expired [Renew]	✗
Email Filtering	Expired [Renew]	✗
<b>FortiCloud</b>		
Account	<a href="#">Activate</a>	
<b>SMS</b>		
Status	Unreachable	✗
<b>FortiToken Mobile</b>		
Registered/Allowed	0 of 0	
<b>FortiClient Software</b>		
Registered/Allowed	[Mac] [Windows]	
Registered/Allowed	0 of 10 [Details]	

You can also view the FortiGuard connection status by going to *System > Config > FortiGuard*.

**Figure 233:**FortiGuard availability

<b>Support Contract</b>		
Registration	Registered (Login ID: [redacted]) [Login Now]	✓
Hardware	8 x 5 support (Expires: 2012-11-26)	✓
Firmware	8 x 5 support (Expires: 2012-11-26)	✓
Enhanced Support	24 x 7 support (Expires: 2012-11-26)	✓
Comprehensive Support	24 x 7 support (Expires: 2012-11-26)	✓
<b>FortiGuard Subscription Services</b>		
AntiVirus	Valid License (Expires 2012-11-26)	✓
AV Definitions	14.00000 (Updated 2011-08-24 via Manual Update) [Update]	
AV Engine	4.00382 (Updated 2011-10-28 via Manual Update)	
Intrusion Protection	Valid License (Expires 2012-11-26)	✓
IPS Definitions	3.00097 (Updated 2011-10-28 via Manual Update) [Update]	
IPS Engine	1.00241 (Updated 2011-10-28 via Manual Update)	
Web Filtering	Not Registered	✗
Email Filtering	Not Registered	✗
Vulnerability Management	Valid License (Expires 2012-11-26)	✓
VCM Plugin	1.00238 (Updated 2011-11-25 via Manual Update) [Update]	
Analysis & Management Service	Expired [Renew] [Update]	✗
<b>FortiToken Seed Server</b>		
Registration	Reachable (0 Tokens Registered)	✓

- ▶ AntiVirus and IPS Options
- ▶ Web Filtering and Email Filtering Options
- ▶ FortiGuard Analysis & Management Service Options

## CLI verification

You can also use the CLI to see what FortiGuard servers are available to your FortiGate unit. Use the following CLI command to ping the FDN for a connection:

```
ping guard.fortinet.net
```

You can also use diagnose command to find out what FortiGuard servers are available:

```
diagnose debug rating
```

From this command, you will see output similar to the following:

```
Locale : english
License : Contract
Expiration : Sun Jul 24 20:00:00 2011
Hostname : service.fortiguard.net

--- Server List (Tue Nov 2 11:12:28 2010) ---

IP Weight RTT Flags TZ Packets Curr Lost Total Lost
69.20.236.180 0 10 -5 77200 0 0 42
69.20.236.179 0 12 -5 52514 0 0 34
66.117.56.42 0 32 -5 34390 0 0 62
80.85.69.38 50 164 0 34430 0 0 11763
208.91.112.194 81 223 D -8 42530 0 0 8129
216.156.209.26 286 241 DI -8 55602 0 0 21555
```

An extensive list of servers are available. Should you see a list of three to five available servers, the FortiGuard servers are responding to DNS replies to service.FortiGuard.net, but the INIT requests are not reaching FDS services on the servers.

The rating flags indicate the server status:

<b>D</b>	Indicates the server was found via the DNS lookup of the hostname. If the hostname returns more than one IP address, all of them will be flagged with 'D' and will be used first for INIT requests before falling back to the other servers.
<b>I</b>	Indicates the server to which the last INIT request was sent
<b>F</b>	The server has not responded to requests and is considered to have failed.
<b>T</b>	The server is currently being timed.

The server list is sorted first by weight and then the server with the smallest RTT is put at the top of the list, regardless of weight. When a packet is lost, it will be resent to the next server in the list.

The weight for each server increases with failed packets and decreases with successful packets. To lower the possibility of using a faraway server, the weight is not allowed to dip below a base weight, which is calculated as the difference in hours between the FortiGate unit and the server multiplied by 10. The further away the server is, the higher its base weight and the lower in the list it will appear.

## Port assignment

FortiGate units contact the FortiGuard Distribution Network (FDN) for the latest list of FDN servers by sending UDP packets with typical source ports of 1027 or 1031, and destination ports of 53 or 8888. The FDN reply packets have a destination port of 1027 or 1031.



If your ISP blocks UDP packets in this port range, the FortiGate unit cannot receive the FDN reply packets. As a result, the FortiGate unit will not receive the complete FDN server list.

If your ISP blocks the lower range of UDP ports (around 1024), you can configure your FortiGate unit to use higher-numbered ports, using the CLI command...

```
config system global
 set ip-src-port-range <start port>-<end port>
end
```

...where the <start port> and <end port> are numbers ranging of 1024 to 25000.

For example, you could configure the FortiGate unit to not use ports lower than 2048 or ports higher than the following range:

```
config system global
 set ip-src-port-range 2048-20000
end
```

Trial and error may be required to select the best source port range. You can also contact your ISP to determine the best range to use. Push updates might be unavailable if:

- there is a NAT device installed between the unit and the FDN
- your unit connects to the Internet using a proxy server.

# FortiCloud

FortiCloud is a hosted security management and log retention service for FortiGate products. It gives you a centralized reporting, traffic analysis, configuration and log retention without the need for additional hardware and software.

## FortiCloud Features

FortiCloud offers a wide range of features:

### **Simplified central management for your FortiGate network**

FortiCloud provides a central web-based management console to manage individual or aggregated FortiGate and FortiWifi devices. Adding a device to the FortiCloud management subscription is straightforward and provides detailed traffic and application visibility across the whole network.

### **Hosted log retention with large default storage allocated**

Log retention is an integral part of any security and compliance program but administering a separate storage system is burdensome. FortiCloud takes care of this automatically and stores the valuable log information in the cloud. Each device is allowed up to 200Gb of log retention storage. Different types of logs can be stored including Traffic, System Events, Web, Applications and Security Events.

### **Monitoring and alerting in real time**

Network availability is critical to a good end-user experience. FortiCloud enables you to monitor your FortiGate network in real time with different alerting mechanisms to pinpoint potential issues. Alerting mechanisms can be delivered via email.

### **Customized or pre-configured reporting and analysis tools**

Reporting and analysis are your eyes and ears into your network's health and security. Pre-configured reports are available, as well as custom reports that can be tailored to your specific reporting and compliance requirements. For example, you may want to look closely at application usage or web site violations. The reports can be emailed as PDFs and can cover different time periods.

### **Maintain important configuration information uniformly**

The correct configuration of the devices within your network is essential to maintaining an optimum performance and security posture. In addition, maintaining the correct firmware (operating system) level allows you to take advantage of the latest features.

### **Service security**

All communication (including log information) between the devices and the clouds is encrypted. Redundant data centers are always used to give the service high availability. Operational

security measures have been put in place to make sure your data is secure — only you can view or retrieve it.

## Registration and Activation

There are five key activation steps. The procedure for each step may vary depending on your model and your FortiOS firmware version, and whether your device (FortiGate or FortiWifi) is brand new.

The steps are:

1. Registering with Support (New devices only)
2. Activating your FortiCloud account
3. Enabling logging to FortiCloud
4. Logging into the FortiCloud portal
5. Upgrading to a 200Gb subscription (Recommended)

### Registering with Support

Registration is very important for new devices, as it allows interaction with the Fortinet back-end systems, such as support. This registration will also allow other services, such as support and data space expansion contracts, to be used with your FortiCloud account.

### Registering and Activating your FortiCloud account

FortiCloud accounts can be registered manually through the FortiCloud website, <https://www.forticloud.com>, but you can easily register and activate your account directly within your FortiGate unit. Your registration process will vary somewhat, depending on which firmware version and device you have.

#### FortiGate 300 and below, all FortiWifi units

1. On your device's dashboard, in the License Information widget, select the green *Activate* button in the FortiCloud section.

The screenshot shows the FortiGate 100D dashboard with the 'License Information' widget expanded. The widget is divided into three main sections:

- Support Contract:**

Registration	Registered (Login: jmaddison@fortinet.com) [Login Now]	✓
Hardware	8 x 5 support (Expires: 2013-09-07)	✓
Firmware	8 x 5 support (Expires: 2013-09-07)	✓
Enhanced Support	24 x 7 support (Expires: 2013-09-07)	✓
Comprehensive Support	24 x 7 support (Expires: 2013-09-07)	✓
- FortiGuard Services:**

AntiVirus	Not Registered [Subscribe]	✗
IPS	Not Registered [Subscribe]	✗
Vulnerability Scan	Not Registered [Subscribe]	✗
Web Filtering	Not Registered [Configure]	✗
Email Filtering	Not Registered [Configure]	✗
- FortiCloud:**

Account	activate	
---------	----------	--

Below the FortiCloud section, there are sections for SMS (Messages Sent: 0 of 5) and softToken.

On the right side of the dashboard, there is an 'Alert Message Console' showing a list of system events, including system restarts and firmware upgrades.

2. A dialogue asking you to register your FortiCloud account will appear. Enter your information, view and accept the Terms and Conditions and select *Create Account*.

3. A second dialogue window will appear, asking you to enter your information to confirm your account. This will send a confirmation email to your registered email. The dashboard widget will update to show that confirmation is required.

Vulnerability Scan	Not Registered [Subscribe]	✕	2012-09-24 0*
Web Filtering	Not Registered [Configure]	✕	2012-09-24 0*
Email Filtering	Not Registered [Configure]	✕	2012-09-21 1*
<b>FortiCloud</b>			
Account	<b>Activation Pending.</b> Please view confirmation email.		
<b>SMS</b>			
Messages Sent	0 of 5		
<b>SoftToken</b>			
Usage	0 of 5000		
<b>Virtual Domain</b>			

4. Open your email, and follow the confirmation link contained in it.

A FortiCloud page will open, stating that your account has been confirmed. The Activation Pending message on the dashboard will change to state the type of account you have ('1Gb Free' or '200Gb Subscription'), and will now provide a link to the FortiCloud portal.

## FortiGate 600 to 800

For 600 through 800, FortiCloud registration must be done through the FortiGate CLI Console. Devices beyond the FortiGate 800 do not support the FortiCloud service.

## Enabling logging to FortiCloud

In order to enable remote logging to the FortiCloud Service, you must first configure the FortiGate's log uploading settings. You must also enable logging in each policy that covers traffic that you want to be logged.

### FortiOS 5.0

FortiOS 5.0 will automatically start logging Traffic and Event logs to FortiCloud upon activation. Logging can be disabled or configured through the FortiGate interface or CLI Console.

### Configuring policies 5.0

After enabling logging functionality, you will need to select which policies will be logged.

1. Open the Policy list.
2. Choose the policy you would like to log, and select *Edit*.
3. Check the box next to *Log all sessions*.
4. Select *OK*.

## Logging into the FortiCloud portal

Once logging has been configured and you have registered your account, you can log into the FortiCloud portal and begin viewing your logging results. There are two methods to reach the FortiCloud portal:

- If you have direct networked access to the FortiGate unit, you can simply open your Dashboard and check the License Information widget. Next to the current FortiCloud connection status will be a link to reach the FortiCloud Portal.
- If you do not currently have access to the FortiGate's interface, you can visit the FortiCloud website (<https://forticloud.com>) and log in remotely, using your email and password. It will ask you to confirm the FortiCloud account you are connecting to and then you will be granted access. Connected devices can be remotely configured using the Scripts page in the Management Tab, useful if an administrator may be away from the unit for a long period of time.

## Upgrading to a 200Gb subscription

Upgrading your subscription is simple but must be done through the FortiGate unit, as the storage contract is allocated based on devices rather than user accounts.

1. Open the FortiGate Dashboard.
2. In the License Information widget, select *Upgrade* next to 'Type' in the FortiCloud section.
3. A new window will open, showing the Fortinet Support portal. Follow the on-screen instructions to register your contract.
4. Wait approximately 10 minutes for the contract to be applied and then visit your Dashboard.

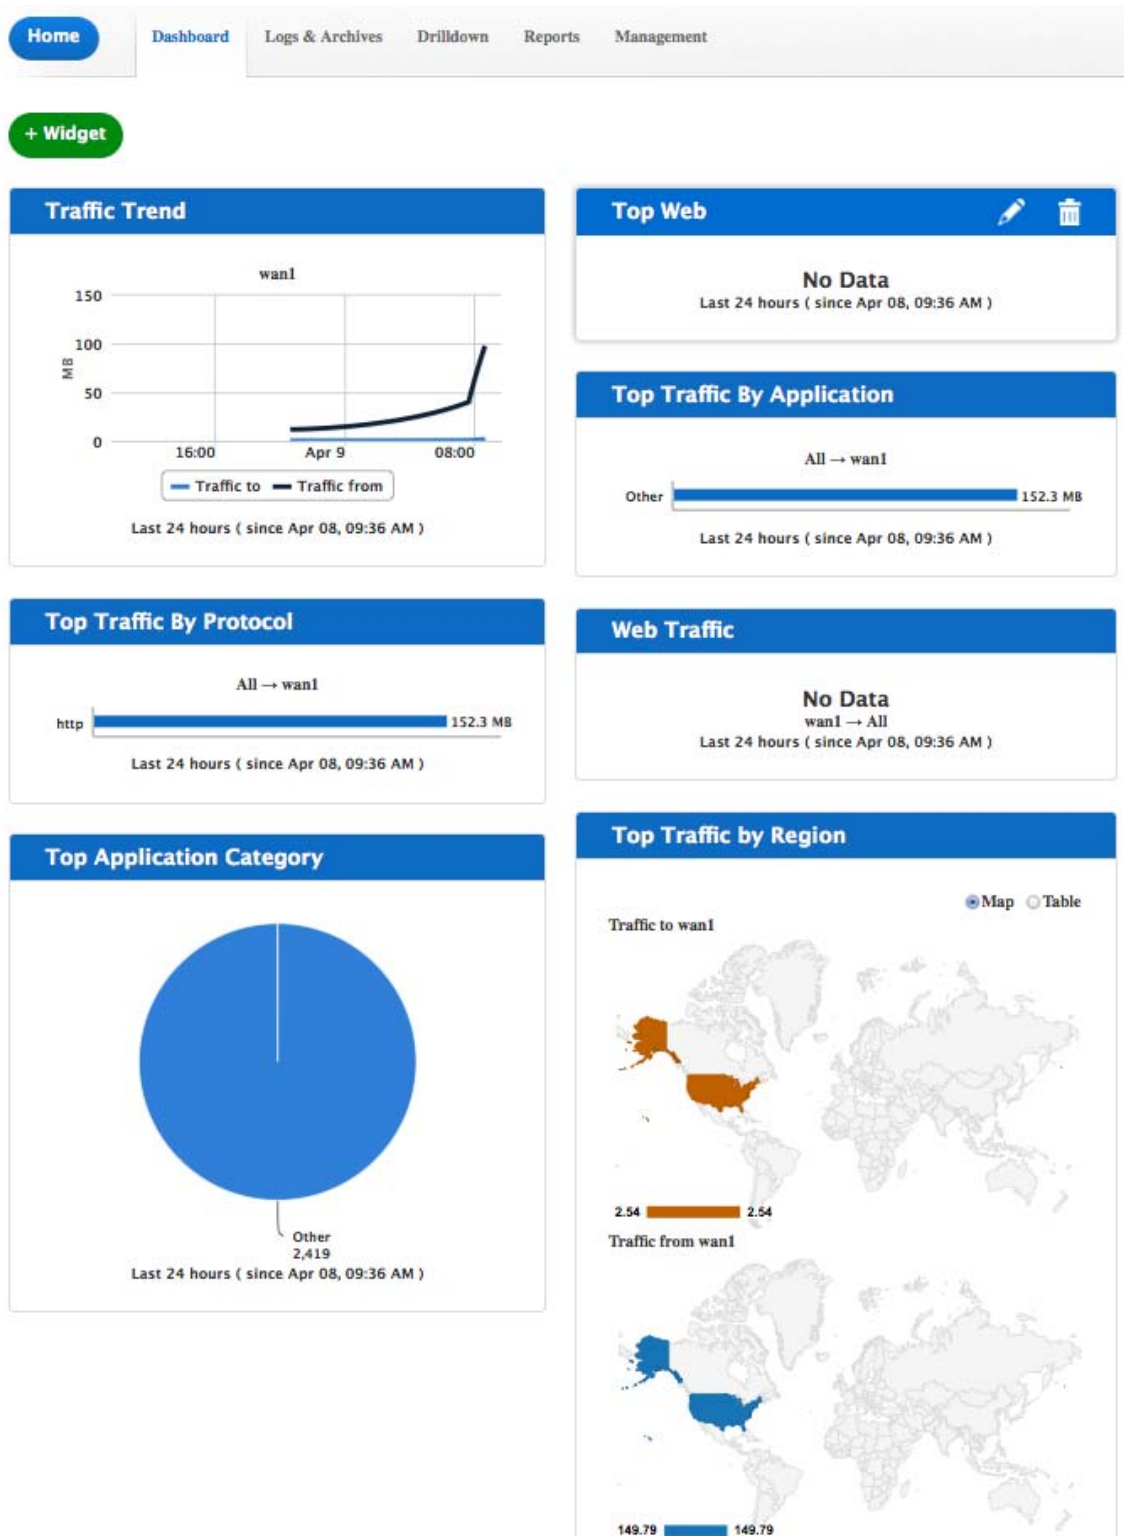
In the License Information widget, Type will have changed from 'Free' to 'Subscribed'. Your maximum listed storage will also have updated.

## The FortiCloud Portal

There are five main tabs in the FortiCloud portal, which allow you to access different features and information. The FortiCloud Settings, Help, and Logout buttons appear in the upper right.

- Dashboards
- Logs & Archives
- Drilldown
- Reports
- Management
- AV Submissions (this tab only appears if sandboxing has occurred, see "[Cloud Sandboxing](#)" on page 1479)

Figure 234: The FortiCloud Portal



## Using FortiCloud

Below is a list of possible tasks that show you how to make the best of the features that FortiCloud has to offer.

**Tasks:**

- Adding a new dashboard with custom charts
- Filtering logs to find specific information
- Downloading logs
- Using drilldown charts to find specific information
- Viewing and printing existing reports
- Generating scheduled and immediate reports
- Creating and configuring a new report with your logo
- Checking the status of your registration contract
- Adding a new user account to a FortiCloud account

For further information about using FortiCloud, please see the [FortiCloud Getting Started Guide](#).

## Cloud Sandboxing

FortiCloud can be used for automated sample tracking, or sandboxing, for files from a FortiGate unit. This allows suspicious files to be sent to be inspected without risking network security. If the file exhibits risky behavior, or is found to contain a virus, a new virus signature is created and added to the FortiGuard antivirus signature database. This feature was formerly known as FortiGuard Analytics.

Cloud sandboxing is configured by going to *System > Config > FortiSandbox*. After enabling FortiSandbox, select *Cloud Sandbox (FortiCloud)*.

Sandboxing results will be shown in a new tab called *AV Submissions* in the FortiCloud portal. This tab will only appear after a file has been sent for sandboxing.

# Interfaces

Interfaces, both physical and virtual, enable traffic to flow to and from the internal network, and the Internet and between internal networks. The FortiGate unit has a number of options for setting up interfaces and groupings of subnetworks that can scale to a company's growing requirements.

This chapter includes:

- [Physical](#)
- [Interface settings](#)
- [Software switch](#)
- [Virtual Switch](#)
- [Loopback interfaces](#)
- [Redundant interfaces](#)
- [One-armed sniffer](#)
- [Aggregate Interfaces](#)
- [DHCP addressing mode on an interface](#)
- [Administrative access](#)
- [Wireless](#)
- [Interface MTU packet size](#)
- [Secondary IP addresses to an interface](#)
- [Virtual domains](#)
- [Virtual LANs](#)
- [Zones](#)

## Physical

FortiGate units have a number of physical ports where you connect ethernet or optical cables. Depending on the model, they can have anywhere from four to 40 physical ports. Some units have a grouping of ports labelled as internal, providing a built-in switch functionality.

In FortiOS, the port names, as labeled on the FortiGate unit, appear in the web-based manager in the *Unit Operation* widget, found on the Dashboard. They also appear when you are configuring the interfaces, by going to *System > Network > Interface*. As shown below, the FortiGate-100D (Generation 2) has 22 interfaces.



Two of the physical ports on the FortiGate-100D (Generation 2) are SFP ports. These ports share the numbers 15 and 16 with RJ-45 ports. Because of this, when SFP port 15 is used, RJ-45 port 15 cannot be used, and vice versa.

These ports also share the same MAC address.

---



Figure 235:FortiGate-100D physical interfaces



Figure 236:FortiGate-100D interfaces on the Dashboard

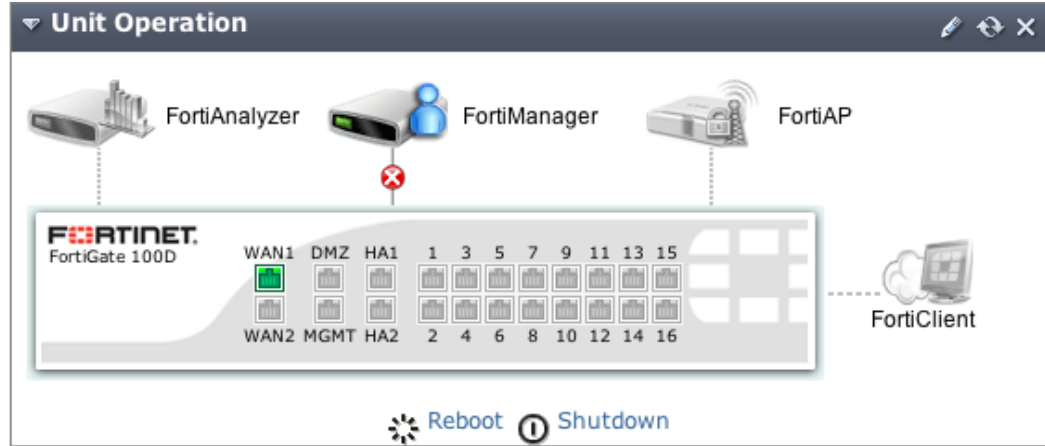


Figure 237:Configuring the FortiGate-100D ports

Name	Type	IP/Netmask	Access	Administrative Status	Link Status	Ref.
<input type="checkbox"/> wan1	Physical	172.20.120.230 / 255.255.255.0	HTTP,HTTPS,PING,SSH	🟢	🟢 100 Mbps/Full Duplex	1
<input type="checkbox"/> dmz	Physical	10.10.10.1 / 255.255.255.0	HTTPS,PING,FMG-Access	🟢	🔴	1
<input type="checkbox"/> modem	Physical	0.0.0.0 / 0.0.0.0		🔴	🔴	0
<input type="checkbox"/> wan2	Physical	0.0.0.0 / 0.0.0.0	PING,FMG-Access	🟢	🔴	1
<input type="checkbox"/> mgmt	Physical	192.168.1.99 / 255.255.255.0	HTTPS,PING,FMG-Access	🟢	🔴	0
<input type="checkbox"/> ha1	Physical	0.0.0.0 / 0.0.0.0		🟢	🔴	0
<input type="checkbox"/> ha2	Physical	0.0.0.0 / 0.0.0.0		🟢	🔴	0
<input type="checkbox"/> internal	Physical	192.168.100.99 / 255.255.255.0	HTTPS,PING,FMG-Access	🟢	🟢 100 Mbps/Full Duplex	2

Normally the internal interface is configured as a single interface shared by all physical interface connections - a switch. The switch mode feature has two states - switch mode and interface mode. Switch mode is the default mode with only one interface and one address for the entire internal switch. Interface mode enables you to configure each of the internal switch physical interface connections separately. This enables you to assign different subnets and netmasks to each of the internal physical interface connections.

The larger FortiGate units can also include Advanced Mezzanine Cards (AMC), which can provide additional interfaces (Ethernet or optical), with throughput enhancements for more efficient handling of specialized traffic. These interfaces appear in FortiOS as port amc/sw1, amc/sw2 and so on. In the following illustration, the FortiGate-3810A has three AMC cards installed: two single-width (amc/sw1, amc/sw2) and one double-width (amc/dw).

Figure 238: FortiGate-3810A AMC card port naming

<input type="checkbox"/>	Name	IP/Netmask	Access	Administrative Status	Link Status
<input type="checkbox"/>	amc-dw2/1	0.0.0.0 / 0.0.0.0		⬆	⬇
<input type="checkbox"/>	amc-dw2/2	0.0.0.0 / 0.0.0.0		⬆	⬇
<input type="checkbox"/>	amc-sw1/1	0.0.0.0 / 0.0.0.0		⬆	⬇
<input type="checkbox"/>	amc-sw1/2	0.0.0.0 / 0.0.0.0		⬆	⬇
<input type="checkbox"/>	amc-sw1/3	0.0.0.0 / 0.0.0.0		⬆	⬇
<input type="checkbox"/>	amc-sw1/4	0.0.0.0 / 0.0.0.0		⬆	⬇
<input type="checkbox"/>	amc-sw2/1	0.0.0.0 / 0.0.0.0		⬆	⬇
<input type="checkbox"/>	amc-sw2/2	0.0.0.0 / 0.0.0.0		⬆	⬇
<input type="checkbox"/>	amc-sw2/3	0.0.0.0 / 0.0.0.0		⬆	⬇
<input type="checkbox"/>	amc-sw2/4	0.0.0.0 / 0.0.0.0		⬆	⬇
<input type="checkbox"/>	april	0.0.0.0 / 0.0.0.0		⬆	
<input type="checkbox"/>	port1	10.21.101.101 / 255.255.255.0	HTTPS,PING,SSH	⬆	⬆
<input type="checkbox"/>	port2	192.168.100.99 / 255.255.255.0	PING	⬆	⬇
<input type="checkbox"/>	port3	0.0.0.0 / 0.0.0.0	PING	⬆	⬇
<input type="checkbox"/>	port4	0.0.0.0 / 0.0.0.0	PING	⬆	⬇
<input type="checkbox"/>	port5	0.0.0.0 / 0.0.0.0	PING	⬆	⬇
<input type="checkbox"/>	port6	0.0.0.0 / 0.0.0.0	PING	⬆	⬇
<input type="checkbox"/>	port7	0.0.0.0 / 0.0.0.0	PING	⬆	⬇
<input type="checkbox"/>	port8	0.0.0.0 / 0.0.0.0	PING	⬆	⬇
<input type="checkbox"/>	port9	0.0.0.0 / 0.0.0.0	PING	⬆	⬇
<input type="checkbox"/>	port10	0.0.0.0 / 0.0.0.0	PING	⬆	⬇

## Interface settings

In *System > Network > Interface*, you configure the interfaces, physical and virtual, for the FortiGate unit. There are different options for configuring interfaces when the FortiGate unit is in NAT mode or transparent mode. On FortiOS Carrier, you can also enable the Gi gatekeeper on each interface for anti-overbilling.

### Interface page

**Create New** Select to add a new interface, zone or, in transparent mode, port pair.

For more information on configuring zones, see [Zones](#).

Depending on the model you can add a VLAN interface, a loopback interface, a IEEE 802.3ad aggregated interface, or a redundant interface.

When VDOMs are enabled, you can also add Inter-VDOM links.

**Name** The names of the physical interfaces on your FortiGate unit. This includes any alias names that have been configured.

When you combine several interfaces into an aggregate or redundant interface, only the aggregate or redundant interface is listed, not the component interfaces.

If you have added VLAN interfaces, they also appear in the name list, below the physical or aggregated interface to which they have been added.

If you have added loopback interfaces, they also appear in the interface list, below the physical interface to which they have been added. If you have software switch interfaces configured, you will be able to view them. For more information, see [“Software switch” on page 1486](#).

If your FortiGate unit supports AMC modules, the interfaces are named amc-sw1/1, amc-dw1/2, and so on.

**Type** The configuration type for the interface.

<b>IP/Netmask</b>	The current IP address and netmask of the interface.  In VDOM mode, when VDOMs are not all in NAT or transparent mode some values may not be available for display and will be displayed as “-”.
<b>Access</b>	The administrative access configuration for the interface.
<b>Administrative Status</b>	Indicates if the interface can be accessed for administrative purposes. If the administrative status is a green arrow, and administrator could connect to the interface using the configured access.  If the administrative status is a red arrow, the interface is administratively down and cannot be accessed for administrative purposes.
<b>Link Status</b>	The status of the interface physical connection. Link status can be either up (green arrow) or down (red arrow). If link status is up the interface is connected to the network and accepting traffic. If link status is down the interface is not connected to the network or there is a problem with the connection. You cannot change link status from the web-based manager, and typically is indicative of an ethernet cable plugged into the interface.  Link status is only displayed for physical interfaces.
<b>MAC</b>	The MAC address of the interface.
<b>Mode</b>	Shows the addressing mode of the interface. The addressing mode can be manual, DHCP, or PPPoE.
<b>Secondary IP</b>	Displays the secondary IP addresses added to the interface.
<b>MTU</b>	The maximum number of bytes per transmission unit (MTU) for the interface.
<b>Virtual Domain</b>	The virtual domain to which the interface belongs. This column is visible when VDOM configuration is enabled.
<b>VLAN ID</b>	The configured VLAN ID for VLAN subinterfaces.

## Interface configuration and settings

To configure an interface, go to *System > Network > Interface* and select *Create New*.

<b>Name</b>	Enter a name of the interface. Physical interface names cannot be changed.
<b>Alias</b>	Enter an alternate name for a physical interface on the FortiGate unit. This field appears when editing an existing physical interface.  The alias can be a maximum of 25 characters. The alias name will not appear in logs.
<b>Link Status</b>	Indicates whether the interface is connected to a network (link status is <i>Up</i> ) or not (link status is <i>Down</i> ). This field appears when editing an existing physical interface.
<b>Type</b>	Select the type of interface that you want to add.  On some models you can set <i>Type</i> to <i>802.3ad Aggregate</i> or <i>Redundant Interface</i> .

<b>Interface</b>	<p>Displayed when <i>Type</i> is set to <i>VLAN</i>.</p> <p>Select the name of the physical interface to which to add a VLAN interface. Once created, the VLAN interface is listed below its physical interface in the Interface list.</p> <p>You cannot change the physical interface of a VLAN interface except when adding a new VLAN interface.</p>
<b>VLAN ID</b>	<p>Displayed when <i>Type</i> is set to <i>VLAN</i>.</p> <p>Enter the VLAN ID. You cannot change the <i>VLAN ID</i> except when adding a new VLAN interface.</p> <p>The VLAN ID can be any number between 1 and 4094 and must match the VLAN ID added by the IEEE 802.1Q-compliant router or switch connected to the VLAN subinterface.</p>
<b>Virtual Domain</b>	<p>Select the virtual domain to add the interface to.</p> <p>Admin accounts with <i>super_admin</i> profile can change the <i>Virtual Domain</i>.</p>
<b>Physical Interface Members</b>	<p>This section has two different forms depending on the interface type:</p> <ul style="list-style-type: none"> <li>• <b>Software switch interface</b> - this section is a display-only field showing the interfaces that belong to the software switch virtual interface.</li> <li>• <b>802.3ad aggregate or Redundant interface</b> - this section includes available interface and selected interface lists to enable adding or removing interfaces from the interface. For more information, see <a href="#">Redundant interfaces</a>.</li> </ul> <p>Select interfaces from this <i>Available Interfaces</i> list and select the right arrow to add an interface to the <i>Selected Interface</i> list.</p>
<b>Addressing mode</b>	<p>Select the addressing mode for the interface.</p> <ul style="list-style-type: none"> <li>• Select <i>Manual</i> and add an <i>IP/Netmask</i> for the interface. If IPv6 configuration is enabled you can add both a IPv4 and an IPv6 IP address.</li> <li>• Select <i>DHCP</i> to get the interface IP address and other network settings from a DHCP server. For more information, see <a href="#">DHCP addressing mode on an interface</a></li> <li>• Select <i>PPPoE</i> to get the interface IP address and other network settings from a PPPoE server. For more information, see <a href="#">PPPoE addressing mode on an interface</a>.</li> <li>• Select <i>One-Arm Sniffer</i> to enable the interface as a means to detect possible traffic threats. This option is available on physical ports not configured for the primary Internet connection. For more information see <a href="#">One-armed sniffer</a>.</li> <li>• Select <i>Dedicate to FortiAP/FortiSwitch</i> to have a FortiAP unit or FortiSwitch unit connect exclusively to the interface. This option is only available when editing a physical interface, and it has a static IP address. When you enter the IP address, the FortiGate unit automatically creates a DHCP server using the subnet entered. This option is not available on the ADSL interface.</li> </ul> <p>The FortiSwitch option is currently only available on the FortiGate-100D.</p>

<b>IP/Netmask</b>	If <i>Addressing Mode</i> is set to <i>Manual</i> , enter an IPv4 address/subnet mask for the interface. FortiGate interfaces cannot have IP addresses on the same subnet.
<b>IPv6 Address</b>	If <i>Addressing Mode</i> is set to <i>Manual</i> and IPv6 support is enabled, enter an IPv6 address/subnet mask for the interface. A single interface can have both an IPv4 and IPv6 address or just one or the other.
<b>Administrative Access</b>	Select the types of administrative access permitted for IPv4 connections to this interface.
<b>HTTPS</b>	Allow secure HTTPS connections to the web-based manager through this interface. If configured, this option will enable automatically when selecting the <i>HTTP</i> option. For information on this setting, see <a href="#">“HTTPS redirect” on page 1438</a> .
<b>PING</b>	Interface responds to pings. Use this setting to verify your installation and for testing.
<b>HTTP</b>	Allow HTTP connections to the web-based manager through this interface. If configured, this option will also enable the <i>HTTPS</i> option. For information on this setting, see <a href="#">“HTTPS redirect” on page 1438</a> .
<b>SSH</b>	Allow SSH connections to the CLI through this interface.
<b>SNMP</b>	Allow a remote SNMP manager to request SNMP information by connecting to this interface.
<b>TELNET</b>	Allow Telnet connections to the CLI through this interface. Telnet connections are not secure and can be intercepted by a third party.
<b>FMG-Access</b>	Allow FortiManager authorization automatically during the communication exchange between the FortiManager and FortiGate units.
<b>FCT-Access</b>	You can configure a FortiGate interface as an interface that will accept FortiClient connections. When configured, the FortiGate unit sends broadcast messages which the FortiClient software running on a end user PC is listening for.
<b>CAPWAP</b>	Allows the FortiGate unit’s wireless controller to manage a wireless access point, such as a FortiAP unit.
<b>IPv6 Administrative Access</b>	Select the types of administrative access permitted for IPv6 connections to this interface. These types are the same as for Administrative Access.
<b>Security Mode</b>	Select a captive portal for the interface. When selected, you can define the portal message and look that the user sees when logging into the interface. You can also define one or more user groups that have access to the interface.
<b>DHCP Server</b>	Select to enable a DHCP server for the interface. For more information on configuring a DHCP server on the interface, see <a href="#">“DHCP servers and relays” on page 1577</a> .
<b>Detect and Identify Devices</b>	Select to enable the interface to be used with BYOD hardware such as iPhones. Define the device definitions by going to <i>User &amp; Device &gt; Device</i> .

<b>Add New Devices to Vulnerability Scan List</b>	This option appears when <i>Detect and Identify Devices</i> is enabled. When enabled, the FortiGate unit performs a network vulnerability scan of any devices detected or seen on the interface. The vulnerability scan occur as configured, either on demand, or as scheduled.
<b>Broadcast Discovery Messages</b>	Available when <i>FCT-Access</i> is enabled for the <i>Administrative Access</i> . Select to enable sends broadcast messages which the FortiClient software running on a end user PC is listening for.  Once enabled, the FortiGate unit broadcasts a discovery message that includes the IP address of the interface and listening port number to the local network. All PCs running FortiClient on that network listen for this discovery message.
<b>Enable Explicit Web Proxy</b>	Available when enabling explicit proxy on the <i>System Information Dashboard (System &gt; Dashboard &gt; Status)</i> .  This option is not available for a VLAN interface selection. Select to enable explicit web proxying on this interface. When enabled, this interface will be displayed on <i>System &gt; Network &gt; Explicit Proxy</i> under <i>Listen on Interfaces</i> and web traffic on this interface will be proxied according to the Web Proxy settings.
<b>Enable STP</b>	With FortiGate units with a switch interface is in switch mode, this option is enabled by default. It enables the single instance MSTP spanning tree protocol.
<b>Listen for RADIUS Accounting Messages</b>	Select to use the interface as a listening port for RADIUS content.
<b>Secondary IP Address</b>	Add additional IPv4 addresses to this interface. Select the Expand Arrow to expand or hide the section.
<b>Comments</b>	Enter a description up to 63 characters to describe the interface.
<b>Administrative Status</b>	Select either <i>Up</i> (green arrow) or <i>Down</i> (red arrow) as the status of this interface.  <i>Up</i> indicates the interface is active and can accept network traffic.  <i>Down</i> indicates the interface is not active and cannot accept traffic.
<b>Gi Gatekeeper (FortiOS Carrier only)</b>	For FortiOS Carrier, enable Gi Gatekeeper to enable the Gi firewall as part of the anti-overbilling configuration. You must also configure <i>Gi Gatekeeper Settings</i> by going to <i>System &gt; Admin &gt; Settings</i> .

## Software switch

A software switch, or soft switch, is a virtual switch that is implemented at the software, or firmware level, rather than the hardware level. A software switch can be used to simplify communication between devices connected to different FortiGate interfaces. For example, using a software switch, you can place the FortiGate interface connected to an internal network on the same subnet as your wireless interfaces. Then devices on the internal network can communicate with devices on the wireless network without any additional configuration such as additional security policies, on the FortiGate unit.

It can also be useful if you require more hardware ports on for the switch on a FortiGate unit. For example, if your FortiGate unit has a 4-port switch, WAN1, WAN2 and DMZ interfaces, and you need one more port, you can create a soft switch that can include the 4-port switch and the DMZ interface all on the same subnet. These types of applications also apply to wireless interfaces and virtual wireless interfaces and physical interfaces such as those with FortiWiFi and FortiAP unit.

Similar to a hardware switch, a software switch functions like a single interface. A software switch has one IP address; all of the interfaces in the software switch are on the same subnet. Traffic between devices connected to each interface are not regulated by security policies, and traffic passing in and out of the switch are affected by the same policy.

There are a few things to consider when setting up a software switch:

- Ensure you create a back up of the configuration.
- Ensure you have at least one port or connection such as the console port to connect to the FortiGate unit. If you accidentally combine too many ports, you will need a way to undo any errors.
- The ports that you include must not have any link or relation to any other aspect of the FortiGate unit. For example, DHCP servers, security policies, and so on.
- For increased security, you can create a captive portal for the switch, allowing only specific user groups access to the resources connected to the switch.

#### **To create a software switch - web-based manager**

1. Go to *System > Network > Interface* and select *Create New*.
2. For *Type*, select *Software Switch*.
3. In the *Physical Interface Members* option, select the interfaces to include.
4. Configure the remaining interface settings
5. Select *OK*.

#### **To create a software switch - CLI**

```
config system switch-interface
 edit <switch-name>
 set type switch
 set member <interface_list>
 end
config system interface
 edit <switch_name>
 set ip <ip_address>
 set allowaccess https ssh ping
 end
```

### **Soft switch example**

For this example, the wireless interface (WiFi) needs to be on the same subnet as the DMZ1 interface to facilitate wireless syncing from an iPhone and a local computer. The syncing

between two subnets is problematic. By putting both interfaces on the same subnet the synching will work. The software switch will accomplish this.



In this example, the soft switch includes a wireless interface. Remember to configure any wireless security before proceeding. If you leave this interface open without any password or other security, it leaves open access to not only the wireless interface but to any other interfaces and devices connected within the software switch.

### Clear the interfaces and back up the configuration

First, ensure that the interfaces are not being used with any other security policy or other use on the FortiGate unit. Check the WiFi and DMZ1 ports to ensure DHCP is not enabled on the interface and there are no other dependencies with these interfaces.

Next, save the current configuration, in the event something doesn't work, recovery can be quick.

### Merge the interfaces

The plan is to merge the WiFi port and DMZ1 port. This will create a software switch with a name of "synchro" with an IP address of 10.10.21.12. The steps will create the switch, add the IP and then set the administrative access for HTTPS, SSH and Ping.

#### To merge the interfaces - CLI

```
config system switch-interface
 edit synchro
 set type switch
 set member dmz1 wifi
 end
config system interface
 edit synchro
 set ip 10.10.21.12
 set allowaccess https ssh ping
 end
```

### Final steps

With the switch set up, you can now add security policies, DHCP servers and any other configuration that you would normally do to configure interfaces on the FortiGate unit.

## Virtual Switch

Virtual switch feature enables you create virtual switches on top of the physical switch(es) with designated interfaces/ports so that a virtual switch can build up its forwarding table through learning and forward traffic accordingly. When traffic is forwarded among interfaces belonging to the same virtual switch, the traffic doesn't need to go up to the software stack, but forwarded directly by the switch. When traffic has to be relayed to interfaces not on the virtual switch, the traffic will go through the normal data path and be offloaded to NP4 when possible.

This feature is only available on mid to high end FortiGate units, including the 100D, 600C, 1000C, and 1240B.



**To enable and configure the virtual switch, enter the CLI commands:**

```
config system virtual-switch
 edit vs1
 set physical-switch sw0
 config port
 edit 1
 set port port1
 set speed xx
 set duplex xx
 set status [up|down]
 edit 2
 set port port2
 set ...
 end
 end
end
```

## Loopback interfaces

A loopback interface is a logical interface that is always up (no physical link dependency) and the attached subnet is always present in the routing table.

The FortiGate's loopback IP address does not depend on one specific external port, and is therefore possible to access it through several physical or VLAN interfaces. Multiple loopback interfaces can be configured in either non-VDOM mode or in each VDOM.

Loopback interfaces still require appropriate firewall policies to allow traffic to and from this type of interface.

A loopback interface can be used with:

- Management access
- BGP (TCP) peering
- PIM RP

Loopback interfaces are a good practice for OSPF. Setting the OSPF router ID the same as loopback IP address troubleshooting OSPF easier, and remembering the management IP addresses (telnet to "router ID").

Dynamic routing protocols can be enabled on loopback interfaces

For black hole static route, use the black hole route type instead of the loopback interface.

## Redundant interfaces

On some models you can combine two or more physical interfaces to provide link redundancy. This feature enables you to connect to two or more switches to ensure connectivity in the event one physical interface or the equipment on that interface fails.

In a redundant interface, traffic is only going over one interface at any time. This differs from an aggregated interface where traffic is going over all interfaces for distribution of increased bandwidth. This difference means redundant interfaces can have more robust configurations with fewer possible points of failure. This is important in a fully-meshed HA configuration.

An interface is available to be in a redundant interface if:

- it is a physical interface, not a VLAN interface
- it is not already part of an aggregated or redundant interface
- it is in the same VDOM as the redundant interface
- it has no defined IP address
- is not configured for DHCP or PPPoE
- it has no DHCP server or relay configured on it
- it does not have any VLAN subinterfaces
- it is not referenced in any security policy, VIP, or multicast policy
- it is not monitored by HA
- it is not one of the FortiGate-5000 series backplane interfaces

When an interface is included in a redundant interface, it is not listed on the *System > Network > Interface* page. You cannot configure the interface individually and it is not available for inclusion in security policies, VIPs, or routing.

## One-armed sniffer

A one-armed sniffer is used to configure a physical interface on the FortiGate unit as a one-arm intrusion detection system (IDS). Traffic sent to the interface is examined for matches to the configured IPS sensor and application control list. Matches are logged and then all received traffic is dropped. Sniffing only reports on attacks. It does not deny or otherwise influence traffic.

Using the one-arm sniffer, you can configure a FortiGate unit to operate as an IDS appliance by sniffing network traffic for attacks without actually processing the packets. To configure one-arm IDS, you enable sniffer mode on a FortiGate interface and connect the interface to a hub or to the SPAN port of a switch that is processing network traffic.

To assign an interface as a sniffer interface, go to *System > Network > Interface*, edit the interface and select *One-Arm Sniffer*.

If the check box is not available, the interface is in use. Ensure that the interface is not selected in any firewall policies, routes, virtual IPs or other features in which a physical interface is specified.

---

### Enable Filters

Select to include filters to define a more granular sniff of network traffic. Select specific addresses, ports, VLANs and protocols.

In all cases, enter a number, or number range, for the filtering type. For Protocol values, standard protocols are:

- UDP - 17
- TCP - 6
- ICMP - 1

---

### Include IPv6 Packets

If your network is running a combination of IPv4 and IPv6 addressing, select to sniff both addressing types. Otherwise, the FortiGate unit will only sniff IPv4 traffic.

---

<b>Include Non-IP Packets</b>	Select for a more intense scan of content in the traffic.
<b>UTM Security Profiles</b>	IPS sensors, and application control lists enable you to select specific sensors and application you want to identify within the traffic.

## Aggregate Interfaces

Link aggregation (IEEE 802.3ad) enables you to bind two or more physical interfaces together to form an aggregated (combined) link. This new link has the bandwidth of all the links combined. If a link in the group fails, traffic is transferred automatically to the remaining interfaces with the only noticeable effect being a reduced bandwidth.

This is similar to redundant interfaces with the major difference being that a redundant interface group only uses one link at a time, where an aggregate link group uses the total bandwidth of the functioning links in the group, up to eight.

Support of the IEEE standard 802.3ad for link aggregation is available on some models.

An interface is available to be an aggregate interface if:

- it is a physical interface, not a VLAN interface or subinterface
- it is not already part of an aggregate or redundant interface
- it is in the same VDOM as the aggregated interface. Aggregate ports cannot span multiple VDOMs.
- it does not have an IP address and is not configured for DHCP or PPPoE
- it is not referenced in any security policy, VIP, IP Pool or multicast policy
- it is not an HA heartbeat interface
- it is not one of the FortiGate-5000 series backplane interfaces

Some models of FortiGate units do not support aggregate interfaces. In this case, the aggregate option is not an option in the web-based manager or CLI. As well, you cannot create aggregate interfaces from the interfaces in a switch port.

To see if a port is being used or has other dependencies, use the following diagnose command:

```
diagnose sys checkused system.interface.name <interface_name>
```

When an interface is included in an aggregate interface, it is not listed on the *System > Network > Interface* page. Interfaces will still appear in the CLI, although configuration for those interfaces will not take effect. You cannot configure the interface individually and it is not available for inclusion in security policies, VIPs, IP pools, or routing.

### Example

This example creates an aggregate interface on a FortiGate-3810A using ports 4-6 with an internal IP address of 10.13.101.100, as well as the administrative access to HTTPS and SSH.

#### To create an aggregate interface - web-based manager

1. Go to *System > Network > Interface* and select *Create New*.
2. Enter the Name as *Aggregate*.
3. For the *Type*, select *802.3ad Aggregate*.

If this option does not appear, your FortiGate unit does not support aggregate interfaces.

4. In the *Available Interfaces* list, select port 4, 5 and 6 and move it to the *Selected Interfaces* list.

5. Select the *Addressing Mode* of *Manual*.
6. Enter the IP address for the port of 10.13.101.100/24.
7. For *Administrative Access* select HTTPS and SSH.
8. Select *OK*.

#### To create aggregate interface - CLI

```

config system interface
 edit Aggregate
 set type aggregate
 set member port4 port5 port6
 set vdom root
 set ip 172.20.120.100/24
 set allowaccess https ssh
 end

```

## DHCP addressing mode on an interface

If you configure an interface to use DHCP, the FortiGate unit automatically broadcasts a DHCP request from the interface. The interface is configured with the IP address and any DNS server addresses and default gateway address that the DHCP server provides.



DHCP IPv6 is similar to DHCP IPv4, however there is:

- no default gateway option defined because a host learns the gateway using router advertisement messages
- there is no WINS servers because it is obsolete.

For more information about DHCP IPv6, see RFC 3315.

Configure DHCP for an interface in *System > Network > Interface* and selecting the interface from the list, and selecting *DHCP* in the *Address Mode*. The table describes the DHCP status information when DHCP is configured for an interface.

#### Addressing mode section of New Interface page for DHCP information

<b>Status</b>	<p>Displays DHCP status messages as the interface connects to the DHCP server and gets addressing information. Select <i>Status</i> to refresh the addressing mode status message.</p> <p>Status can be one of:</p> <ul style="list-style-type: none"> <li>• <b>initializing</b> - No activity.</li> <li>• <b>connecting</b> - interface attempts to connect to the DHCP server.</li> <li>• <b>connected</b> - interface retrieves an IP address, netmask, and other settings from the DHCP server.</li> <li>• <b>failed</b> - interface was unable to retrieve an IP address and other settings from the DHCP server.</li> </ul>
<b>Obtained IP/Netmask</b>	The IP address and netmask leased from the DHCP server. Only displayed if <i>Status</i> is <i>connected</i> .
<b>Renew</b>	Select to renew the DHCP license for this interface. Only displayed if <i>Status</i> is <i>connected</i> .

<b>Expiry Date</b>	The time and date when the leased IP address and netmask is no longer valid for the interface. The IP address is returned to the pool to be allocated to the next user request for an IP address. Only displayed if <i>Status</i> is <i>connected</i> .
<b>Default Gateway</b>	The IP address of the gateway defined by the DHCP server. Only displayed if <i>Status</i> is <i>connected</i> , and if <i>Receive default gateway from server</i> is selected.
<b>Distance</b>	Enter the administrative distance for the default gateway retrieved from the DHCP server. The administrative distance, an integer from 1-255, specifies the relative priority of a route when there are multiple routes to the same destination. A lower administrative distance indicates a more preferred route.
<b>Retrieve default gateway from server</b>	Enable to retrieve a default gateway IP address from the DHCP server. The default gateway is added to the static routing table.
<b>Override internal DNS</b>	Enable to use the DNS addresses retrieved from the DHCP server instead of the DNS server IP addresses on the DNS page.  When VDOMs are enabled, you can override the internal DNS only on the management VDOM.

## PPPoE addressing mode on an interface

If you configure the interface to use PPPoE, the FortiGate unit automatically broadcasts a PPPoE request from the interface.

The FortiGate units support many PPPoE RFC features (RFC 2516) including unnumbered IPs, initial discovery timeout and PPPoE Active Discovery Terminate (PADT).

PPPoE is only configurable in the web-based manager on desktop FortiGate units. 1U FortiGates and up must be configured in the CLI using the commands:

```
config system interface
 edit <port_name>
 set mode pppoe
 set username <ISP_username>
 set password <ISP_password>
 set idle-timeout <seconds>
 set distance <integer>
 set ipunnumbered <unnumbered-IP>
 set disc-retry-timeout <seconds>
 set padt-retry-timeout <seconds>
 set lcp-echo-interval <seconds>
 set dns-server-override {enable | disable}
 end
```

Configure PPPoE on an interface in *System > Network > Interface*. The table describes the PPPoE status information when PPPoE is configured for an interface.

---

### Addressing mode section of New Interface page

<b>Status</b>	<p>Displays PPPoE status messages as the FortiGate unit connects to the PPPoE server and gets addressing information. Select Status to refresh the addressing mode status message.</p> <p>The status is only displayed if you selected <i>Edit</i>.</p> <p>Status can be any one of the following 4 messages.</p>
<b>Initializing</b>	No activity.
<b>Connecting</b>	The interface is attempting to connect to the PPPoE server.
<b>Connected</b>	<p>The interface retrieves an IP address, netmask, and other settings from the PPPoE server.</p> <p>When the status is connected, PPPoE connection information is displayed.</p>
<b>Failed</b>	The interface was unable to retrieve an IP address and other information from the PPPoE server.
<b>Reconnect</b>	<p>Select to reconnect to the PPPoE server.</p> <p>Only displayed if Status is connected.</p>
<b>User Name</b>	The PPPoE account user name.
<b>Password</b>	The PPPoE account password.
<b>Unnumbered IP</b>	Specify the IP address for the interface. If your ISP has assigned you a block of IP addresses, use one of them. Otherwise, this IP address can be the same as the IP address of another interface or can be any IP address.
<b>Initial Disc Timeout</b>	Enter Initial discovery timeout. Enter the time to wait before starting to retry a PPPoE discovery.
<b>Initial PADT timeout</b>	Enter Initial PPPoE Active Discovery Terminate (PADT) timeout in seconds. Use this timeout to shut down the PPPoE session if it is idle for this number of seconds. PADT must be supported by your ISP. Set initial PADT timeout to 0 to disable.
<b>Distance</b>	Enter the administrative distance for the default gateway retrieved from the PPPoE server. The administrative distance, an integer from 1-255, specifies the relative priority of a route when there are multiple routes to the same destination. A lower administrative distance indicates a more preferred route. The default distance for the default gateway is 1.
<b>Retrieve default gateway from server</b>	Enable to retrieve a default gateway IP address from a PPPoE server. The default gateway is added to the static routing table.
<b>Override internal DNS</b>	<p>Enable to replace the DNS server IP addresses on the System DNS page with the DNS addresses retrieved from the PPPoE server.</p> <p>When VDOMs are enabled, you can override the internal DNS only on the management VDOM.</p>

## Administrative access

Interfaces, especially the public-facing ports can be potentially accessed by those who you may not want access to the FortiGate unit. When setting up the FortiGate unit, you can set the type of protocol an administrator must use to access the FortiGate unit. The options include:

- HTTPS
- HTTP
- SSH
- TELNET
- SNMP
- PING
- FortiManager Access (FMG-Access)
- FortiClient Access (FCT-Access)

You can select as many, or as few, even none, that are accessible by an administrator.

This example adds an IPv4 address 172.20.120.100 to the WAN1 interface as well as the administrative access to HTTPS and SSH. As a good practice, set the administrative access when you are setting the IP address for the port.

### To add an IP address on the WAN1 interface - web-based manager

1. Go to *System > Network > Interface*.
2. Select the WAN1 interface row and select *Edit*.
3. Select the *Addressing Mode* of *Manual*.
4. Enter the IP address for the port of 172.20.120.100/24.
5. For *Administrative Access*, select *HTTPS* and *SSH*.
6. Select *OK*.

### To create IP address on the WAN1 interface - CLI

```
config system interface
 edit wan1
 set ip 172.20.120.100/24
 set allowaccess https ssh
 end
```



When adding to, or removing a protocol, you must type the entire list again. For example, if you have an access list of HTTPS and SSH, and you want to add PING, typing:

```
set allowaccess ping
```

...only PING will be set. In this case, you must type...

```
set allowaccess https ssh ping
```

## Wireless

A wireless interface is similar to a physical interface only it does not include a physical connection. The FortiWiFi units enables you to add multiple wireless interfaces that can be available at the same time (the FortiWiFi-30B can only have one wireless interface). On FortiWiFi units, you can configure the device to be either an access point, or a wireless client. As an access point, the FortiWiFi unit can have up to four separate SSIDs, each on their own subnet for wireless access. In client mode, the FortiWiFi only has one SSID, and is used as a receiver, to enable remote users to connect to the existing network using wireless protocols.

Wireless interfaces also require additional security measures to ensure the signal does not get hijacked and data tampered or stolen.

For more information on configuring wireless interfaces see the [Deploying Wireless Networks Guide](#).

## Interface MTU packet size

You can change the maximum transmission unit (MTU) of the packets that the FortiGate unit transmits to improve network performance. Ideally, the MTU should be the same as the smallest MTU of all the networks between the FortiGate unit and the destination of the packets. If the packets that the FortiGate unit sends are larger than the smallest MTU, they are broken up or fragmented, which slows down transmission. You can easily experiment by lowering the MTU to find an MTU size for optimum network performance.

To change the MTU, select Override default MTU value (1500) and enter the MTU size based on the addressing mode of the interface

- 68 to 1 500 bytes for static mode
- 576 to 1 500 bytes for DHCP mode
- 576 to 1 492 bytes for PPPoE mode
- larger frame sizes if supported by the FortiGate model

Only available on physical interfaces. Virtual interfaces associated with a physical interface inherit the physical interface MTU size.

Interfaces on some models support frames larger than the traditional 1500 bytes. Jumbo frames are supported on FortiGate models that have either a SOC2 or NP4lite, except for the FortiGate-30D, as well as on FortiGate-100D series models (for information about your FortiGate unit's hardware, see the [Hardware Acceleration](#) guide). For other models, please contact Fortinet Customer Support for the maximum frame size that is supported.

If you need to enable sending larger frames over a route, you need all Ethernet devices on that route to support that larger frame size, otherwise your larger frames will not be recognized and are dropped.

If you have standard size and larger size frame traffic on the same interface, routing alone cannot route them to different routes based only on frame size. However, you can use VLANs to make sure the larger frame traffic is routed over network devices that support that larger size. VLANs will inherit the MTU size from the parent interface. You will need to configure the VLAN to include both ends of the route as well as all switches and routers along the route.

MTU packet size is changed in the CLI. If you select an MTU size larger than your FortiGate unit supports, an error message will indicate this. In this situation, try a smaller MTU size until the value is supported.



In Transparent mode, if you change the MTU of an interface, you must change the MTU of all interfaces on the FortiGate unit to match the new MTU.

---



To change the MTU size, use the following CLI commands:

```
config system interface
 edit <interface_name>
 set mtu-override enable
 set mtu <byte_size>
 end
```

## Secondary IP addresses to an interface

If an interface is configured with a manual or static IP address, you can also add secondary static IP addresses to the interface. Adding secondary IP addresses effectively adds multiple IP addresses to the interface. Secondary IP addresses cannot be assigned using DHCP or PPPoE.

All of the IP addresses added to an interface are associated with the single MAC address of the physical interface and all secondary IP addresses are in the same VDOM as the interface that are added to. You configure interface status detection for gateway load balancing separately for each secondary IP addresses. As with all other interface IP addresses, secondary IP addresses cannot be on the same subnet as any other primary or secondary IP address assigned to a FortiGate interface unless they are in separate VDOMs.

To configure a secondary IP, go to *System > Network > Interface*, select *Edit* or *Create New* and select the *Secondary IP Address* check box.

## Virtual domains

Virtual domains (VDOMs) are a method of dividing a FortiGate unit into two or more virtual units that function as multiple independent units. A single FortiGate unit is then flexible enough to serve multiple departments of an organization, separate organizations, or to act as the basis for a service provider's managed security service.

VDOMs provide separate security domains that allow separate zones, user authentication, security policies, routing, and VPN configurations. By default, each FortiGate unit has a VDOM named *root*. This VDOM includes all of the FortiGate physical interfaces, modem, VLAN subinterfaces, zones, security policies, routing settings, and VPN settings.

When a packet enters a VDOM, it is confined to that VDOM. In a VDOM, you can create security policies for connections between Virtual LAN (VLAN) subinterfaces or zones in the VDOM. Packets do not cross the virtual domain border internally. To travel between VDOMs, a packet must pass through a firewall on a physical interface. The packet then arrives at another VDOM on a different interface, but it must pass through another firewall before entering the VDOM. Both VDOMs are on the same FortiGate unit. Inter-VDOMs change this behavior in that they are internal interfaces; however their packets go through all the same security measures as on physical interfaces.

This example shows how to enable VDOMs on the FortiGate unit and the basic and create a VDOM accounting on the DMZ2 port and assign an administrator to maintain the VDOM. First enable Virtual Domains on the FortiGate unit. When you enable VDOMs, the FortiGate unit will log you out.

For desktop and low-end FortiGate units, VDOMs are enabled using the CLI. On larger FortiGate units, you can enable on the web-based manager or the CLI. Once enabled all further configuration can be made in the web-based manager or CLI.

### To enable VDOMs - web-based manager

1. Go to *System > Dashboard > Status*.

2. In the *System Information* widget, select *Enable for Virtual Domain*.

The FortiGate unit logs you out. Once you log back in, you will notice that the menu structure has changed. This reflects the global settings for all Virtual Domains.

#### To enable VDOMs - CLI

```
config system global
 set vdom-admin enable
end
```

Next, add the VDOM called accounting.

#### To add a VDOM - web-based manager

1. Go to *System > VDOM > VDOM*, and select *Create New*.
2. Enter the VDOM name *accounting*.
3. Select *OK*.

To add a VDOM - CLI

```
config vdom
 edit <new_vdom_name>
end
```

With the Virtual Domain created, you can assign a physical interface to it, and assign it an IP address.

#### To assign physical interface to the accounting Virtual Domain - web-based manager

1. Go to *System > Network > Interface*.
2. Select the DMZ2 port row and select *Edit*.
3. For the *Virtual Domain* drop-down list, select *accounting*.
4. Select the *Addressing Mode* of *Manual*.
5. Enter the IP address for the port of 10.13.101.100/24.
6. Set the *Administrative Access* to *HTTPS* and *SSH*.
7. Select *OK*.

#### To assign physical interface to the accounting Virtual Domain - CLI

```
config global
 config system interface
 edit dmz2
 set vdom accounting
 set ip 10.13.101.100/24
 set allowaccess https ssh
 next
 end
```

## Virtual LANs

The term VLAN subinterface correctly implies the VLAN interface is not a complete interface by itself. You add a VLAN subinterface to the physical interface that receives VLAN-tagged packets. The physical interface can belong to a different VDOM than the VLAN, but it must be connected to a network route that is configured for this VLAN. Without that route, the VLAN will not be connected to the network, and VLAN traffic will not be able to access this interface. The traffic on the VLAN is separate from any other traffic on the physical interface.

FortiGate unit interfaces cannot have overlapping IP addresses, the IP addresses of all interfaces must be on different subnets. This rule applies to both physical interfaces and to virtual interfaces such as VLAN subinterfaces. Each VLAN subinterface must be configured with its own IP address and netmask. This rule helps prevent a broadcast storm or other similar network problems.

Any FortiGate unit, with or without VDOMs enabled, can have a maximum of 255 interfaces in Transparent operating mode. In NAT/Route operating mode, the number can range from 255 to 8192 interfaces per VDOM, depending on the FortiGate model. These numbers include VLANs, other virtual interfaces, and physical interfaces. To have more than 255 interfaces configured in Transparent operating mode, you need to configure multiple VDOMs with many interfaces on each VDOM.

This example shows how to add a VLAN, `vlan_accounting` on the FortiGate unit internal interface with an IP address of 10.13.101.101.

### To add a VLAN - web-based manager

1. Go to *System > Network > Interface* and select *Create New*.

The *Type* is by default set to VLAN.

2. Enter a name for the VLAN to `vlan_accounting`.
3. Select the *Internal* interface.
4. Enter the *VLAN ID*.

The VLAN ID is a number between 1 and 4094 that allow groups of IP addresses with the same VLAN ID to be associated together.

5. Select the *Addressing Mode* of *Manual*.
6. Enter the IP address for the port of 10.13.101.101/24.
7. Set the *Administrative Access* to *HTTPS* and *SSH*.
8. Select *OK*.

### To add a VLAN - CLI

```
config system interface
 edit VLAN_1
 set interface internal
 set type vlan
 set vlanid 100
 set ip 10.13.101.101/24
 set allowaccess https ssh
 next
end
```

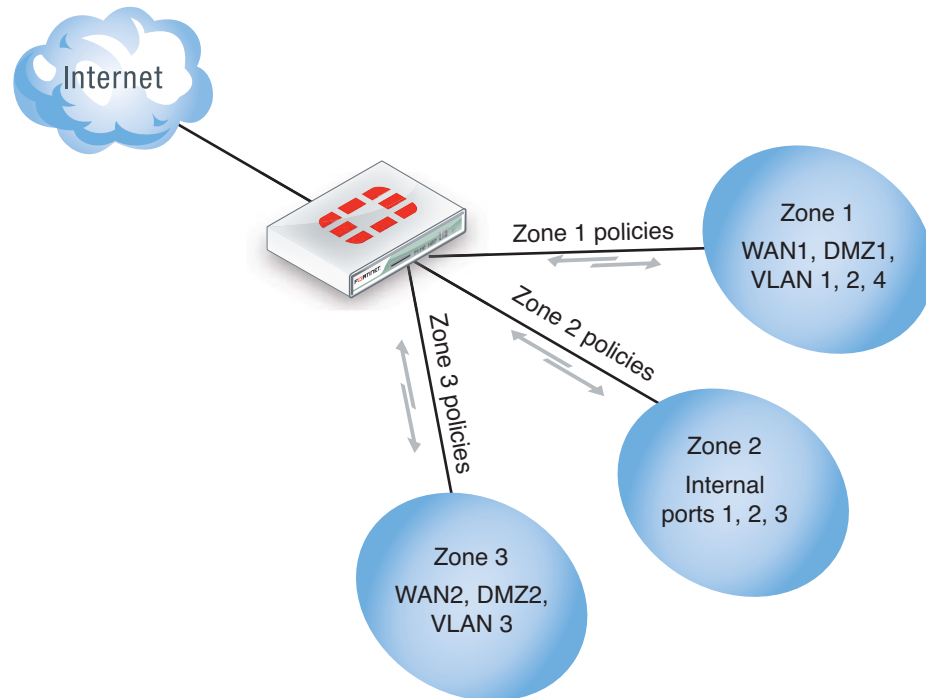
## Zones

Zones are a group of one or more FortiGate interfaces, both physical and virtual, that you can apply security policies to control inbound and outbound traffic. Grouping interfaces and VLAN subinterfaces into zones simplifies the creation of security policies where a number of network segments can use the same policy settings and protection profiles. When you add a zone, you select the names of the interfaces and VLAN subinterfaces to add to the zone. Each interface still has its own address and routing is still done between interfaces, that is, routing is not affected by zones. Security policies can also be created to control the flow of intra-zone traffic.

For example, in the illustration below, the network includes three separate groups of users representing different entities on the company network. While each group has its own set of

port and VLANs, in each area, they can all use the same security policy and protection profiles to access the Internet. Rather than the administrator making nine separate security policies, he can add the required interfaces to a zone, and create three policies, making administration simpler.

**Figure 239:**Network zones



You can configure policies for connections to and from a zone, but not between interfaces in a zone. Using the above example, you can create a security policy to go between zone 1 and zone 3, but not between WAN2 and WAN1, or WAN1 and DMZ1.

This example explains how to set up a zone to include the Internal interface and a VLAN.

**To create a zone - web-based manager**

1. Go to *System > Network > Interface*.
2. Select the arrow on the *Create New* button and select *Zone*.
3. Enter a zone name of *Zone\_1*.
4. Select the Internal interface and the virtual LAN interface *vlan\_accounting* created previously.
5. Select *OK*.

**To create a zone - CLI**

```
config system zone
 edit Zone_1
 set interface internal VLAN_1
 end
```

## Probing Interfaces

Server probes can be used on interfaces. In order for this to occur, the probe response must first be enabled and configured, then the probe response must be allowed administrative access on the interface. Both steps must be done through the CLI.

### **Enabling and configuring the probe**

```
config system probe-response
 set http-probe-port <port>
 set http-probe enable
end
```

### **Allowing the probe response to have administrative access to the interface**

```
config system interface
 edit <port>
 set allowaccess probe-response
 end
end
```

# Central management

This chapter describes the basics of using FortiManager as an administration tool for multiple FortiGate units. It describes the basics of setting up a FortiGate unit in FortiManager and some key management features you can use within FortiManager to manage the FortiGate unit.

This section includes the topics:

- [Adding a FortiGate to FortiManager](#)
- [Configuration through FortiManager](#)
- [Firmware updates](#)
- [FortiGuard](#)
- [Backup and restore configurations](#)
- [Administrative domains](#)



In order for the FortiGate unit and FortiManager unit to properly connect, both units must have compatible firmware. To find out if your firmware is compatible, refer to the FortiOS or FortiManager Release Notes.

## Adding a FortiGate to FortiManager

Before you can maintain a FortiGate unit using a FortiManager unit, you need to add it to the FortiManager. To do this requires configuration on both the FortiGate and FortiManager. This section describes the basics to configure management using a FortiManager device. For more information on the interaction of FortiManager with the FortiGate unit, see the FortiManager documentation.

### FortiGate configuration

These steps ensure that the FortiGate unit will be able to receive updated antivirus and IPS updates and allow remote management through the FortiManager system. You can add a FortiGate unit whether it is running in either NAT mode or transparent mode. The FortiManager unit provides remote management of a FortiGate unit over TCP port 541.

If you have not already done so, register the FortiGate unit by visiting <http://support.fortinet.com> and select *Product Registration*. By registering your Fortinet unit, you will receive updates to threat detection and prevention databases (Antivirus, Intrusion Detection, etc.) and will also ensure your access to technical support.

You must enable the FortiGate management option so the FortiGate unit can accept management updates to firmware, antivirus signatures, and IPS signatures.

#### **To configure the FortiGate unit - web-based manager**

1. Log in to the FortiGate unit.
2. Go to *System > Admin > Settings*.
3. Enter the *IP address* for the FortiManager unit.
4. Select *Send Request*.

The FortiManager ID now appears in the Trusted FortiManager table.

As an additional security measure, you can also select *Registration Password* and enter a password to connect to the FortiManager.

### To configure the FortiGate unit - CLI

```
config system central-management
 set fmg <ip_address>
end
```

To use the registration password enter:

```
execute central-mgmt register-device
 <fmg-serial-no><fmg-register-password><fgt-username><fgt-password>
```

## Configuring an SSL connection

An SSL connection can be configured between the two devices and an encryption level selected. Use the following CLI commands in the FortiGate CLI to configure the connection:

```
config system central-management
 set status enable
 set enc-algorithm {default* | high | low}
end
```

The default encryption automatically sets high and medium encryption algorithms. Algorithms used for high, medium, and low follows openssl definitions:

- **High** - Key lengths larger than 128 bits, and some cipher suites with 128-bit keys.  
Algorithms are: DHE-RSA-AES256-SHA:AES256-SHA: EDH-RSA-DES-CBC3-SHA: DES-CBC3-SHA:DES-CBC3-MD5:DHE-RSA-AES128-SHA:AES128-SHA
- **Medium** - Key strengths of 128 bit encryption.  
Algorithms are: RC4-SHA:RC4-MD5:RC4-MD
- **Low** - Key strengths of 64 or 56 bit encryption algorithms but excluding export cipher suites  
Algorithms are: EDH-RSA-DES-CBC-SHA; DES-CBC-SHA; DES-CBC-MD5.

## FortiManager configuration

Once the connection between the FortiGate unit and the FortiManager unit has been configured, you can add the FortiGate to the Device Manager in the FortiManager unit's web-based manager.

## Configuration through FortiManager

With the FortiManager system, you can monitor and configure multiple FortiGate units from one location and log in. Using the FortiManager's Device Manager, you can view the FortiGate units and make the usual configuration updates and changes, without having to log in and out of multiple FortiGate units.

FortiManager enables you to complete the configuration, by going to the Device Manager, selecting the FortiGate unit and using the same menu structure and pages as you would see in the FortiGate web-based manager. All changes to the FortiGate configuration are stored locally on the FortiManager unit until you synchronize with the FortiGate unit.

When a FortiGate unit is under control of a FortiManager system, administrators will not be able to change the configuration using the FortiGate. When trying to change options, the unit

displays a message that it is configured through FortiManager, and any changes may be reverted.

## Global objects

If you are maintaining a number of FortiGate units within a network, many of the policies and configuration elements will be the same across the corporation. In these instances, the adding and editing of many of the same policies will be come a tedious and error-prone activity. With FortiManager global objects, this level of configuration is simplified.

A global object is an object that is not associated specifically with one device or group. Global objects includes security policies, a DNS server, VPN, and IP pools.

The Global Objects window is where you can configure global objects and copy the configurations to the FortiManager device database for a selected device or a group of devices. You can also import configurations from the FortiManager device database for a selected device and modify the configuration as required.

When configuring or creating a global policy object the interface, prompts, and fields are the same as creating the same object on a FortiGate unit using the FortiGate web-based manager.

## Locking the FortiGate web-based manager

When you use the FortiManager to manager multiple FortiGate units, a local FortiGate unit becomes locked from any configuration changes using the web-based manager for most administrators. The super\_admin will still be able to make changes to the configuration; however, this is not recommended as it may cause conflicts with the FortiManager.

## Firmware updates

A FortiManager unit can also perform firmware updates for multiple FortiGate units, saving time rather than upgrading each FortiGate unit individually.

The FortiManager unit stores local copies of firmware images, either by downloading images from the Fortinet Distribution Network (FDN) or by accepting firmware images that are uploaded from the management computer.

If you are using the FortiManager unit to download firmware images, the FDN first validates device licenses and support contracts and then provides a list of currently available firmware images. For devices with valid Fortinet Technical Support contracts, you can download new firmware images from the FDN and the firmware release notes.

After firmware images have been either downloaded from the FDN or imported to the firmware list, you can either schedule or immediately upgrade/downgrade a device or group of device's firmware.

## FortiGuard

FortiManager can also connect to the FortiGuard Distribution Network (FDN) to receive push updates for IPS signatures and antivirus definitions. These updates can then be used to update multiple FortiGate units throughout an organization. By using the FortiManager as the host for updates, bandwidth use is minimized as updates are downloaded to one source instead of many.

To receive IPS and antivirus updates from FortiManager, indicate an alternate IP address on the FortiGate unit.



### To configure updates from FortiManager

1. Go to *System > Config > FortiGuard*.
2. Select *AntiVirus and IPS Options* to expand the options.
3. Enable both *Allow Push Update* and *Use override push IP*.
4. Enter the IP address of the FortiManager unit.
5. Select *Apply*.

## Backup and restore configurations

FortiManager stores configuration files for backup and restore purposes. FortiManager also enables you to save revisions of configuration files. Configuration backups occur automatically when the administrator logs out, the administrator login session expires, or the FortiGate restarts. Administrators can also start a backup manually.

FortiManager also enables you to view differences between different configurations to view where changes have been made.

Configure the FortiGate as follows to support backing up the configuration to FortiManager:

```
config system central-management
 set mode backup
 set fortimanager-fds-override enable
 set fmg "192.168.206.26"
end
```

On the FortiManager site, the ADOM that includes the FortiGate must be set to backup.

Enabling `fortimanager-fds-override` means that the FortiGate must use the FortiManager for FortiGuard updates and FortiGuard web filtering lookups.

## Administrative domains

FortiManager administrative domains enable the `super_admin` to create groupings of devices for configured administrators to monitor and manage. FortiManager can manage a large number of Fortinet appliances. This enables administrators to maintain managed devices specific to their geographic location or business division. This also includes FortiGate units with multiple configured VDOMs.

Each administrator is tied to an administrative domain (ADOM). When that particular administrator logs in, they see only those devices or VDOMs configured for that administrator and ADOM. The one exception is the `super_admin` account that can see and maintain all administrative domains and the devices within those domains.

Administrative domains are not enabled by default and enabling and configuring the domains can only be performed by the `super_admin`.

The maximum number of administrative domains you can add depends on the FortiManager model.

# Monitoring

With network administration, the first step is installing and configuring the FortiGate unit to be the protector of the internal network. Once the system is running efficiently, the next step is to monitor the system and network traffic, making configuration changes as necessary when a threat or vulnerability is discovered.

This chapter discusses the various methods of monitoring both the FortiGate unit and the network traffic through a range of different tools available within FortiOS.

This section includes the topics:

- [Dashboard](#)
- [sFlow](#)
- [Monitor menus](#)
- [Logging](#)
- [Alert email](#)
- [SNMP](#)

## Dashboard

The FortiOS dashboard provides a location to view real-time system information. By default, the dashboard displays the key statistics of the FortiGate unit itself, providing the memory and CPU status, as well as the health of the ports, whether they are up or down and their throughput.

## Widgets

Within the dashboard is a number of smaller windows, called widgets, that provide this status information. Beyond what is visible by default, you can add a number of other widgets that display other key traffic information including application use, traffic per IP address, top attacks, traffic history and logging statistics.

You can add multiple dashboards to reflect what data you want to monitor, and add the widgets accordingly. Dashboard configuration is only available through the web-based manager. Administrators must have read and write privileges to customize and add widgets when in either menu. Administrators must have read privileges if they want to view the information.

### To add a dashboard and widgets

1. Go to *System > Dashboard*.
2. Select the *Dashboard* menu at the top of the window and select *Add Dashboard*.
3. Enter a name for the widget.
4. Select the *Widget* menu at the top of the window.
5. From the screen, select the type of information you want to add.
6. When done, select the X in the top right of the widget.

Dashboard widgets provide an excellent method to view real-time data about the events occurring on the FortiGate unit and the network. For example, by adding the Network Protocol Usage widget, you can monitor the activity of various protocols over a selected span of time. Based on that information you can add or adjust traffic shaping and/or security policies to control traffic.

## FortiClient software

The *License Information* widget includes information for the FortiClient connections. It displays the number of FortiClient connections allowed and the number of users connecting. By selecting the *Details* link for the number of connections, you can view more information about the connecting user, including IP address, user name, and type of operating system the user is connecting with.

Included with this information is a link for Mac and Windows. Selecting these links automatically downloads the FortiClient install file (.dmg or .exe) to the management computer.

## sFlow

sFlow is a method of monitoring the traffic on your network to identify areas on the network that may impact performance and throughput. FortiOS implements sFlow version 5.

sFlow uses packet sampling to monitor network traffic. The sFlow Agent captures packet information at defined intervals and sends them to an sFlow Collector for analysis, providing real-time data analysis. The information sent is only a sampling of the data for minimal impact on network throughput and performance.

The sFlow Agent is embedded in the FortiGate unit. Once configured, the FortiGate unit sends sFlow datagrams of the sampled traffic to the sFlow Collector, also called an sFlow Analyzer. The sFlow Collector receives the datagrams, and provides real-time analysis and graphing to indicate where potential traffic issues are occurring. sFlow Collector software is available from a number of third party software vendors.

sFlow data captures only a sampling of network traffic, not all traffic like the traffic logs on the FortiGate unit. Sampling works by the sFlow Agent looking at traffic packets when they arrive on an interface. A decision is made whether the packet is dropped and allowed to be to its destination or if a copy is forwarded to the sFlow Collector. The sample used and its frequency are determined during configuration.

sFlow is not supported on virtual interfaces such as vdom link, ipsec, ssl.<vdom> or gre.

The sFlow datagram sent to the Collector contains the information:

- Packet header (e.g. MAC,IPv4,IPv6,IPX,AppleTalk,TCP,UDP, ICMP)
- Sample process parameters (rate, pool etc.)
- Input/output ports
- Priority (802.1p and TOS)
- VLAN (802.1Q)
- Source/destination prefix
- Next hop address
- Source AS, Source Peer AS
- Destination AS Path
- Communities, local preference
- User IDs (TACACS/RADIUS) for source/destination
- URL associated with source/destination
- Interface statistics (RFC 1573, RFC 2233, and RFC 2358)

sFlow agents can be added to any type of FortiGate interface. sFlow isn't supported on some virtual interfaces such as VDOM link, IPsec, gre, and ssl.<vdom>.

For more information on sFlow, Collector software and sFlow MIBs, visit [www.sflow.org](http://www.sflow.org).

## Configuration

sFlow configuration is available only from the CLI. Configuration requires two steps: enabling the sFlow Agent and configuring the interface for the sampling information.

### Enable sFlow

```
config system sflow
 set collector-ip <ip_address>
 set collector-port <port_number>
end
```

The default port for sFlow is UDP 6343. To configure in VDOM, use the commands:

```
config system vdom-sflow
 set vdom-sflow enable
 set collector-ip <ip_address>
 set collector-port <port_number>
end
```

Configure sFlow agents per interface.

```
config system interface
 edit <interface_name>
 set sflow-sampler enable
 set sample-rate <every_n_packets>
 set sample-direction [tx | rx | both]
 set polling-interval <seconds>
 end
```

## Monitor menus

The *Monitor* menus enable you to view session and policy information and other activity occurring on your FortiGate unit. The monitors provide the details of user activity, traffic and policy usage to show live activity. Monitors are available for DHCP, routing, security policies, traffic shaping, load balancing, security features, VPN, users, WiFi, and logging.

## Logging

FortiOS provides a robust logging environment that enables you to monitor, store, and report traffic information and FortiGate events, including attempted log ins and hardware status. Depending on your requirements, you can log to a number of different hosts.

To configure logging in the web-based manager, go to *Log & Report > Log Config > Log Settings*.

To configure logging in the CLI use the commands `config log <log_location>`.

For details on configuring logging see the [Logging and Reporting Guide](#).

If you will be using several FortiGate units, you can also use a FortiAnalyzer unit for logging. For more information, see the [FortiAnalyzer Administration Guide](#).

## FortiCloud

The FortiCloud is a subscription-based hosted service. With this service, you can have centralized management, logging, and reporting capabilities available in FortiAnalyzer and FortiManager platforms, without any additional hardware to purchase, install or maintain. In most cases, FortiCloud is the recommended location for saving and viewing logs.

This service includes a full range of reporting, analysis and logging, firmware management and configuration revision history. It is hosted within the Fortinet global FortiGuard Network for maximum reliability and performance, and includes reporting, and drill-down analysis widgets makes it easy to develop custom views of network and security events.

The FortiGate unit sends log messages to the FortiCloud using TCP port 443. Configuration is available once a user account has been set up and confirmed. To enable the account on the FortiGate unit, go to *System > Dashboard > Status*, select *Activate*, and enter the account ID.

For FortiCloud traffic, you can identify a specific port/IP address for logging traffic. Configuration of these services is performed in the CLI, using the command `set source-ip`. When configured, this becomes the dedicated port to send this traffic over.

For example, to set the source IP of the FortiCloud server to be on the DMZ1 port with an IP of 192.168.4.5, the commands are:

```
config log fortiguard setting
 set status enable
 set source-ip 192.168.4.5
end
```

From the FortiGate unit, you can configure the connection and sending of log messages to be sent over an SSL tunnel to ensure log messages are sent securely. To do this, use the CLI commands to enable the encrypted connection and define the level of encryption.

```
config log fortiguard setting
 set status enable
 set enc-algorithm {default | high | low | disable}
end
```

For more information on each encryption level see [“Configuring an SSL connection” on page 1511](#).

## FortiGate memory

Logs are saved to the internal memory by default. Inexpensive yet volatile, for basic event logs or verifying traffic, AV or spam patterns, logging to memory is a simple option. However, because logs are stored in the limited space of the internal memory, only a small amount is available for logs. As such logs can fill up and be overridden with new entries, negating the use of recursive data. This is especially true for traffic logs. Also, should the FortiGate unit be shut down or rebooted, all log information will be lost.

## FortiGate hard disk

For those FortiGate units with an internal hard disk or SDHC card, you can store logs to this location. Efficient and local, the hard disk provides a convenient storage location. If you choose to store logs in this manner, remember to backup the log data regularly.

Configure log disk settings is performed in the CLI using the commands:

```
config log disk setting
 set status enable
end
```

Further options are available when enabled to configure log file sizes, and uploading/backup events.

As well, note that the write speeds of hard disks compared to the logging of ongoing traffic may cause the dropping such, it is recommended that traffic logging be sent to a FortiAnalyzer or other device meant to handle large volumes of data.

## Syslog server

An industry standard for collecting log messages, for off-site storage. In the web-based manager, you are able to send logs to a single syslog server, however in the CLI you can configure up to three syslog servers where you can also use multiple configuration options. For example, send traffic logs to one server, antivirus logs to another. The FortiGate unit sends Syslog traffic over UDP port 514. Note that if a secure tunnel is configured for communication to a FortiAnalyzer unit, then Syslog traffic will be sent over an IPSec connection, using UDP 500/4500, protocol IP/50.

To configure a Syslog server in the web-based manager, go to *Log & Report > Log Config > Log Settings*. In the CLI use the commands:

```
config log syslogd setting
 set status enable
end
```

Further options are available when enabled to configure a different port, facility and server IP address.

For Syslog traffic, you can identify a specific port/IP address for logging traffic. Configuration of these services is performed in the CLI, using the command `set source-ip`. When configured, this becomes the dedicated port to send this traffic over.

For example, to set the source IP of a Syslog server to be on the DMZ1 port with an IP of 192.168.4.5, the commands are:

```
config log syslogd setting
 set status enable
 set source-ip 192.168.4.5
end
```

## FortiAnalyzer

The FortiAnalyzer family of logging, analyzing, and reporting appliances securely aggregate log data from Fortinet devices and other syslog-compatible devices. Using a comprehensive suite of easily-customized reports, users can filter and review records, including traffic, event, virus, attack, Web content, and email data, mining the data to determine your security stance and assure regulatory compliance. FortiAnalyzer also provides advanced security management functions such as quarantined file archiving, event correlation, vulnerability assessments, traffic analysis, and archiving of email, Web access, instant messaging and file transfer content.

The FortiGate unit sends log messages over UDP port 514 or OFTP (TCP 514). If a secure connection has been configured, log traffic is sent over UDP port 500/4500, Protocol IP/50. For more information on configuring a secure connection see [“Sending logs using a secure connection” on page 1511](#).

For FortiAnalyzer traffic, you can identify a specific port/IP address for logging traffic. Configuration of these services is performed in the CLI, using the command `set source-ip`. When configured, this becomes the dedicated port to send this traffic over.

For example, to set the source IP of a FortiAnalyzer unit to be on port 3 with an IP of 192.168.21.12, the commands are:

```
config log fortiguard setting
 set status enable
 set source-ip 192.168.21.12
end
```

## Sending logs using a secure connection

From the FortiGate unit, you can configure the connection and sending of log messages over an SSL tunnel to ensure log messages are sent securely. To do this, use the CLI commands below to enable the encrypted connection and define the level of encryption.



You must configure the secure tunnel on **both** ends of the tunnel, the FortiGate unit and the FortiAnalyzer unit.

This configuration is for FortiAnalyzer OS version 4.0 MR2 or lower. For version 4.0 MR3, see “Configuring an SSL connection” on page 1511.

### To configure a secure connection to the FortiAnalyzer unit

On the FortiAnalyzer unit, enter the commands:

```
config log device
 edit <device_name>
 set secure psk
 set psk <name_of_IPSec_tunnel>
 set id <fortigate_device_name_on_the_fortianalyzer>
 end
```

### To configure a secure connection on the FortiGate unit

On the FortiGate CLI, enter the commands:

```
config log fortianalyzer setting
 set status enable
 set server <ip_address>
 set local
 set localid <name_of_IPSec_tunnel>
end
```

## Configuring an SSL connection

An SSL connection can be configured between the two devices, and an encryption level selected.

Use the CLI commands to configure the encryption connection:

```
config log fortianalyzer setting
 set status enable
 set enc-algorithm {default* | high | low | disable}
end
```

The default encryption automatically sets high and medium encryption algorithms. Algorithms used for high, medium, and low follows openssl definitions:

- **High** - Key lengths larger than 128 bits, and some cipher suites with 128-bit keys.  
Algorithms are: DHE-RSA-AES256-SHA:AES256-SHA: EDH-RSA-DES-CBC3-SHA: DES-CBC3-SHA:DES-CBC3-MD5:DHE-RSA-AES128-SHA:AES128-SHA
- **Medium** - Key strengths of 128 bit encryption.  
Algorithms are: RC4-SHA:RC4-MD5:RC4-MD
- **Low** - Key strengths of 64 or 56 bit encryption algorithms but excluding export cipher suites  
Algorithms are: EDH-RSA-DES-CBC-SHA; DES-CBC-SHA; DES-CBC-MD5.

If you want to use an IPSec tunnel to connect to the FortiAnalyzer unit, you need to first disable the enc-algorithm:

```
config log fortianalyzer setting
 set status enable
 set enc-algorithm disable
```

Then set the IPSec encryption:

```
set encrypt enable
 set psksecret <preshared_IPSec_tunnel_key>
end
```

## Packet Capture

When troubleshooting networks, it helps to look inside the header of the packets. This helps to determine if the packets, route, and destination are all what you expect. Packet capture can also be called a network tap, packet sniffing, or logic analyzing.

### To use the packet capture.

1. Go to *System > Network > Packet Capture*.
2. Select the interface to monitor and select the number of packets to keep.
3. Select *Enable Filters*.
4. Enter the information you want to gather from the packet capture.
5. Select *OK*.

To run the capture, select the play button in the progress column in the packet capture list. If not active, *Not Running* will also appear in the column cell. The progress bar will indicate the status of the capture. You can stop and restart it at any time.

When the capture is complete, select the *Download* icon to save the packet capture file to your hard disk for further analysis.

Packet capture tells you what is happening on the network at a low level. This can be very useful for troubleshooting problems, such as:

- finding missing traffic
- seeing if sessions are setting up properly
- locating ARP problems such as broadcast storm sources and causes
- confirming which address a computer is using on the network if they have multiple addresses or are on multiple networks
- confirming routing is working as you expect
- wireless client connection problems



- intermittent missing PING packets
- a particular type of packet is having problems, such as UDP, which is commonly used for streaming video

If you are running a constant traffic application such as ping, packet capture can tell you if the traffic is reaching the destination, how the port enters and exits the FortiGate unit, if the ARP resolution is correct, and if the traffic is returning to the source as expected. You can also use packet switching to verify that NAT or other configuration is translating addresses or routing traffic the way that you want it to.

Before you start capturing packets, you need to have a good idea of what you are looking for. Capture is used to confirm or deny your ideas about what is happening on the network. If you try capture without a plan to narrow your search, you could end up with too much data to effectively analyze. On the other hand, you need to capture enough packets to really understand all of the patterns and behavior that you are looking for.

## Alert email

As an administrator, you want to be certain you can respond quickly to issues occurring on your network or on the FortiGate unit. Alert emails provide an efficient and direct method of notifying an administrator of events. By configuring alert messages, you can define the threshold when a problem becomes critical and needs attention. When this threshold is reached, the FortiGate unit will send an email to one or more individuals, notifying them of the issue.

In the following example, the FortiGate unit is configured to send email to two administrators (admin1 and admin2) when multiple intrusions are detected every two minutes. The FortiGate unit has its own email address on the mail server.

### To configure the email service

1. Go to *System > Config > Messaging Servers*.
2. Complete the following and select *Apply*:

<b>SMTP Server</b>	Enter the address or name of the email server. For example, <code>smtp.example.com</code> .
<b>Default Reply To</b>	Enter an email address to associate with the alert email. This field is optional. If you enter an email address here, it overrides the email address entered when configuring alert email in <i>Log &amp; Report &gt; Alert E-mail</i> .
<b>Authentication</b>	Enable authentication if required by the email server.
<b>SMTP User</b>	FortiGate
<b>Password</b>	*****

### To configure alert email - web-based manager

1. Go to *Log & Report > Log Config > Alert E-mail*.
2. Enter the information:

<b>Email from</b>	fortigate@example.com
<b>Email to</b>	admin1@example.com admin2@example.com

3. For the *Interval Time*, enter 2.
4. Select *Intrusion Detected*.
5. Select *Apply*.

#### To configure alert email - CLI

```
config system email-server
 set port 25
 set server smtp.example.com
 set authenticate enable
 set username FortiGate
 set password *****
end
config alertemail setting
 set username fortigate@example.com
 set mailto1 admin1@example.com
 set mailto2 admin2@example.com
 set filter category
 set IPS-logs enable
end
```

## SNMP

Simple Network Management Protocol (SNMP) enables you to monitor hardware on your network. You can configure the hardware, such as the FortiGate SNMP agent, to report system information and send traps (alarms or event messages) to SNMP managers. An SNMP manager, or host, is typically a computer running an application that can read the incoming trap and event messages from the agent and send out SNMP queries to the SNMP agents. A FortiManager unit can act as an SNMP manager to one or more FortiGate units. FortiOS supports SNMP using IPv4 and IPv6 addressing.

By using an SNMP manager, you can access SNMP traps and data from any FortiGate interface or VLAN subinterface configured for SNMP management access. Part of configuring an SNMP manager is to list it as a host in a community on the FortiGate unit it will be monitoring. Otherwise, the SNMP monitor will not receive any traps from that FortiGate unit or be able to query that unit.

The FortiGate SNMP implementation is read-only. SNMP v1, v2c, and v3 compliant SNMP managers have read-only access to FortiGate system information through queries and can receive trap messages from the FortiGate unit.

To monitor FortiGate system information and receive FortiGate traps, you must first compile the Fortinet and FortiGate Management Information Base (MIB) files. A MIB is a text file that describes a list of SNMP data objects that are used by the SNMP manager. These MIBs provide information the SNMP manager needs to interpret the SNMP trap, event, and query messages sent by the FortiGate unit SNMP agent.

FortiGate core MIB files are available for download by going to *System > Config > SNMP* and selecting the download link on the page.

The Fortinet implementation of SNMP includes support for most of RFC 2665 (Ethernet-like MIB) and most of RFC 1213 (MIB II). For more information, see [“Fortinet MIBs” on page 1520](#). RFC support for SNMP v3 includes Architecture for SNMP Frameworks (RFC 3411), and partial support of User-based Security Model (RFC 3414).

SNMP traps alert you to events that occur such as an a full log disk or a virus detected.

SNMP fields contain information about the FortiGate unit, such as CPU usage percentage or the number of sessions. This information is useful for monitoring the condition of the unit on an ongoing basis and to provide more information when a trap occurs.

The FortiGate SNMP v3 implementation includes support for queries, traps, authentication, and privacy. Authentication and encryption are configured in the CLI. See the `system snmp user` command in the [FortiGate CLI Reference](#).

## SNMP configuration settings

Before a remote SNMP manager can connect to the FortiGate agent, you must configure one or more FortiGate interfaces to accept SNMP connections by going to *System > Network > Interface*. Select the interface and, in the *Administrative Access*, select *SNMP*.

For VDOMS, SNMP traps can only be sent on interfaces in the management VDOM. Traps cannot be sent over other interfaces outside the management VDOM.

To configure SNMP settings, go to *System > Config > SNMP*.

<b>SNMP Agent</b>	Select to enable SNMP communication.
<b>Description</b>	Enter descriptive information about the FortiGate unit. The description can be up to 35 characters.
<b>Location</b>	Enter the physical location of the FortiGate unit. The system location description can be up to 35 characters long.
<b>Contact</b>	Enter the contact information for the person responsible for this FortiGate unit. The contact information can be up to 35 characters.

### SNMP v1/v2c section

To create a new SNMP community, see [New SNMP Community page](#).

<b>Community Name</b>	The name to identify the community.
<b>Queries</b>	Indicates whether queries protocols (v1 and v2c) are enabled or disabled. A green check mark indicates queries are enabled; a gray x indicates queries are disabled. If one query is disabled and another one enabled, there will still be a green check mark.
<b>Traps</b>	Indicates whether trap protocols (v1 and v2c) are enabled or disabled. A green check mark indicates traps are enabled; a gray x indicates traps are disabled. If one query is disabled and another one enabled, there will still be a green check mark.
<b>Enable</b>	Select the check box to enable or disable the community.

### SNMP v3 section

To create a new SNMP community, see [Create New SNMP V3 User](#).

<b>User Name</b>	The name of the SNMPv3 user.
<b>Security Level</b>	The security level of the user.
<b>Notification Host</b>	The IP address or addresses of the host.
<b>Queries</b>	Indicates whether queries are enabled or disabled. A green check mark indicates queries are enabled; a gray x indicates queries are disabled

---

### ***New SNMP Community page***

---

**Community Name** Enter a name to identify the SNMP community

---

#### **Hosts (section)**

---

**IP Address** Enter the IP address and Identify the SNMP managers that can use the settings in this SNMP community to monitor the FortiGate unit.

You can also set the IP address to 0.0.0.0 to so that any SNMP manager can use this SNMP community.

---

**Interface** Optionally, select the name of the interface that this SNMP manager uses to connect to the FortiGate unit. You only have to select the interface if the SNMP manager is not on the same subnet as the FortiGate unit. This can occur if the SNMP manager is on the Internet or behind a router.

In virtual domain mode, the interface must belong to the management VDOM to be able to pass SNMP traps.

---

**Delete** Removes an SNMP manager from the list within the *Hosts* section.

---

**Add** Select to add a blank line to the Hosts list. You can add up to eight SNMP managers to a single community.

---

#### **Queries (section)**

---

**Protocol** The SNMP protocol. In the v1 row, this means that the settings are for SNMP v1. In the v2c row, this means that the settings are for SNMP v2c.

---

**Port** Enter the port number (161 by default) that the SNMP managers in this community use for SNMP v1 and SNMP v2c queries to receive configuration information from the FortiGate unit. Select the *Enable* check box to activate queries for each SNMP version.

**Note:** The SNMP client software and the FortiGate unit must use the same port for queries.

---

**Enable** Select to enable that SNMP protocol.

---

#### **Traps (section)**

---

**Protocol** The SNMP protocol. In the v1 row, this means that the settings are for SNMP v1. In the v2c row, this means that the settings are for SNMP v2c.

---

**Local** Enter the remote port numbers (port 162 for each by default) that the FortiGate unit uses to send SNMP v1 or SNMP v2c traps to the SNMP managers in this community. Select the *Enable* check box to activate traps for each SNMP version.

**Note:** The SNMP client software and the FortiGate unit must use the same port for traps.

---

**Remote** Enter the remote port number (port 162 is default) that the FortiGate unit uses to send SNMP v1 or v2c traps to the SNMP managers in this community.

**Note:** The SNMP client software and the FortiGate unit must use the same port for queries.

---

<b>Enable</b>	Select to activate traps for each SNMP version.
<b>SNMP Event</b>	<p>Enable each SNMP event for which the FortiGate unit should send traps to the SNMP managers in this community.</p> <p><i>CPU Over usage</i> traps sensitivity is slightly reduced, by spreading values out over 8 polling cycles. This prevents sharp spikes due to CPU intensive short-term events such as changing a policy.</p> <p><i>Power Supply Failure</i> event trap is available only on some models.</p> <p><i>AMC interfaces enter bypass mode</i> event trap is available only on models that support AMC modules.</p>
<b>Enable</b>	Select to enable the SNMP event.
<b>Create New SNMP V3 User</b>	
<b>User Name</b>	Enter the name of the user.
<b>Security Level</b>	Select the type of security level the user will have.
<b>Notification Host</b>	Enter the IP address of the notification host. If you want to add more than one host, after entering the IP address of the first host, select the plus sign to add another host.
<b>Enable Query</b>	Select to enable or disable the query. By default, the query is enabled.
<b>Port</b>	Enter the port number in the field.
<b>Events</b>	Select the SNMP events that will be associated with that user.

## Gigabit interfaces

When determining the interface speed of a FortiGate unit with a 10G interface, the IF-MIB.ifSpeed may not return the correct value. IF-MIB.ifSpeed is a 32-bit gauge used to report interface speeds in bits/second and cannot convert to a 64-bit value. The 32-bit counter wrap the output too fast to be accurate.

In this case, you can use the value ifHighSpeed. It reports interface speeds in megabits/second. This ensures that 10Gb interfaces report the correct value.

## SNMP agent

You need to first enter information and enable the FortiGate SNMP Agent. Enter information about the FortiGate unit to identify it so that when your SNMP manager receives traps from the FortiGate unit, you will know which unit sent the information.

### To configure the SNMP agent - web-based manager

1. Go to *System > Config > SNMP*.
2. Select *Enable* for the *SNMP Agent*.
3. Enter a descriptive name for the agent.
4. Enter the location of the FortiGate unit.
5. Enter a contact or administrator for the SNMP Agent or FortiGate unit.
6. Select *Apply*.

### To configure SNMP agent - CLI

```
config system snmp sysinfo
 set status enable
 set contact-info <contact_information>
 set description <description_of_FortiGate>
 set location <FortiGate_location>
end
```

## SNMP community

An SNMP community is a grouping of devices for network administration purposes. Within that SNMP community, devices can communicate by sending and receiving traps and other information. One device can belong to multiple communities, such as one administrator terminal monitoring both a firewall SNMP and a printer SNMP community.

Add SNMP communities to your FortiGate unit so that SNMP managers can connect to view system information and receive SNMP traps.

You can add up to three SNMP communities. Each community can have a different configuration for SNMP queries and traps. Each community can be configured to monitor the FortiGate unit for a different set of events. You can also add the IP addresses of up to 8 SNMP managers to each community.

When the FortiGate unit is in virtual domain mode, SNMP traps can only be sent on interfaces in the management virtual domain. Traps cannot be sent over other interfaces.

### To add an SNMP v1/v2c community - web-based manager

1. Go to *System > Config > SNMP*.
2. In the *SNMP v1/v2c* area, select *Create New*.
3. Enter a *Community Name*.
4. Enter the IP address and identify the SNMP managers that can use the settings in this SNMP community to monitor the FortiGate unit.
5. Select the interface if the SNMP manager is not on the same subnet as the FortiGate unit.
6. Enter the *Port* number that the SNMP managers in this community use for SNMP v1 and SNMP v2c queries to receive configuration information from the FortiGate unit. Select the *Enable* check box to activate queries for each SNMP version.
7. Enter the Local and Remote port numbers that the FortiGate unit uses to send SNMP v1 and SNMP v2c traps to the SNMP managers in this community.
8. Select the *Enable* check box to activate traps for each SNMP version.
9. Select *OK*.

### To add an SNMP v1/v2c community - CLI

```
config system snmp community
 edit <index_number>
 set events <events_list>
 set name <community_name>
 set query-v1-port <port_number>
 set query-v1-status {enable | disable}
 set query-v2c-port <port_number>
 set query-v2c-status {enable | disable}
 set status {enable | disable}
 set trap-v1-lport <port_number>
 set trap-v1-rport <port_number>
 set trap-v1-status {enable | disable}
 set trap-v2c-lport <port_number>
 set trap-v2c-rport <port_number>
 set trap-v2c-status {enable | disable}
 end
```

### To add an SNMP v3 community - web-based manager

1. Go to *System > Config > SNMP*.
2. In the *SNMP v3* area, select *Create New*.
3. Enter a *User Name*.
4. Select a *Security Level* and associated authorization algorithms.
5. Enter the IP address of the *Notification Host* SNMP managers that can use the settings in this SNMP community to monitor the FortiGate unit.
6. Enter the *Port* number that the SNMP managers in this community use to receive configuration information from the FortiGate unit. Select the *Enable* check box to activate queries for each SNMP version.
7. Select the *Enable* check box to activate traps.
8. Select *OK*.

### To add an SNMP v3 community - CLI

```
config system snmp user
 edit <index_number>
 set security-level [auth-priv | auth-no-priv | no-auth-no-priv]
 set queries enable
 set query-port <port_number>
 set notify-hosts <ip_address>
 set events <event_selections>
 end
```

## Enabling on the interface

Before a remote SNMP manager can connect to the FortiGate agent, you must configure one or more FortiGate interfaces to accept SNMP connections.

### To configure SNMP access - web-based manager

1. Go to *System > Network > Interface*.
2. Choose an interface that an SNMP manager connects to and select *Edit*.

3. In *Administrative Access*, select *SNMP*.
4. Select *OK*.

#### To configure SNMP access - CLI

```
config system interface
 edit <interface_name>
 set allowaccess snmp
 end
```



If the interface you are configuring already has protocols that are allowed access, use the command `append allowaccess snmp` instead, or else the other protocols will be replaced. For more information, see [“Adding and removing options from lists” on page 1417](#).

---

## Fortinet MIBs

The FortiGate SNMP agent supports Fortinet proprietary MIBs as well as standard RFC 1213 and RFC 2665 MIBs. RFC support includes support for the parts of RFC 2665 (Ethernet-like MIB) and the parts of RFC 1213 (MIB II) that apply to FortiGate unit configuration.

There are two MIB files for FortiGate units - the Fortinet MIB, and the FortiGate MIB. The Fortinet MIB contains traps, fields and information that is common to all Fortinet products. The FortiGate MIB contains traps, fields and information that is specific to FortiGate units. Each Fortinet product has its own MIB. If you use other Fortinet products you will need to download their MIB files as well. Both MIB files are used for FortiOS and FortiOS Carrier; there are no additional traps for the Carrier version of the operating system.

The Fortinet MIB and FortiGate MIB along with the two RFC MIBs are listed in tables in this section. You can download the two FortiGate MIB files from Fortinet Customer Support. The Fortinet MIB contains information for Fortinet products in general. the Fortinet FortiGate MIB includes the system information for The FortiGate unit and version of FortiOS. Both files are required for proper SNMP data collection.

To download the MIB files, go to *System > Config > SNMP* and select a MIB link in the *FortiGate SNMP MIB* section.

Your SNMP manager may already include standard and private MIBs in a compiled database that is ready to use. You must add the Fortinet proprietary MIB to this database to have access to the Fortinet specific information.

---



There were major changes to the MIB files between v3.0 and v4.0. You need to use the new MIBs for v4.0 or you may mistakenly access the wrong traps and fields.

MIB files are updated for each version of FortiOS. When upgrading the firmware ensure that you updated the Fortinet FortiGate MIB file as well.

---



**Table 68:** Fortinet MIBs

MIB file name or RFC	Description
<b>FORTINET-CORE-MIB.mib</b>	<p>The Fortinet MIB includes all system configuration information and trap information that is common to all Fortinet products.</p> <p>Your SNMP manager requires this information to monitor FortiGate unit configuration settings and receive traps from the FortiGate SNMP agent.</p>
<b>FORTINET-FORTIGATE-MIB.mib</b>	<p>The FortiGate MIB includes all system configuration information and trap information that is specific to FortiGate units.</p> <p>Your SNMP manager requires this information to monitor FortiGate configuration settings and receive traps from the FortiGate SNMP agent. FortiManager systems require this MIB to monitor FortiGate units.</p>
<b>RFC-1213 (MIB II)</b>	<p>The FortiGate SNMP agent supports MIB II groups with these exceptions.</p> <ul style="list-style-type: none"> <li>• No support for the EGP group from MIB II (RFC 1213, section 3.11 and 6.10).</li> <li>• Protocol statistics returned for MIB II groups (IP/ICMP/TCP/UDP/etc.) do not accurately capture all FortiGate traffic activity. More accurate information can be obtained from the information reported by the Fortinet MIB.</li> </ul>
<b>RFC-2665 (Ethernet-like MIB)</b>	<p>The FortiGate SNMP agent supports Ethernet-like MIB information. FortiGate SNMP does not support for the dot3Tests and dot3Errors groups.</p>

## SNMP get command syntax

Normally, to get configuration and status information for a FortiGate unit, an SNMP manager would use an SNMP get commands to get the information in a MIB field. The SNMP get command syntax would be similar to:

```
snmpget -v2c -c <community_name> <address_ipv4> {<OID> | <MIB_field>}
```

...where...

<community\_name> is an SNMP community name added to the FortiGate configuration. You can add more than one community name to a FortiGate SNMP configuration. The most commonly used community name is `public`.

<address\_ipv4> is the IP address of the FortiGate interface that the SNMP manager connects to.

{<OID> | <MIB\_field>} is the object identifier (OID) for the MIB field or the MIB field name itself.

The SNMP `get` command gets firmware version running on the FortiGate unit. The community name is `public`. The IP address of the interface configured for SNMP management access is `10.10.10.1`. The firmware version MIB field is `fgSysVersion` and the OID for this MIB field is

1.3.6.1.4.1.12356.101.4.1.1 The first command uses the MIB field name and the second uses the OID:

```
snmpget -v2c -c public 10.10.10.1 fgSysVersion.0
```

```
snmpget -v2c -c public 10.10.10.1 1.3.6.1.4.1.12356.101.4.1.1.0
```



The OIDs and object names used in these examples are dependent on the version of MIB and are subject to change.

---

# VLANs

Virtual Local Area Networks (VLANs) multiply the capabilities of your FortiGate unit, and can also provide added network security. Virtual LANs (VLANs) use ID tags to logically separate devices on a network into smaller broadcast domains. These smaller domains forward packets only to devices that are part of that VLAN domain. This reduces traffic and increases network security.

A Local Area Network (LAN) is a group of connected computers and devices that are arranged into network broadcast domains. A LAN broadcast domain includes all the computers that receive a packet broadcast from any computer in that broadcast domain. A switch will automatically forward the packets to all of its ports; in contrast, routers do not automatically forward network broadcast packets. This means routers separate broadcast domains. If a network has only switches and no routers, that network is considered one broadcast domain, no matter how large or small it is. Smaller broadcast domains are more efficient because fewer devices receive unnecessary packets. They are more secure as well because a hacker reading traffic on the network will have access to only a small portion of the network instead of the entire network's traffic.

Virtual LANs (VLANs) use ID tags to logically separate a LAN into smaller broadcast domains. Each VLAN is its own broadcast domain. Smaller broadcast domains reduce traffic and increase network security. The IEEE 802.1Q standard defines VLANs. All layer-2 and layer-3 devices along a route must be 802.1Q-compliant to support VLANs along that route. For more information, see [“VLAN switching and routing” on page 1524](#) and [“VLAN layer-3 routing” on page 1527](#).

VLANs reduce the size of the broadcast domains by only forwarding packets to interfaces that are part of that VLAN or part of a VLAN trunk link. Trunk links form switch-to-switch or switch-to-router connections, and forward traffic for all VLANs. This enables a VLAN to include devices that are part of the same broadcast domain, but physically distant from each other.

VLAN ID tags consist of a 4-byte frame extension that switches and routers apply to every packet sent and received in the VLAN. Workstations and desktop computers, which are commonly originators or destinations of network traffic, are not an active part of the VLAN process. All the VLAN tagging and tag removal is done after the packet has left the computer. For more information, see [“VLAN ID rules” on page 1524](#).

Any FortiGate unit without VDOMs enabled can have a maximum of 255 interfaces in transparent operating mode. The same is true for any single VDOM. In NAT mode, the number can range from 255 to 8192 interfaces per VDOM, depending on the FortiGate model. These numbers include VLANs, other virtual interfaces, and physical interfaces. To have more than 255 interfaces configured in transparent operating mode, you need to configure multiple VDOMs that enable you to divide the total number of interfaces over all the VDOMs.

One example of an application of VLANs is a company's accounting department. Accounting computers may be located at both main and branch offices. However, accounting computers need to communicate with each other frequently and require increased security. VLANs allow the accounting network traffic to be sent only to accounting computers and to connect accounting computers in different locations as if they were on the same physical subnet.

This guide uses the term “packet” to refer to both layer-2 frames and layer-3 packets.

## VLAN ID rules

Layer-2 switches and layer-3 devices add VLAN ID tags to the traffic as it arrives and remove them before they deliver the traffic to its final destination. Devices such as PCs and servers on the network do not require any special configuration for VLANs. Twelve bits of the 4-byte VLAN tag are reserved for the VLAN ID number. Valid VLAN ID numbers are from 1 to 4094, while 0 is used for high priority frames, and 4095 is reserved.

On a layer-2 switch, you can have only one VLAN subinterface per physical interface, unless that interface is configured as a trunk link. Trunk links can transport traffic for multiple VLANs to other parts of the network.

On a FortiGate unit, you can add multiple VLANs to the same physical interface. However, VLAN subinterfaces added to the same physical interface cannot have the same VLAN ID or have IP addresses on the same subnet. You can add VLAN subinterfaces with the same VLAN ID to different physical interfaces.

Creating VLAN subinterfaces with the same VLAN ID does not create any internal connection between them. For example a VLAN ID of 300 on port1 and VLAN ID of 300 on port2 are allowed, but they are not connected. Their relationship is the same as between any two FortiGate network interfaces.

## VLAN switching and routing

VLAN switching takes place on the OSI model layer-2, just like other network switching. VLAN routing takes place on the OSI model layer-3. The difference between them is that during VLAN switching, VLAN packets are simply forwarded to their destination. This is different from VLAN routing where devices can open the VLAN packets and change their VLAN ID tags to route the packets to a new destination.

### VLAN layer-2 switching

Ethernet switches are layer-2 devices, and generally are 802.1Q compliant. Layer 2 refers to the second layer of the seven layer Open Systems Interconnect (OSI) basic networking model; the Data Link layer. FortiGate units act as layer-2 switches or bridges when they are in transparent mode. The units simply tag and forward the VLAN traffic or receive and remove the tags from the packets. A layer-2 device does not inspect incoming packets or change their contents; it only adds or removes tags and routes the packet.

A VLAN can have any number of physical interfaces assigned to it. Multiple VLANs can be assigned to the same physical interface. Typically two or more physical interfaces are assigned to a VLAN, one for incoming and one for outgoing traffic. Multiple VLANs can be configured on one FortiGate unit, including trunk links.

### Layer-2 VLAN example


To better understand VLAN operation, this example shows what happens to a data frame on a network that uses VLANs.

The network topology consists of two 8-port switches that are configured to support VLANs on a network. Both switches are connected through port 8 using an 802.1Q trunk link. Subnet 1 is connected to switch A, and subnet 2 is connected to switch B. The ports on the switches are configured as follows.

**Table 69:** How ports and VLANs are used on Switch A and B


Switch	Ports	VLAN
A	1 - 4	100
A	5 - 7	200
A & B	8	Trunk link
B	4 - 5	100
B	6	200

In this example, switch A is connected to the Branch Office and switch B to the Main Office.




1. A computer on port 1 of switch A sends a data frame over the network.
2. Switch A tags the data frame with a VLAN 100 ID tag upon arrival because port 1 is part of VLAN 100.
3. Switch A forwards the tagged data frame to the other VLAN 100 ports — ports 2 through 4. Switch A also forwards the data frame to the 802.1Q trunk link (port 8) so other parts of the network that may contain VLAN 100 groups will receive VLAN 100 traffic.

This data frame is not forwarded to the other ports on switch A because they are not part of VLAN 100. This increases security and decreases network traffic.



4. Switch B receives the data frame over the trunk link (port 8).
5. Because there are VLAN 100 ports on switch B (ports 4 and 5), the data frame is forwarded to those ports. As with switch A, the data frame is not delivered to VLAN 200.

If there were no VLAN 100 ports on switch B, the switch would not forward the data frame and it would stop there.



6. The switch removes the VLAN 100 ID tag before it forwards the data frame to an end destination.

The sending and receiving computers are not aware of any VLAN tagging on the data frames that are being transmitted. When any computer receives that data frame, it appears as a normal data frame.

## VLAN layer-3 routing

Routers are layer-3 devices. Layer 3 refers to the third layer of the OSI networking model, the Network layer. FortiGate units in NAT mode act as layer-3 devices. As with layer 2, FortiGate units acting as layer-3 devices are 802.1Q-compliant.

The main difference between layer-2 and layer-3 devices is how they process VLAN tags. Layer-2 switches just add, read and remove the tags. They do not alter the tags or do any other high-level actions. Layer-3 routers not only add, read and remove tags but also analyze the data frame and its contents. This analysis allows layer-3 routers to change the VLAN tag if it is appropriate and send the data frame out on a different VLAN.

In a layer-3 environment, the 802.1Q-compliant router receives the data frame and assigns a VLAN ID. The router then forwards the data frame to other members of the same VLAN broadcast domain. The broadcast domain can include local ports, layer-2 devices and layer-3 devices such as routers and firewalls. When a layer-3 device receives the data frame, the device removes the VLAN tag and examines its contents to decide what to do with the data frame. The layer-3 device considers:

- source and destination addresses
- protocol
- port number.


The data frame may be forwarded to another VLAN, sent to a regular non-VLAN-tagged network or just forwarded to the same VLAN as a layer-2 switch would do. Or, the data frame may be discarded if the proper security policy has been configured to do so.

### Layer-3 VLAN example


In this example, switch A is connected to the Branch Office subnet, the same as subnet 1 in the layer-2 example. In the Main Office subnet, VLAN 300 is on port 5 of switch B. The FortiGate unit is connected to switch B on port 1 and the trunk link connects the FortiGate unit's port 3 to switch A. The other ports on switch B are unassigned.

This example explains how traffic can change VLANs originating on VLAN 100 and arriving at a destination on VLAN 300. Layer-2 switches alone cannot accomplish this, but a layer-3 router can.

1. The VLAN 100 computer at the Branch Office sends the data frame to switch A, where the VLAN 100 tag is added.



2. Switch A forwards the tagged data frame to the FortiGate unit over the 802.1Q trunk link, and to the VLAN 100 interfaces on Switch A. Up to this point everything is the same as in the layer-2 example.




3. The FortiGate unit removes the VLAN 100 tag, and inspects the content of the data frame. The FortiGate unit uses the content to select the correct security policy and routing options.
4. The FortiGate unit's security policy allows the data frame to go to VLAN 300 in this example. The data frame will be sent to all VLAN 300 interfaces, but in the example there is only port 1




on the FortiGate unit. Before the data frame leaves, the FortiGate unit adds the VLAN ID 300 tag to the data frame.

This is the step that layer 2 cannot do. Only layer 3 can retag a data frame as a different VLAN.



5. Switch B receives the data frame, and removes the VLAN ID 300 tag, because this is the last hop, and forwards the data frame to the computer on port 5.



In this example, a data frame arrived at the FortiGate unit tagged as VLAN 100. After checking its content, the FortiGate unit retagged the data frame for VLAN 300. It is this change from

VLAN 100 to VLAN 300 that requires a layer-3 routing device, in this case the FortiGate unit. Layer-2 switches cannot perform this change.

## VLANs in NAT mode

In NAT mode the FortiGate unit functions as a layer-3 device. In this mode, the FortiGate unit controls the flow of packets between VLANs, but can also remove VLAN tags from incoming VLAN packets. The FortiGate unit can also forward untagged packets to other networks, such as the Internet.

In NAT mode, the FortiGate unit supports VLAN trunk links with IEEE 802.1Q-compliant switches, or routers. The trunk link transports VLAN-tagged packets between physical subnets or networks. When you add VLAN sub-interfaces to the FortiGate unit physical interfaces, the VLANs have IDs that match the VLAN IDs of packets on the trunk link. The FortiGate unit directs packets with VLAN IDs to sub-interfaces with matching IDs.

You can define VLAN sub-interfaces on all FortiGate physical interfaces. However, if multiple virtual domains are configured on the FortiGate unit, you will have access to only the physical interfaces on your virtual domain. The FortiGate unit can tag packets leaving on a VLAN subinterface. It can also remove VLAN tags from incoming packets and add a different VLAN tag to outgoing packets.

Normally in VLAN configurations, the FortiGate unit's internal interface is connected to a VLAN trunk, and the external interface connects to an Internet router that is not configured for VLANs. In this configuration the FortiGate unit can apply different policies for traffic on each VLAN interface connected to the internal interface, which results in less network traffic and better security.

## Adding VLAN subinterfaces

A VLAN subinterface, also called a VLAN, is a virtual interface on a physical interface. The subinterface allows routing of VLAN tagged packets using that physical interface, but it is separate from any other traffic on the physical interface.

Adding a VLAN subinterface includes configuring:

- [Physical interface](#)
- [IP address and netmask](#)
- [VLAN ID](#)
- [VDOM](#)

### Physical interface

The term VLAN subinterface correctly implies the VLAN interface is not a complete interface by itself. You add a VLAN subinterface to the physical interface that receives VLAN-tagged packets. The physical interface can belong to a different VDOM than the VLAN, but it must be connected to a network router that is configured for this VLAN. Without that router, the VLAN will not be connected to the network, and VLAN traffic will not be able to access this interface. The traffic on the VLAN is separate from any other traffic on the physical interface.

When you are working with interfaces on your FortiGate unit, use the *Column Settings* on the Interface display to make sure the information you need is displayed. When working with VLANs, it is useful to position the *VLAN ID* column close to the IP address. If you are working with VDOMs, including the *Virtual Domain* column as well will help you troubleshoot problems more quickly.

To view the Interface display, go to *System > Network > Interface*.

## IP address and netmask

FortiGate unit interfaces cannot have overlapping IP addresses. The IP addresses of all interfaces must be on different subnets. This rule applies to both physical interfaces and to virtual interfaces such as VLAN subinterfaces. Each VLAN subinterface must be configured with its own IP address and netmask pair. This rule helps prevent a broadcast storm or other similar network problems.



If you are unable to change your existing configurations to prevent IP overlap, enter the CLI command `config system global and set ip-overlap enable` to allow IP address overlap. If you enter this command, multiple VLAN interfaces can have an IP address that is part of a subnet used by another interface. This command is recommended for advanced users only.

## VLAN ID

The VLAN ID is part of the VLAN tag added to the packets by VLAN switches and routers. The VLAN ID is a number between 1 and 4094 that allow groups of IP addresses with the same VLAN ID to be associated together. VLAN ID 0 is used only for high priority frames, and 4095 is reserved.

All devices along a route must support the VLAN ID of the traffic along that route. Otherwise, the traffic will be discarded before reaching its destination. For example, if your computer is part of VLAN\_100 and a co-worker on a different floor of your building is also on the same VLAN\_100, you can communicate with each other over VLAN\_100, only if all the switches and routers support VLANs and are configured to pass along VLAN\_100 traffic properly. Otherwise, any traffic you send your co-worker will be blocked or not delivered.

## VDOM

If VDOMs are enabled, each VLAN subinterface must belong to a VDOM. This rule also applies for physical interfaces.



Interface-related CLI commands require a VDOM to be specified, regardless of whether the FortiGate unit has VDOMs enabled.

VLAN subinterfaces on separate VDOMs cannot communicate directly with each other. In this situation, the VLAN traffic must exit the FortiGate unit and re-enter the unit again, passing through firewalls in both directions. This situation is the same for physical interfaces.

A VLAN subinterface can belong to a different VDOM than the physical interface it is part of. This is because the traffic on the VLAN is handled separately from the other traffic on that interface. This is one of the main strengths of VLANs.

The following procedure will add a VLAN subinterface called `VLAN_100` to the FortiGate internal interface with a VLAN ID of 100. It will have an IP address and netmask of `172.100.1.1/255.255.255.0`, and allow HTTPS, PING, and Telnet administrative access. Note that in the CLI, you must enter “`set type vlan`” before setting the `vlanid`, and that the `allowaccess` protocols are lower case.

### To add a VLAN subinterface in NAT mode - web-based manager

1. If *Current VDOM* appears at the bottom left of the screen, select *Global* from the list of VDOMs.
2. Go to *System > Network > Interface*.

3. Select *Create New* to add a VLAN subinterface.
4. Enter the following:

<b>VLAN Name</b>	VLAN_100
<b>Type</b>	VLAN
<b>Interface</b>	internal
<b>VLAN ID</b>	100
<b>Addressing Mod</b>	Manual
<b>IP/Netmask</b>	172.100.1.1/255.255.255.0
<b>Administrative Access</b>	HTTPS, PING, TELNET

5. Select *OK*.

To view the new VLAN subinterface, select the expand arrow next to the parent physical interface (the internal interface). This will expand the display to show all VLAN subinterfaces on this physical interface. If there is no expand arrow displayed, there are no subinterfaces configured on that physical interface.

For each VLAN, the list displays the name of the VLAN, and, depending on column settings, its IP address, the Administrative access you selected for it, the VLAN ID number, and which VDOM it belongs to if VDOMs are enabled.

#### To add a VLAN subinterface in NAT mode - CLI

```
config system interface
 edit VLAN_100
 set interface internal
 set type vlan
 set vlanid 100
 set ip 172.100.1.1 255.255.255.0
 set allowaccess https ping telnet
 end
```

## Configuring security policies and routing

Once you have created a VLAN subinterface on the FortiGate unit, you need to configure security policies and routing for that VLAN. Without these, the FortiGate unit will not pass VLAN traffic to its intended destination. Security policies direct traffic through the FortiGate unit between interfaces. Routing directs traffic across the network.

### Configuring security policies

Security policies permit communication between the FortiGate unit's network interfaces based on source and destination IP addresses. Interfaces that communicate with the VLAN interface need security policies to permit traffic to pass between them and the VLAN interface.

Each VLAN needs a security policy for each of the following connections the VLAN will be using:

- from this VLAN to an external network
- from an external network to this VLAN
- from this VLAN to another VLAN in the same virtual domain on the FortiGate unit
- from another VLAN to this VLAN in the same virtual domain on the FortiGate unit.

The packets on each VLAN are subject to antivirus scans and other UTM measures as they pass through the FortiGate unit.

## Configuring routing

As a minimum, you need to configure a default static route to a gateway with access to an external network for outbound packets. In more complex cases, you will have to configure different static or dynamic routes based on packet source and destination addresses.

As with firewalls, you need to configure routes for VLAN traffic. VLANs need routing and a gateway configured to send and receive packets outside their local subnet just as physical interfaces do. The type of routing you configure, static or dynamic, will depend on the routing used by the subnet and interfaces you are connecting to. Dynamic routing can be routing information protocol (RIP), border gateway protocol (BGP), open shortest path first (OSPF), or multicast.

If you enable SSH, PING, Telnet, HTTPS and HTTP on the VLAN, you can use those protocols to troubleshoot your routing and test that it is properly configured. Enabling logging on the interfaces and using CLI diagnose commands such as `diagnose sniff packet <interface_name>` can also help locate any possible configuration or hardware issues.


## Example VLAN configuration in NAT mode

In this example two different internal VLAN networks share one interface on the FortiGate unit, and share the connection to the Internet. This example shows that two networks can have separate traffic streams while sharing a single interface. This configuration could apply to two departments in a single company, or to different companies.

There are two different internal network VLANs in this example. VLAN\_100 is on the 10.1.1.0/255.255.255.0 subnet, and VLAN\_200 is on the 10.1.2.0/255.255.255.0 subnet. These VLANs are connected to the VLAN switch, such as a Cisco 2950 Catalyst switch.

The FortiGate internal interface connects to the VLAN switch through an 802.1Q trunk. The internal interface has an IP address of 192.168.110.126 and is configured with two VLAN subinterfaces (VLAN\_100 and VLAN\_200). The external interface has an IP address of 172.16.21.2 and connects to the Internet. The external interface has no VLAN subinterfaces.

**Figure 240:**FortiGate unit with VLANs in NAT mode



When the VLAN switch receives packets from VLAN\_100 and VLAN\_200, it applies VLAN ID tags and forwards the packets of each VLAN both to local ports and to the FortiGate unit across the trunk link. The FortiGate unit has policies that allow traffic to flow between the VLANs, and from the VLANs to the external network.

This section describes how to configure a FortiGate unit and a Cisco Catalyst 2950 switch for this example network topology. The Cisco configuration commands used in this section are IOS commands.

It is assumed that both the FortiGate unit and the Cisco 2950 switch are installed and connected and that basic configuration has been completed. On the switch, you will need to be able to access the CLI to enter commands. Refer to the manual for your FortiGate model as well as the manual for the switch you select for more information.

It is also assumed that no VDOMs are enabled.

## General configuration steps

The following steps provide an overview of configuring and testing the hardware used in this example. For best results in this configuration, follow the procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

1. Configure the FortiGate unit
  - Configure the external interface
  - Add two VLAN subinterfaces to the internal network interface
  - Add firewall addresses and address ranges for the internal and external networks
  - Add security policies to allow:
    - the VLAN networks to access each other
    - the VLAN networks to access the external network.
2. Configure the VLAN switch

## Configure the FortiGate unit

Configuring the FortiGate unit includes:

- [Configure the external interface](#)
- [Add VLAN subinterfaces](#)
- [Add the firewall addresses](#)
- [Add the security policies](#)

### Configure the external interface

The FortiGate unit's external interface will provide access to the Internet for all internal networks, including the two VLANs.

#### To configure the external interface - web-based manager

1. Go to *System > Network > Interface*.
2. Select *Edit* for the external interface.
3. Enter the following information and select *OK*:

---

<b>Addressing mode</b>	Manual
<b>IP/Network Mask</b>	172.16.21.2/255.255.255.0

---

#### To configure the external interface - CLI

```
config system interface
 edit external
 set mode static
 set ip 172.16.21.2 255.255.255.0
 end
```

### Add VLAN subinterfaces

This step creates the VLANs on the FortiGate unit internal physical interface. The IP address of the internal interface does not matter to us, as long as it does not overlap with the subnets of the VLAN subinterfaces we are configuring on it.

The rest of this example shows how to configure the VLAN behavior on the FortiGate unit, configure the switches to direct VLAN traffic the same as the FortiGate unit, and test that the configuration is correct.

Adding VLAN subinterfaces can be completed through the web-based manager, or the CLI.

#### To add VLAN subinterfaces - web-based manager

1. Go to *System > Network > Interface*.
2. Select *Create New*.
3. Enter the following information and select *OK*:

---

<b>Name</b>	VLAN_100
<b>Interface</b>	internal
<b>VLAN ID</b>	100
<b>Addressing mode</b>	Manual

---

<b>IP/Network Mask</b>	10.1.1.1/255.255.255.0
<b>Administrative Access</b>	HTTPS, PING, TELNET

4. Select *Create New*.
5. Enter the following information and select *OK*:

<b>Name</b>	VLAN_200
<b>Interface</b>	internal
<b>VLAN ID</b>	200
<b>Addressing mode</b>	Manual
<b>IP/Network Mask</b>	10.1.2.1/255.255.255.0
<b>Administrative Access</b>	HTTPS, PING, TELNET

#### To add VLAN subinterfaces - CLI

```

config system interface
 edit VLAN_100
 set vdom root
 set interface internal
 set type vlan
 set vlanid 100
 set mode static
 set ip 10.1.1.1 255.255.255.0
 set allowaccess https ping telnet
 next
 edit VLAN_200
 set vdom root
 set interface internal
 set type vlan
 set vlanid 200
 set mode static
 set ip 10.1.2.1 255.255.255.0
 set allowaccess https ping telnet
end

```

#### Add the firewall addresses

You need to define the addresses of the VLAN subnets for use in security policies. The FortiGate unit provides one default address, “all”, that you can use when a security policy applies to all addresses as a source or destination of a packet. However, using “all” is less secure and should be avoided when possible.

In this example, the “\_Net” part of the address name indicates a range of addresses instead of a unique address. When choosing firewall address names, use informative and unique names.

#### To add the firewall addresses - web-based manager

1. Go to *Firewall Objects > Address > Addresses*.
2. Select *Create New*.



3. Enter the following information and select *OK*:

<b>Name</b>	VLAN_100_Net
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	10.1.1.0/255.255.255.0

4. Select *Create New*.
5. Enter the following information and select *OK*:

<b>Name</b>	VLAN_200_Net
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	10.1.2.0/255.255.255.0

### To add the firewall addresses - CLI

```
config firewall address
 edit VLAN_100_Net
 set type ipmask
 set subnet 10.1.1.0 255.255.255.0
 next
 edit VLAN_200_Net
 set type ipmask
 set subnet 10.1.2.0 255.255.255.0
end
```

### Add the security policies

Once you have assigned addresses to the VLANs, you need to configure security policies for them to allow valid packets to pass from one VLAN to another and to the Internet.



You can customize the Security Policy display by including some or all columns, and customize the column order onscreen. Due to this feature, security policy screenshots may not appear the same as on your screen.

If you do not want to allow all services on a VLAN, you can create a security policy for each service you want to allow. This example allows all services.

### To add the security policies - web-based manager

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and the *Policy Subtype* as *Address*.
3. Enter the following information and select *OK*:

<b>Incoming Interface</b>	VLAN_100
<b>Source Address</b>	VLAN_100_Net
<b>Outgoing Interface</b>	VLAN_200
<b>Destination Address</b>	VLAN_200_Net

<b>Schedule</b>	Always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Enable

4. Select *Create New*.
5. Leave the *Policy Type* as *Firewall* and the *Policy Subtype* as *Address*.
6. Enter the following information and select *OK*:

<b>Incoming Interface</b>	VLAN_200
<b>Source Address</b>	VLAN_200_Net
<b>Outgoing Interface</b>	VLAN_100
<b>Destination Address</b>	VLAN_100_Net
<b>Schedule</b>	Always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Enable

7. Select *Create New*.
8. Leave the *Policy Type* as *Firewall* and the *Policy Subtype* as *Address*.
9. Enter the following information and select *OK*:

<b>Incoming Interface</b>	VLAN_100
<b>Source Address</b>	VLAN_100_Net
<b>Outgoing Interface</b>	external
<b>Destination Address</b>	all
<b>Schedule</b>	Always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Enable

10. Select *Create New*.
11. Leave the *Policy Type* as *Firewall* and the *Policy Subtype* as *Address*.
12. Enter the following information and select *OK*:

<b>Incoming Interface</b>	VLAN_200
<b>Source Address</b>	VLAN_200_Net

<b>Outgoing Interface</b>	external
<b>Destination Address</b>	all
<b>Schedule</b>	Always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Enable

### To add the security policies - CLI

```

config firewall policy
 edit 1
 set srcintf VLAN_100
 set srcaddr VLAN_100_Net
 set dstintf VLAN_200
 set dstaddr VLAN_200_Net
 set schedule always
 set service ALL
 set action accept
 set nat enable
 set status enable
 next
 edit 2
 set srcintf VLAN_200
 set srcaddr VLAN_200_Net
 set dstintf VLAN_100
 set dstaddr VLAN_100_Net
 set schedule always
 set service ALL
 set action accept
 set nat enable
 set status enable
 next
 edit 3
 set srcintf VLAN_100
 set srcaddr VLAN_100_Net
 set dstintf external
 set dstaddr all
 set schedule always
 set service ALL
 set action accept
 set nat enable
 set status enable
 next
 edit 4
 set srcintf VLAN_200
 set srcaddr VLAN_200_Net

```

```
set dstintf external
set dstaddr all
set schedule always
set service ALL
set action accept
set nat enable
set status enable
end
```

## Configure the VLAN switch

On the Cisco Catalyst 2950 Catalyst VLAN switch, you need to define VLANs 100 and 200 in the VLAN database, and then add a configuration file to define the VLAN subinterfaces and the 802.1Q trunk interface.

One method to configure a Cisco switch is to connect over a serial connection to the console port on the switch, and enter the commands at the CLI. Another method is to designate one interface on the switch as the management interface and use a web browser to connect to the switch's graphical interface. For details on connecting and configuring your Cisco switch, refer to the installation and configuration manuals for the switch.

The switch used in this example is a Cisco Catalyst 2950 switch. The commands used are IOS commands. Refer to the switch manual for help with these commands.

### To configure the VLAN subinterfaces and the trunk interfaces

Add this file to the Cisco switch:

```
!
interface FastEthernet0/3
switchport access vlan 100
!
interface FastEthernet0/9
switchport access vlan 200
!
interface FastEthernet0/24
switchport trunk encapsulation dot1q
switchport mode trunk
!
```

The switch has the configuration:

<b>Port 0/3</b>	VLAN ID 100
<b>Port 0/9</b>	VLAN ID 200
<b>Port 0/24</b>	802.1Q trunk



To complete the setup, configure devices on VLAN\_100 and VLAN\_200 with default gateways. The default gateway for VLAN\_100 is the FortiGate VLAN\_100 subinterface. The default gateway for VLAN\_200 is the FortiGate VLAN\_200 subinterface.

## Test the configuration

Use diagnostic commands, such as `tracert`, to test traffic routed through the FortiGate unit and the Cisco switch.

### Testing traffic from VLAN\_100 to VLAN\_200

In this example, a route is traced between the two internal networks. The route target is a host on VLAN\_200.

Access a command prompt on a Windows computer on the VLAN\_100 network, and enter the following command:

```
C:\>tracert 10.1.2.2
Tracing route to 10.1.2.2 over a maximum of 30 hops:
 1 <10 ms <10 ms <10 ms 10.1.1.1
 2 <10 ms <10 ms <10 ms 10.1.2.2
Trace complete.
```

### Testing traffic from VLAN\_200 to the external network

In this example, a route is traced from an internal network to the external network. The route target is the external network interface of the FortiGate-800 unit.

From VLAN\_200, access a command prompt and enter this command:

```
C:\>tracert 172.16.21.2
Tracing route to 172.16.21.2 over a maximum of 30 hops:
 1 <10 ms <10 ms <10 ms 10.1.2.1
 2 <10 ms <10 ms <10 ms 172.16.21.2
Trace complete.
```

## VLANs in transparent mode

In transparent mode, the FortiGate unit behaves like a layer-2 bridge but can still provide services such as antivirus scanning, web filtering, spam filtering and intrusion protection to traffic. There are some limitations in transparent mode in that you cannot use SSL VPN, PPTP/L2TP VPN, DHCP server, or easily perform NAT on traffic. The limits in transparent mode apply to IEEE 802.1Q VLAN trunks passing through the unit.

### VLANs and transparent mode

You can insert the FortiGate unit operating in transparent mode into the VLAN trunk without making changes to your network. In a typical configuration, the FortiGate unit internal interface accepts VLAN packets on a VLAN trunk from a VLAN switch or router connected to internal network VLANs. The FortiGate external interface forwards VLAN-tagged packets through another VLAN trunk to an external VLAN switch or router and on to external networks such as the Internet. You can configure the unit to apply different policies for traffic on each VLAN in the trunk.

To pass VLAN traffic through the FortiGate unit, you add two VLAN subinterfaces with the same VLAN ID, one to the internal interface and the other to the external interface. You then create a security policy to permit packets to flow from the internal VLAN interface to the external VLAN interface. If required, you create another security policy to permit packets to flow from the external VLAN interface to the internal VLAN interface. Typically in transparent mode, you do not permit packets to move between different VLANs. Network protection features, such as

spam filtering, web filtering and anti-virus scanning, are applied through the UTM profiles specified in each security policy, enabling very detailed control over traffic.

When the FortiGate unit receives a VLAN-tagged packet at a physical interface, it directs the packet to the VLAN subinterface with the matching VLAN ID. The VLAN tag is removed from the packet, and the FortiGate unit then applies security policies using the same method it uses for non-VLAN packets. If the packet exits the FortiGate unit through a VLAN subinterface, the VLAN ID for that subinterface is added to the packet and the packet is sent to the corresponding physical interface. For a configuration example, see [“Example of VLANs in transparent mode” on page 1544](#).

There are two essential steps to configure your FortiGate unit to work with VLANs in transparent mode:

- [Add VLAN subinterfaces](#)
- [Create security policies](#)

You can also configure the protection profiles that manage antivirus scanning, web filtering and spam filtering. For more information on UTM profiles, see [“Unified Threat Management for FortiOS 5.0” on page 2018](#).

### Add VLAN subinterfaces

The VLAN ID of each VLAN subinterface must match the VLAN ID added by the IEEE 802.1Q-compliant router or switch. The VLAN ID can be any number between 1 and 4094, with 0 being used only for high priority frames and 4095 being reserved. You add VLAN subinterfaces to the physical interface that receives VLAN-tagged packets.

For this example, we are creating a VLAN called `internal_v225` on the internal interface, with a VLAN ID of 225. Administrative access is enabled for HTTPS and SSH. VDOMs are not enabled.

#### To add VLAN subinterfaces in transparent mode - web-based manager

1. Go to *System > Network > Interface*.
2. Select *Create New*.
3. Enter the following information and select *OK*.

<b>Name</b>	internal_v225
<b>Type</b>	VLAN
<b>Interface</b>	internal
<b>VLAN ID</b>	225
<b>Administrative Access</b>	Enable HTTPS, and SSH. These are very secure access methods.
<b>Comments</b>	VLAN 225 on internal interface

The FortiGate unit adds the new subinterface to the interface that you selected.

Repeat steps 2 and 3 to add additional VLANs. You will need to change the *VLAN ID*, *Name*, and possibly *Interface* when adding additional VLANs.

## To add VLAN subinterfaces in transparent mode - CLI

```
config system interface
 edit internal_v225
 set interface internal
 set vlanid 225
 set allowaccess HTTPS SSH
 set description "VLAN 225 on internal interface"
 set vdom root
 end
```

## Create security policies

In transparent mode, the FortiGate unit performs antivirus and antispam scanning on each VLAN's packets as they pass through the unit. You need security policies to permit packets to pass from the VLAN interface where they enter the unit to the VLAN interface where they exit the unit. If there are no security policies configured, no packets will be allowed to pass from one interface to another.

## To add security policies for VLAN subinterfaces - web based manager

1. Go to *Firewall Objects > Address > Addresses*.
2. Select *Create New* to add firewall addresses that match the source and destination IP addresses of VLAN packets.
3. Go to *Policy > Policy > Policy* and select *Create New*.
4. Leave the *Policy Type* as *Firewall* and the *Policy Subtype* as *Address*.
5. From the *Incoming Interface/Zone* list, select the VLAN interface where packets enter the unit.
6. From the *Outgoing Interface/Zone* list, select the VLAN interface where packets exit the unit.
7. Select the *Source* and *Destination Address* names that you added in step 2.
8. Select *OK*.

## To add security policies for VLAN subinterfaces - CLI

```
config firewall address
 edit incoming_VLAN_address
 set associated-interface <incoming_VLAN_interface>
 set type ipmask
 set subnet <IPv4_address_mask>
 next
 edit outgoing_VLAN_address
 set associated-interface <outgoing_VLAN_interface>
 set type ipmask
 set subnet <IPv4_address_mask>
 next
end
config firewall policy
 edit <unused_policy_number>
 set srcintf <incoming_VLAN_interface>
 set srcaddr incoming_VLAN_address
 set destintf <outgoing_VLAN_interface>
```

```

set destaddr outgoing_VLAN_address
set service <protocol_to_allow_on_VLAN>
set action ACCEPT
next
end

```

## Example of VLANs in transparent mode


In this example, the FortiGate unit is operating in transparent mode and is configured with two VLANs: one with an ID of 100 and the other with ID 200. The internal and external physical interfaces each have two VLAN subinterfaces, one for VLAN\_100 and one for VLAN\_200.

The IP range for the internal VLAN\_100 network is 10.100.0.0/255.255.0.0, and for the internal VLAN\_200 network is 10.200.0.0/255.255.0.0.

The internal networks are connected to a Cisco 2950 VLAN switch, which combines traffic from the two VLANs onto one the FortiGate unit internal interface. The VLAN traffic leaves the FortiGate unit on the external network interface, goes on to the VLAN switch, and on to the Internet. When the FortiGate unit receives a tagged packet, it directs it from the incoming VLAN subinterface to the outgoing VLAN subinterface for that VLAN.

This section describes how to configure a FortiGate-800 unit, Cisco switch, and Cisco router in the network topology shown in [Figure 180](#).

**Figure 241:** VLAN transparent network topology



## General configuration steps

The following steps summarize the configuration for this example. For best results, follow the procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.



1. Configure the FortiGate unit which includes
  - Adding VLAN subinterfaces
  - Adding the security policies
2. Configure the Cisco switch and router

## Configure the FortiGate unit

The FortiGate unit must be configured with the VLAN subinterfaces and the proper security policies to enable traffic to flow through the FortiGate unit.

### Add VLAN subinterfaces

For each VLAN, you need to create a VLAN subinterface on the internal interface and another one on the external interface, both with the same VLAN ID.

#### To add VLAN subinterfaces - web-based manager

1. Go to *System > Network > Interface*.
2. Select *Create New*.
3. Enter the following information and select *OK*:

<b>Name</b>	VLAN_100_int
<b>Interface</b>	internal
<b>VLAN ID</b>	100

4. Select *Create New*.
5. Enter the following information and select *OK*:

<b>Name</b>	VLAN_100_ext
<b>Interface</b>	external
<b>VLAN ID</b>	100

6. Select *Create New*.
7. Enter the following information and select *OK*:

<b>Name</b>	VLAN_200_int
<b>Interface</b>	internal
<b>VLAN ID</b>	200

8. Select *Create New*.
9. Enter the following information and select *OK*:

<b>Name</b>	VLAN_200_ext
<b>Interface</b>	external
<b>VLAN ID</b>	200

### To add VLAN subinterfaces - CLI

```
config system interface
 edit VLAN_100_int
 set status down
 set type vlan
 set interface internal
 set vlanid 100
 next
 edit VLAN_100_ext
 set status down
 set type vlan
 set interface external
 set vlanid 100
 next
 edit VLAN_200_int
 set status down
 set type vlan
 set interface internal
 set vlanid 200
 next
 edit VLAN_200_ext
 set status down
 set type vlan
 set interface external
 set vlanid 200
end
```

### Add the security policies

Security policies allow packets to travel between the VLAN\_100\_int interface and the VLAN\_100\_ext interface. Two policies are required; one for each direction of traffic. The same is required between the VLAN\_200\_int interface and the VLAN\_200\_ext interface, for a total of four required security policies.

#### To add the security policies - web-based manager

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and the *Policy Subtype* as *Address*.
3. Enter the following information and select *OK*:

<b>Incoming Interface</b>	VLAN_100_int
<b>Source Address</b>	all
<b>Outgoing Interface</b>	VLAN_100_ext
<b>Destination Address</b>	all
<b>Schedule</b>	Always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT

4. Select *Create New*.
5. Leave the *Policy Type* as *Firewall* and the *Policy Subtype* as *Address*.
6. Enter the following information and select *OK*:

<b>Incoming Interface</b>	VLAN_100_ext
<b>Source Address</b>	all
<b>Outgoing Interface</b>	VLAN_100_int
<b>Destination Address</b>	all
<b>Schedule</b>	Always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT

7. Go to *Policy > Policy > Policy* and select *Create New*.
8. Leave the *Policy Type* as *Firewall* and the *Policy Subtype* as *Address*.
9. Enter the following information and select *OK*:

<b>Incoming Interface</b>	VLAN_200_int
<b>Source Address</b>	all
<b>Outgoing Interface</b>	VLAN_200_ext
<b>Destination Address</b>	all
<b>Schedule</b>	Always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Enable

10. Select *Create New*.
11. Leave the *Policy Type* as *Firewall* and the *Policy Subtype* as *Address*.
12. Enter the following information and select *OK*:

<b>Incoming Interface</b>	VLAN_200_ext
<b>Source Address</b>	all
<b>Outgoing Interface</b>	VLAN_200_int
<b>Destination Address</b>	all
<b>Schedule</b>	Always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT

## To add the security policies - CLI

```
config firewall policy
 edit 1
 set srcintf VLAN_100_int
 set srcaddr all
 set dstintf VLAN_100_ext
 set dstaddr all
 set action accept
 set schedule always
 set service ALL
 next
 edit 2
 set srcintf VLAN_100_ext
 set srcaddr all
 set dstintf VLAN_100_int
 set dstaddr all
 set action accept
 set schedule always
 set service ALL
 next
 edit 3
 set srcintf VLAN_200_int
 set srcaddr all
 set dstintf VLAN_200_ext
 set dstaddr all
 set action accept
 set schedule always
 set service ALL
 next
 edit 4
 set srcintf VLAN_200_ext
 set srcaddr all
 set dstintf VLAN_200_int
 set dstaddr all
 set action accept
 set schedule always
 set service ALL
end
```

## Configure the Cisco switch and router

This example includes configuration for the Cisco Catalyst 2900 ethernet switch, and for the Cisco Multiservice 2620 ethernet router. If you have access to a different VLAN enabled switch or VLAN router you can use them instead, however their configuration is not included in this document.

### Configure the Cisco switch

On the VLAN switch, you need to define VLAN\_100 and VLAN\_200 in the VLAN database and then add a configuration file to define the VLAN subinterfaces and the 802.1Q trunk interface.

Add this file to the Cisco switch:

```
interface FastEthernet0/3
 switchport access vlan 100
!
interface FastEthernet0/9
 switchport access vlan 200
!
interface FastEthernet0/24
 switchport trunk encapsulation dot1q
 switchport mode trunk
!
```

The switch has the following configuration:

<b>Port 0/3</b>	VLAN ID 100
<b>Port 0/9</b>	VLAN ID 200
<b>Port 0/24</b>	802.1Q trunk

### Configure the Cisco router

You need to add a configuration file to the Cisco Multiservice 2620 ethernet router. The file defines the VLAN subinterfaces and the 802.1Q trunk interface on the router. The 802.1Q trunk is the physical interface on the router.

The IP address for each VLAN on the router is the gateway for that VLAN. For example, all devices on the internal VLAN\_100 network will have 10.100.0.1 as their gateway.

Add this file to the Cisco router:

```
!
interface FastEthernet0/0
!
interface FastEthernet0/0.1
 encapsulation dot1Q 100
 ip address 10.100.0.1 255.255.255.0
!
interface FastEthernet0/0.2
 encapsulation dot1Q 200
 ip address 10.200.0.1 255.255.255.0
!
```

The router has the following configuration:

<b>Port 0/0.1</b>	VLAN ID 100
<b>Port 0/0.2</b>	VLAN ID 200
<b>Port 0/0</b>	802.1Q trunk

## Test the configuration

Use diagnostic network commands such as traceroute (`tracert`) and ping to test traffic routed through the network.

### Testing traffic from VLAN\_100 to VLAN\_200

In this example, a route is traced between the two internal networks. The route target is a host on VLAN\_200. The Windows traceroute command `tracert` is used.

From VLAN\_100, access a Windows command prompt and enter this command:

```
C:\>tracert 10.1.2.2
Tracing route to 10.1.2.2 over a maximum of 30 hops:
 1 <10 ms <10 ms <10 ms 10.1.1.1
 2 <10 ms <10 ms <10 ms 10.1.2.2
Trace complete.
```

## Troubleshooting VLAN issues

Several problems can occur with your VLANs. Since VLANs are interfaces with IP addresses, they behave as interfaces and can have similar problems that you can diagnose with tools such as ping, traceroute, packet sniffing, and diag debug.

### Asymmetric routing

You might discover unexpectedly that hosts on some networks are unable to reach certain other networks. This occurs when request and response packets follow different paths. If the FortiGate unit recognizes the response packets, but not the requests, it blocks the packets as invalid. Also, if the FortiGate unit recognizes the same packets repeated on multiple interfaces, it blocks the session as a potential attack.

This is asymmetric routing. By default, the FortiGate unit blocks packets or drops the session when this happens. You can configure the FortiGate unit to permit asymmetric routing by using the following CLI commands:

```
config vdom
 edit <vdom_name>
 config system settings
 set asymroute enable
 end
 end
```

If VDOMs are enabled, this command is per VDOM. You must set it for each VDOM that has the problem. If this solves your blocked traffic issue, you know that asymmetric routing is the cause. But allowing asymmetric routing is not the best solution, because it reduces the security of your network.

For a long-term solution, it is better to change your routing configuration or change how your FortiGate unit connects to your network. The [Asymmetric Routing and Other FortiGate Layer-2](#)

[Installation Issues](#) technical note provides detailed examples of asymmetric routing situations and possible solutions.



If you enable asymmetric routing, antivirus and intrusion prevention systems will not be effective. Your FortiGate unit will be unaware of connections and treat each packet individually. It will become a stateless firewall.

## Layer-2 and Arp traffic

By default, FortiGate units do not pass layer-2 traffic. If there are layer-2 protocols such as IPX, PPTP or L2TP in use on your network, you need to configure your FortiGate unit interfaces to pass these protocols without blocking. Another type of layer-2 traffic is ARP traffic. For more information on ARP traffic, see [“ARP traffic” on page 1551](#).

You can allow these layer-2 protocols using the CLI command:

```
config vdom
 edit <vdom_name>
 config system interface
 edit <name_str>
 set l2forward enable
 end
 end
 end
```

where `<name_str>` is the name of an interface.

If VDOMs are enabled, this command is per VDOM. You must set it for each VDOM that has the problem. If you enable layer-2 traffic, you may experience a problem if packets are allowed to repeatedly loop through the network. This repeated looping, very similar to a broadcast storm, occurs when you have more than one layer-2 path to a destination. Traffic may overflow and bring your network to a halt. You can break the loop by enabling Spanning Tree Protocol (STP) on your network’s switches and routers. For more information, see [“STP forwarding” on page 1262](#).

### ARP traffic

Address Resolution Protocol (ARP) packets are vital to communication on a network, and ARP support is enabled on FortiGate unit interfaces by default. Normally you want ARP packets to pass through the FortiGate unit, especially if it is sitting between a client and a server or between a client and a router.

ARP traffic can cause problems, especially in transparent mode where ARP packets arriving on one interface are sent to all other interfaces including VLAN subinterfaces. Some layer-2 switches become unstable when they detect the same MAC address originating on more than one switch interface or from more than one VLAN. This instability can occur if the layer-2 switch does not maintain separate MAC address tables for each VLAN. Unstable switches may reset and cause network traffic to slow down considerably.

The default ARP timeout value is 5 minutes (300 seconds). This timeout is not configurable.

Usually ARP entries are removed after 5 minutes. However, some conditions can cause ARP entries to remain for a longer time. Enter the `get system arp` CLI command to view the entries in the ARP list.

## Multiple VDOMs solution

By default, physical interfaces are in the root domain. If you do not configure any of your VLANs in the root VDOM, it will not matter how many interfaces are in the root VDOM.

The multiple VDOMs solution is to configure multiple VDOMs on the FortiGate unit, one for each VLAN. In this solution, you configure one inbound and one outbound VLAN interface in each VDOM. ARP packets are not forwarded between VDOMs. This configuration limits the VLANs in a VDOM and correspondingly reduces the administration needed per VDOM.

As a result of this configuration, the switches do not receive multiple ARP packets with duplicate MACs. Instead, the switches receive ARP packets with different VLAN IDs and different MACs. Your switches are stable.

However, you should **not** use the multiple VDOMs solution under any of the following conditions:

- you have more VLANs than licensed VDOMs
- you do not have enough physical interfaces

Instead, use one of two possible solutions, depending on which operation mode you are using:

- In NAT mode, you can use the `vlan forward` CLI command.
- In transparent mode, you can use the `forward-domain` CLI command. But you still need to be careful in some rare configurations.

## Vlanforward solution

If you are using NAT mode, the solution is to use the `vlanforward` CLI command for the interface in question. By default, this command is enabled and will forward VLAN traffic to all VLANs on this interface. When disabled, each VLAN on this physical interface can send traffic only to the same VLAN. There is no cross-talk between VLANs, and ARP packets are forced to take one path along the network which prevents the multiple paths problem.

In the following example, `vlanforward` is disabled on `port1`. All VLANs configured on `port1` will be separate and will not forward any traffic to each other.

```
config system interface
 edit port1
 set vlanforward disable
 end
```

## Forward-domain solution

If you are using transparent mode, the solution is to use the `forward-domain` CLI command. This command tags VLAN traffic as belonging to a particular collision group, and only VLANs tagged as part of that collision group receive that traffic. It is like an additional set of VLANs. By default, all interfaces and VLANs are part of forward-domain collision group 0. The many benefits of this solution include reduced administration, the need for fewer physical interfaces, and the availability of more flexible network solutions.

In the following example, forward-domain collision group 340 includes VLAN 340 traffic on `port1` and untagged traffic on port 2. Forward-domain collision group 341 includes VLAN 341 traffic on port 1 and untagged traffic on port 3. All other interfaces are part of forward-domain collision group 0 by default. This configuration separates VLANs 340 and 341 from each other on port 1, and prevents the ARP packet problems from before.

Use these CLI commands:

```
config system interface
 edit port1
```



```

next
edit port2
 set forward_domain 340
next
edit port3
 set forward_domain 341
next
edit port1-340
 set forward_domain 340
 set interface port1
 set vlanid 340
next
edit port1-341
 set forward_domain 341
 set interface port1
 set vlanid 341
end

```

You may experience connection issues with layer-2 traffic, such as ping, if your network configuration has:

- packets going through the FortiGate unit in transparent mode more than once
- more than one forwarding domain (such as incoming on one forwarding domain and outgoing on another)
- IPS and AV enabled.

Now IPS and AV is applied the first time packets go through the FortiGate unit, but not on subsequent passes. Only applying IPS and AV to this first pass fixes the network layer-2 related connection issues.

## NetBIOS

Computers running Microsoft Windows operating systems that are connected through a network rely on a WINS server to resolve host names to IP addresses. The hosts communicate with the WINS server by using the NetBIOS protocol.

To support this type of network, you need to enable the forwarding of NetBIOS requests to a WINS server. The following example will forward NetBIOS requests on the internal interface for the WINS server located at an IP address of 192.168.111.222.

```

config system interface
 edit internal
 set netbios_forward enable
 set wins-ip 192.168.111.222
 end

```

These commands apply only in NAT mode. If VDOMs are enabled, these commands are per VDOM. You must set them for each VDOM that has the problem.

## STP forwarding

The FortiGate unit does not participate in the Spanning Tree Protocol (STP). STP is an IEEE 802.1 protocol that ensures there are no layer-2 loops on the network. Loops are created when there is more than one route for traffic to take and that traffic is broadcast back to the original switch. This loop floods the network with traffic, reducing available bandwidth to nothing.

If you use your FortiGate unit in a network topology that relies on STP for network loop protection, you need to make changes to your FortiGate configuration. Otherwise, STP recognizes your FortiGate unit as a blocked link and forwards the data to another path. By default, your FortiGate unit blocks STP as well as other non-IP protocol traffic.

Using the CLI, you can enable forwarding of STP and other layer-2 protocols through the interface. In this example, layer-2 forwarding is enabled on the external interface:

```
config system interface
 edit external
 set l2forward enable
 set stpforward enable
 end
```

By substituting different commands for `stpforward enable`, you can also allow layer-2 protocols such as IPX, PPTP or L2TP to be used on the network. For more information, see [“Layer-2 and Arp traffic” on page 1551](#).

## Too many VLAN interfaces

Any virtual domain can have a maximum of 255 interfaces in transparent mode. This includes VLANs, other virtual interfaces, and physical interfaces. NAT mode supports from 255 to 8192 depending on the FortiGate model. This total number of interfaces includes VLANs, other virtual interfaces, and physical interfaces.

Your FortiGate unit may allow you to configure more interfaces than this. However, if you configure more than 255 interfaces, your system will become unstable and, over time, will not work properly. As all interfaces are used, they will overflow the routing table that stores the interface information, and connections will fail. When you try to add more interfaces, an error message will state that the maximum limit has already been reached.

If you see this error message, chances are you already have too many VLANs on your system and your routing has become unstable. To verify, delete a VLAN and try to add it back. If you have too many, you will not be able to add it back on to the system. In this case, you will need to remove enough interfaces (including VLANs) so that the total number of interfaces drops to 255 or less. After doing this, you should also reboot your FortiGate unit to clean up its memory and buffers, or you will continue to experience unstable behavior.

To configure more than 255 interfaces on your FortiGate unit in transparent mode, you have to configure multiple VDOMs, each with many VLANs. However, if you want to create more than the default 10 VDOMs (or a maximum of 2550 interfaces), you must buy a license for additional VDOMs.

With these extra licenses, you can configure up to 500 VDOMs, with each VDOM containing up to 255 VLANs in transparent mode. This is a theoretical maximum of over 127 500 interfaces. However, system resources will quickly get used up before reaching that theoretical maximum. To achieve the maximum number of VDOMs, you need to have top-end hardware with the most resources possible.

In NAT mode, if you have a top-end model, the maximum interfaces per VDOM can be as high as 8192, enough for all the VLANs in your configuration.



Your FortiGate unit has limited resources, such as CPU load and memory, that are divided between all configured VDOMs. When running 250 or more VDOMs, you may need to monitor the system resources to ensure there is enough to support the configured traffic processing.

---

# PPTP and L2TP

A virtual private network (VPN) is a way to use a public network, such as the Internet, as a vehicle to provide remote offices or individual users with secure access to private networks. FortiOS supports the Point-to-Point Tunneling Protocol (PPTP), which enables interoperability between FortiGate units and Windows or Linux PPTP clients. Because FortiGate units support industry standard PPTP VPN technologies, you can configure a PPTP VPN between a FortiGate unit and most third-party PPTP VPN peers.

This section describes how to configure PPTP and L2TP VPNs as well as PPTP passthrough.

This section includes the topics:

- [How PPTP VPNs work](#)
- [FortiGate unit as a PPTP server](#)
- [Configuring the FortiGate unit for PPTP VPN](#)
- [Configuring the FortiGate unit for PPTP pass through](#)
- [Testing PPTP VPN connections](#)
- [Logging VPN events](#)
- [Configuring L2TP VPNs](#)
- [L2TP configuration overview](#)

## How PPTP VPNs work

The Point-to-Point Tunneling Protocol enables you to create a VPN between a remote client and your internal network. Because it is a Microsoft Windows standard, PPTP does not require third-party software on the client computer. As long as the ISP supports PPTP on its servers, you can create a secure connection by making relatively simple configuration changes to the client computer and the FortiGate unit.

PPTP uses Point-to-Point protocol (PPP) authentication protocols so that standard PPP software can operate on tunneled PPP links. PPTP packages data in PPP packets and then encapsulates the PPP packets within IP packets for transmission through a VPN tunnel.

When the FortiGate unit acts as a PPTP server, a PPTP session and tunnel is created as soon as the PPTP client connects to the FortiGate unit. More than one PPTP session can be supported on the same tunnel. FortiGate units support PAP, CHAP, and plain text authentication. PPTP clients are authenticated as members of a user group.


Traffic from one PPTP peer is encrypted using PPP before it is encapsulated using Generic Routing Encapsulation (GRE) and routed to the other PPTP peer through an ISP network. PPP packets from the remote client are addressed to a computer on the private network behind the FortiGate unit. PPTP packets from the remote client are addressed to the public interface of the FortiGate unit. See [Figure 242 on page 1556](#)



PPTP control channel messages are not authenticated, and their integrity is not protected. Furthermore, encapsulated PPP packets are not cryptographically protected and may be read or modified unless appropriate encryption software such as Secure Shell (SSH) or Secure File Transfer Protocol (SFTP) is used to transfer data after the tunnel has been established.

As an alternative, you can use encryption software such as Microsoft Point-to-Point Encryption (MPPE) to secure the channel. MPPE is built into Microsoft Windows clients and can be installed on Linux clients. FortiGate units support MPPE.

**Figure 242:**Packet encapsulation



In [Figure 242](#), traffic from the remote client is addressed to a computer on the network behind the FortiGate unit. When the PPTP tunnel is established, packets from the remote client are encapsulated and addressed to the FortiGate unit. The FortiGate unit forwards disassembled packets to the computer on the internal network.

When the remote PPTP client connects, the FortiGate unit assigns an IP address from a reserved range of IP addresses to the client PPTP interface. The PPTP client uses the assigned IP address as its source address for the duration of the connection.

When the FortiGate unit receives a PPTP packet, the unit disassembles the PPTP packet and forwards the packet to the correct computer on the internal network. The security policy and protection profiles on the FortiGate unit ensure that inbound traffic is screened and processed securely.




PPTP clients must be authenticated before a tunnel is established. The authentication process relies on FortiGate user group definitions, which can optionally use established authentication mechanisms such as RADIUS or LDAP to authenticate PPTP clients. All PPTP clients are challenged when a connection attempt is made.

## FortiGate unit as a PPTP server

In the most common Internet scenario, the PPTP client connects to an ISP that offers PPP connections with dynamically-assigned IP addresses. The ISP forwards PPTP packets to the Internet, where they are routed to the FortiGate unit.

**Figure 243:**FortiGate unit as a PPTP server



If the FortiGate unit will act as a PPTP server, there are a number of steps to complete:

- Configure user authentication for PPTP clients.
- Enable PPTP.
- Specify the range of addresses that are assigned to PPTP clients when connecting
- Configure the security policy.

### Configuring user authentication for PPTP clients

To enable authentication for PPTP clients, you must create user accounts and a user group to identify the PPTP clients that need access to the network behind the FortiGate unit. Within the user group, you must add a user for each PPTP client.

You can choose to use a plain text password for authentication or forward authentication requests to an external RADIUS, LDAP, or TACACS+ server. If password protection will be provided through a RADIUS, LDAP, or TACACS+ server, you must configure the FortiGate unit to forward authentication requests to the authentication server.

This example creates a basic user/password combination.

#### Configuring a user account

##### To add a local user - web-based manager

1. Go to *User & Device > User > User Definition* and select *Create New*.
2. Enter a *User Name*.
3. Enter a *Password* for the user. The password should be at least six characters.
4. Select *OK*.

### To add a local user - CLI

```
config user local
 edit <username>
 set type password
 set passwd <password>
 end
```

## Configuring a user group

To ease configuration, create user groups that contain users in similar categories or departments.

### To create a user group - web-based manager

1. Go to *User & Device > User > User Group* and select *Create New*.
2. Enter a *Name* for the group.
3. Select the *Type of Firewall*.
4. From the *Available Users* list, select the required users and select the right-facing arrow to add them to the *Members* list.
5. Select *OK*.

### To create a user group - CLI

```
config user group
 edit <group_name>
 set group-type firewall
 set members <user_names>
 end
```

## Enabling PPTP and specifying the PPTP IP address range

The PPTP address range specifies the range of addresses reserved for remote PPTP clients. When a PPTP client connects to the FortiGate unit, the client is assigned an IP address from this range. Afterward, the FortiGate unit uses the assigned address to communicate with the PPTP client.

The address range that you reserve can be associated with private or routable IP addresses. If you specify a private address range that matches a network behind the FortiGate unit, the assigned address will make the PPTP client appear to be part of the internal network.

PPTP requires two IP addresses, one for each end of the tunnel. The PPTP address range is the range of addresses reserved for remote PPTP clients. When the remote PPTP client establishes a connection, the FortiGate unit assigns an IP address from the reserved range of IP addresses to the client PPTP interface or retrieves the assigned IP address from the PPTP user group. If you use the PPTP user group, you must also define the FortiGate end of the tunnel by entering the IP address of the unit in *Local IP* (web-based manager) or `local-ip` (CLI). The PPTP client uses the assigned IP address as its source address for the duration of the connection.

PPTP configuration is only available through the CLI. In the example below, PPTP is enabled with the use of an IP range of 192.168.1.1 to 192.168.1.10 for addressing.



The start and end IPs in the PPTP address range must be in the same 24-bit subnet, for example, 192.168.1.1 - 192.168.1.254.

```

config vpn pptp
 set status enable
 set ip-mode range
 set eip 192.168.1.10
 set sip 192.168.1.1
end

```

In this example, PPTP is enabled with the use of a user group for addressing, where the IP address of the PPTP server is 192.168.1.2 and the user group is hr\_admin.

```

config vpn pptp
 set status enable
 set ip-mode range
 set local-ip 192.168.2.1
 set usrgrp hr_admin
end

```

## Adding the security policy

The security policy specifies the source and destination addresses that can generate traffic inside the PPTP tunnel and defines the scope of services permitted through the tunnel. If a selection of services are required, define a service group.

### To configure the firewall for the PPTP tunnel - web-based manager

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and the *Policy Subtype* as *Address*.
3. Complete the following and select *OK*:

<b>Incoming Interface</b>	The FortiGate interface connected to the Internet.
<b>Source Address</b>	Select the name that corresponds to the range of addresses that you reserved for PPTP clients.
<b>Outgoing Interface</b>	The FortiGate interface connected to the internal network.
<b>Destination Address</b>	Select the name that corresponds to the IP addresses behind the FortiGate unit.
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT

Do not select identity-based policy, as this will cause the PPTP access to fail. Authentication is configured in the PPTP configuration setup

## To configure the firewall for the PPTP tunnel - CLI

```
config firewall policy
 edit 1
 set srcintf <interface to internet>
 set dstintf <interface to internal network>
 set srcaddr <reserved_range>
 set dstaddr <internal_addresses>
 set action accept
 set schedule always
 set service ALL
 end
```

## Configuring the FortiGate unit for PPTP VPN

To arrange for PPTP packets to pass through the FortiGate unit to an external PPTP server, perform the following tasks in the order given:

- Configure user authentication for PPTP clients.
- Enable PPTP on the FortiGate unit and specify the range of addresses that can be assigned to PPTP clients when they connect.
- Configure PPTP pass through on the FortiGate unit.

## Configuring the FortiGate unit for PPTP pass through

To forward PPTP packets to a PPTP server on the network behind the FortiGate unit, you need to perform the following configuration tasks on the FortiGate unit:

- Define a virtual IP address that points to the PPTP server.
- Create a security policy that allows incoming PPTP packets to pass through to the PPTP server.



The address range is the external (public) ip address range which requires access to the internal PPTP server through the FortiGate virtual port-forwarding firewall.

IP addresses used in this document are fictional and follow the technical documentation guidelines specific to Fortinet. Real external IP addresses are not used.

---

### Configuring a virtual IP address

The virtual IP address will be the address of the PPTP server host.

#### To define a virtual IP for PPTP pass through - web-based manager

1. Go to *Firewall Objects > Virtual IP > Virtual IP*.
2. Select *Create New*.
3. Enter the name of the VIP, for example, *PPTP\_Server*.
4. Select the *External Interface* where the packets will be received for the PPTP server.
5. Enter the *External IP Address* for the VIP.
6. Select *Port Forwarding*.
7. Set the *Protocol to TCP*.



8. Enter the *External Service Port* of 1723, the default for PPTP.
9. Enter the *Map to Port* to 1723.
10. Select *OK*.

#### To define a virtual IP for PPTP pass through - web-based manager

```
config firewall vip
 edit PPTP_Server
 set extintf <interface>
 set extip <ip_address>
 set portforward enable
 set protocol tcp
 set extport 1723
 set mappedport 1723
end
```

### Configuring a port-forwarding security policy

To create a port-forwarding security policy for PPTP pass through you must first create an address range reserved for the PPTP clients.

#### To create an address range - web-based manager

1. Go to *Firewall Objects > Address > Addresses* and select *Create New*.
2. Enter a *Name* for the range, for example, *External\_PPTP*.
3. Select a *Type* of *Subnet/IP Range*.
4. Enter the IP address range.
5. Select the *Interface* to the Internet.
6. Select *OK*.

#### To create an address range - CLI

```
config firewall address
 edit External_PPTP
 set iprange <ip_range>
 set start-ip <ip_address>
 set end-ip <ip_address>
 set associated-interface <internet_interface>
end
```

With the address set, you can add the security policy.

#### To add the security policy - web-based manager

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and the *Policy Subtype* as *Address*.
3. Complete the following and select *OK*:

<b>Incoming Interface</b>	The FortiGate interface connected to the Internet.
<b>Source Address</b>	Select the address range created in the previous step.
<b>Outgoing Interface</b>	The FortiGate interface connected to the PPTP server.
<b>Destination Address</b>	Select the VIP address created in the previous steps.

<b>Schedule</b>	always
<b>Service</b>	PPTP
<b>Action</b>	ACCEPT

#### To add the security policy - CLI

```
config firewall policy
 edit <policy_number>
 set srcintf <interface to internet>
 set dstintf <interface to PPTP server>
 set srcaddr <address_range>
 set dstaddr <PPTP_server_address>
 set action accept
 set schedule always
 set service PPTP
 end
```

## Testing PPTP VPN connections

To confirm that a PPTP VPN between a local network and a dialup client has been configured correctly, at the dialup client, issue a ping command to test the connection to the local network. The PPTP VPN tunnel initializes when the dialup client attempts to connect.

## Logging VPN events

PPTP VPN, activity is logged when enabling VPN logging. The FortiGate unit connection events and tunnel status (up/down) are logged.

#### To log VPN events

1. Go to *Log & Report > Log Config > Log Settings*.
2. Enable the storage of log messages to one or more locations.
3. Select *VPN activity event*.
4. Select *Apply*.

#### To view event logs

1. Go to *Log & Report > Event Log > VPN*.
2. If the option is available from the Log Type list, select the log file from disk or memory.

## Configuring L2TP VPNs

This section describes how to configure a FortiGate unit to establish a Layer Two Tunneling Protocol (L2TP) tunnel with a remote dialup client. The FortiGate implementation of L2TP enables a remote dialup client to establish an L2TP tunnel with the FortiGate unit directly.

According to RFC 2661, an Access Concentrator (LAC) can establish an L2TP tunnel with an L2TP Network Server (LNS). In a typical scenario, the LAC is managed by an ISP and located on the ISP premises; the LNS is the gateway to a private network. When a remote dialup client connects to the Internet through the ISP, the ISP uses a local database to establish the identity

of the caller and determine whether the caller needs access to an LNS through an L2TP tunnel. If the services registered to the caller indicate that an L2TP connection to the LNS is required, the ISP LAC attempts to establish an L2TP tunnel with the LNS.

A FortiGate unit can be configured to act as an LNS. The FortiGate implementation of L2TP enables a remote dialup client to establish an L2TP tunnel with the FortiGate unit directly, bypassing any LAC managed by an ISP. The ISP must configure its network access server to forward L2TP traffic from the remote client to the FortiGate unit directly whenever the remote client requires an L2TP connection to the FortiGate unit.

When the FortiGate unit acts as an LNS, an L2TP session and tunnel is created as soon as the remote client connects to the FortiGate unit. The FortiGate unit assigns an IP address to the client from a reserved range of IP addresses. The remote client uses the assigned IP address as its source address for the duration of the connection.

More than one L2TP session can be supported on the same tunnel. FortiGate units can be configured to authenticate remote clients using a plain text user name and password, or authentication can be forwarded to an external RADIUS or LDAP server. L2TP clients are authenticated as members of a user group.




FortiGate units support L2TP with Microsoft Point-to-Point Encryption (MPPE) encryption only. Later implementations of Microsoft L2TP for Windows use IPsec and require certificates for authentication and encryption. If you want to use Microsoft L2TP with IPsec to connect to a FortiGate unit, the IPsec and certificate elements must be disabled on the remote client

Traffic from the remote client must be encrypted using MPPE before it is encapsulated and routed to the FortiGate unit. Packets originating at the remote client are addressed to a computer on the private network behind the FortiGate unit. Encapsulated packets are addressed to the public interface of the FortiGate unit. See [Figure 244](#).

When the FortiGate unit receives an L2TP packet, the unit disassembles the packet and forwards the packet to the correct computer on the internal network. The security policy and protection profiles on the FortiGate unit ensure that inbound traffic is screened and processed securely.

**Figure 244:**L2TP encapsulation




FortiGate units cannot deliver non-IP traffic such as Frame Relay or ATM frames encapsulated in L2TP packets — FortiGate units support the IPv4 and IPv6 addressing schemes only

## Network topology

The remote client connects to an ISP that determines whether the client requires an L2TP connection to the FortiGate unit. If an L2TP connection is required, the connection request is forwarded to the FortiGate unit directly.

**Figure 245:**Example L2TP configuration



## L2TP infrastructure requirements

- The FortiGate unit must be operating in NAT mode and have a static public IP address.
- The ISP must configure its network access server to forward L2TP traffic from remote clients to the FortiGate unit directly.
- The remote client must not generate non-IP traffic (Frame Relay or ATM frames).
- The remote client includes L2TP support with MPPE encryption. If the remote client includes Microsoft L2TP with IPSec, the IPSec and certificate components must be disabled.

## L2TP configuration overview

To configure a FortiGate unit to act as an LNS, you perform the following tasks:

- Create an L2TP user group containing one user for each remote client.
- Enable L2TP on the FortiGate unit and specify the range of addresses that can be assigned to remote clients when they connect.
- Define firewall source and destination addresses to indicate where packets transported through the L2TP tunnel will originate and be delivered.
- Create the security policy and define the scope of permitted services between the source and destination addresses.
- Configure the remote clients.

## Authenticating L2TP clients

L2TP clients must be authenticated before a tunnel is established. The authentication process relies on FortiGate user group definitions, which can optionally use established authentication mechanisms such as RADIUS or LDAP to authenticate L2TP clients. All L2TP clients are challenged when a connection attempt is made.

To enable authentication, you must create user accounts and a user group to identify the L2TP clients that need access to the network behind the FortiGate unit.

You can choose to use a plain text password for authentication or forward authentication requests to an external RADIUS or LDAP server. If password protection will be provided through a RADIUS or LDAP server, you must configure the FortiGate unit to forward authentication requests to the authentication server.

## Enabling L2TP and specifying an address range

The L2TP address range specifies the range of addresses reserved for remote clients. When a remote client connects to the FortiGate unit, the client is assigned an IP address from this range. Afterward, the FortiGate unit uses the assigned address to communicate with the remote client.

The address range that you reserve can be associated with private or routable IP addresses. If you specify a private address range that matches a network behind the FortiGate unit, the assigned address will make the remote client appear to be part of the internal network.

To enable L2TP and specify the L2TP address range, use the `config vpn l2tp` CLI command.

The following example shows how to enable L2TP and set the L2TP address range using a starting address of `192.168.10.80` and an ending address of `192.168.10.100` for an existing group of L2TP users named `L2TP_users`:

```
config vpn l2tp
 set sip 192.168.10.80
 set eip 192.168.10.100
 set status enable
 set usrgrp L2TP_users
end
```

## Defining firewall source and destination addresses

Before you define the security policy, you must define the source and destination addresses of packets that are to be transported through the L2TP tunnel:

- For the source address, enter the range of addresses that you reserved for remote L2TP clients (for example `192.168.10.[80-100]`).
- For the destination address, enter the IP addresses of the computers that the L2TP clients need to access on the private network behind the FortiGate unit (for example, `172.16.5.0/24` for a subnet, or `172.16.5.1` for a server or host, or `192.168.10.[10-15]` for an IP address range).

### To define the firewall source address

1. Go to *Firewall Objects > Address > Addresses* and select *Create New*.
2. In the *Address Name* field, type a name that represents the range of addresses that you reserved for remote clients (for example, `Ext_L2TPrange`).
3. In *Type*, select *Subnet / IP Range*.

4. In the *Subnet / IP Range* field, type the corresponding IP address range.
5. In *Interface*, select the FortiGate interface that connects to the clients.
6. This is usually the interface that connects to the Internet.
7. Select *OK*.

**To define the firewall destination address**

1. Go to *Firewall Objects > Address > Addresses* and select *Create New*.
2. In the *Address Name* field, type a name that represents a range of IP addresses on the network behind the FortiGate unit (for example, *Int\_L2TPaccess*).
3. In *Type*, select *Subnet / IP Range*.
4. In the *Subnet / IP Range* field, type the corresponding IP address range.
5. In *Interface*, select the FortiGate interface that connects to the network behind the FortiGate unit.
6. Select *OK*.

## Adding the security policy

The security policy specifies the source and destination addresses that can generate traffic inside the L2TP tunnel and defines the scope of services permitted through the tunnel. If a selection of services are required, define a service group.

**To define the traffic and services permitted inside the L2TP tunnel**

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and the *Policy Subtype* as *Address*.
3. Enter these settings:

<b>Incoming Interface</b>	Select the FortiGate interface to the Internet.
<b>Source Address</b>	Select the name that corresponds to the address range that reserved for L2TP clients (for example, <i>Ext_L2TPrange</i> ).
<b>Outgoing Interface</b>	Select the FortiGate interface to the internal (private) network.
<b>Destination Address</b>	Select the name that corresponds to the IP addresses behind the FortiGate unit (for example, <i>Int_L2TPaccess</i> ).
<b>Service</b>	Select ALL, or if selected services are required instead, select the service group that you defined previously.
<b>Action</b>	ACCEPT

4. Select *OK*.

## Configuring a Linux client

This procedure outlines how to install L2TP client software and run an L2TP tunnel on a Linux computer. Obtain an L2TP client package that meets your requirements (for example, *rp-l2tp*). If needed to encrypt traffic, obtain L2TP client software that supports encryption using MPPE.

To establish an L2TP tunnel with a FortiGate unit that has been set up to accept L2TP connections, you can obtain and install the client software following these guidelines:

1. If encryption is required but MPPE support is not already present in the kernel, download and install an MPPE kernel module and reboot your computer.
2. Download and install the L2TP client package.
3. Configure an L2TP connection to run the L2TP program.
4. Configure routes to determine whether all or some of your network traffic will be sent through the tunnel. You must define a route to the remote network over the L2TP link and a host route to the FortiGate unit.
5. Run `l2tpd` to start the tunnel.

Follow the software supplier's documentation to complete the steps.

To configure the system, you need to know the public IP address of the FortiGate unit, and the user name and password that has been set up on the FortiGate unit to authenticate L2TP clients. Contact the FortiGate administrator if required to obtain this information.

## Monitoring L2TP sessions

You can display a list of all active sessions and view activity by port number. By default, port 1701 is used for L2TP VPN-related communications. If required, active sessions can be stopped from this view. Use the Top Sessions Dashboard Widget.

## Testing L2TP VPN connections

To confirm that a VPN between a local network and a dialup client has been configured correctly, at the dialup client, issue a ping command to test the connection to the local network. The VPN tunnel initializes when the dialup client attempts to connect.

## Logging L2TP VPN events

You can configure the FortiGate unit to log VPN events. For L2TP VPNs, connection events and tunnel status (up/down) are logged.

### To log VPN events - web-based manager

1. Go to *Log & Report > Log Config > Log Settings*.
2. Enable the storage of log messages to one or more locations.
3. Select *Enable*, and then select *VPN activity event*.
4. Select *Apply*.

### To log VPN events - CLI

```
config log memory setting
 set diskfull overright
 set status enable
end
config log eventfilter
 set ppp
end
```

# Advanced concepts

This chapter provides configuration concepts and techniques to enhance your network security.

This section includes the topics:

- [Dual internet connections \(redundant Internet connections\)](#)
- [Single firewall vs. multiple virtual domains](#)
- [Modem](#)
- [DHCP servers and relays](#)
- [Assigning IP address by MAC address](#)
- [DNS services](#)
- [Dynamic DNS](#)
- [FortiClient discovery and registration](#)
- [IP addresses for self-originated traffic](#)
- [Administration for schools](#)
- [Tag management](#)
- [Replacement messages list](#)
- [Disk](#)
- [CLI Scripts](#)
- [Rejecting PING requests](#)
- [Opening TCP 113](#)
- [Obfuscate HTTP responses](#)

## Dual internet connections (redundant Internet connections)

Dual internet connection, dual WAN, or redundant internet connection refers to using two FortiGate interfaces to connect to the Internet. Dual internet connections can be used in three ways:


- redundant interfaces, should one interface go down, the second automatically becomes the main internet connection
- for load sharing to ensure better throughput.
- a combination of redundancy and load sharing.

### Redundant interfaces

Redundant interfaces, ensures that should your internet access be no longer available through a certain port, the FortiGate unit will use an alternate port to connect to the Internet.



**Figure 246:**Configuring redundant interfaces



In this scenario, two interfaces, WAN1 and WAN2 are connected to the Internet using two different ISPs. WAN1 is the primary connection. In an event of a failure of WAN1, WAN2 automatically becomes the connection to the Internet. For this configuration to function correctly, you need to configure three specific settings:

- configure a ping server to determine when the primary interface (WAN1) is down and when the connection returns
- configure a default route for each interface.
- configure security policies to allow traffic through each interface to the internal network.

### Ping server

Adding a ping server is required for routing fail over traffic. A ping server will confirm the connectivity of the device's interface

#### To add a ping server - web-based manager

1. Go to *Router > Static > Settings* and select *Create New*.  
For low-end FortiGate units, go to *System > Network > Routing* and select *Create New*.
2. Select the *Interface* that will send ping requests.
3. For the *Ping Server* field, enter the IP address of a server that the FortiGate unit will send ping requests to. This is typically a next hop router or gateway device.
4. Select the *Detect Protocol* type.
5. For the *Ping Interval*, enter the number of seconds to send ping requests.
6. For the *Failover Threshold*, enter the number of lost pings is acceptable before the port is determined to be down.
7. Select *OK*.

### To add a ping server - CLI

```
config router gwdetect
 edit wan1
 set server <ISP_IP_address>
 set failtime <failure_count>
 set interval <seconds>
 end
```

## Routing

You need to configure a default route for each interface and indicate which route is preferred by specifying the distance. The lower distance is declared active and placed higher in the routing table.



When you have dual WAN interfaces that are configured to provide fail over, you might not be able to connect to the backup WAN interface because the FortiGate unit may not route traffic (even responses) out of the backup interface. The FortiGate unit performs a reverse path lookup to prevent spoofed traffic. If no entry can be found in the routing table which sends the return traffic out the same interface, then the incoming traffic is dropped.

### To configure the routing of the two interfaces - web-based manager

1. Go to *Router > Static > Static Routes* and select *Create New*.  
For low-end FortiGate units, go to *System > Network > Routing* and select *Create New*.
2. Set the *Destination IP/Mask* to the address and netmask to 0.0.0.0/0.0.0.0.
3. Select the *Device* to the primary connection, *WAN1*.
4. Enter the *Gateway* address.
5. Select *Advanced*.
6. Set the *Distance* to 10.
7. Select *OK*.
8. Repeat steps 1 through 7 setting the *Device* to *WAN2* and a *Distance* of 20.

### To configure the routing of the two interfaces - CLI

```
config router static
 edit 1
 set dst 0.0.0.0 0.0.0.0
 set device WAN1
 set gateway 0.0.0.0 0.0.0.0
 set distance 10
 next
 edit 1
 set dst <ISP_Address>
 set device WAN2
 set gateway <gateway_address>
 set distance 20
 next
end
```

## Security policies

When creating security policies, you need to configure duplicate policies to ensure that after traffic fails over WAN1, regular traffic will be allowed to pass through WAN2 as it did with WAN1. This ensures that fail-over will occur with minimal affect to users. For more information on creating security policies see the [Firewall Guide](#).

## Load sharing

Load sharing enables you to use both connections to the internet at the same time, but do not provide fail over support. When configuring for load sharing, you need to ensure routing is configured for both external ports, for example, WAN1 and WAN2, have static routes with the same distance and priority.

Further configuration can be done using Equal Cost Multiple Path (ECMP). For more information on ECMP and load sharing, see the [Advanced Routing Guide](#).

## Link redundancy and load sharing

In this scenario, both links are available to distribute Internet traffic over both links. Should one of the interfaces fail, the FortiGate unit will continue to send traffic over the other active interface. Configuration is similar to the [Redundant interfaces](#) configuration, with the main difference being that the configured routes should have equal distance settings.

This means both routes will remain active in the routing table. To make one interface the preferred interface, use a default policy route to indicate the interface that is preferred for accessing the Internet. If traffic matches the security policy, the policy overrides all entries in the routing table, including connected routes. You may need to add a specific policy routes that override these default policy routes.

To redirect traffic over the secondary interface, create policy routes to direct some traffic onto it rather than the primary interface. When adding the policy route, only define the outgoing interface and leave the gateway blank. This ensures that the policy route will not be active when the link is down.

## Single firewall vs. multiple virtual domains

A typical FortiGate setup, with a small to mid-range appliance, enables you to include a number of subnets on your network using the available ports and switch interfaces. This can potentially provide a means of having three or more mini networks for the various groups in a company. Within this infrastructure, multiple network administrators have access to the FortiGate to maintain security policies.

However, the FortiGate unit may not have enough interfaces to match the number of departments in the organization. If the FortiGate unit is running in transparent mode however, there is only one interface, and multiple network branches through the FortiGate are not possible.


A FortiGate unit with Virtual Domains (VDMs) enabled, provides a means to provide the same functionality in transparent mode as a FortiGate in NAT mode. VDMs are a method of dividing a FortiGate unit into two or more virtual units that function as multiple independent units. VDMs can provide separate security policies and, in NAT mode, completely separate configurations for routing and VPN services for each connected network. For administration, an administrator can be assigned to each VDM, minimizing the possibility of error or fouling network communications.

By default, your FortiGate unit supports a maximum of 10 VDMs. For select FortiGate models you can purchase a license key to increase the number of VDMs.


The FortiGate-20C and 30B and FortiWifi-20C and 30B do not support VDOMs.

## Single firewall vs. vdoms

When VDOMs are not enabled, and the FortiGate unit is in transparent mode, all the interfaces on your unit become broadcast interfaces. The problem is there are no interfaces free for additional network segments.



A FortiGate with three interfaces means only limited network segments are possible without purchasing more FortiGate devices.



With multiple VDOMs you can have one of them configured in transparent mode, and the rest in NAT mode. In this configuration, you have an available transparent mode FortiGate unit you can drop into your network for troubleshooting, and you also have the standard.

This example shows how to enable VDOMs on the FortiGate unit and the basic and create a VDOM accounting on the DMZ2 port and assign an administrator to maintain the VDOM. First enable Virtual Domains on the FortiGate unit.

### To enable VDOMs - web-based manager

1. Go to *System > Dashboard > Status*.
2. In the *System Information* widget, select *Enable* for *Virtual Domain*.

Note that on FortiGate-60 series and lower models, you need to enable VDOMs in the CLI only.

The FortiGate unit logs you out. Once you log back in, you will notice that the menu structure has changed. This reflects the global settings for all Virtual Domains.

### To enable VDOMs - CLI

```
config system global
 set vdom-admin enable
end
```

Next, add the VDOM called accounting.

### To add a VDOM - web-based manager

1. Go to *Global > VDOM > VDOM*, and select *Create New*.
2. Enter the VDOM name *accounting*.
3. Select *OK*.

### To add a VDOM - CLI

```
config vdom
 edit <new_vdom_name>
end
```

With the Virtual Domain created, you can assign a physical interface to it, and assign it an IP address.

### To assign physical interface to the accounting Virtual Domain - web-based manager

1. Go to *Global > Network > Interface*.
2. Select the DMZ2 port row and select *Edit*.
3. For the *Virtual Domain* drop-down list, select *accounting*.
4. Select the *Addressing Mode* of *Manual*.
5. Enter the IP address for the port of 10.13.101.100/24.
6. Set the *Administrative Access* to *HTTPS* and *SSH*.
7. Select *OK*.

### To assign physical interface to the accounting Virtual Domain - CLI

```
config global
 config system interface
 edit dmz2
 set vdom accounting
 set ip 10.13.101.100/24
 set allowaccess https ssh
 next
 end
```

## Modem

FortiGate units support the use of wireless, 3G and 4G modems connected using the USB port or, if available, the express card slot. Modem access provides either primary or secondary (redundant) access to the Internet. For FortiGate units that do not include an internal modem (those units with an “M” designation), the modem interface will not appear in the web-based manager until enabled in the CLI. To enable the modem interface enter the CLI commands:

```
config system modem
 set status enable
end
```

You will need to log out of the FortiGate and log back in to see the modem configuration page at *System > Network > Modem*. Once enabled, modem options become available by going to *System > Network > Interface*.

Note that the modem interface is only available when the FortiGate unit is in NAT mode.

To configure modem settings, go to *System > Network > Modem*.

Configuring the modem settings is a matter of entering the ISP phone number, user name and password. Depending on the modem, additional information may need to be supplied such as product identifiers, and initialization strings.

The FortiGate unit includes a number of common modems within its internal database. You can view these by selecting the *Configure Modem* link on the *Modem Settings* page. If your modem is not on the list, select *Create New* to add the information. This information is stored on the device, and will remain after a reboot.

Fortinet has an online database of modem models and configuration settings through FortiGuard. A subscription to the FortiGuard services is not required to access the information. As models are added, you can select the *Configure Modem* link and select *Update Now* to download new configurations.

## USB modem port

Each USB modem has a specific dial-out ttyusb port. This will be indicated with the documentation for your modem. To enable the correct USB port, use the CLI commands:

```
config system modem
 set wireless-port {ttyusb0 | ttyusb1 | ttyusb2}
end
```

To test the port, use the diagnose command:

```
diagnose sys modem com /ttyusb1
```

The ttyusb1 will be the value of your USB port selected. The response will be:

```
Serial port: /dev/ttyusb1
Press Ctrl+W to exit.
```

If the port does not respond the output will be:

```
Can not open modem device '/dev/ttyusb1' : Broken pipe
```

## Modes

The FortiGate unit allows for two modes of operation for the modem; stand alone and redundant. In stand alone mode, the modem connects to a dialup ISP account to provide the connection to the Internet. In redundant mode, the modem acts as a backup method of connecting to the Internet, should the primary port for this function fails.

Configuring either stand alone or redundant modes are very similar. The primary difference is the selection of the interface that the modem will replace in the event of it failing, and the configuration of a PING server to monitor the chosen interface.

### Configuring stand alone mode

Configuring stand alone mode is a matter of configuring the modem information and the dialing mode. The dial mode is either *Always Connect* or *Dial on demand*. Selecting *Always Connect* ensures that once the modem has connected, it remains connected to the ISP. Selecting *Dial on Demand*, the modem only calls the ISP if packets are routed to the modem interface. Once sent, the modem will disconnect after a specified amount of time.

#### To configure standalone mode as needed - web-based manager

1. Go to *System > Network > Modem*.
2. Select the *Mode* of *Standalone*.
3. Select the *Dial Mode* of *Dial on Demand*.
4. Enter the *Idle Timeout* of 2 minutes.
5. Select the number of redials the modem attempts if connection fails to 5.
6. Select *Apply*.

#### To configure standalone mode as needed- CLI

```
config system modem
 set mode standalone
 set auto-dial enable
 set idle-timer 2
 set redial 5
end
```

### Configuring redundant mode

Redundant mode provides a backup to an interface, typically to the Internet. If that interface fails or disconnects, the modem automatically dials the configured phone number(s). Once connected, the FortiGate unit routes all traffic to the modem interface until the monitored interface is up again. The FortiGate unit pings the connection to determine when it is back online.

For the FortiGate to verify when the interface is back up, you need to configure a Ping server for that interface. You will also need to configure security policies between the modem interface and the other interfaces of the FortiGate unit to ensure traffic flow.

#### To configure redundant mode as needed - web-based manager

1. Go to *System > Network > Modem*.
2. Select the *Mode* of *Redundant*.
3. Select the interface the modem takes over from if it fails.
4. Select the *Dial Mode* of *Dial on Demand*.
5. Enter the *Idle Timeout* of 2 minutes.
6. Select the number of redials the modem attempts if connection fails to 5.
7. Select *Apply*.

### To configure standalone mode as needed- CLI

```
config system modem
 set mode redundant
 set interface wan1
 set auto-dial enable
 set idle-timer 2
 set redial 5
end
```

### Ping server

Adding a ping server is required for routing fail over traffic. A ping server will confirm the connectivity of the device's interface.

For low-end FortiGate units, go to *System > Admin > Settings* and enable *Dynamic Routing* before continuing.

#### To add a ping server - web-based manager

1. Go to *Router > Static > Settings* and select *Create New*.
2. Select the *Interface* that will send ping requests.
3. For the *Ping Server* field, enter the IP address of a server that the FortiGate unit will send ping requests to. This is typically a next hop router or gateway device.
4. Select the *Detect Protocol* type *ICMP Ping*.
5. For the *Ping Interval*, enter the number of seconds to send ping requests.
6. For the *Failover Threshold*, enter the number of lost pings is acceptable before the port is determined to be down.
7. Select *OK*.

#### To add a ping server - CLI

```
config router gwdetect
 edit wan1
 set server <ISP_IP_address>
 set failtime <failure_count>
 set interval <seconds>
 end
```

### Additional modem configuration

The CLI provides additional configuration options when setting up the modem options including adding multiple ISP dialing and initialization options and routing. For more information, see the [CLI Reference](#).

### Modem interface routing

The modem interface can be used in FortiOS as a dedicated interface. Once enabled and configured, you can use it in security policies and define static and dynamic routing. Within the CLI commands for the modem, you can configure the distance and priority of routes involving the modem interface. The CLI commands are:

```
config sysetm modem
 set distance <route_distance>
 set priority <priority_value>
end
```



For more information on the routing configuration in the CLI, see the [CLI Reference](#). For more information on routing and configuring routing, see the [Advanced Routing Guide](#).

## DHCP servers and relays

Note that DHCP server options are not available in transparent mode.

A DHCP server provides an address to a client on the network, when requested, from a defined address range.

An interface cannot provide both a server and a relay for connections of the same type (regular or IPSec). However, you can configure a Regular DHCP server on an interface only if the interface is a physical interface with a static IP address. You can configure an IPSec DHCP server on an interface that has either a static or a dynamic IP address.

You can configure one or more DHCP servers on any FortiGate interface. A DHCP server dynamically assigns IP addresses to hosts on the network connected to the interface. The host computers must be configured to obtain their IP addresses using DHCP.

If an interface is connected to multiple networks via routers, you can add a DHCP server for each network. The IP range of each DHCP server must match the network address range. The routers must be configured for DHCP relay.

You can configure a FortiGate interface as a DHCP relay. The interface forwards DHCP requests from DHCP clients to an external DHCP server and returns the responses to the DHCP clients. The DHCP server must have appropriate routing so that its response packets to the DHCP clients arrive at the unit.

### DHCP Server configuration

To add a DHCP server, go to *System > Network > Interface*. Edit the interface, and select *Enable* for the *DHCP Server* row.

<b>DHCP Server IP</b>	This appears only when <i>Mode</i> is <i>Relay</i> . Enter the IP address of the DHCP server where the FortiGate unit obtains the requested IP address.
<b>Address Range</b>	By default, the FortiGate unit assigns an address range based on the address of the interface for the complete scope of the address. For example, if the interface address is 172.20.120.230, the default range created is 172.20.120.231 to 172.20.120.254. Select the range and select <i>Edit</i> to adjust the range as needed, or select <i>Create New</i> to add a different range.
<b>Netmask</b>	Enter the netmask of the addresses that the DHCP server assigns.
<b>Default Gateway</b>	Select to either use the same IP as the interface or select <i>Specify</i> and enter the IP address of the default gateway that the DHCP server assigns to DHCP clients.
<b>DNS Server</b>	Select to use the system's DNS settings or select <i>Specify</i> and enter the IP address of the DNS server.
<b>Advanced</b>	
<b>Mode</b>	Select the type of DHCP server the FortiGate unit will be. By default, it is a server. Select <i>Relay</i> if needed. When <i>Relay</i> is selected, the above configuration is replaced by a field to enter the <i>DHCP Server IP</i> address.

<b>Type</b>	Select to use the DHCP in regular or IPsec mode.
<b>MAC Address Access Control List</b>	<p>Select to match an IP address from the DHCP server to a specific client or device using its MAC address.</p> <p>In a typical situation, an IP address is assigned ad hoc to a client, and that assignment times out after a specific time of inactivity from the client, known as the lease time. To ensure a client or device always has the same IP address, that is, there is no lease time, use IP reservation.</p> <p>For more information, see <a href="#">“Assigning IP address by MAC address” on page 1580</a>.</p>
<b>Add from DHCP Client List</b>	If the client is currently connected and using an IP address from the DHCP server, you can select this option to select the client from the list.

## DHCP in IPv6

You can use DHCP with IPv6 using the CLI. To configure DHCP, ensure IPv6 is enabled by going to *System > Admin > Settings* and enable *IPv6*. Use the CLI command

```
config system dhcp6.
```

For more information on the configuration options, see the [CLI Reference](#).

## Service

On low-end FortiGate units, a DHCP server is configured, by default on the Internal interface:

<b>IP Range</b>	192.168.1.110 to 192.168.1.210
<b>Netmask</b>	255.255.255.0
<b>Default gateway</b>	192.168.1.99
<b>Lease time</b>	7 days
<b>DNS Server 1</b>	192.168.1.99

These settings are appropriate for the default Internal interface IP address of 192.168.1.99. If you change this address to a different network, you need to change the DHCP server settings to match.

Alternatively, after the FortiGate unit assigns an address, you can go to *System > Monitor > DHCP Monitor*, locate the particular user. Select the check box for the user and select *Add to Reserved*.

## Lease time

The lease time determines the length of time an IP address remains assigned to a client. Once the lease expires, the address is released for allocation to the next client request for an IP

address The default lease time is seven days. To change the lease time, use the following CLI commands:

```
config system dhcp server
 edit <server_entry_number>
 set lease-time <seconds>
 end
```

To have an unlimited lease time, set the value to zero.

## DHCP options

When adding a DHCP server, you have the ability to include DHCP codes and options. The DHCP options are BOOTP vendor information fields that provide additional vendor-independent configuration parameters to manage the DHCP server. For example, you may need to configure a FortiGate DHCP server that gives out a separate option as well as an IP address. For example, an environment that needs to support PXE boot with Windows images.

The option numbers and codes are specific to the particular application. The documentation for the application will indicate the values to use. Option codes are represented in a option value/HEX value pairs. The option is a value 1 and 255.

You can add up to three DHCP code/option pairs per DHCP server.

### To configure option 252 with value <http://192.168.1.1/wpad.dat> - CLI

```
config system dhcp server
 edit <server_entry_number>
 set option1 252
 687474703a2f2f3139322e3136382e312e312f777061642e646174
 end
```

For detailed information about DHCP options, see [RFC 2132](#), DHCP Options and BOOTP Vendor Extensions.

## Exclude addresses in DHCP a range

If you have a large address range for the DHCP server, you can block a range of addresses that will not be included in the available addresses for the connecting users. To do this, go to the CLI and enter the commands:

```
config system dhcp server
 edit <server_entry_number>
 config exclude-range
 edit <sequence_number>
 set start-ip <address>
 set end-ip <address>
 end
 end
 end
```

## DHCP Monitor

To view information about DHCP server connections, go to *System > Monitor > DHCP Monitor*. On this page, you can also add IP address to the reserved IP address list.

## Breaking a address lease

Should you need to end an IP address lease, you can break the lease using the CLI. This is useful if you have limited addresses, longer lease times where leases are no longer necessary. For example, with corporate visitors.

**To break a lease enter the CLI command:**

```
execute dhcp lease-clear <ip_address>
```

## Assigning IP address by MAC address

To prevent users in the from changing their IP addresses and causing IP address conflicts or unauthorized use of IP addresses, you can bind an IP address to a specific MAC address using DHCP.

Use the CLI to reserve an IP address for a particular client identified by its device MAC address and type of connection. The DHCP server then always assigns the reserved IP address to the client. The number of reserved addresses that you can define ranges from 10 to 200 depending on the FortiGate model.

After setting up a DHCP server on an interface by going to *System > Network > Interface*, select the blue arrow next to *Advanced* to expand the options. If you know the MAC address of the system select *Create New* to add it, or if the system has already connected, locate it in the list, select its check box and select *Add from DHCP Client List*.

You can also match an address to a MAC address in the CLI. In the example below, the IP address 10.10.10.55 for User1 is assigned to MAC address 00:09:0F:30:CA:4F.

```
config system dhcp reserved-address
 edit User1
 set ip 10.10.10.55
 set mac 00:09:0F:30:CA:4F
 set type regular
 end
```

## DNS services

A DNS server is a public service that converts symbolic node names to IP addresses. A Domain Name System (DNS) server implements the protocol. In simple terms, it acts as a phone book for the Internet. A DNS server matches domain names with the computer IP address. This enables you to use readable locations, such as fortinet.com when browsing the Internet. FortiOS supports DNS configuration for both IPv4 and IPv6 addressing.

The FortiGate unit includes default DNS server addresses. However, these should be changed to those provided by your Internet Service Provider. The defaults are DNS proxies and are not as reliable as those from your ISP.

Within FortiOS, there are two DNS configuration options; each provide a specific service, and can work together to provide a complete DNS solution.

## DNS settings

Basic DNS queries are configured on interfaces that connect to the Internet. When a web site is requested, for example, the FortiGate unit will look to the configured DNS servers to provide the IP address to know which server to contact to complete the transaction.

DNS server addresses are configured by going to *System > Network > DNS*. Here you specify the DNS server addresses. Typically, these addresses are supplied by your ISP. An additional option is available if you have local Microsoft domains on the network, by entering a domain name in the *Local Domain Name* field.

In a situation where all three fields are configured, the FortiGate unit will first look to the local domain. If no match is found, a request is sent to the external DNS servers.

If virtual domains are enabled, you create a DNS database in each VDOM. All of the interfaces in a VDOM share the DNS database in that VDOM.

## Additional DNS CLI configuration

Further options are available from the CLI with the command `config system dns`. Within this command you can set the following commands:

- `dns-cache-limit` - enables you to set how many DNS entries are stored in the cache. Entries that remain in the cache provide a quicker response to requests than going out to the Internet to get the same information.
- `dns-cache-ttl` - enables you to set how long entries remain in the cache in seconds, between 60 and 86,400 (24 hours).
- `cache-notfound-responses` - when enabled, any DNS requests that are returned with NOTFOUND can be stored in the cache.
- `source-ip` - enables you to define a dedicated IP address for communications with the DNS server.

## DNS server

You can also create local DNS servers for your network. Depending on your requirements, you can manually maintain your entries (master DNS server), or use it as a jumping point, where the server refers to an outside source (slave DNS server). A local master DNS server works similarly to the DNS server addresses configured in *System > Network > DNS*, but all entries must be added manually. This enables you to add a local DNS server to include specific URL/IP address combinations.

The DNS server options are not visible in the web-based manager by default. To enable the server, go to *System > Admin > Settings* and select *DNS Database*.

While a master DNS server is an easy method of including regularly used addresses to save on going to an outside DNS server, it is not recommended to make it the authoritative DNS server. IP addresses may change, and maintaining any type of list can quickly become labor-intensive.

A FortiGate master DNS server is best set for local services. For example, if your company has a web server on the DMZ that is accessed by internal employees as well as external users, such as customers or remote users. In this situation, the internal users when accessing the site would send a request for `website.example.com`, that would go out to the DNS server on the web, to return an IP address or virtual IP. With an internal DNS, the same site request is resolved internally to the internal web server IP address, minimizing inbound/outbound traffic and access time.

As a slave, DNS server, the FortiGate server refers to an external or alternate source as way to obtain the url/IP combination. This useful if there is a master DNS server for a large company where a list is maintained. Satellite offices can then connect to the master DNS server to obtain the correct addressing.

The DNS server entries does not allow CNAME entries, as per [rfc 1912](#), section 2.4.

### To configure a master DNS server - web-based manager

1. Go to *System > Network > DNS Server*, and select *Create New*.
2. Select the *Type* of *Master*.
3. Select the *View* as *Shadow*.
4. The view is the accessibility of the DNS server. Selecting *Public*, external users can access, or use, the DNS server. Selecting *Shadow*, only internal users can use it.
5. Enter the *DNS Zone*, for example, *WebServer*.
6. Enter the domain name for the zone, for example *example.com*.
7. Enter the hostname of the DNS server, for example, *Corporate*.
8. Enter the contact address for the administrator, for example, *admin@example.com*.
9. Set *Authoritative* to *Disable*.
10. Select *OK*.
11. Enter the DNS entries for the server by selecting *Create New*.
12. Select the *Type*, for example, *Address (A)*.
13. Enter the *Hostname*, for example *web.example.com*.
14. Enter the remaining information, which varies depending on the *Type* selected.
15. Select *OK*.

### To configure a DNS server - CLI

```
config system dns-database
 edit WebServer
 set domain example.com
 set type master
 set view shadow
 set ttl 86400
 set primary-name corporate
 set contact admin@exmple.com
 set authoritative disable
 config dns-entry
 edit 1
 set hostname web.example.com
 set type A
 set ip 192.168.21.12
 set status enable
 end
 end
end
```

## Recursive DNS

You can set an option to ensure these types of DNS server is not the authoritative server. When configured, the FortiGate unit will check its internal DNS server (Master or Slave). If the request cannot be fulfilled, it will look to the external DNS servers. This is known as a split DNS configuration.

You can also have the FortiGate unit look to an internal server should the Master or Slave not fulfill the request by using the CLI commands:

```
config system dns-database
 edit example.com
 ...
 set view shadow
end
```

For this behavior to work completely, for the external port, you must set the DNS query for the external interface to be recursive. This option is configured in the CLI only.

### To set the DNS query

```
config system dns-server
 edit wan1
 set mode recursive
end
```

## Dynamic DNS

If your ISP changes your external IP address on a regular basis, and you have a static domain name, you can configure the external interface to use a dynamic DNS service to ensure external users and/or customers can always connect to your company firewall.

If you have a FortiGuard subscription, you can use FortiGuard as your DDNS server. To configure dynamic DNS in the web-based manager, go to *System > Network > DNS*, select *Enable FortiGuard DDNS*, and enter the relevant information for the interface communicating to the server, and which server to use, and relevant information.

If you do not have a FortiGuard subscription, or want to use an alternate server, you can configure dynamic DNS in the CLI use the commands below. Within the CLI you can configure a DDNS for each interface. Only the first configured port appears in the web-based manager. Additional commands vary with the DDNS server you select.

```
config system ddns
 edit <instance_value>
 set monitor-interface <external_interface>
 set ddns-server <ddns_server_selection>
end
```

You can also use FortiGuard (when subscribed) as a DDNS as well. To configure, use the CLI commands:

```
config system fortiguard
 set ddns-server-ip
 set ddns-server-port
end
```

## FortiClient discovery and registration

FortiOS provides a means of allowing users running FortiClient Endpoint Control software to connect to specific interfaces when connecting to the FortiGate unit. As well as ensuring that remote or local users have FortiClient Endpoint Control software installed on their PC or mobile device.

## FortiClient discovery

You can configure a FortiGate interface as an interface that will accept FortiClient connections. When configured, the FortiGate unit sends broadcast messages which the FortiClient software running on an end user PC is listening for.

### To enable the broadcast message

1. Go to *System > Network > Interface*.
2. Edit the interface to send the broadcast messages.
3. Select *FCT-Access*.
4. In *Device Management*, select *Broadcast Discovery Messages*.
5. Select *OK*.

Once enabled, the FortiGate unit broadcasts a discovery message that includes the IP address of the interface and listening port number to the local network. All PCs running FortiClient on that network listen for this discovery message.

You also have the option of including a registration key. When the FortiClient discovers the FortiGate unit, it is prompted to enter a registration key, defined by the administrator.

### To add a registration key

1. Go to *System > Config > Advanced*.
2. Select *Enable Registration Key for FortiClient*, and enter the key.
3. Select *Apply*.

Ensure you distribute the key to the users that need to connect to the FortiGate unit.

## FortiClient Registration

On the end user side, if FortiClient has not been registered with the FortiGate unit, it is continually listening for the FortiGate discovery message. When this message is detected the un-registered client will pop-up a FortiGate Detected message. The user can choose to either register or ignore the message.

Clients that have registered with that FortiGate unit will not be listening for these messages and will not display the message again.

If you enabled the registration key, the user is prompted to enter the key before a connection can be completed.

For more information on FortiGate registration, see the [FortiClient Administration Guide](#).

## IP addresses for self-originated traffic

On the FortiGate unit, there are a number of protocols and traffic that is specific to the internal workings of FortiOS. For many of these traffic sources, you can identify a specific port/IP



address for this self-originating traffic. The following traffic can be configured to a specific port/IP address:

- SNMP
- Syslog
- alert email
- FortiManager connection IP
- FortiGuard services
- FortiAnalyzer logging
- NTP
- DNS
- Authorization requests such as RADIUS
- FSAE

Configuration of these services is performed in the CLI. In each instance, there is a command `set source-ip`. For example, to set the source IP of NTP to be on the DMZ1 port with an IP of 192.168.4.5, the commands are:

```
config system ntp
 set ntpsyn enable
 set syncinterval 5
 set source-ip 192.168.4.5
end
```

To see which services are configured with source-ip settings, use the `get` command:

```
get system source-ip status
```

The output will appear similar to the sample below:

```
NTP: x.x.x.x
DNS: x.x.x.x
SNMP: x.x.x.x
Central Management: x.x.x.x
FortiGuard Updates (AV/IPS): x.x.x.x
FortiGuard Queries (WebFilter/SpamFilter): x.x.x.x
```

## Administration for schools

For system administrator in the school system it is particularly difficult to maintain a network and access to the Internet. There are potential legal liabilities if content is not properly filtered and children are allowed to view pornography and other non-productive and potentially dangerous content. For a school, too much filtering is better than too little. This section describes some basic practices administrators can employ to help maintain control without being too draconian for access to the internet.

### Security policies

The default security policies in FortiOS allow all traffic on all ports and all IP addresses. Not the most secure. While applying UTM profiles can help to block viruses, detect attacks and prevent spam, this doesn't provide a solid overall security option. The best approach is a layered approach; the first layer being the security policy.

When creating outbound security policies, you need to know the answer to the question “What are the students allowed to do?” The answer is surf the web, connect to FTP sites, send/receive email, and so on.

Once you know what the students need to do, you can research the software used and determine the ports the applications use. For example, if the students only require web surfing, then there are only two ports (80 - HTTP and 443 - HTTPS) needed to complete their tasks. Setting the security policies to only allow traffic through two ports (rather than all 65,000), this will significantly lower any possible exploits. By restricting the ports to known services, means stopping the use of proxy servers, as many of them operate on a non-standard port to hide their traffic from URL filtering or HTTP inspection.

## DNS

Students should not be allowed to use whatever DNS they want. This opens another port for them to use and potentially smuggle traffic on. The best approach is to point to an internal DNS server and only allow those devices out on port 53. It's the same approach one would use for SMTP. Only allow the mail server to use port 25 since nothing else should be sending email.

If there is no internal DNS server, then the list of allowed DNS servers they can use should be restrictive. One possible exploit would be for them to set up their own DNS server at home that serves different IPs for known hosts, such as having Google.com sent back the IP for playboy.com.

## Encrypted traffic (HTTPS)

Generally speaking, students should not be allowed to access encrypted web sites. Encrypted traffic cannot be sniffed, and therefore, cannot be monitored. HTTPS traffic should only be allowed when necessary. Most web sites a student needs to access are HTTP, not HTTPS. Due to the nature of HTTPS protocol, and the fact that encryption is an inherent security risk to your network, its use should be restricted.

Adding a security policy that encompasses a list of allowed secure sites will ensure that any HTTPS sites that are required are the only sites a student can go to.

## FTP

For the most part, students should not be using FTP. FTP is not HTTP or HTTPS so you cannot use URL flitting to restrict where they go. This can be controlled with destination IPs in the security policy. With a policy that specifically outlines which FTP addresses are allowed, all other will be blocked.

## Example security policies

Given these requirements, an example set of security policies could look like the following illustration. In a large setup, all the IPs for the students are treated by one of these four policies.

**Figure 247:**Simple security policy setup

<input type="checkbox"/>	Seq. No.	ID	Source	Destination	Schedule	Service	Action	Status
<input type="checkbox"/>	1	2	Student PCs	Allowed Websites	always	HTTPS	✓	✓
<input type="checkbox"/>	2	3	Student PCs	all	always	HTTP	✓	✓
<input type="checkbox"/>	3	4	Student PCs	Allowed DNS	always	DNS	✓	✓
<input type="checkbox"/>	4	5	Student PCs	Allowed FTP	always	FTP	✓	✓
<input type="checkbox"/>	5		all	all	always	ANY	⊘	Implicit

The last policy in the list, included by default, is a deny policy. This adds to the potential of error that could end up allowing unwanted traffic to pass. The deny policy ensures that any traffic making it to this point is stopped. It can also help in further troubleshooting by viewing the logs for denied traffic.

With these policies in place, even before packet inspection occurs, the FortiGate, and the network are fairly secure. Should any of the UTM profiles fail, there is still a basic level of security.

## UTM security profiles

### Antivirus profiles

Antivirus screening should be enabled for any service you have enabled in the security policies. In the case above, HTTP, FTP, as well as POP3 and SMTP (assuming there is email access for students). There is not a virus scan option for HTTPS, because the content is encrypted. Generally speaking, most of the network traffic will be students surfing the web.

To configure antivirus profiles in the web-based manager, go to *UTM Security Profiles > Antivirus > Profile*, or use the CLI commands under `config antivirus profile`.

### Web filtering

The actual filtering of URLs - sites and content - should be performed by FortiGuard. It is easier and web sites are constantly being monitored, and new ones reviewed and added to the FortiGuard databases every day. The FortiGuard categories provide an extensive list of offensive, and non-productive sites.

As well, there are additional settings to include in a web filtering profile to best contain a student's web browsing.

- Web URL filtering should be enabled to set up exemptions for web sites that are blocked or reasons other than category filtering. It also prevents the use of IP addresses to get around web filtering.
- Block invalid URLs - HTTPS only. This option inspects the HTTPS certificate and looks at the URL to ensure it's valid. It is common for proxy sites to create an HTTPS certificate with a garbage URL. If the site is legitimate, it should be set up correctly. If the site approach to security is to ignore it, then their security policy puts your network at risk and the site should be blocked.

Web filtering options are configured in the web-based manager by going to *UTM Security Profiles > Web filter > Profile*, or in the CLI under `config webfilter profile`.

### Advanced options

There are a few Advanced options to consider for a web filtering profile:

- Enable *Provide details for blocked HTTP 4xx and 5xx errors*. Under normal circumstances there are exploits that can be used with 400 and 500 series messages to access the web

site. While most students probably won't know how to do this, there is no harm in being cautious. It only takes one.

- Enable *Rate Images by URL*. This option only works with Google images. It examines the URL that the image is stored at to get a rating on it, then blocks or allows the image based on the rating of the originating URL. It does not inspect the image contents. Most image search engines to a perfect and pass the images directly to the browser.
- Enable *Block HTTP redirects by rating*. An HTTP redirect is one method of getting around ratings. Go to one web site that has an allowed rating, and it redirects to another web site that may want blocked.

## Categories and Classifications

For the selection of what FortiGuard categories and classifications that should be blocked, that is purely based on the school system and its Internet information policy.

## Email Filtering

Other than specific teacher-led email inboxes, there is no reason why a student should be able to access, read or send personal email. Ports for POP3, SMTP and IMAP should not be opened in a security policies.

## IPS

The intrusion protection profiles should be used to ensure the student PCs are not vulnerable to attacks, nor do you want students making attacks. As well, IPS can do more than simple vulnerability scans. With a FortiGuard subscription, IPS signatures are pushed to the FortiGate unit. New signatures are released constantly for various intrusions as they are discovered.

FortiOS includes a number of predefined IPS sensors that you can enable by default. Selecting the all\_default signature is a good place to start as it includes the major signatures.

To configure IPS sensors in the web-based manager, go to *UTM Security Profiles > Intrusion Protection > IPS Sensor*, on the CLI use commands under `config ips sensor`.

## Application control

Application control uses IPS signatures to limit the use of instant messaging and peer-to-peer applications which can lead to possible infections on a student's PC. FortiOS includes a number of pre-defined application categories. To configure and maintain application control profiles in the web-based manager, go to *UTM Security Profiles > Application Control > Application Sensor*. In the CLI use commands under `config application list`.

Some applications to consider include proxies, botnets, toolbars and P2P applications.

## Logging

Turn on all logging - every option in this section should be enabled. This is not where you decide what you are going to log. It is simply defining what the UTM profiles can log.

Logging everything is a way to monitor traffic on the network, see what student's are utilizing the most, and locate any potential holes in your security plan. As well, keeping this information may help to prove negligence later in necessary.

## Tag management

Tag management provide a method of categorizing, or labelling objects within FortiOS using keywords. You can give the following elements a “tag”, similar to a keyword:

- IPS signature
- application signature
- security policy
- firewall address

Tagging is way to organize the various elements, especially if you have a large number of addresses, security policies to manage and keep track of. Tagging enables you to break these elements into groups, but each element can belong to more than one group. Tags help you find elements which have something in common, be it a group, user or location. This is very similar to tagging found on photo sharing sites.

To use tagging, you need to enable it for 1U FortiGate units. It is enabled by default on all 2U FortiGate units and blades.

### To enable tagging - web-based manager

1. Go to *System > Admin > Settings*.
2. Select *Object Tagging and Coloring*.
3. Select *Apply*.

### To enable tagging - CLI

```
config system global
 set gui-object-tags enable
end
```

## Adding and removing tags

You add and remove tags when you create the various elements. For example, when adding a firewall address, a section below the Interface selection enables you to add tags for that element, such as the department, region, or really, anything to help identify the element. When editing, applied tags appear as well. To add a tag, right-click on the element you want to add a tag to.

**Figure 248:**Adding tags to a new address.

The screenshot shows the 'New Address' configuration window in FortiOS. The fields are as follows:

- Address Name: User\_1
- Color: [Change]
- Type: Subnet / IP Range
- Subnet / IP Range: 172.20.120.12
- Interface: dmz2
- Tags: Applied tags list contains 'accounting'. The 'Add tags' input field contains 'west coast'.

To remove a tag, in the element, click the tag in the Applied Tags list.

## Reviewing tags

Tags can be reviewed in one location by going to *System > Config > Tag Management*. In this screen, all tags used appear. The visual size of the tag name indicates the usage; the bigger the size, the more it is used. By hovering over the keyword, a fly out indicates how many times it has been used.

To see where it was used, click the keyword. An *Object Usage* window displays all the reference categories where the keyword was used, and the number of times. Selecting the expand arrow further details its use.

Further, for security policies for example, you can select the *View* icon and see the details of the particular element. If need be, select the *Edit* icon to modify the element.

**Figure 249:** Viewing the address information for a tagged object



## Tagging guidelines

Given the ease that tags can be added to elements in FortiOS, it makes sense to jump right in and begin applying tags to elements and object. However, this type of methodology will lead to problems down the road as new elements are added.

A methodology should be considered and developed before applying tags. This doesn't mean you need to develop an entire thesaurus or reference guide for all possibilities of tags. However, taking some time to develop a methodology for the keywords you intend to use will benefit later when new security policies, addresses, and so on are added. Some things to consider when developing a tag list:

- the hierarchy used for the organization such as region, city location, building location
- department names and if short forms or long forms are used
- will acronyms be used or terms spelled out.
- how granular will the tagging be

As tags are added, previously used tags appear so there is an opportunity to use previously used tags. However, you want to avoid a situation where both accounting and acct are both

options. This is also important if there are multiple administrators in different locations to ensure consistency.

At any time, you can change or even remove tags. It is best to do a bit of planning ahead of time to avoid unnecessary work later on.

## Replacement messages list

The replacement message list in *System > Config > Replacement Messages*.

The replacement messages list enables you to view and customize replacement messages. Use the expand arrow beside each type to display the replacement messages for that category. Select the *Edit* icon beside each replacement message to customize that message for your requirements.

Should you make a major error to the code, you can select the *Restore Default* to return to the original message and code base.

If you are viewing the replacement messages list in a VDOM, any messages that have been customized for that VDOM are displayed with a Reset icon that you can use to reset the replacement message to the global version.

For connections requiring authentication, the FortiGate unit uses HTTP to send an authentication disclaimer page for the user to accept before a security policy is in effect. Therefore, the user must initiate HTTP traffic first in order to trigger the authentication disclaimer page. Once the disclaimer is accepted, the user can send whatever traffic is allowed by the security policy.

## Replacement message images

You can add images to replacement messages to:

- disclaimer pages
- login pages
- declined disclaimer pages
- login failed page
- login challenge pages
- keepalive pages

Image embedding is also available to the endpoint NAC download portal and recommendation portal replacement messages, as well as HTTP replacement messages.

Supported image formats are GIF, JPEG, TIFF and PNG. The maximum file size supported is 6000 bytes.

## Adding images to replacement messages

### To upload an image for use in a message

1. Go to *System > Config > Replacement Messages*.
2. Select *Manage Images* at the top of the page.
3. Select *Create New*.
4. Enter a *Name* for the image.
5. Select the *Content Type*.
6. Select *Browse* to locate the file and select *OK*.

The image that you include in a replacement message, must have the following html:

```
<img src=%%IMAGE: <config_image_name>%% size=<bytes> >
```

For example:

```

```

## Modifying replacement messages

Replacement messages can be modified to include a message or content that suits your organization.

Use the expand arrows to view the replacement message list for a given category. Messages are in HTML format. For descriptions of the replacement message tags, see [Replacement message tags](#).

To change a replacement message, go to *System > Config > Replacement Messages* select the replacement message that you want to modify. At the bottom pane of the window, you can the message on one side and the HTML code on the other side. The message view changes in real-time as you change the content.

A list of common replacement messages appears in the main window. To see the entire list and all categories of replacement messages, in the upper-right corner of the window, select *Extended View*.

## Replacement message tags

Replacement messages can include replacement message tags, or variables. When users receive the message, the message tag is replaced with content relevant to the message. The table lists the replacement message tags that you can use.

**Table 70:** Replacement message tags

Tag	Description
%%AUTH_LOGOUT%%	The URL that will immediately delete the current policy and close the session. Used on the auth-keepalive page.
%%AUTH_REDIR_URL%%	The auth-keepalive page can prompt the user to open a new window which links to this tag.
%%CATEGORY%%	The name of the content category of the web site.
%%DEST_IP%%	The IP address of the request destination from which a virus was received. For email this is the IP address of the email server that sent the email containing the virus. This tag only works with alert email replacement messages.
%%DURATION%% (FortiOS Carrier only)	The amount of time in the reporting period. This is user defined in the protection profile.
%%EMAIL_FROM%%	The email address of the sender of the message from which the file was removed.
%%EMAIL_TO%%	The email address of the intended receiver of the message from which the file was removed.
%%FAILED_MESSAGE%%	The failed to login message displayed on the auth-login-failed page.



**Table 70:** Replacement message tags (continued)

Tag	Description
%%FILE%%	The name of a file that has been removed from a content stream. This could be a file that contained a virus or was blocked by antivirus file blocking. %%FILE%% can be used in virus and file block messages.
%%FORTIGUARD_WF%%	The FortiGuard - Web Filtering logo.
%%FORTINET%%	The Fortinet logo.
%%LINK%%	The link to the FortiClient Host Security installs download for the Endpoint Control feature.
%%HTTP_ERR_CODE%%	The HTTP error code. "404" for example.
%%HTTP_ERR_DESC%%	The HTTP error description.
%%KEEPALIVEURL%% (FortiOS Carrier only)	auth-keepalive-page automatically connects to this URL every %%TIMEOUT%% seconds to renew the connection policy.
%%MMS_SENDER%% (FortiOS Carrier only)	Senders MSISDN from message header.
%%MMS_RECIPIENT%% (FortiOS Carrier only)	Recipients MSISDN from message header.
%%MMS_SUBJECT%% (FortiOS Carrier only)	MMS Subject line to help with message identity.
%%MMS_HASH_CHECKSUM%%	Value derived from hash calculation - will only be shown on duplicate message alerts.
%%MMS_THRESH%%	Mass MMS alert threshold that triggered this alert.
%%NIDSEVENT%%	The IPS attack message. %%NIDSEVENT%% is added to alert email intrusion messages.
%%NUM_MSG%% (FortiOS Carrier only)	The number of time the device tried to send the message with banned content within the reporting period.
%%OVERRIDE%%	The link to the FortiGuard Web Filtering override form. This is visible only if the user belongs to a group that is permitted to create FortiGuard web filtering overrides.
%%OVRD_FORM%%	The FortiGuard web filter block override form. This tag must be present in the FortiGuard Web Filtering override form and should not be used in other replacement messages.

**Table 70:** Replacement message tags (continued)

Tag	Description
%%PROTOCOL%%	The protocol (http, ftp, pop3, imap, or smtp) in which a virus was detected. %%PROTOCOL%% is added to alert email virus messages.
%%QUARFILENAME%%	The name of a file that has been removed from a content stream and added to the quarantine. This could be a file that contained a virus or was blocked by antivirus file blocking. %%QUARFILENAME%% can be used in virus and file block messages. Quarantining is only available on FortiGate units with a local disk.
%%QUOTA_INFO%%	Display information about the traffic shaping quota setting that is blocking the user. Used in traffic quota control replacement messages.
%%QUESTION%%	Authentication challenge question on auth-challenge page. Prompt to enter username and password on auth-login page.
%%SERVICE%%	The name of the web filtering service.
%%SOURCE_IP%%	The IP address of the request originator who would have received the blocked file. For email this is the IP address of the user's computer that attempted to download the message from which the file was removed. This tag only works with alert email replacement messages.
%%TIMEOUT%%	Configured number of seconds between authentication keepalive connections. Used on the auth-keepalive page.
%%URL%%	The URL of a web page. This can be a web page that is blocked by web filter content or URL blocking. %%URL%% can also be used in http virus and file block messages to be the URL of the web page from which a user attempted to download a file that is blocked.
%%VIRUS%%	The name of a virus that was found in a file by the antivirus system. %%VIRUS%% can be used in virus messages

## Administration replacement message

If you enter the following CLI command the FortiGate unit displays the *Administration Login Disclaimer* whenever an administrator logs into the FortiGate unit's web-based manager or CLI.

```
config system global
 set access-banner enable
end
```

The web-based manager administrator login disclaimer contains the text of the Login Disclaimer replacement message as well as Accept and Decline buttons. The administrator must select accept to login.

## Alert Mail replacement messages

The FortiGate unit adds the alert mail replacement messages listed in the following table to alert email messages sent to administrators. If you enable the option *Send alert email for logs based on severity*, whether or not replacement messages are sent by alert email depends on how you set the alert email in *Minimum log level*.

## Authentication replacement messages

The FortiGate unit uses the text of the authentication replacement messages for various user authentication HTML pages that are displayed when a user is required to authenticate because a security policy includes at least one identity-based policy that requires firewall users to authenticate.

These replacement message pages are for authentication using HTTP and HTTPS. You cannot customize the firewall authentication messages for FTP and Telnet.

The authentication login page and the authentication disclaimer include replacement tags and controls not found on other replacement messages.

Users see the authentication login page when they use a VPN or a security policy that requires authentication. You can customize this page in the same way as you modify other replacement messages.

There are some unique requirements for these replacement messages:

- The login page must be an HTML page containing a form with ACTION="/" and METHOD="POST"
- The form must contain the following hidden controls:
  - `<INPUT TYPE="hidden" NAME="%%MAGICID%%" VALUE="%%MAGICVAL%%">`
  - `<INPUT TYPE="hidden" NAME="%%STATEID%%" VALUE="%%STATEVAL%%">`
  - `<INPUT TYPE="hidden" NAME="%%REDIRID%%" VALUE="%%PROTURI%%">`
- The form must contain the following visible controls:
  - `<INPUT TYPE="text" NAME="%%USERNAMEID%%" size=25>`
  - `<INPUT TYPE="password" NAME="%%PASSWORDID%%" size=25>`

### Example

The following is an example of a simple authentication page that meets the requirements listed above.

```
<HTML><HEAD><TITLE>Firewall Authentication</TITLE></HEAD>
 <BODY><H4>You must authenticate to use this service.</H4>
<FORM ACTION="/" method="post">
 <INPUT NAME="%%MAGICID%%" VALUE="%%MAGICVAL%%" TYPE="hidden">
<TABLE ALIGN="center" BGCOLOR="#00cccc" BORDER="0"
 CELLPADDING="15" CELLSPACING="0" WIDTH="320"><TBODY>
<TR><TH>Username:</TH>
 <TD><INPUT NAME="%%USERNAMEID%%" SIZE="25" TYPE="text"> </TD></TR>
<TR><TH>Password:</TH>
 <TD><INPUT NAME="%%PASSWORDID%%" SIZE="25" TYPE="password">
 </TD></TR>
<TR><TD COLSPAN="2" ALIGN="center" BGCOLOR="#00cccc">
 <INPUT NAME="%%STATEID%%" VALUE="%%STATEVAL%%" TYPE="hidden">
```

```
<INPUT NAME="%%REDIRID%%" VALUE="%%PROTURI%%" TYPE="hidden">
<INPUT VALUE="Continue" TYPE="submit"> </TD></TR>
</TBODY></TABLE></FORM></BODY></HTML>
```

## Captive Portal Default replacement messages

The Captive Portal Default replacement messages are used for wireless authentication only. You must have a VAP interface with the security set as captive portal to trigger these replacement messages.

## Device Detection Portal replacement message

The FortiGate unit displays the replacement message when the FortiGate unit cannot determine the type of BYOD or handheld device is used to connect the network.

## Email replacement messages

The FortiGate unit sends the mail replacement messages to email clients using IMAP, POP3, or SMTP when an event occurs such as antivirus blocking a file attached to an email that contains a virus. Email replacement messages are text messages.

If the FortiGate unit supports SSL content scanning and inspection these replacement messages can also be added to IMAPS, POP3S, and SMTPS email messages.

## Endpoint Control replacement message

The FortiGate unit displays the replacement message when the FortiClient Endpoint Security software is not installed or registered correctly with the FortiGate unit.

## FTP replacement messages

The FortiGate unit sends the FTP replacement messages listed in the table below to FTP clients when an event occurs such as antivirus blocking a file that contains a virus in an FTP session. FTP replacement messages are text messages.

## FortiGuard Web Filtering replacement messages

The FortiGate unit sends the FortiGuard Web Filtering replacement messages listed in the table to web browsers using the HTTP protocol when FortiGuard web filtering blocks a URL, provides details about blocked HTTP 4xx and 5xx errors, and for FortiGuard overrides. FortiGuard Web Filtering replacement messages are HTTP pages.

If the FortiGate unit supports SSL content scanning and inspection and if *Protocol Recognition > HTTPS Content Filtering Mode* is set to Deep Scan in the antivirus profile, these replacement messages can also replace web pages downloaded using the HTTPS protocol.

## HTTP replacement messages

The FortiGate unit sends the HTTP replacement messages listed in the following table to web browsers using the HTTP protocol when an event occurs such as antivirus blocking a file that contains a virus in an HTTP session. HTTP replacement messages are HTML pages.

If the FortiGate unit supports SSL content scanning and inspection, and if under HTTPS in the protocol option list has Enable Deep Scan enabled, these replacement messages can also replace web pages downloaded using the HTTPS protocol.

## IM replacement messages

The FortiGate unit sends the IM replacement messages listed in to IM clients using AIM, ICQ, MSN, or Yahoo! Messenger when an event occurs such as antivirus blocking a file attached to an email that contains a virus. IM replacement messages are text messages.

## NNTP replacement messages

The FortiGate unit sends the NNTP replacement messages listed in the following table to NNTP clients when an event occurs such as antivirus blocking a file attached to an NNTP message that contains a virus. NNTP replacement messages are text messages.

## Spam replacement messages

The FortiGate unit adds the Spam replacement messages listed in the following table to SMTP server responses if the email message is identified as spam and the spam action is discard. If the FortiGate unit supports SSL content scanning and inspection these replacement messages can also be added to SMTPS server responses.

## NAC quarantine replacement messages

The page that is displayed for the user depends on whether NAC quarantine blocked the user because a virus was found, a DoS sensor detected an attack, an IPS sensor detected an attack, or a DLP rule with action set to *Quarantine IP address* or *Quarantine Interface* matched a session from the user.

The default messages inform the user of why they are seeing this page and recommend they contact the system administrator. You can customize the pages as required, for example to include an email address or other contact information or if applicable a note about how long the user can expect to be blocked.

## SSL VPN replacement message

The SSL VPN login replacement message is an HTML replacement message that formats the FortiGate SSL VPN portal login page. You can customize this replacement message according to your organization's needs. The page is linked to FortiGate functionality and you must construct it according to the following guidelines to ensure that it will work.

- The login page must be an HTML page containing a form with `ACTION="%%SSL_ACT%%"` and `METHOD="%%SSL_METHOD%%"`
- The form must contain the `%%SSL_LOGIN%%` tag to provide the login form.
- The form must contain the `%%SSL_HIDDEN%%` tag.

## Web Proxy replacement messages

The FortiGate unit sends Web Proxy replacement messages listed in the table below when a web proxy event occurs that is detected and matches the web proxy configuration. These replacement messages are web pages that appear within your web browser.

The following web proxy replacement messages require an identity-based security policy so that the web proxy is successful. You can also enable FTP-over-HTTP by selecting the *FTP* option in *System > Network > Explicit Proxy*.

## Traffic quota control replacement messages

When user traffic is going through the FortiGate unit and it is blocked by traffic shaping quota controls, users see the *Traffic shaper block message* or the *Per IP traffic shaper block message* when they attempt to connect through the FortiGate unit using HTTP.

The traffic quota HTTP pages should contain the `%%QUOTA_INFO%%` tag to display information about the traffic shaping quota setting that is blocking the user.

## MM1 replacement messages

MM1 replacement messages are sent when, during MMS content scanning, FortiOS Carrier detects, for example a virus, using the MMS profile.

You must have *Remove Blocked* selected within the MMS profile if you want to remove the content that is intercepted during MMS scanning on the FortiGate unit.

## MM3 replacement messages

MM3 replacement messages are sent when, during MMS content scanning, FortiOS Carrier detects, for example a virus, using the MMS profile.

You must have *Remove Blocked* selected within the MMS profile if you want to remove the content that is intercepted during MMS scanning on the unit.

## MM4 replacement messages

MM4 replacement messages are sent when, during MMS content scanning, FortiOS Carrier detects, for example a virus, using the MMS profile.

## MM7 replacement messages

MM7 replacement messages are sent when, during MMS content scanning, FortiOS Carrier detects, for example a virus, using the MMS profile.

## MMS replacement messages

The MMS replacement message is sent when a section of an MMS message has been replaced because it contains a blocked file. This replacement message is in HTML format.

The message text is:

```
<HTML><BODY>This section of the message has been replaced because it
contained a blocked file</BODY></HTML>
```

## Replacement message groups

Replacement message groups enable you to view common messages in groups for large carriers. To view grouped replacement messages, go to *System > Admin > Settings* and select *Replacement Message Groups* in the *Display Options on GUI* section.

Message groups can be configured by going to *Config > Replacement Message Group*.

Using the defined groups, you can manage specific replacement messages from a single location, rather than searching through the entire replacement message list.

If you enable virtual domains (VDOMs) on the FortiGate unit, replacement message groups are configured separately for each virtual domain. Each virtual domain has its own default

replacement message group, configured from *System > Config > Replacement Messages Group*.

When you modify a message in a replacement message group, a Reset icon appears beside the message in the group. You can select this Reset icon to reset the message in the replacement message group to the default version.

All MM1/4/7 notification messages for FortiOS Carrier (and MM1 retrieve-conf messages) can contain a SMIL layer and all MM4 notification messages can contain an HTML layer in the message. These layers can be used to brand messages by using logos uploaded to the FortiGate unit via the 'Manage Images' link found on the replacement message group configuration page.

## Disk

To view the status and storage information of the local disk on your FortiGate unit, go to *System > Config > Advanced*. The *Disk* menu appears only on FortiGate units with an internal hard or flash disk.

### Formatting the disk

The internal disk of the FortiGate unit (if available) can be formatted by going to *System > Config > Disk* and selecting *Format*.

Formatting the disk will erase all data on it, including databases for antivirus and IPS; logs, quarantine files, and WAN optimization caches. The FortiGate unit requires a reboot once the disk has been formatted.

### Setting space quotas

If the FortiGate unit has an internal hard or flash disk, you can allocate the space on the disk for specific logging and archiving, and WAN optimization. By default, the space is used on an as required basis. As such, a disk can fill up with basic disk logging, leaving less potential space for quarantine.

By going to *System > Config > Disk*, you can select the *Edit* icon for *Logging and Archiving* and *WAN Optimization & Web Cache* and define the amount of space each log, archive and WAN optimization has on the disk.

## CLI Scripts

To upload bulk CLI commands and scripts, go to *System > Config > Advanced*.

Scripts are text files containing CLI command sequences. Scripts can be used to deploy identical configurations to many devices. For example, if all of your devices use identical security policies, you can enter the commands required to create the security policies in a script, and then deploy the script to all the devices which should use those same settings.

Use a text editor such as Notepad or other application that creates simple text files. Enter the commands in sequence, with each line as one command, similar to examples throughout the FortiOS documentation set.

If you are using a FortiGate unit that is not remotely managed by a FortiManager unit or the FortiGuard Analysis and Management Service, the scripts you upload are executed and discarded. If you want to execute a script more than once, you must keep a copy on your management PC.

If your FortiGate unit is configured to use a FortiManager unit, you can upload your scripts to the FortiManager unit, and run them from any FortiGate unit configured to use the FortiManager unit. If you upload a script directly to a FortiGate unit, it is executed and discarded.

If your FortiGate unit is configured to use FortiGuard Analysis and Management Service, scripts you upload are executed and stored. You can run uploaded scripts from any FortiGate unit configured with your FortiGuard Analysis and Management Service account. The uploaded script files appear on the FortiGuard Analysis and Management Service portal web site.

## Uploading script files

After you have created a script file, you can then upload it through *System > Config > Advanced*. When a script is uploaded, it is automatically executed.

Commands that require the FortiGate unit to reboot when entered in the command line will also force a reboot if included in a script.

### To execute a script

1. Go to *System > Config > Advanced*.
2. Verify that *Upload Bulk CLI Command File* is selected.
3. Select *Browse* to locate the script file.
4. Select *Apply*.

If the FortiGate unit is not configured for remote management, or if it is configured to use a FortiManager unit, uploaded scripts are discarded after execution. Save script files to your management PC if you want to execute them again later.

If the FortiGate unit is configured to use the FortiGuard Analysis and Management Service, the script file is saved to the remote server for later reuse. You can view the script or run it from the FortiGuard Analysis and Management Service portal web site.

## Rejecting PING requests

The factory default configuration of your FortiGate unit allows the default external interface to respond to ping requests. Depending on the model of your FortiGate unit the actual name of this interface will vary. For the most secure operation, you should change the configuration of the external interface so that it does not respond to ping requests. Not responding to ping requests makes it more difficult for a potential attacker to detect your FortiGate unit from the Internet. One such potential threat are Denial of Service (DoS) attacks.

A FortiGate unit responds to ping requests if ping administrative access is enabled for that interface.

### To disable ping administrative access - web-based manager

1. Go to *System > Network > Interface*.
2. Choose the external interface and select *Edit*.
3. Clear the *Ping Administrative Access* check box.
4. Select *OK*.

In the CLI, when setting the *allowaccess* settings, by selecting the access types and not including the PING option, that option is then not selected. In this example, only HTTPS is selected.



### To disable ping administrative access - CLI

```
config system interface
 edit external
 set allowaccess https
 end
```

## Opening TCP 113

Although seemingly contrary to conventional wisdom of closing ports from hackers, this port, which is used for ident requests, should be opened.

Port 113 initially was used as an authentication port, and later defined as an identification port (see RFC 1413). Some servers may still use this port to help in identifying users or other servers and establish a connection. Because port 113 receives a lot of unsolicited traffic, many routers, including on the FortiGate unit, close this port.

The issue arises in that unsolicited requests are stopped by the FortiGate unit, which will send a response saying that the port is closed. In doing so, it also lets the requesting server know there is a device at the given address, and thus announcing its presence. By enabling traffic on port 113, requests will travel to this port, and will most likely, be ignored and never responded to.

By default, the ident port is closed. To open it, use the following CLI commands:

```
config system interface
 edit <port_name>
 set ident_accept enable
 end
```

You could also further use port forwarding to send the traffic to a non-existent IP address and thus never have a response packet sent.

## Obfuscate HTTP responses

The FortiGate unit can obfuscate the HTTP responses from the FortiGate admin GUI and SSL VPN servers. By default this option is not enabled. To obfuscate HTTP headers, use the following CLI command:

```
config system global
 set http-obfuscate {none | header-only | modified | no-error}
end
```

Where:

`none` — do not hide the FortiGate web server identity.

`header-only` — hides the HTTP server banner.

`modified` — provides modified error responses.

`no-error` — suppresses error responses.

# Session helpers

The FortiOS firewall can analyze most TCP/IP protocol traffic by comparing packet header information to security policies. This comparison determines whether to accept or deny the packet and the session that the packet belongs to.

Some protocols include information in the packet body (or payload) that must be analyzed to successfully process sessions for this protocol. For example, the SIP VoIP protocol uses TCP control packets with a standard destination port to set up SIP calls. But the packets that carry the actual conversation can use a variety of UDP protocols with a variety of source and destination port numbers. The information about the protocols and port numbers used for a SIP call is contained in the body of the SIP TCP control packets. To successfully process SIP VoIP calls, FortiOS must be able to extract information from the body of the SIP packet and use this information to allow the voice-carrying packets through the firewall.

FortiOS uses session helpers to analyze the data in the packet bodies of some protocols and adjust the firewall to allow those protocols to send packets through the firewall.

This section includes the topics:

- [Viewing the session helper configuration](#)
- [Changing the session helper configuration](#)
- [DCE-RPC session helper \(dcerpc\)](#)
- [DNS session helpers \(dns-tcp and dns-udp\)](#)
- [File transfer protocol \(FTP\) session helper \(ftp\)](#)
- [H.245 session helpers \(h245l and h245O\)](#)
- [H.323 and RAS session helpers \(h323 and ras\)](#)
- [Media Gateway Controller Protocol \(MGCP\) session helper \(mgcp\)](#)
- [ONC-RPC portmapper session helper \(pmap\)](#)
- [PPTP session helper for PPTP traffic \(pptp\)](#)
- [Remote shell session helper \(rsh\)](#)
- [Real-Time Streaming Protocol \(RTSP\) session helper \(rtsp\)](#)
- [Session Initiation Protocol \(SIP\) session helper \(sip\)](#)
- [Trivial File Transfer Protocol \(TFTP\) session helper \(tftp\)](#)
- [Oracle TNS listener session helper \(tns\)](#)

## Viewing the session helper configuration

You can view the session helpers enabled on your FortiGate unit in the CLI using the commands below. The following output shows the first two session helpers. The number of session helpers can vary to around 20.

```
show system session-helper
config system session-helper
edit 1
 set name pptp
 set port 1723
 set protocol 6
```

```

end
next
 set name h323
 set port 1720
 set protocol 6
next
end
.
.

```

The configuration for each session helper includes the name of the session helper and the port and protocol number on which the session helper listens for sessions. Session helpers listed on protocol number 6 (TCP) or 17 (UDP). For a complete list of protocol numbers see: [Assigned Internet Protocol Numbers](#).

For example, the output above shows that FortiOS listens for PPTP packets on TCP port 1723 and H.323 packets on port TCP port 1720.

If a session helper listens on more than one port or protocol the more than one entry for the session helper appears in the `config system session-helper` list. For example, the `pmap` session helper appears twice because it listens on TCP port 111 and UDP port 111. The `rsh` session helper appears twice because it listens on TCP ports 514 and 512.

## Changing the session helper configuration

Normally you will not need to change the configuration of the session helpers. However in some cases you may need to change the protocol or port the session helper listens on.

### Changing the protocol or port that a session helper listens on

Most session helpers are configured to listen for their sessions on the port and protocol that they typically use. If your FortiGate unit receives sessions that should be handled by a session helper on a non-standard port or protocol you can use the following procedure to change the port and protocol used by a session helper. The following example shows how to change the port that the `pmap` session helper listens on for Sun RPC portmapper TCP sessions. By default `pmap` listens on TCP port 111.

#### To change the port that the `pmap` session helper listens on to TCP port 112

1. Confirm that the TCP `pmap` session helper entry is 11 in the session-helper list:

```

show system session-helper 11
config system session-helper
 edit 11
 set name pmap
 set port 111
 set protocol 6
 next
end

```

2. Enter the following command to change the TCP port to 112.

```

config system session-helper
 edit 11
 set port 112
 end

```

3. The pmap session helper also listens on UDP port 111. Confirm that the UDP pmap session helper entry is 12 in the session-helper list:

```
show system session-helper 12
 config system session-helper
 edit 12
 set name pmap
 set port 111
 set protocol 17
 next
 end
```

4. Enter the following command to change the UDP port to 112.

```
config system session-helper
 edit 12
 set port 112
 end
end
```

Use the following command to set the h323 session helper to listen for ports on the UDP protocol.

#### **To change the protocol that the h323 session helper listens on**

1. Confirm that the h323 session helper entry is 2 in the session-helper list:

```
show system session-helper 2
 config system session-helper
 edit 2
 set name h323
 set port 1720
 set protocol 6
 next
 end
```

2. Enter the following command to change the protocol to UDP.

```
config system session-helper
 edit 2
 set protocol 17
 end
end
```

If a session helper listens on more than one port or protocol, then multiple entries for the session helper must be added to the session helper list, one for each port and protocol combination. For example, the rtsp session helper listens on TCP ports 554, 7070, and 8554 so there are three rtsp entries in the session-helper list. If your FortiGate unit receives rtsp packets on a different TCP port (for example, 6677) you can use the following command to configure the rtsp session helper to listen on TCP port 6677.

#### **To configure a session helper to listen on a new port and protocol**

```
config system session-helper
 edit 0
 set name rtsp
 set port 6677
 set protocol 6
 end
```

## Disabling a session helper

In some cases you may need to disable a session helper. Disabling a session helper just means removing it from the session-helper list so that the session helper is not listening on a port. You can completely disable a session helper by deleting all of its entries from the session helper list. If there are multiple entries for a session helper on the list you can delete one of the entries to prevent the session helper from listening on that port.

### To disable the mgcp session helper from listening on UDP port 2427

1. Enter the following command to find the mgcp session helper entry that listens on UDP port 2427:

```
show system session-helper
.
.
.
edit 19
 set name mgcp
 set port 2427
 set protocol 17
next
.
.
.
```

2. Enter the following command to delete session-helper list entry number 19 to disable the mgcp session helper from listening on UDP port 2427:

```
config system session-helper
 delete 19
```

By default the mgcp session helper listens on UDP ports 2427 and 2727. The previous procedure shows how to disable the mgcp protocol from listening on port 2427. The following procedure completely disables the mgcp session helper by also disabling it from listening on UDP port 2727.

### To completely disable the mgcp session helper

1. Enter the following command to find the mgcp session helper entry that listens on UDP port 2727:

```
show system session-helper
.
.
.
edit 20
 set name mgcp
 set port 2727
 set protocol 17
next
.
.
.
```

2. Enter the following command to delete session-helper list entry number 20 to disable the mgcp session helper from listening on UDP port 2727:

```
config system session-helper
 delete 20
```

## DCE-RPC session helper (dcerpc)

Distributed Computing Environment Remote Procedure Call (DCE-RPC) provides a way for a program running on one host to call procedures in a program running on another host. DCE-RPC (also called MS RPC for Microsoft RPC) is similar to ONC-RPC. Because of the large number of RPC services, for example, MAPI, the transport address of an RPC service is dynamically negotiated based on the service program's universal unique identifier (UUID). The Endpoint Mapper (EPM) binding protocol in FortiOS maps the specific UUID to a transport address.

To accept DCE-RPC sessions you must add a security policy with service set to any or to the DCE-RPC pre-defined service (which listens on TCP and UDP ports 135). The dcerpc session helper also listens on TCP and UDP ports 135.

The session allows FortiOS to handle DCE-RPC dynamic transport address negotiation and to ensure UUID-based security policy enforcement. You can define a security policy to permit all RPC requests or to permit by specific UUID number.

In addition, because a TCP segment in a DCE-RPC stream might be fragmented, it might not include an intact RPC PDU. This fragmentation occurs in the RPC layer; so FortiOS does not support parsing fragmented packets.

## DNS session helpers (dns-tcp and dns-udp)

FortiOS includes two DNS session helpers, dns-tcp, a session helper for DNS over TCP, and dns-udp, a session helper for DNS over UDP.

To accept DNS sessions you must add a security policy with service set to any or to the DNS pre-defined service (which listens on TCP and UDP ports 53). The dns-udp session helper also listens on UDP port 53. By default the dns-tcp session helper is disabled. If needed you can use the following command to enable the dns-tcp session helper to listen for DNS sessions on TCP port 53:

```
config system session-helper
 edit 0
 set name dns-tcp
 set port 53
 set protocol 6
 end
```

## File transfer protocol (FTP) session helper (ftp)

The FTP session helper monitors PORT, PASV and 227 commands and NATs the IP addresses and port numbers in the body of the FTP packets and opens ports on the FortiGate unit as required.

To accept FTP sessions you must add a security policy with service set to any or to the FTP, FTP\_Put, and FTP\_GET pre-defined services (which all listen on TCP port 21).

## H.245 session helpers (h245I and h245O)

H.245 is a control channel protocol used for H.323 and other similar communication sessions. H.245 sessions transmit non-telephone signals. H.245 sessions carry information needed for multimedia communication, such as encryption, flow control jitter management and others.

FortiOS includes two H.245 sessions helpers, h245I which is for H.245 call in and h245O which is for H.245 call out sessions. There is no standard port for H.245. By default the H.245 sessions helpers are disabled. You can enable them as you would any other session helper. When you enable them, you should specify the port and protocol on which the FortiGate unit receives H.245 sessions.

## H.323 and RAS session helpers (h323 and ras)

The H.323 session helper supports secure H.323 voice over IP (VoIP) sessions between terminal endpoints such as IP phones and multimedia devices. In H.323 VoIP networks, gatekeeper devices manage call registration, admission, and call status for VoIP calls. The FortiOS h323 session helper supports gatekeepers installed on two different networks or on the same network.

To accept H.323 sessions you must add a security policy with service set to any or to the H323 pre-defined service (which listens on TCP port numbers 1720 and 1503 and on UDP port number 1719). The h323 session helper listens on TCP port 1720.

The ras session helper is used with the h323 session helper for H.323 Registration, Admission, and Status (RAS) services. The ras session helper listens on UDP port 1719.

### Alternate H.323 gatekeepers

The h323 session helper supports using H.323 alternate gatekeepers. All the H.323 end points must register with a gatekeeper through the Registration, Admission, and Status (RAS) protocol before they make calls. During the registration process, the primary gatekeeper sends Gatekeeper Confirm (GCF) and Registration Confirm (RCF) messages to the H.323 end points that contain the list of available alternate gatekeepers.

The alternate gatekeeper provides redundancy and scalability for the H.323 end points. If the primary gatekeeper fails the H.323 end points that have registered with that gatekeeper are automatically registered with the alternate gatekeeper. To use the H.323 alternate gatekeeper, you need to configure security policies that allow H.323 end points to reach the alternate gatekeeper.

## Media Gateway Controller Protocol (MGCP) session helper (mgcp)

The Media Gateway Control Protocol (MGCP) is a text-based application layer protocol used for VoIP call setup and control. MGCP uses a master-slave call control architecture in which the media gateway controller uses a call agent to maintain call control intelligence, while the media gateways perform the instructions of the call agent.

To accept MGCP sessions you must add a security policy with service set to any or to the MGCP pre-defined service (which listens on UDP port numbers 2427 and 2727). The h323 session helper also listens on UDP port numbers 2427 and 2727.

The MGCP session helper does the following:

- VoIP signalling payload inspection. The payload of the incoming VoIP signalling packet is inspected and malformed packets are blocked.
- Signaling packet body inspection. The payload of the incoming MGCP signaling packet is inspected according to RFC 3435. Malformed packets are blocked.
- Stateful processing of MGCP sessions. State machines are invoked to process the parsed information. Any out-of-state or out-of-transaction packet is identified and properly handled.
- MGCP Network Address Translation (NAT). Embedded IP addresses and ports in packet bodies is properly translated based on current routing information and network topology, and is replaced with the translated IP address and port number, if necessary.
- Manages pinholes for VoIP traffic. To keep the VoIP network secure, the IP address and port information used for media or signalling is identified by the session helper, and pinholes are dynamically created and closed during call setup.

## ONC-RPC portmapper session helper (pmap)

Open Network Computing Remote Procedure Call (ONC-RPC) is a widely deployed remote procedure call system. Also called Sun RPC, ONC-RPC allows a program running on one host to call a program running on another. The transport address of an ONC-RPC service is dynamically negotiated based on the service's program number and version number. Several binding protocols are defined for mapping the RPC program number and version number to a transport address.

To accept ONC-RPC sessions you must add a security policy with service set to any or to the ONC-RPC pre-defined service (which listens on TCP and UDP port number 111). The RPC portmapper session helper (called pmap) handles the dynamic transport address negotiation mechanisms of ONC-RPC.

## PPTP session helper for PPTP traffic (pptp)

The PPTP session help supports port address translation (PAT) for PPTP traffic. PPTP provides IP security at the Network Layer. PPTP consists of a control session and a data tunnel. The control session runs over TCP and helps in establishing and disconnecting the data tunnel. The data tunnel handles encapsulated Point-to-Point Protocol (PPP) packets carried over IP.

To accept PPTP sessions that pass through the FortiGate unit you must add a security policy with service set to any or to the PPTP pre-defined service (which listens on IP port 47 and TCP port 1723). The pptp session helper listens on TCP port 1723.

PPTP uses TCP port 1723 for control sessions and Generic Routing Encapsulation (GRE) (IP protocol 47) for tunneling the encapsulated PPP data. The GRE traffic carries no port number, making it difficult to distinguish between two clients with the same public IP address. PPTP uses the source IP address and the Call ID field in the GRE header to identify a tunnel. When multiple clients sharing the same IP address establish tunnels with the same PPTP server, they may get the same Call ID. The call ID value can be translated in both the control message and the data traffic, but only when the client is in a private network and the server is in a public network.

PPTP clients can either directly connect to the Internet or dial into a network access server to reach the Internet. A FortiGate unit that protects PPTP clients can translate the clients' private IP addresses to a pool of public IP addresses using NAT port translation (NAT-PT). Because the GRE traffic carries no port number for address translation, the pptp session helper treats the Call ID field as a port number as a way of distinguishing multiple clients.



After the PPTP establishing a TCP connection with the PPTP server, the client sends a start control connection request message to establish a control connection. The server replies with a start control connection reply message. The client then sends a request to establish a call and sends an outgoing call request message. FortiOS assigns a Call ID (bytes 12-13 of the control message) that is unique to each PPTP tunnel. The server replies with an outgoing call reply message that carries its own Call ID in bytes 12-13 and the client's call ID in bytes 14-15. The pptp session helper parses the control connection messages for the Call ID to identify the call to which a specific PPP packet belongs. The session helper also identifies an outgoing call request message using the control message type field (bytes 8-9) with the value 7. When the session helper receives this message, it parses the control message for the call ID field (bytes 12-13). FortiOS translates the call ID so that it is unique across multiple calls from the same translated client IP. After receiving outgoing call response message, the session helper holds this message and opens a port that accepts GRE traffic that the PPTP server sends. An outgoing call request message contains the following parts:

- The protocol used for the outgoing call request message (usually GRE)
- Source IP address (PPTP server IP)
- Destination IP address (translated client IP)
- Destination port number (translated client call ID)

The session helper identifies an outgoing call reply message using the control message type field (bytes 8-9) with the value 8. The session helper parses these control messages for the call ID field (bytes 12-13) and the client's call ID (bytes 14-15). The session helper then uses the client's call ID value to find the mapping created for the other direction, and then opens a pinhole to accept the GRE traffic that the client sends.

An outgoing call reply message contains the following parts:

- Protocol used for the outgoing call reply message (usually GRE)
- Source IP address (PPTP client IP)
- Destination IP address (PPTP server IP)
- Destination port number (PPTP server Call ID)

Each port that the session opens creates a session for data traffic arriving in that direction. The session helper opens the following two data sessions for each tunnel:

- Traffic from the PPTP client to the server, using the server's call ID as the destination port
- Traffic from the PPTP server to the client, using the client's translated call ID as the destination port

The default timeout value of the control connection is 30 minutes. The session helper closes the pinhole when the data session exceeds the timeout value or is idle for an extended period.

## Remote shell session helper (rsh)

Using the remote shell program (RSH), authenticated users can run shell commands on remote hosts. RSH sessions most often use TCP port 514. To accept RSH sessions you must add a security policy with service set to any or to the RSH pre-defined service (which listens on TCP port number 514).

FortiOS automatically invokes the rsh session helper to process all RSH sessions on TCP port 514. The rsh session helper opens ports required for the RSH service to operate through a FortiGate unit running NAT or transparent and supports port translation of RSH traffic.

## Real-Time Streaming Protocol (RTSP) session helper (rtsp)

The Real-Time Streaming Protocol (RTSP) is an application layer protocol often used by SIP to control the delivery of multiple synchronized multimedia streams, for example, related audio and video streams. Although RTSP is capable of delivering the data streams itself it is usually used like a network remote control for multimedia servers. The protocol is intended for selecting delivery channels (like UDP, multicast UDP, and TCP) and for selecting a delivery mechanism based on the Real-Time Protocol (RTP). RTSP may also use the SIP Session Description Protocol (SDP) as a means of providing information to clients for aggregate control of a presentation consisting of streams from one or more servers, and non-aggregate control of a presentation consisting of multiple streams from a single server.

To accept RTSP sessions you must add a security policy with service set to any or to the RTSP pre-defined service (which listens on TCP ports 554, 770, and 8554 and on UDP port 554). The rtsp session helper listens on TCP ports 554, 770, and 8554.

The rtsp session help is required because RTSP uses dynamically assigned port numbers that are communicated in the packet body when end points establish a control connection. The session helper keeps track of the port numbers and opens pinholes as required. In Network Address Translation (NAT) mode, the session helper translates IP addresses and port numbers as necessary.

In a typical RTSP session the client starts the session (for example, when the user selects the Play button on a media player application) and establishes a TCP connection to the RTSP server on port 554. The client then sends an OPTIONS message to find out what audio and video features the server supports. The server responds to the OPTIONS message by specifying the name and version of the server, and a session identifier, for example, 24256-1.

The client then sends the DESCRIBE message with the URL of the actual media file the client wants to play. The server responds to the DESCRIBE message with a description of the media in the form of SDP code. The client then sends the SETUP message, which specifies the transport mechanisms acceptable to the client for streamed media, for example RTP/RTCP or RDT, and the ports on which it receives the media.

In a NAT configuration the rtsp session helper keeps track of these ports and addresses translates them as necessary. The server responds to the SETUP message and selects one of the transport protocols. When both client and server agree on a mechanism for media transport the client sends the PLAY message, and the server begins streaming the media.

## Session Initiation Protocol (SIP) session helper (sip)

The sip session helper is described in [“The SIP session helper”](#) on page 2505.

## Trivial File Transfer Protocol (TFTP) session helper (tftp)

To accept TFTP sessions you must add a security policy with service set to any or to the TFTP pre-defined service (which listens on UDP port number 69). The TFTP session helper also listens on UTP port number 69.

TFTP initiates transfers on UDP port 69, but the actual data transfer ports are selected by the server and client during initialization of the connection. The tftp session helper reads the transfer ports selected by the TFTP client and server during negotiation and opens these ports on the firewall so that the TFTP data transfer can be completed. When the transfer is complete the tftp session helper closes the open ports.

## Oracle TNS listener session helper (tns)

The Oracle Transparent Network Substrate (TNS) listener listens on port TCP port 1521 for network requests to be passed to a database instance. The Oracle TNS listener session helper (tns) listens for TNS sessions on TCP port 1521. TNS is a foundation technology built into the Oracle Net foundation layer and used by SQLNET.

# Chapter 11 IPsec VPN for FortiOS 5.0

This FortiOS Handbook chapter contains the following sections:

[IPsec VPN concepts](#) explains the basic concepts that you need to understand about virtual private networks (VPNs).

[IPsec VPN Overview](#) provides a brief overview of IPsec technology and includes general information about how to configure IPsec VPNs using this guide.

[IPsec VPN in the web-based manager](#) describes the IPsec VPN menu of the web-based manager interface.

[Gateway-to-gateway configurations](#) explains how to set up a basic gateway-to-gateway (site-to-site) IPsec VPN. In a gateway-to-gateway configuration, two FortiGate units create a VPN tunnel between two separate private networks.

[Hub-and-spoke configurations](#) describes how to set up hub-and-spoke IPsec VPNs. In a hub-and-spoke configuration, connections to a number of remote peers and/or clients radiate from a single, central FortiGate hub.

[Dynamic DNS configuration](#) describes how to configure a site-to-site VPN, in which one FortiGate unit has a static IP address and the other FortiGate unit has a dynamic IP address and a domain name.

[FortiClient dialup-client configurations](#) guides you through configuring a FortiClient dialup-client IPsec VPN. In a FortiClient dialup-client configuration, the FortiGate unit acts as a dialup server and VPN client functionality is provided by the FortiClient Endpoint Security application installed on a remote host.

[FortiGate dialup-client configurations](#) explains how to set up a FortiGate dialup-client IPsec VPN. In a FortiGate dialup-client configuration, a FortiGate unit with a static IP address acts as a dialup server and a FortiGate unit with a dynamic IP address initiates a VPN tunnel with the FortiGate dialup server.

[Supporting IKE Mode config clients](#) explains how to set up a FortiGate unit as either an IKE Mode Config server or client. IKE Mode Config is an alternative to DHCP over IPsec.

[Internet-browsing configuration](#) explains how to support secure web browsing performed by dialup VPN clients, and hosts behind a remote VPN peer. Remote users can access the private network behind the local FortiGate unit and browse the Internet securely. All traffic generated remotely is subject to the security policy that controls traffic on the private network behind the local FortiGate unit.

[Redundant VPN configurations](#) discusses the options for supporting redundant and partially redundant tunnels in an IPsec VPN configuration. A FortiGate unit can be configured to support redundant tunnels to the same remote peer if the FortiGate unit has more than one interface to the Internet.

[Transparent mode VPNs](#) describes two FortiGate units that create a VPN tunnel between two separate private networks transparently. In transparent mode, all FortiGate unit interfaces except the management interface are invisible at the network layer.

[IPv6 IPsec VPNs](#) describes FortiGate unit VPN capabilities for networks based on IPv6 addressing. This includes IPv4-over-IPv6 and IPv6-over-IPv4 tunnelling configurations. IPv6 IPsec VPNs are available in FortiOS 3.0 MR5 and later.

[L2TP and IPsec \(Microsoft VPN\)](#) explains how to support Microsoft Windows native VPN clients.

[GRE over IPsec \(Cisco VPN\)](#) explains how to interoperate with Cisco VPNs that use Generic Routing Encapsulation (GRE) protocol with IPsec.

[Protecting OSPF with IPsec](#) provides an example of protecting OSPF links with IPsec.

[Auto Key phase 1 parameters](#) provides detailed step-by-step procedures for configuring a FortiGate unit to accept a connection from a remote peer or dialup client. The basic phase 1 parameters identify the remote peer or clients and support authentication through preshared keys or digital certificates. You can increase VPN connection security further using methods such as extended authentication (XAuth).

[Phase 2 parameters](#) provides detailed step-by-step procedures for configuring an IPsec VPN tunnel. During phase 2, the specific IPsec security associations needed to implement security services are selected and a tunnel is established.

[Defining VPN security policies](#) explains how to specify the source and destination IP addresses of traffic transmitted through an IPsec VPN tunnel, and how to define a security encryption policy. Security policies control all IP traffic passing between a source address and a destination address.

[Hardware offloading and acceleration](#) explains how to make use of FortiASIC network processor IPsec accelerated processing capabilities.

[Monitoring and troubleshooting](#) provides VPN monitoring and testing procedures

# IPsec VPN concepts

Virtual Private Network (VPN) technology enables remote users to connect to private computer networks to gain access to their resources in a secure way. For example, an employee traveling or working from home can use a VPN to securely access the office network through the Internet.

Instead of remotely logging on to a private network using an unencrypted and unsecure Internet connection, the use of a VPN ensures that unauthorized parties cannot access the office network and cannot intercept any of the information that is exchanged between the employee and the office. It is also common to use a VPN to connect the private networks of two or more offices.

Fortinet offers VPN capabilities in the FortiGate Unified Threat Management (UTM) appliance and in the FortiClient Endpoint Security suite of applications. A FortiGate unit can be installed on a private network, and FortiClient software can be installed on the user's computer. It is also possible to use a FortiGate unit to connect to the private network instead of using FortiClient software.

This chapter discusses VPN terms and concepts including:


- [VPN tunnels](#)
- [VPN gateways](#)
- [Clients, servers, and peers](#)
- [Encryption](#)
- [Authentication](#)
- [Phase 1 and Phase 2 settings](#)
- [Security Association](#)

## VPN tunnels

The data path between a user's computer and a private network through a VPN is referred to as a tunnel. Like a physical tunnel, the data path is accessible only at both ends. In the telecommuting scenario, the tunnel runs between the FortiClient application on the user's PC, or a FortiGate unit or other network device and the FortiGate unit on the office private network.

Encapsulation makes this possible. IPsec packets pass from one end of the tunnel to the other and contain data packets that are exchanged between the local user and the remote private network. Encryption of the data packets ensures that any third-party who intercepts the IPsec packets can not access the data.

**Figure 250:** Encoded data going through a VPN tunnel



You can create a VPN tunnel between:

- a PC equipped with the FortiClient application and a FortiGate unit
- two FortiGate units
- third-party VPN software and a FortiGate unit

Third-party VPN software is not covered in this document. Refer to the [Fortinet Knowledge Base](#) for more information on this topic.

## VPN gateways


A gateway is a router that connects the local network to other networks. The default gateway setting in your computer's TCP/IP properties specifies the gateway for your local network.

A VPN gateway functions as one end of a VPN tunnel. It receives incoming IPsec packets, decrypts the encapsulated data packets and passes the data packets to the local network. Also, it encrypts data packets destined for the other end of the VPN tunnel, encapsulates them, and sends the IPsec packets to the other VPN gateway. The VPN gateway is a FortiGate unit because the private network behind it is protected, ensuring the security of the unencrypted VPN data. The gateway can also be FortiClient software running on a PC since the unencrypted data is secure on the PC.

The IP address of a VPN gateway is usually the IP address of the network interface that connects to the Internet. Optionally, you can define a secondary IP address for the interface and use that address as the local VPN gateway address. The benefit of doing this is that your existing setup is not affected by the VPN settings.

The following diagram shows a VPN connection between two private networks with FortiGate units acting as the VPN gateways. This configuration is commonly referred to as Gateway-to-Gateway IPsec VPN.

**Figure 251:**VPN tunnel between two private networks



Although the IPsec traffic may actually pass through many Internet routers, you can visualize the VPN tunnel as a simple secure connection between the two FortiGate units.

Users on the two private networks do not need to be aware of the VPN tunnel. The applications on their computers generate packets with the appropriate source and destination addresses, as they normally do. The FortiGate units manage all the details of encrypting, encapsulating and sending the packets to the remote VPN gateway.


The data is encapsulated in IPsec packets only in the VPN tunnel between the two VPN gateways. Between the user's computer and the gateway, the data is on the secure private network and it is in regular IP packets.

For example User1 on the Site A network, at IP address 10.10.1.7, sends packets with destination IP address 192.168.10.8, the address of User2 on the Site B network. The Site A FortiGate unit is configured to send packets with destinations on the 192.168.10.0 network through the VPN, encrypted and encapsulated. Similarly, the Site B FortiGate unit is configured to send packets with destinations on the 10.10.1.0 network through the VPN tunnel to the Site A VPN gateway.

In the site-to-site, or gateway-to-gateway VPN shown in [Figure 251](#), the FortiGate units have static (fixed) IP addresses and either unit can initiate communication.

You can also create a VPN tunnel between an individual PC running FortiClient and a FortiGate unit, as shown below. This is commonly referred to as Client-to-Gateway IPsec VPN.

**Figure 252:**VPN tunnel between a FortiClient PC and a FortiGate unit



On the PC, the FortiClient application acts as the local VPN gateway. Packets destined for the office network are encrypted, encapsulated into IPsec packets, and sent through the VPN tunnel to the FortiGate unit. Packets for other destinations are routed to the Internet as usual. IPsec packets arriving through the tunnel are decrypted to recover the original IP packets.

## Clients, servers, and peers

A FortiGate unit in a VPN can have one of the following roles:

- **server** — responds to a request to establish a VPN tunnel.
- **client** — contacts a remote VPN gateway and requests a VPN tunnel.
- **peer** — brings up a VPN tunnel or responds to a request to do so.

The site-to-site VPN shown in [Figure 251](#) is a peer-to-peer relationship. Either FortiGate unit VPN gateway can establish the tunnel and initiate communications. The FortiClient-to-FortiGate VPN shown in [Figure 252](#) is a client-server relationship. The FortiGate unit establishes a tunnel when the FortiClient PC requests one.

A FortiGate unit cannot be a VPN server if it has a dynamically-assigned IP address. VPN clients need to be configured with a static IP address for the server. A FortiGate unit acts as a server only when the remote VPN gateway has a dynamic IP address or is a client-only device or application, such as FortiClient.



As a VPN server, a FortiGate unit can also offer automatic configuration for FortiClient PCs. The user needs to know only the IP address of the FortiGate VPN server and a valid user name/password. FortiClient downloads the VPN configuration settings from the FortiGate VPN server. For information about configuring a FortiGate unit as a VPN server, see the [FortiClient Administration Guide](#).

## Encryption

Encryption mathematically transforms data to appear as meaningless random numbers. The original data is called plaintext and the encrypted data is called ciphertext. The opposite process, called decryption, performs the inverse operation to recover the original plaintext from the ciphertext.

The process by which the plaintext is transformed to ciphertext and back again is called an algorithm. All algorithms use a small piece of information, a key, in the arithmetic process of converted plaintext to ciphertext, or vice-versa. IPsec uses symmetrical algorithms, in which the same key is used to both encrypt and decrypt the data.

The security of an encryption algorithm is determined by the length of the key that it uses. FortiGate IPsec VPNs offer the following encryption algorithms, in descending order of security:

<b>AES256</b>	A 128-bit block algorithm that uses a 256-bit key.
<b>AES192</b>	A 128-bit block algorithm that uses a 192-bit key.
<b>AES128</b>	A 128-bit block algorithm that uses a 128-bit key.
<b>3DES</b>	Triple-DES, in which plain text is DES-encrypted three times by three keys.
<b>DES</b>	Digital Encryption Standard, a 64-bit block algorithm that uses a 56-bit key

The default encryption algorithms provided on FortiGate units make recovery of encrypted data almost impossible without the proper encryption keys

There is a human factor in the security of encryption. The key must be kept secret, known only to the sender and receiver of the messages. Also, the key must not be something that unauthorized parties might easily guess, such as the sender's name, birthday or simple sequence such as 123456.

## Authentication

In addition to protecting data through encryption, a VPN must ensure that only authorized users can access the private network. You must use either a preshared key on both VPN gateways or RSA X.509 security certificates. The examples in this guide use only preshared key authentication. Refer to the [Fortinet Knowledge Base](#) for articles on RSA X.509 security certificates.

### Preshared keys

A preshared key contains at least six random alphanumeric characters. Users of the VPN must obtain the preshared key from the person who manages the VPN server and add the preshared key to their VPN client configuration.

Although it looks like a password, the preshared key, also known as a shared secret, is never sent by either gateway. The preshared key is used in the calculations at each end that generate

the encryption keys. As soon as the VPN peers attempt to exchange encrypted data, preshared keys that do not match will cause the process to fail.

## Additional authentication

To increase security, you can require additional means of authentication from users:

- an identifier, called a peer ID or a local ID
- extended authentication (XAUTH) which imposes an additional user name/password requirement

A Local ID is an alphanumeric value assigned in the Phase 1 configuration. The Local ID of a peer is called a Peer ID.

## Phase 1 and Phase 2 settings

A VPN tunnel is established in two phases: Phase 1 and Phase 2. Several parameters determine how this is done. Except for IP addresses, the settings simply need to match at both VPN gateways. There are defaults that are appropriate for most cases.

FortiClient distinguishes between Phase 1 and Phase 2 only in the VPN Advanced settings and uses different terms. Phase 1 is called the IKE Policy. Phase 2 is called the IPsec Policy.

### Phase 1

In Phase 1, the two VPN gateways exchange information about the encryption algorithms that they support and then establish a temporary secure connection to exchange authentication information.

When you configure your FortiGate unit or FortiClient application, you must specify the following settings for Phase 1:

<b>Remote Gateway</b>	The remote VPN gateway's address.  FortiGate units also have the option of operating only as a server by selecting the "Dialup User" option.
<b>Preshared key</b>	This must be the same at both ends. It is used to encrypt phase 1 authentication information.
<b>Local interface</b>	The network interface that connects to the other VPN gateway. This applies on a FortiGate unit only.

All other Phase 1 settings have default values. These settings mainly configure the types of encryption to be used. The default settings on FortiGate units and in the FortiClient application are compatible. The examples in this guide use these defaults.

For more detailed information about Phase 1 settings, see the ["Auto Key phase 1 parameters" on page 1637](#).

### Phase 2

Similar to the Phase 1 process, the two VPN gateways exchange information about the encryption algorithms that they support for Phase 2. You may choose different encryption for Phase 1 and Phase 2. If both gateways have at least one encryption algorithm in common, a VPN tunnel can be established. Keep in mind that more algorithms each phase does not share

with the other gateway, the longer negotiations will take. In extreme cases this may cause timeouts during negotiations.

To configure default Phase 2 settings on a FortiGate unit, you need only select the name of the corresponding Phase 1 configuration. In FortiClient, no action is required to enable default Phase 2 settings.

For more detailed information about Phase 2 settings, see [“Phase 2 parameters” on page 1653](#).

## Security Association

The establishment of a Security Association (SA) is the successful outcome of Phase 1 negotiations. Each peer maintains a database of information about VPN connections. The information in each SA can include cryptographic algorithms and keys, keylife, and the current packet sequence number. This information is kept synchronized as the VPN operates. Each SA has a Security Parameter Index (SPI) that is provided to the remote peer at the time the SA is established. Subsequent IPsec packets from the peer always reference the relevant SPI. It is possible for peers to have multiple VPNs active simultaneously, and correspondingly multiple SPIs.

# IPsec VPN Overview

This section provides a brief overview of IPsec technology and includes general information about how to configure IPsec VPNs using this guide.

The following topics are included in this section:

- [Types of VPNs](#)
- [Planning your VPN](#)
- [General preparation steps](#)
- [How to use this guide to configure an IPsec VPN](#)

VPN configurations interact with the firewall component of the FortiGate unit. There must be a security policy in place to permit traffic to pass between the private network and the VPN tunnel.

Security policies for VPNs specify:

- the FortiGate interface that provides the physical connection to the remote VPN gateway, usually an interface connected to the Internet
- the FortiGate interface that connects to the private network
- IP addresses associated with data that has to be encrypted and decrypted
- optionally, a schedule that restricts when the VPN can operate
- optionally, the services (types of data) that can be sent

When the first packet of data that meets all of the conditions of the security policy arrives at the FortiGate unit, a VPN tunnel may be initiated and the encryption or decryption of data is performed automatically afterward. For more information, see [“Defining VPN security policies” on page 1659](#).

## Types of VPNs

FortiGate unit VPNs can be policy-based or route-based. There is little difference between the two types. In both cases, you specify phase 1 and phase 2 settings. However there is a difference in implementation. A route-based VPN creates a virtual IPsec network interface that applies encryption or decryption as needed to any traffic that it carries. That is why route-based VPNs are also known as interface-based VPNs. A policy-based VPN is implemented through a special security policy that applies the encryption you specified in the phase 1 and phase 2 settings.

### Route-based VPNs

For a route-based VPN, you create two security policies between the virtual IPsec interface and the interface that connects to the private network. In one policy the virtual interface is the source. In the other policy the virtual interface is the destination. The Action for both policies is Accept. This creates bidirectional policies that ensure traffic will flow in both directions over the VPN.

## Policy-based VPNs

For a policy-based VPN, one security policy enables communication in both directions. You must select IPSEC as the Action and then select the VPN tunnel you defined in the phase 1 settings. You can then enable inbound and outbound traffic as needed within that policy, or create multiple policies of this type to handle different types of traffic differently. For example HTTPS traffic may not require the same level of scanning as FTP traffic.

## Comparing policy-based or route-based VPNs

For both VPN types you create phase 1 and phase 2 configurations. Both types are handled in the stateful inspection security layer, assuming there is no IPS or AV. For more information on the security layers, see [“Life of a Packet” on page 2271](#).

The main difference is in the security policy.

You create a policy-based VPN by defining an IPSEC security policy between two network interfaces and associating it with the VPN tunnel (phase 1) configuration.

You create a route-based VPN by enabling IPsec interface mode in the VPN phase 1 configuration. This creates a virtual IPsec interface. You then define a regular ACCEPT security policy to permit traffic to flow between the virtual IPsec interface and another network interface. And lastly, configure a static route to allow traffic over the VPN.

Where possible, you should create route-based VPNs. Generally, route-based VPNs are more flexible and easier to configure than policy-based VPNs — by default they are treated as interfaces. However, these two VPN types have different requirements that limit where they can be used.

**Table 71:** Comparison of policy-based and route-based VPNs

Features	Policy-based	Route-based
<ul style="list-style-type: none"><li>Both NAT and transparent modes available</li></ul>	<ul style="list-style-type: none"><li>Yes</li></ul>	<ul style="list-style-type: none"><li>NAT mode only</li></ul>
<ul style="list-style-type: none"><li>L2TP-over-IPsec supported</li></ul>	<ul style="list-style-type: none"><li>Yes</li></ul>	<ul style="list-style-type: none"><li>No</li></ul>
<ul style="list-style-type: none"><li>GRE-over-IPsec supported</li></ul>	<ul style="list-style-type: none"><li>No</li></ul>	<ul style="list-style-type: none"><li>Yes</li></ul>
<ul style="list-style-type: none"><li>security policy requirements</li></ul>	<ul style="list-style-type: none"><li>Requires a security policy with IPSEC action that specifies the VPN tunnel</li></ul>	<ul style="list-style-type: none"><li>Requires only a simple security policy with ACCEPT action</li></ul>
<ul style="list-style-type: none"><li>Number of policies per VPN</li></ul>	<ul style="list-style-type: none"><li>One policy controls connections in both directions</li></ul>	<ul style="list-style-type: none"><li>A separate policy is required for connections in each direction</li></ul>

## Planning your VPN

It is a good idea to plan the VPN configuration ahead of time. This will save time later and help you configure your VPN correctly.

All VPN configurations comprise a number of required and optional parameters. Before you begin, you need to determine:

- where does the IP traffic originate and where does it need to be delivered
- which hosts, servers, or networks to include in the VPN
- which VPN devices to include in the configuration
- through which interfaces the VPN devices communicate
- through which interfaces do private networks access the VPN gateways

Once you have this information, you can select a VPN topology that meets the requirements of your situation.

## Network topologies

The topology of your network will determine how remote peers and clients connect to the VPN and how VPN traffic is routed. You can read about various network topologies and find the high-level procedures needed to configure IPsec VPNs in one of these sections.

**Table 72:** VPN network topologies and brief descriptions

Topology	Description
<a href="#">Gateway-to-gateway configurations</a>	Standard one-to-one VPN between two FortiGate units. See <a href="#">“Gateway-to-gateway configurations”</a> on page 1665.
<a href="#">Hub-and-spoke configurations</a>	One central FortiGate unit has multiple VPNs to other remote FortiGate units. See <a href="#">“Hub-and-spoke configurations”</a> on page 1679.
<a href="#">Dynamic DNS configuration</a>	One end of the VPN tunnel has a changing IP address and the other end must go to a dynamic DNS server for the current IP address before establishing a tunnel. See <a href="#">“Dynamic DNS configuration”</a> on page 1695.
<a href="#">FortiClient dialup-client configurations</a>	Typically remote FortiClient dialup-clients use dynamic IP addresses through NAT devices. The FortiGate unit acts as a dialup server allowing dialup VPN connections from multiple sources. See <a href="#">“FortiClient dialup-client configurations”</a> on page 1709.
<a href="#">FortiGate dialup-client configurations</a>	Similar to FortiClient dialup-client configurations but with more gateway-to-gateway settings such as unique user authentication for multiple users on a single VPN tunnel. See <a href="#">“FortiGate dialup-client configurations”</a> on page 1724.
<a href="#">Internet-browsing configuration</a>	Secure web browsing performed by dialup VPN clients, and/or hosts behind a remote VPN peer. See <a href="#">“Internet-browsing configuration”</a> on page 1737.
<a href="#">Redundant VPN configurations</a>	Options for supporting redundant and partially redundant IPsec VPNs, using route-based approaches. See <a href="#">“Redundant VPN configurations”</a> on page 1741.
<a href="#">Transparent mode VPNs</a>	In transparent mode, the FortiGate acts as a bridge with all incoming traffic being broadcast back out on all other interfaces. Routing and NAT must be performed on external routers. See <a href="#">“Transparent mode VPNs”</a> on page 1766.

**Table 72:** VPN network topologies and brief descriptions

Topology	Description
<a href="#">Manual-key configurations</a>	Manually define cryptographic keys to establish an IPsec VPN, either policy-based or route-based. See <a href="#">“Manual-key configurations”</a> on page 1770.
<a href="#">L2TP and IPsec (Microsoft VPN)</a>	Configure VPN for Microsoft Windows dialup clients using the built in L2TP software. Users do not have to install any See <a href="#">“L2TP and IPsec (Microsoft VPN)”</a> on page 1786.

These sections contain high-level configuration guidelines with cross-references to detailed configuration procedures. If you need more detail to complete a step, select the cross-reference in the step to drill-down to more detail. Return to the original procedure to complete the procedure. For a general overview of how to configure a VPN, see [“General preparation steps”](#) below.

## General preparation steps

A VPN configuration defines relationships between the VPN devices and the private hosts, servers, or networks making up the VPN. Configuring a VPN involves gathering and recording the following information. You will need this information to configure the VPN.

- **The private IP addresses of participating hosts, servers, and/or networks.** These IP addresses represent the source addresses of traffic that is permitted to pass through the VPN. A IP source address can be an individual IP address, an address range, or a subnet address.
- **The public IP addresses of the VPN end-point interfaces.** The VPN devices establish tunnels with each other through these interfaces.
- **The private IP addresses associated with the VPN-device interfaces to the private networks.** Computers on the private networks behind the VPN gateways will connect to their VPN gateways through these interfaces.

## How to use this guide to configure an IPsec VPN

This guide uses a task-based approach to provide all of the procedures needed to create different types of VPN configurations. Follow the step-by-step configuration procedures in this guide to set up the VPN.

The following configuration procedures are common to all IPsec VPNs:

1. Define the phase 1 parameters that the FortiGate unit needs to authenticate remote peers or clients and establish a secure a connection. See [“Auto Key phase 1 parameters”](#) on page 1637.
2. Define the phase 2 parameters that the FortiGate unit needs to create a VPN tunnel with a remote peer or dialup client. See [“Phase 2 parameters”](#) on page 1653.
3. Specify the source and destination addresses of IP packets that are to be transported through the VPN tunnel. See [“Defining policy addresses”](#) on page 1659.
4. Create an IPsec security policy to define the scope of permitted services between the IP source and destination addresses. See [“Defining VPN security policies”](#) on page 1660.

# IPsec VPN in the web-based manager

The IPsec VPN menu in FortiOS provides settings to configure an IPsec VPN. IPsec VPNs that are configured by using the general procedure below. With these steps, your FortiGate unit will automatically generate unique IPsec encryption and authentication keys.

1. Define phase 1 parameters to authenticate remote peers and clients for a secure connection. See [“Phase 1 configuration” on page 1625](#).
2. Define phase 2 parameters to create a VPN tunnel with a remote peer or dialup client. See [“Phase 2 configuration” on page 1629](#).
3. Create a security policy to permit communication between your private network and the VPN. Policy-based VPNs have an action of IPSEC, where for interface-based VPNs the security policy action is ACCEPT. See [“Defining VPN security policies” on page 1659](#).

The FortiGate unit implements the Encapsulated Security Payload (ESP) protocol. Internet Key Exchange (IKE) is performed automatically based on pre-shared keys or X.509 digital certificates. Interface mode, supported in NAT mode only, creates a virtual interface for the local end of a VPN tunnel.

This topic contains the following:

- [Auto Key \(IKE\)](#)
- [Manual Key](#)
- [Concentrator](#)

## Auto Key (IKE)

You can configure VPN peers (or a FortiGate dialup server and a VPN client) to generate unique Internet Key Exchange (IKE) keys automatically during the IPsec phase 1 and phase 2 exchanges.

When you define phase 2 parameters, you can choose any set of phase 1 parameters to set up a secure connection for the tunnel and authenticate the remote peer. Auto Key configuration applies to both tunnel-mode and interface-mode VPNs.

To configure VPN peers go to *VPN > IPsec > Auto Key (IKE)*.

<b>Create Phase 1</b>	Creates a new phase 1 tunnel configuration. For more information, see <a href="#">“Phase 1 configuration” on page 1625</a> .
<b>Create Phase 2</b>	Creates a new phase 2 configuration. For more information, see <a href="#">“Phase 2 configuration” on page 1629</a> .
<b>Create FortiClient VPN</b>	Creates a new FortiClient VPN. For more information, see <a href="#">“FortiClient VPN” on page 1632</a> .

If you want to control how the IKE negotiation process controls traffic when there is no traffic, as well as the length of time the FortiGate unit waits for negotiations to occur, use the `negotiation-timeout` and `auto-negotiation` commands in the CLI.



## Phase 1 configuration

The basic phase 1 settings associate IPsec phase 1 parameters with a remote gateway, if a pre-shared key or digital certificate will be used, and if a special identifier will be used to identify the remote VPN peer or client.

<b>Name</b>	<p>Type a name for the phase 1 definition. The maximum name length is 15 characters for an interface mode VPN, 35 characters for a policy-based VPN. If <i>Remote Gateway</i> is <i>Dialup User</i>, the maximum name length is further reduced depending on the number of dialup tunnels that can be established: by 2 for up to 9 tunnels, by 3 for up to 99 tunnels, 4 for up to 999 tunnels, and so on.</p> <p>For a tunnel mode VPN, the name normally reflects where the remote connection originates. For a route-based tunnel, the FortiGate unit also uses the name for the virtual IPsec interface that it creates automatically.</p>
<b>Remote Gateway</b>	<p>Select the category of the remote connection:</p> <ul style="list-style-type: none"><li>• <i>Static IP Address</i> — If the remote peer has a static IP address.</li><li>• <i>Dialup User</i> — If one or more FortiClient or FortiGate dialup clients with dynamic IP addresses will connect to the FortiGate unit.</li><li>• <i>Dynamic DNS</i> — If a remote peer that has a domain name and subscribes to a dynamic DNS service will connect to the FortiGate unit.</li></ul>
<b>IP Address</b>	<p>If you selected <i>Static IP Address</i>, enter the IP address of the remote peer.</p>
<b>Dynamic DNS</b>	<p>If you selected <i>Dynamic DNS</i>, enter the domain name of the remote peer.</p>
<b>Local Interface</b>	<p>This option is available in NAT mode only. Select the name of the interface through which remote peers or dialup clients connect to the FortiGate unit.</p> <p>By default, the local VPN gateway IP address is the IP address of the interface that you selected. Optionally, you can specify a unique IP address for the VPN gateway in the <i>Advanced</i> settings.</p>
<b>Mode</b>	<ul style="list-style-type: none"><li>• <i>Main mode</i> — the phase 1 parameters are exchanged in multiple rounds with encrypted authentication information.</li><li>• <i>Aggressive mode</i> — the phase 1 parameters are exchanged in single message with authentication information that is not encrypted.</li></ul> <p>When the remote VPN peer has a dynamic IP address and is authenticated by a pre-shared key, you must select Aggressive mode if there is more than one dialup phase1 configuration for the interface IP address.</p> <p>When the remote VPN peer has a dynamic IP address and is authenticated by a certificate, you must select Aggressive mode if there is more than one phase 1 configuration for the interface IP address and these phase 1 configurations use different proposals.</p>
<b>Authentication Method</b>	<p>Select <i>Preshared Key</i> or <i>RSA Signature</i>.</p>

<b>Pre-shared Key</b>	If you selected <i>Pre-shared Key</i> , enter the pre-shared key that the FortiGate unit will use to authenticate itself to the remote peer or dialup client during phase 1 negotiations. You must define the same key at the remote peer or client. The key must contain at least 6 printable characters. For optimum protection against currently known attacks, the key must consist of a minimum of 16 randomly chosen alphanumeric characters.
<b>Certificate Name</b>	If you selected <i>RSA Signature</i> , select the name of the server certificate that the FortiGate unit will use to authenticate itself to the remote peer or dialup client during phase 1 negotiations. For information about obtaining and loading the required server certificate, see the <a href="#">FortiOS User Authentication guide</a> .
<b>Peer Options</b>	Peer options are available to authenticate VPN peers or clients, depending on the <i>Remote Gateway</i> and <i>Authentication Method</i> settings.
<b>Accept any peer ID</b>	Accept the local ID of any remote VPN peer or client. The FortiGate unit does not check identifiers (local IDs). You can set <i>Mode</i> to <i>Aggressive</i> or <i>Main</i> .  You can use this option with RSA Signature authentication. But, for highest security, configure a PKI user/group for the peer and set <i>Peer Options</i> to <i>Accept this peer certificate only</i> .
<b>Accept this peer ID</b>	This option is available when <i>Aggressive Mode</i> is enabled. Enter the identifier that is used to authenticate the remote peer. This identifier must match the Local ID that the remote peer's administrator has configured.  If the remote peer is a FortiGate unit, the identifier is specified in the <i>Local ID</i> field of the Advanced phase 1 configuration.  If the remote peer is a FortiClient user, the identifier is specified in the <i>Local ID</i> field, accessed by selecting <i>Config</i> in the <i>Policy</i> section of the VPN connection's <i>Advanced Settings</i> .
<b>Accept peer ID in dialup group</b>	Authenticate multiple FortiGate or FortiClient dialup clients that use unique identifiers and unique pre-shared keys (or unique pre-shared keys only) through the same VPN tunnel.  You must create a dialup user group for authentication purposes. Select the group from the list next to the <i>Accept peer ID in dialup group</i> option.  You must set <i>Mode</i> to <i>Aggressive</i> when the dialup clients use unique identifiers and unique pre-shared keys. If the dialup clients use unique pre-shared keys only, you can set <i>Mode</i> to <i>Main</i> if there is only one dialup phase 1 configuration for this interface IP address.
<b>Advanced</b>	Defines advanced phase 1 parameters. For more information, see <a href="#">Phase 1 advanced configuration settings</a> .

## Phase 1 advanced configuration settings

You use the advanced parameters to select the encryption and authentication algorithms that the FortiGate unit uses to generate keys for the IKE exchange. You can also select these advanced settings to ensure the smooth operation of phase 1 negotiations.

To configure Phase 1 settings, go to *VPN > Auto Key (IKE)* and select *Create Phase 1*.

---

<b>Enable IPsec Interface Mode</b>	<p>This is available in NAT mode only.</p> <p>Create a virtual interface for the local end of the VPN tunnel. Select this option to create a route-based VPN, clear it to create a policy-based VPN.</p>
<b>IKE Version</b>	<p>Select the version of IKE to use. This is available only if <i>IPsec Interface Mode</i> is enabled. For more information about IKE v2, refer to RFC 4306.</p> <p>IKE v2 is not available if <i>Mode</i> is <i>Aggressive</i>.</p> <p>When <i>IKE Version</i> is 2, <i>Mode</i> and <i>XAUTH</i> are not available.</p>
<b>IPv6 Version</b>	<p>Select if you want to use IPv6 addresses for the remote gateway and interface IP addresses. This is available only when <i>Enable IPsec Interface Mode</i> is selected and IPv6 Support is enabled in the administrative settings (<i>System &gt; Admin &gt; Settings</i>).</p>
<b>Local Gateway IP</b>	<p>If you selected <i>Enable IPsec Interface Mode</i>, specify an IP address for the local end of the VPN tunnel. Select one of the following:</p> <ul style="list-style-type: none"><li>• <i>Main Interface IP</i> — The FortiGate unit obtains the IP address of the interface from the network interface settings.</li><li>• <i>Specify</i> — Enter a secondary address of the interface selected in the phase 1 <i>Local Interface</i> field.</li></ul> <p>You cannot configure Interface mode in a transparent mode VDOM.</p>
<b>P1 Proposal</b>	<p>Select the encryption and authentication algorithms used to generate keys for protecting negotiations and add encryption and authentication algorithms as required.</p> <p>You need to select a minimum of one and a maximum of three combinations. The remote peer or client must be configured to use at least one of the proposals that you define.</p> <p>Select one of the following symmetric-key encryption algorithms:</p> <ul style="list-style-type: none"><li>• <i>DES</i> — Digital Encryption Standard, a 64-bit block algorithm that uses a 56-bit key.</li><li>• <i>3DES</i> — Triple-DES, in which plain text is encrypted three times by three keys.</li><li>• <i>AES128</i> — a 128-bit block Cipher Block Chaining (CBC) algorithm that uses a 128-bit key.</li><li>• <i>AES192</i> — a 128-bit block Cipher Block Chaining (CBC) algorithm that uses a 192-bit key.</li><li>• <i>AES256</i> — a 128-bit block Cipher Block Chaining (CBC) algorithm that uses a 256-bit key.</li></ul>

---

---

	<p>Select either of the following authentication message digests to check the authenticity of messages during phase 1 negotiations:</p> <ul style="list-style-type: none"> <li>• <i>MD5</i> — Message Digest 5, the hash algorithm developed by RSA Data Security.</li> <li>• <i>SHA1</i> — Secure Hash Algorithm 1, which produces a 160-bit message digest.</li> <li>• <i>SHA256</i> — Secure Hash Algorithm 2, which produces a 256-bit message digest.</li> </ul> <p>To specify a third combination, use the <i>Add</i> button beside the fields for the second combination.</p>
<b>DH Group</b>	<p>Select one or more Diffie-Hellman groups from DH group 1, 2, 5 and 14. At least one of the <i>DH Group</i> settings on the remote peer or client must match one the selections on the FortiGate unit. Failure to match one or more DH groups will result in failed negotiations.</p>
<b>Keylife</b>	<p>Enter the time (in seconds) that must pass before the IKE encryption key expires. When the key expires, a new key is generated without interrupting service. The keylife can be from 120 to 172 800 seconds.</p>
<b>Local ID</b>	<p>If the FortiGate unit will act as a VPN client and you are using peer IDs for authentication purposes, enter the identifier that the FortiGate unit will supply to the VPN server during the phase 1 exchange.</p> <p>If the FortiGate unit will act as a VPN client, and you are using security certificates for authentication, select the distinguished name (DN) of the local server certificate that the FortiGate unit will use for authentication purposes.</p> <p>If the FortiGate unit is a dialup client and will not be sharing a tunnel with other dialup clients (that is, the tunnel will be dedicated to this Fortinet dialup client), set <i>Mode</i> to <i>Aggressive</i>.</p> <p>Note that this Local ID value must match the peer ID value given for the remote VPN peer's Peer Options.</p>
<b>XAuth</b>	<p>This option supports the authentication of dialup clients. It is available for IKE v1 only.</p> <ul style="list-style-type: none"> <li>• <i>Disable</i> — Select if you do not use XAuth.</li> <li>• <i>Enable as Client</i> — If the FortiGate unit is a dialup client, enter the user name and password that the FortiGate unit will need to authenticate itself to the remote XAuth server.</li> <li>• <i>Enable as Server</i> — This is available only if <i>Remote Gateway</i> is set to <i>Dialup User</i>. Dialup clients authenticate as members of a dialup user group. You must first create a user group for the dialup clients that need access to the network behind the FortiGate unit.</li> </ul> <p>You must also configure the FortiGate unit to forward authentication requests to an external RADIUS or LDAP authentication server.</p> <p>Select a <i>Server Type</i> setting to determine the type of encryption method to use between the FortiGate unit, the XAuth client and the external authentication server, and then select the user group from the User Group list.</p>
<b>Username</b>	<p>Enter the user name that is used for authentication.</p>

---

<b>Password</b>	Enter the password that is used for authentication.
<b>NAT Traversal</b>	Select the check box if a NAT device exists between the local FortiGate unit and the VPN peer or client. The local FortiGate unit and the VPN peer or client must have the same NAT traversal setting (both selected or both cleared) to connect reliably.
<b>Keepalive Frequency</b>	If you enabled <i>NAT-traversal</i> , enter a keepalive frequency setting.
<b>Dead Peer Detection</b>	<p>Select this check box to reestablish VPN tunnels on idle connections and clean up dead IKE peers if required. You can use this option to receive notification whenever a tunnel goes up or down, or to keep the tunnel connection open when no traffic is being generated inside the tunnel. For example, in scenarios where a dialup client or dynamic DNS peer connects from an IP address that changes periodically, traffic may be suspended while the IP address changes.</p> <p>With <i>Dead Peer Detection</i> selected, you can use the <code>config vpn ipsec phase1 (tunnel mode)</code> or <code>config vpn ipsec phase1-interface (interface mode)</code> CLI command to optionally specify a retry count and a retry interval.</p>

## Phase 2 configuration

After IPsec phase 1 negotiations end successfully, you begin phase 2. You configure the phase 2 parameters to define the algorithms that the FortiGate unit may use to encrypt and transfer data for the remainder of the session. During phase 2, you select specific IPsec security associations needed to implement security services and establish a tunnel.

The basic phase 2 settings associate IPsec phase 2 parameters with the phase 1 configuration that specifies the remote end point of the VPN tunnel. In most cases, you need to configure only basic phase 2 settings.

To configure Phase 2 settings, go to *VPN > Auto Key (IKE)* and select *Create Phase 2*.

<b>Name</b>	Type a name to identify the phase 2 configuration.
<b>Phase 1</b>	Select the phase 1 tunnel configuration. For more information on configuring phase 1, see <a href="#">“Phase 1 configuration” on page 1625</a> . The phase 1 configuration describes how remote VPN peers or clients will be authenticated on this tunnel, and how the connection to the remote peer or client will be secured.
<b>Advanced</b>	Define advanced phase 2 parameters. For more information, see <a href="#">“Phase 2 advanced configuration settings” on page 1629</a> .

## Phase 2 advanced configuration settings

In phase 2, the FortiGate unit and the VPN peer or client exchange keys again to establish a secure communication channel between them. You select the encryption and authentication algorithms needed to generate keys for protecting the implementation details of Security Associations (SAs). These are called P2 Proposal parameters. The keys are generated automatically using a Diffie-Hellman algorithm.

You can use a number of additional advanced phase 2 settings to enhance the operation of the tunnel.

---

<b>P2 Proposal</b>	<p>Select the encryption and authentication algorithms that will be proposed to the remote VPN peer. You can specify up to three proposals. To establish a VPN connection, at least one of the proposals that you specify must match configuration on the remote peer.</p> <p>Initially there are two proposals. <i>Add</i> and <i>Delete</i> icons are next to the second <i>Authentication</i> field.</p> <p>It is invalid to set both <i>Encryption</i> and <i>Authentication</i> to NULL.</p>
<b>Encryption</b>	<p>Select one of the following symmetric-key algorithms:</p> <ul style="list-style-type: none"><li>• <i>NULL</i> — Do not use an encryption algorithm.</li><li>• <i>DES</i> — Digital Encryption Standard, a 64-bit block algorithm that uses a 56-bit key.</li><li>• <i>3DES</i> — Triple-DES, in which plain text is encrypted three times by three keys.</li><li>• <i>AES128</i> — a 128-bit block Cipher Block Chaining (CBC) algorithm that uses a 128-bit key.</li><li>• <i>AES192</i> — a 128-bit block CBC algorithm that uses a 192-bit key.</li><li>• <i>AES256</i> — a 128-bit block CBC algorithm that uses a 256-bit key.</li></ul>
<b>Authentication</b>	<p>Select one of the following message digests to check the authenticity of messages during an encrypted session:</p> <ul style="list-style-type: none"><li>• <i>NULL</i> — Do not use a message digest.</li><li>• <i>MD5</i> — Message Digest 5, the hash algorithm developed by RSA Data Security.</li><li>• <i>SHA1</i> — Secure Hash Algorithm 1, which produces a 160-bit message digest.</li><li>• <i>SHA256</i> — Secure Hash Algorithm 2, which produces a 256-bit message digest.</li><li>• <i>SHA384</i> — Secure Hash Algorithm 2, which produces a 384-bit message digest.</li><li>• <i>SHA512</i> — Secure Hash Algorithm 2, which produces a 512-bit message digest.</li></ul>
<b>Enable replay detection</b>	<p>Replay attacks occur when an unauthorized party intercepts a series of IPsec packets and replays them back into the tunnel.</p>
<b>Enable perfect forward secrecy (PFS)</b>	<p>Perfect forward secrecy (PFS) improves security by forcing a new Diffie-Hellman exchange whenever keylife expires.</p>
<b>DH Group</b>	<p>Select one Diffie-Hellman group (1, 2, 5 or 14). This must match the DH Group that the remote peer or dialup client uses.</p>
<b>Keylife</b>	<p>Select the method for determining when the phase 2 key expires: <i>Seconds</i>, <i>KBytes</i>, or <i>Both</i>. If you select <i>Both</i>, the key expires when either the time has passed or the number of KB have been processed.</p>
<b>Autokey Keep Alive</b>	<p>Select the check box if you want the tunnel to remain active when no data is being processed.</p>

---

<b>DHCP-IPSec</b>	<p>Provide IP addresses dynamically to VPN clients. This is available for phase 2 configurations associated with a dialup phase 1 configuration.</p> <p>You also need configure a DHCP server or relay on the private network interface. You must configure the DHCP parameters separately.</p> <p>If you configure the DHCP server to assign IP addresses based on RADIUS user group attributes, you must also set the Phase 1 <i>Peer Options</i> to <i>Accept peer ID in dialup group</i> and select the appropriate user group. See <a href="#">“Phase 1 configuration” on page 1625</a>.</p> <p>If the FortiGate unit acts as a dialup server and you manually assigned FortiClient dialup clients VIP addresses that match the network behind the dialup server, selecting the check box will cause the FortiGate unit to act as a proxy for the dialup clients.</p>
<b>Quick Mode Selector</b>	<p>Specify the source and destination IP addresses to be used as selectors for IKE negotiations. If the FortiGate unit is a dialup server, keep the default value of 0.0.0.0/0 unless you need to circumvent problems caused by ambiguous IP addresses between one or more of the private networks making up the VPN. You can specify a single host IP address, an IP address range, or a network address. You may optionally specify source and destination port numbers and a protocol number.</p> <p>If you are editing an existing phase 2 configuration, the <i>Source address</i> and <i>Destination address</i> fields are unavailable if the tunnel has been configured to use firewall addresses as selectors. This option exists only in the CLI.</p>
<b>Source address</b>	<p>If the FortiGate unit is a dialup server, enter the source IP address that corresponds to the local senders or network behind the local VPN peer (for example, 172.16.5.0/24 or 172.16.5.0/255.255.255.0 for a subnet, or 172.16.5.1/32 or 172.16.5.1/255.255.255.255 for a server or host, or 192.168.10.[80-100] or 192.168.10.80-192.168.10.100 for an address range). A value of 0.0.0.0/0 means all IP addresses behind the local VPN peer.</p> <p>If the FortiGate unit is a dialup client, source address must refer to the private network behind the Fortinet dialup client.</p>
<b>Source port</b>	<p>Enter the port number that the local VPN peer uses to transport traffic related to the specified service (protocol number). The range is from 0 to 65535. To specify all ports, type 0.</p>
<b>Destination address</b>	<p>Enter the destination IP address that corresponds to the recipients or network behind the remote VPN peer (for example, 192.168.20.0/24 for a subnet, or 172.16.5.1/32 for a server or host, or 192.168.10.[80-100] for an address range). A value of 0.0.0.0/0 means all IP addresses behind the remote VPN peer.</p>
<b>Destination port</b>	<p>Enter the port number that the remote VPN peer uses to transport traffic related to the specified service (protocol number). To specify all ports, enter 0.</p>
<b>Protocol</b>	<p>Enter the IP protocol number of the service. To specify all services, enter 0.</p>

## FortiClient VPN

Use the FortiClient VPN configuration settings when configuring an IPsec VPN for remote users to connect to the VPN tunnel using FortiClient.

To create a FortiClient VPN tunnel, go to *VPN > IPsec > Auto Key (IKE)* and select *Create FortiClient VPN* at the top of the screen.

When configuring a FortiClient VPN connection, the settings for phase 1 and phase 2 settings are automatically configured by the FortiGate unit. They are set to:

- Remote Gateway — Dialup User
- Mode — Aggressive
- IPsec Interface Mode — Enabled
- Default settings for P1 and P2 Proposal
- XAUTH Enable as Server (Auto)
- IKE mode-config will be enabled
- Peer Option — “Accept any peer ID”

The remainder of the settings use the current FortiGate defaults. Note that FortiClient settings need to match these FortiGate defaults. If you need to configure advanced settings for the FortiClient VPN, select *Edit* on the Auto Key (IKE) page (Go to *VPN > IPsec > Auto Key (IKE)*) and configure the peer options or advanced options.

<b>Name</b>	Enter a name for the FortiClient VPN.
<b>Local Outgoing Interface</b>	Select the local outgoing interface for the VPN.
<b>Authentication Method</b>	Select the type of authentication used when logging in to the VPN.
<b>Preshared Key</b>	If <i>Pre-shared Key</i> was selected in <i>Authentication Method</i> , enter the pre-shared key in the field provided.
<b>User Group</b>	Select a user group. You can also create a user group from the drop-down list by selecting <i>Create New</i> .
<b>Address Range Start IP</b>	Enter the start IP address for the DHCP address range for the client.
<b>Address Range End IP</b>	Enter the end IP address for the address range.
<b>Subnet Mask</b>	Enter the subnet mask.
<b>Enable IPv4 Split Tunnel</b>	Enabled by default, this option enables the FortiClient user to use the VPN to access internal resources while other Internet access is not sent over the VPN, alleviating potential traffic bottlenecks in the VPN connection. Disable this option to have all traffic sent through the VPN tunnel.
<b>Accessible Networks</b>	Select from a list of internal networks that the FortiClient user can access.



<b>Client Options</b>	<p>These options affect how the FortiClient application behaves when connected to the FortiGate VPN tunnel. When enabled, a check box for the corresponding option appears on the VPN login screen in FortiClient, and is not enabled by default.</p> <p><b>Save Password</b> - When enabled, if the user selects this option, their password is stored on the user's computer and will automatically populate each time they connect to the VPN.</p> <p><b>Auto Connect</b> - When enabled, if the user selects this option, when the FortiClient application is launched, for example after a reboot or system startup, FortiClient will automatically attempt to connect to the VPN tunnel.</p> <p><b>Always Up (Keep Alive)</b> - When enabled, if the user selects this option, the FortiClient connection will not shut down. When not selected, during periods of inactivity, FortiClient will attempt to stay connected every three minutes for a maximum of 10 minutes.</p>
<b>Endpoint Registration</b>	<p>When selected, the FortiGate unit requests a registration key from FortiClient before a connection can be established. A registration key is defined by going to <i>System &gt; Config &gt; Advanced</i>.</p> <p>For more information on FortiClient VPN connections to a FortiGate unit, see the <i>FortiClient Administration Guide</i>.</p>
<b>DNS Server</b>	<p>Select which DNS server to use for this VPN:</p> <ul style="list-style-type: none"> <li>• <i>Use System DNS</i> — Use the same DNS servers as the FortiGate unit. These are configured at <i>System &gt; Interface &gt; DNS</i>. This is the default option.</li> <li>• <i>Specify</i> — Specify the IP address of a different DNS server.</li> </ul>

## Manual Key

Use manual keys only if it is unavoidable. There are potential difficulties in keeping keys confidential and in propagating changed keys to remote VPN peers securely.

If required, you can manually define cryptographic keys for establishing an IPsec VPN tunnel. You would define manual keys in situations where:

- you require prior knowledge of the encryption or authentication key (that is, one of the VPN peers requires a specific IPsec encryption or authentication key).
- you need to disable encryption and authentication.

In both cases, you do not specify IPsec phase 1 and phase 2 parameters; you define manual keys from the CLI using the `config vpn ipsec manualkey` command.

### Manual key configuration settings

If one of the VPN devices is manually keyed, the other VPN device must also be manually keyed with the identical authentication and encryption keys. In addition, it is essential that both VPN devices be configured with complementary Security Parameter Index (SPI) settings. The administrators of the devices need to cooperate to achieve this.

If you are not familiar with the security policies, SAs, selectors, and SA databases for your particular installation, do not attempt these procedures without qualified assistance.

Each SPI identifies a Security Association (SA). The value is placed in ESP datagrams to link the datagrams to the SA. When an ESP datagram is received, the recipient refers to the SPI to determine which SA applies to the datagram. You must manually specify an SPI for each SA. There is an SA for each direction, so for each VPN you must specify two SPIs, a local SPI and a remote SPI, to cover bidirectional communications between two VPN devices.

---

<b>Name</b>	Type a name for the VPN tunnel. The maximum name length is 15 characters for an interface mode VPN, 35 characters for a policy-based VPN.
<b>Local SPI</b>	Type a hexadecimal number (up to 8 characters, 0-9, a-f) that represents the SA that handles outbound traffic on the local FortiGate unit. The valid range is from 0x100 to 0xffffffff. This value must match the Remote SPI value in the manual key configuration at the remote peer.
<b>Remote SPI</b>	Type a hexadecimal number (up to 8 characters, 0-9, a-f) that represents the SA that handles inbound traffic on the local FortiGate unit. The valid range is from 0x100 to 0xffffffff. This value must match the Local SPI value in the manual key configuration at the remote peer.
<b>Remote Gateway</b>	Enter the IP address of the public interface to the remote peer. The address identifies the recipient of ESP datagrams.
<b>Local Interface</b>	This option is available in NAT mode only. Select the name of the interface to which the IPsec tunnel will be bound. The FortiGate unit obtains the IP address of the interface from the network interface settings.
<b>Encryption Algorithm</b>	Select one of the following symmetric-key encryption algorithms: <ul style="list-style-type: none"><li>• <i>NULL</i> — Do not use an encryption algorithm.</li><li>• <i>DES</i> — Digital Encryption Standard, a 64-bit block algorithm that uses a 56-bit key.</li><li>• <i>3DES</i> — Triple-DES, where plain text is encrypted three times by three keys.</li><li>• <i>AES128</i> — a 128-bit block Cipher Block Chaining algorithm that uses a 128-bit key.</li><li>• <i>AES192</i> — a 128-bit block Cipher Block Chaining ) algorithm that uses a 192-bit key.</li><li>• <i>AES256</i> — a 128-bit block Cipher Block Chaining algorithm that uses a 256-bit key.</li></ul> <p><b>Note:</b> The algorithms for encryption and authentication cannot both be NULL.</p>

---

<b>Authentication Algorithm</b>	<p>Select one of the following message digests:</p> <ul style="list-style-type: none"> <li>• <i>NULL</i> — Do not use a message digest.</li> <li>• <i>MD5</i> — Message Digest 5 algorithm, which produces a 128-bit message digest.</li> <li>• <i>SHA1</i> — Secure Hash Algorithm 1, which produces a 160-bit message digest.</li> <li>• <i>SHA256</i> — Secure Hash Algorithm 2, which produces a 256-bit message digest.</li> <li>• <i>SHA384</i> — Secure Hash Algorithm 2, which produces a 384-bit message digest.</li> <li>• <i>SHA512</i> — Secure Has Algorithm 2, which produces a 512-bit message digest.</li> </ul> <p><b>Note:</b> The Algorithms for encryption and authentication cannot both be NULL.</p>
<b>IPsec Interface Mode</b>	<p>Create a virtual interface for the local end of the VPN tunnel. Select this check box to create a route-based VPN, clear it to create a policy-based VPN.</p> <p>This is available only in NAT mode.</p>

## Concentrator

In a hub-and-spoke configuration, policy-based VPN connections to a number of remote peers radiate from a single, central FortiGate unit. Site-to-site connections between the remote peers do not exist; however, you can establish VPN tunnels between any two of the remote peers through the FortiGate unit's "hub".

In a hub-and-spoke network, all VPN tunnels terminate at the hub. The peers that connect to the hub are known as "spokes". The hub functions as a concentrator on the network, managing all VPN connections between the spokes. VPN traffic passes from one tunnel to the other through the hub.

You define a concentrator to include spokes in the hub-and-spoke configuration. You create the concentrator in *VPN > IPsec > Concentrator* and select *Create New*. A concentrator configuration specifies which spokes to include in an IPsec hub-and-spoke configuration.

<b>Concentrator Name</b>	Type a name for the concentrator.
<b>Available Tunnels</b>	A list of defined IPsec VPN tunnels. Select a tunnel from the list and then select the right arrow.
<b>Members</b>	A list of tunnels that are members of the concentrator. To remove a tunnel from the concentrator, select the tunnel and select the left arrow.

## IPsec Monitor

You can use the IPsec Monitor to view activity on IPsec VPN tunnels and start or stop those tunnels. The display provides a list of addresses, proxy IDs, and timeout information for all active tunnels, including tunnel mode and route-based (interface mode) tunnels.

To view the IPsec monitor, go to *VPN > Monitor > IPsec Monitor*.

For dialup VPNs, the list provides status information about the VPN tunnels established by dialup clients, and their IP addresses.

For static IP or dynamic DNS VPNs, the list provides status and IP addressing information about VPN tunnels, active or not, to remote peers that have static IP addresses or domain names. You can also start and stop individual tunnels from the list.

# Auto Key phase 1 parameters

This chapter provides detailed step-by-step procedures for configuring a FortiGate unit to accept a connection from a remote peer or dialup client. The phase 1 parameters identify the remote peer or clients and support authentication through preshared keys or digital certificates. You can increase access security further using peer identifiers, certificate distinguished names, group names, or the FortiGate extended authentication (XAuth) option for authentication purposes.

For more information on phase 1 parameters in the web-based manager, see [“Phase 1 configuration” on page 1625](#).

The following topics are included in this section:

- [Overview](#)
- [Defining the tunnel ends](#)
- [Choosing main mode or aggressive mode](#)
- [Authenticating the FortiGate unit](#)
- [Authenticating remote peers and clients](#)
- [Defining IKE negotiation parameters](#)
- [Using XAuth authentication](#)

## Overview

To configure IPsec phase 1 settings, go to *VPN > IPsec > Auto Key (IKE)* and select *Create Phase 1*. IPsec phase 1 settings define:

- the remote and local ends of the IPsec tunnel
- if phase 1 parameters are exchanged in multiple rounds with encrypted authentication information (main mode) or in a single message with authentication information that is not encrypted (aggressive mode)
- if a preshared key or digital certificates will be used to authenticate the FortiGate unit to the VPN peer or dialup client
- if the VPN peer or dialup client is required to authenticate to the FortiGate unit. A remote peer or dialup client can authenticate by peer ID or, if the FortiGate unit authenticates by certificate, it can authenticate by peer certificate.
- the IKE negotiation proposals for encryption and authentication
- optional XAuth authentication, which requires the remote user to enter a user name and password. A FortiGate VPN server can act as an XAuth server to authenticate dialup users. A FortiGate unit that is a dialup client can also be configured as an XAuth client to authenticate itself to the VPN server.

For all the phase 1 web-based manager fields, see [“Phase 1 configuration” on page 1625](#).

If you want to control how the IKE negotiation process controls traffic when there is no traffic, as well as the length of time the unit waits for negotiations to occur, use the `negotiation-timeout` and `auto-negotiation` commands in the CLI.

## Defining the tunnel ends

To begin defining the phase 1 configuration, go to *VPN > IPsec > Auto Key (IKE)* and select *Create Phase 1*. Enter a descriptive name for the VPN tunnel. This is particularly important if you will create several tunnels.

The phase 1 configuration mainly defines the ends of the IPsec tunnel. The remote end is the remote gateway with which the FortiGate unit exchanges IPsec packets. The local end is the FortiGate interface that sends and receives IPsec packets.

The remote gateway can be:

- a static IP address
- a domain name with a dynamic IP address
- a dialup client

A statically addressed remote gateway is the simplest to configure. You specify the IP address. Unless restricted in the security policy, either the remote peer or a peer on the network behind the FortiGate unit can bring up the tunnel.

If the remote peer has a domain name and subscribes to a dynamic DNS service, you need to specify only the domain name. The FortiGate unit performs a DNS query to determine the appropriate IP address. Unless restricted in the security policy, either the remote peer or a peer on the network behind the FortiGate unit can bring up the tunnel.

If the remote peer is a dialup client, only the dialup client can bring up the tunnel. The IP address of the client is not known until it connects to the FortiGate unit. This configuration is a typical way to provide a VPN for client PCs running VPN client software such as the FortiClient Endpoint Security application.

The local end of the VPN tunnel, the Local Interface, is the FortiGate interface that sends and receives the IPsec packets. This is usually the public interface of the FortiGate unit that is connected to the Internet (typically the WAN1 port). Packets from this interface pass to the private network through a security policy.

By default, the local VPN gateway is the IP address of the selected Local Interface. If you are configuring an interface mode VPN, you can optionally use a secondary IP address of the Local Interface as the local gateway.

## Choosing main mode or aggressive mode

The FortiGate unit and the remote peer or dialup client exchange phase 1 parameters in either Main mode or Aggressive mode. This choice does not apply if you use IKE version 2, which is available only for route-based configurations.

- In Main mode, the phase 1 parameters are exchanged in multiple rounds with encrypted authentication information
- In Aggressive mode, the phase 1 parameters are exchanged in single message with authentication information that is not encrypted.

Although Main mode is more secure, you must select Aggressive mode if there is more than one dialup phase 1 configuration for the interface IP address, and the remote VPN peer or client is authenticated using an identifier local ID). Descriptions of the peer options in this guide indicate whether Main or Aggressive mode is required.

## Choosing the IKE version

If you create a route-based VPN, you have the option of selecting IKE version 2. Otherwise, IKE version 1 is used.

IKEv2, defined in RFC 4306, simplifies the negotiation process that creates the security association (SA).

If you select IKEv2:

- There is no choice in Phase 1 of Aggressive or Main mode.
- FortiOS does not support Peer Options or Local ID.
- Extended Authentication (XAUTH) is not available.
- You can select only one DH Group.

## Authenticating the FortiGate unit

The FortiGate unit can authenticate itself to remote peers or dialup clients using either a pre-shared key or an RSA Signature (certificate).

### Authenticating the FortiGate unit with digital certificates

To authenticate the FortiGate unit using digital certificates, you must have the required certificates installed on the remote peer and on the FortiGate unit. The signed server certificate on one peer is validated by the presence of the root certificate installed on the other peer. If you use certificates to authenticate the FortiGate unit, you can also require the remote peers or dialup clients to authenticate using certificates.

For more information about obtaining and installing certificates, see the [FortiOS User Authentication guide](#).

#### To authenticate the FortiGate unit using digital certificates

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. Create a new phase 1 configuration or edit an existing phase 1 configuration.
3. Include appropriate entries as follows:

<b>Name</b>	Enter a name that reflects the origination of the remote connection. For interface mode, the name can be up to 15 characters long.
<b>Remote Gateway</b>	Select the nature of the remote connection. Each option changes the available fields you must configure. For more information, see <a href="#">“Defining the tunnel ends” on page 1638</a> .
<b>Local Interface</b>	Select the interface that is the local end of the IPsec tunnel. For more information, see <a href="#">“Defining the tunnel ends” on page 1638</a> . The local interface is typically the WAN1 port.

<b>Mode</b>	<p>Select a mode. It is easier to use aggressive mode.</p> <ul style="list-style-type: none"> <li>In Main mode, parameters are exchanged in multiple encrypted rounds.</li> <li>In Aggressive mode, parameters are exchanged in a single unencrypted message.</li> </ul> <p>Aggressive mode must be used when the remote VPN peer or client has a dynamic IP address, or the remote VPN peer or client will be authenticated using an identifier (local ID).</p> <p>For more information, see <a href="#">“Choosing main mode or aggressive mode” on page 1638</a>.</p>
<b>Authentication Method</b>	Select <i>RSA Signature</i> .
<b>Certificate Name</b>	<p>Select the name of the server certificate that the FortiGate unit will use to authenticate itself to the remote peer or dialup client during phase 1 negotiations.</p> <p>You must obtain and load the required server certificate before this selection. See the <a href="#">FortiOS User Authentication guide</a>. If you have not loaded any certificates, use the certificate named <i>Fortinet_Factory</i>.</p>
<b>Peer Options</b>	<p>Peer options define the authentication requirements for remote peers or dialup clients. They are not for your FortiGate unit itself.</p> <p>See <a href="#">“Authenticating remote peers and clients” on page 1642</a>.</p>
<b>Advanced</b>	<p>You can use the default settings for most phase 1 configurations. Changes are required only if your network requires them. These settings includes IKE version, DNS server, P1 proposal encryption and authentication settings, and XAuth settings. See <a href="#">“Defining IKE negotiation parameters” on page 1646</a>.</p>

- If you are configuring authentication parameters for a dialup user group, optionally define extended authentication (XAuth) parameters in the Advanced section. See [“Using the FortiGate unit as an XAuth server” on page 1651](#).
- Select *OK*.

## Authenticating the FortiGate unit with a pre-shared key

The simplest way to authenticate a FortiGate unit to its remote peers or dialup clients is by means of a pre-shared key. This is less secure than using certificates, especially if it is used alone, without requiring peer IDs or extended authentication (XAuth). Also, you need to have a secure way to distribute the pre-shared key to the peers.

If you use pre-shared key authentication alone, all remote peers and dialup clients must be configured with the same pre-shared key. Optionally, you can configure remote peers and dialup clients with unique pre-shared keys. On the FortiGate unit, these are configured in user accounts, not in the phase\_1 settings. For more information, see [“Enabling VPN access with user accounts and pre-shared keys” on page 1645](#).

The pre-shared key must contain at least 6 printable characters and best practices dictate that it be known only to network administrators. For optimum protection against currently known attacks, the key must consist of a minimum of 16 randomly chosen alphanumeric characters.

If you authenticate the FortiGate unit using a pre-shared key, you can require remote peers or dialup clients to authenticate using peer IDs, but not client certificates.



## To authenticate the FortiGate unit with a pre-shared key

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. Create a new phase 1 configuration or edit an existing phase 1 configuration.
3. Include appropriate entries as follows:

<b>Name</b>	Enter a name that reflects the origination of the remote connection.
<b>Remote Gateway</b>	Select the nature of the remote connection. For more information, see <a href="#">“Defining the tunnel ends” on page 1638</a> .
<b>Local Interface</b>	Select the interface that is the local end of the IPsec tunnel. For more information, see <a href="#">“Defining the tunnel ends” on page 1638</a> . The local interface is typically the WAN1 port.
<b>Mode</b>	<p>Select Main or Aggressive mode.</p> <ul style="list-style-type: none"><li>• In Main mode, the phase 1 parameters are exchanged in multiple rounds with encrypted authentication information.</li><li>• In Aggressive mode, the phase 1 parameters are exchanged in single message with authentication information that is not encrypted.</li></ul> <p>When the remote VPN peer or client has a dynamic IP address, or the remote VPN peer or client will be authenticated using an identifier (local ID), you must select Aggressive mode if there is more than one dialup phase 1 configuration for the interface IP address.</p> <p>For more information, see <a href="#">“Choosing main mode or aggressive mode” on page 1638</a>.</p>
<b>Authentication Method</b>	Select <i>Pre-shared Key</i> .
<b>Pre-shared Key</b>	Enter the preshared key that the FortiGate unit will use to authenticate itself to the remote peer or dialup client during phase 1 negotiations. You must define the same value at the remote peer or client. The key must contain at least 6 printable characters and best practices dictate that it only be known by network administrators. For optimum protection against currently known attacks, the key must consist of a minimum of 16 randomly chosen alphanumeric characters.
<b>Peer options</b>	Peer options define the authentication requirements for remote peers or dialup clients, not for the FortiGate unit itself. You can require the use of peer IDs, but not client certificates. For more information, see <a href="#">“Authenticating remote peers and clients” on page 1642</a> .
<b>Advanced</b>	You can retain the default settings unless changes are needed to meet your specific requirements. See <a href="#">“Defining IKE negotiation parameters” on page 1646</a> .

4. If you are configuring authentication parameters for a dialup user group, optionally define extended authentication (XAuth) parameters. See [“Using the FortiGate unit as an XAuth server” on page 1651](#).
5. Select *OK*.

## Authenticating remote peers and clients

Certificates or pre-shared keys restrict who can access the VPN tunnel, but they do not identify or authenticate the remote peers or dialup clients. You have the following options for authentication:

**Table 73:** Methods of authenticating remote VPN peers

Certificates or Pre-shared key	Local ID	User account pre-shared keys	Reference
Certificates			See <a href="#">“Enabling VPN access for specific certificate holders”</a> on page 1642
Either	<b>X</b>		See <a href="#">“Enabling VPN access by peer identifier”</a> on page 1644
Pre-shared key		<b>X</b>	See <a href="#">“Enabling VPN access with user accounts and pre-shared keys”</a> on page 1645
Pre-shared key	<b>X</b>	<b>X</b>	See <a href="#">“Enabling VPN access with user accounts and pre-shared keys”</a> on page 1645

For authentication of users of the remote peer or dialup client device, see [“Using XAuth authentication”](#) on page 1650.

### Enabling VPN access for specific certificate holders

When a VPN peer or dialup client is configured to authenticate using digital certificates, it sends the DN of its certificate to the FortiGate unit. This DN can be used to allow VPN access for the certificate holder. That is, a FortiGate unit can be configured to deny connections to all remote peers and dialup clients except the one having the specified DN.

#### Before you begin

The following procedures assume that you already have an existing phase 1 configuration (see [“Authenticating the FortiGate unit with digital certificates”](#) on page 1639). Follow the procedures below to add certificate-based authentication parameters to the existing configuration.

Before you begin, you must obtain the certificate DN of the remote peer or dialup client. If you are using the FortiClient application as a dialup client, refer to [FortiClient online Help](#) for information about how to view the certificate DN. To view the certificate DN of a FortiGate unit, see [“To view server certificate information and obtain the local DN”](#) on page 1643.

Use the `config user peer` CLI command to load the DN value into the FortiGate configuration. For example, if a remote VPN peer uses server certificates issued by your own organization, you would enter information similar to the following:

```
config user peer
 edit DN_FG1000
 set cn 192.168.2.160
 set cn-type ipv4
 end
```

The value that you specify to identify the entry (for example, DN\_FG1000) is displayed in the Accept this peer certificate only list in the IPsec phase 1 configuration when you return to the web-based manager.

If the remote VPN peer has a CA-issued certificate to support a higher level of credibility, you would enter information similar to the following:

```
config user peer
 edit CA_FG1000
 set ca CA_Cert_1
 set subject FG1000_at_site1
 end
```

The value that you specify to identify the entry (for example, CA\_FG1000) is displayed in the Accept this peer certificate only list in the IPsec phase 1 configuration when you return to the web-based manager. For more information about these CLI commands, see the “user” chapter of the [FortiGate CLI Reference](#).

A group of certificate holders can be created based on existing user accounts for dialup clients. To create the user accounts for dialup clients, see the “User” chapter of the [FortiGate Administration Guide](#). To create the certificate group afterward, use the `config user peergrp` CLI command. See the “user” chapter of the [FortiGate CLI Reference](#).

#### To view server certificate information and obtain the local DN

1. Go to *System > Certificates > Local Certificates*.
2. Note the CN value in the *Subject* field (for example, CN = 172.16.10.125, CN = info@fortinet.com, or CN = www.example.com).

#### To view CA root certificate information and obtain the CA certificate name

1. Go to *System > Certificates > CA Certificates*.
2. Note the value in the *Name* column (for example, CA\_Cert\_1).

### Configuring certificate authentication for a VPN

With peer certificates loaded, peer users and peer groups defined, you can configure your VPN to authenticate users by certificate.

#### To enable access for a specific certificate holder or a group of certificate holders

1. At the FortiGate VPN server, go to *VPN > IPsec > Auto Key (IKE)*.
2. In the list of defined configurations, select the phase 1 configuration and edit it.
3. From the *Authentication Method* list, select *RSA Signature*.
4. From the *Certificate Name* list, select the name of the server certificate that the FortiGate unit will use to authenticate itself to the remote peer or dialup client
5. Under *Peer Options*, select one of these options:
  - To accept a specific certificate holder, select *Accept this peer certificate only* and select the name of the certificate that belongs to the remote peer or dialup client. The certificate DN must be added to the FortiGate configuration through CLI commands before it can be selected here. See “Before you begin” on page 1642.
  - To accept dialup clients who are members of a certificate group, select *Accept this peer certificate group only* and select the name of the group. The group must be added to the FortiGate configuration through CLI commands before it can be selected here. See “Before you begin” on page 1642.
6. If you want the FortiGate VPN server to supply the DN of a local server certificate for authentication purposes, select *Advanced* and then from the *Local ID* list, select the DN of the certificate that the FortiGate VPN server is to use.
7. Select *OK*.

## Enabling VPN access by peer identifier

Whether you use certificates or pre-shared keys to authenticate the FortiGate unit, you can require that remote peers or clients have a particular peer ID. This adds another piece of information that is required to gain access to the VPN. More than one FortiGate/FortiClient dialup client may connect through the same VPN tunnel when the dialup clients share a preshared key and assume the same identifier.

A peer ID, also called local ID, can be up to 63 characters long containing standard regular expression characters. Local ID is set in phase1 Aggressive Mode configuration.

You cannot require a peer ID for a remote peer or client that uses a pre-shared key and has a static IP address.

### To authenticate remote peers or dialup clients using one peer ID

1. At the FortiGate VPN server, go to *VPN > IPsec > Auto Key (IKE)*.
2. In the list, select a phase 1 configuration and edit its parameters.
3. Select *Aggressive* mode in any of the following cases:
  - the FortiGate VPN server authenticates a FortiGate dialup client that uses a dedicated tunnel
  - a FortiGate unit has a dynamic IP address and subscribes to a dynamic DNS service
  - FortiGate/FortiClient dialup clients sharing the same preshared key and local ID connect through the same VPN tunnel
4. Select *Accept this peer ID* and type the identifier into the corresponding field.
5. Select *OK*.

### To assign an identifier (local ID) to a FortiGate unit

Use this procedure to assign a peer ID to a FortiGate unit that acts as a remote peer or dialup client.

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. In the list, select a phase 1 configuration and edit its parameters.
3. Select *Advanced*.
4. In the *Local ID* field, type the identifier that the FortiGate unit will use to identify itself.
5. Set *Mode* to *Aggressive* if any of the following conditions apply:
  - The FortiGate unit is a dialup client that will use a unique ID to connect to a FortiGate dialup server through a dedicated tunnel.
  - The FortiGate unit has a dynamic IP address, subscribes to a dynamic DNS service, and will use a unique ID to connect to the remote VPN peer through a dedicated tunnel.
  - The FortiGate unit is a dialup client that shares the specified ID with multiple dialup clients to connect to a FortiGate dialup server through the same tunnel.
6. Select *OK*.

### To configure the FortiClient application

Follow this procedure to add a peer ID to an existing FortiClient configuration:

1. Start the FortiClient application.
2. Go to *VPN > Connections*, select the existing configuration.
3. Select *Advanced > Edit > Advanced*.
4. Under *Policy*, select *Config*.

5. In the *Local ID* field, type the identifier that will be shared by all dialup clients. This value must match the *Accept this peer ID* value that you specified previously in the phase 1 gateway configuration on the FortiGate unit.
6. Select *OK* to close all dialog boxes.
7. Configure all dialup clients the same way using the same preshared key and local ID.

## Enabling VPN access with user accounts and pre-shared keys

You can permit access only to remote peers or dialup clients that have pre-shared keys and/or peer IDs configured in user accounts on the FortiGate unit.

If you want two VPN peers (or a FortiGate unit and a dialup client) to accept reciprocal connections based on peer IDs, you must enable the exchange of their identifiers when you define the phase 1 parameters.

The following procedures assume that you already have an existing phase 1 configuration (see [“Authenticating the FortiGate unit with digital certificates” on page 1639](#)). Follow the procedures below to add ID checking to the existing configuration.

Before you begin, you must obtain the identifier (local ID) of the remote peer or dialup client. If you are using the FortiClient Endpoint Security application as a dialup client, refer to the [Authenticating FortiClient Dialup Clients Technical Note](#) to view or assign an identifier. To assign an identifier to a FortiGate dialup client or a FortiGate unit that has a dynamic IP address and subscribes to a dynamic DNS service, see [“To assign an identifier \(local ID\) to a FortiGate unit” on page 1644](#).

If required, a dialup user group can be created from existing user accounts for dialup clients. To create the user accounts and user groups, see the [User Authentication](#) chapter of The Handbook.

The following procedure supports FortiGate/FortiClient dialup clients that use unique preshared keys and/or peer IDs. The client must have an account on the FortiGate unit and be a member of the dialup user group.

The dialup user group must be added to the FortiGate configuration before it can be selected. For more information, see the [User Authentication](#) chapter of The Handbook.

The FortiGate dialup server compares the local ID that you specify at each dialup client to the FortiGate user-account user name. The dialup-client preshared key is compared to a FortiGate user-account password.

### To authenticate dialup clients using unique preshared keys and/or peer IDs

1. At the FortiGate VPN server, go to *VPN > IPsec > Auto Key (IKE)*.
2. In the list, select the *Edit* icon of a phase 1 configuration to edit its parameters.
3. If the clients have unique peer IDs, set *Mode* to *Aggressive*.
4. Clear the *Pre-shared Key* field.  
The user account password will be used as the preshared key.
5. Select *Accept peer ID in dialup group* and then select the group name from the list of user groups.
6. Select *OK*.

Follow this procedure to add a unique pre-shared key and unique peer ID to an existing FortiClient configuration.

### To configure FortiClient - pre-shared key and peer ID

1. Start the FortiClient Endpoint Security application.
2. Go to *VPN > Connections*, select the existing configuration.

3. Select *Advanced > Edit*.
4. In the *Preshared Key* field, type the FortiGate password that belongs to the dialup client (for example, 1234546).  
The user account password will be used as the preshared key.
5. Select *Advanced*.
6. Under *Policy*, select *Config*.
7. In the *Local ID* field, type the FortiGate user name that you assigned previously to the dialup client (for example, FortiClient1).
8. Select *OK* to close all dialog boxes.

Configure all FortiClient dialup clients this way using unique preshared keys and local IDs.

Follow this procedure to add a unique pre-shared key to an existing FortiClient configuration.

#### **To configure FortiClient - preshared key only**

1. Start the FortiClient Endpoint Security application.
2. Go to *VPN > Connections*, select the existing configuration
3. Select *Advanced > Edit*.
4. In the *Preshared Key* field, type the user name, followed by a “+” sign, followed by the password that you specified previously in the user account settings on the FortiGate unit (for example, FC2+1FG6LK)
5. Select *OK* to close all dialog boxes.

Configure all the FortiClient dialup clients this way using their unique peer ID and pre-shared key values.

## **Defining IKE negotiation parameters**

In phase 1, the two peers exchange keys to establish a secure communication channel between them. As part of the phase 1 process, the two peers authenticate each other and negotiate a way to encrypt further communications for the duration of the session. For more information see [“Authenticating remote peers and clients” on page 1642](#). The P1 Proposal parameters select the encryption and authentication algorithms that are used to generate keys for protecting negotiations.

The IKE negotiation parameters determine:

- which encryption algorithms may be applied for converting messages into a form that only the intended recipient can read
- which authentication hash may be used for creating a keyed hash from a preshared or private key
- which Diffie-Hellman group (DH Group) will be used to generate a secret session key

Phase 1 negotiations (in main mode or aggressive mode) begin as soon as a remote VPN peer or client attempts to establish a connection with the FortiGate unit. Initially, the remote peer or dialup client sends the FortiGate unit a list of potential cryptographic parameters along with a session ID. The FortiGate unit compares those parameters to its own list of advanced phase 1 parameters and responds with its choice of matching parameters to use for authenticating and encrypting packets. The two peers handle the exchange of encryption keys between them, and authenticate the exchange through a preshared key or a digital signature.

## Generating keys to authenticate an exchange

The FortiGate unit supports the generation of secret session keys automatically using a Diffie-Hellman algorithm. These algorithms are defined in RFC 2409. The *Keylife* setting in the *P1 Proposal* area determines the amount of time before the phase 1 key expires. Phase 1 negotiations are rekeyed automatically when there is an active security association. See [“Dead peer detection” on page 1650](#).

You can enable or disable automatic rekeying between IKE peers through the `phase1-rekey` attribute of the `config system global` CLI command. For more information, see the “system” chapter of the [FortiGate CLI Reference](#).



When in FIPS-CC mode, the FortiGate unit requires DH key exchange to use values at least 3072 bits long. However most browsers need the key size set to 1024. You can set the minimum size of the DH keys in the CLI.

```
config system global
 set dh-params 3072
end
```

---

When you use a preshared key (shared secret) to set up two-party authentication, the remote VPN peer or client and the FortiGate unit must both be configured with the same preshared key. Each party uses a session key derived from the Diffie-Hellman exchange to create an authentication key, which is used to sign a known combination of inputs using an authentication algorithm (such as HMAC-MD5, HMAC-SHA-1, or HMAC-SHA-256). Hash-based Message Authentication Code (HMAC) is a method for calculating an authentication code using a hash function plus a secret key, and is defined in RFC 2104. Each party signs a different combination of inputs and the other party verifies that the same result can be computed.



SHA-256, SHA-384 and SHA-512 are not accelerated by some FortiASIC processors (including FortiASIC network processors and security processors). As a result, using SHA-256, SHA-384 and SHA-512 may reduce the performance of the FortiGate unit more significantly than SHA-1 which is accelerated by all FortiASIC processors.

---

When you use preshared keys to authenticate VPN peers or clients, you must distribute matching information to all VPN peers and/or clients whenever the preshared key changes.

As an alternative, the remote peer or dialup client and FortiGate unit can exchange digital signatures to validate each other’s identity with respect to their public keys. In this case, the required digital certificates must be installed on the remote peer and on the FortiGate unit. By exchanging certificate DNs, the signed server certificate on one peer is validated by the presence of the root certificate installed on the other peer.

The following procedure assumes that you already have a phase 1 definition that describes how remote VPN peers and clients will be authenticated when they attempt to connect to a local FortiGate unit. For information about the Local ID and XAuth options, see [“Enabling VPN access with user accounts and pre-shared keys” on page 1645](#) and [“Using the FortiGate unit as an XAuth server” on page 1651](#). Follow this procedure to add IKE negotiation parameters to the existing definition.

## Defining IKE negotiation parameters

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. In the list, select the *Edit* button to edit the phase 1 parameters for a particular remote gateway.

3. Select *Advanced* and include appropriate entries and select *OK*:

---

**P1 Proposal**

Select the encryption and authentication algorithms that will be used to generate keys for protecting negotiations.

Add or delete encryption and authentication algorithms as required. Select a minimum of one and a maximum of three combinations. The remote peer must be configured to use at least one of the proposals that you define.

You can select any of these symmetric-key algorithms:

- DES-Digital Encryption Standard, a 64-bit block algorithm that uses a 56-bit key.
- 3DES-Triple-DES, in which plain text is encrypted three times by three keys.
- AES128-A 128-bit block algorithm that uses a 128-bit key.
- AES192-A 128-bit block algorithm that uses a 192-bit key.
- AES256-A 128-bit block algorithm that uses a 256-bit key.

You can select one of the following message digests to check the authenticity of messages during phase 1 negotiations:

- MD5-Message Digest 5, the hash algorithm developed by RSA Data Security.
- SHA1-Secure Hash Algorithm 1, which produces a 160-bit message digest.
- SHA-256 Secure Hash Algorithm 256, which produces a 256-bit message digest
- SHA-384 Secure Hash Algorithm 384, which produces a 384-bit message digest
- SHA-512 Secure Hash Algorithm 512, which produces a 512-bit message digest

To specify a third combination, use the add button beside the fields for the second combination.

SHA-256, SHA-384 and SHA-512 are not accelerated by some FortiASIC processors (including FortiASIC network processors and security processors). As a result, using SHA-256, SHA-384 and SHA-512 may reduce the performance of the FortiGate unit more significantly than SHA-1 which is accelerated by all FortiASIC processors.

---



<b>DH Group</b>	<p>Select one or more Diffie-Hellman groups from DH group 1, 2, and 5. When using aggressive mode, DH groups cannot be negotiated.</p> <p>If both VPN peers (or a VPN server and its client) have static IP addresses and use aggressive mode, select a single DH group. The setting on the FortiGate unit must be identical to the setting on the remote peer or dialup client.</p> <p>When the remote VPN peer or client has a dynamic IP address and uses aggressive mode, select up to three DH groups on the FortiGate unit and one DH group on the remote peer or dialup client. The setting on the remote peer or dialup client must be identical to one of the selections on the FortiGate unit.</p> <p>If the VPN peer or client employs main mode, you can select multiple DH groups. At least one of the settings on the remote peer or dialup client must be identical to the selections on the FortiGate unit.</p>
<b>Keylife</b>	<p>Type the amount of time (in seconds) that will be allowed to pass before the IKE encryption key expires. When the key expires, a new key is generated without interrupting service. The keylife can be from 120 to 172800 seconds.</p>
<b>Nat-traversal</b>	<p>Enable this option if a NAT device exists between the local FortiGate unit and the VPN peer or client. The local FortiGate unit and the VPN peer or client must have the same NAT traversal setting (both selected or both cleared). When in doubt, enable NAT-traversal. See <a href="#">“NAT traversal” on page 1649</a>.</p>
<b>Keepalive Frequency</b>	<p>If you enabled NAT traversal, enter a keepalive frequency setting. The value represents an interval from 0 to 900 seconds where the connection will be maintained with no activity. For additional security this value must be as low as possible. See <a href="#">“NAT keepalive frequency” on page 1650</a>.</p>
<b>Dead Peer Detection</b>	<p>Enable this option to reestablish VPN tunnels on idle connections and clean up dead IKE peers if required. This feature minimizes the traffic required to check if a VPN peer is available or unavailable (dead). See <a href="#">“Dead peer detection” on page 1650</a>.</p>

## NAT traversal

Network Address Translation (NAT) is a way to convert private IP addresses to publicly routable Internet addresses and vice versa. When an IP packet passes through a NAT device, the source or destination address in the IP header is modified. FortiGate units support NAT version 1 (encapsulate on port 500 with non-IKE marker), version 3 (encapsulate on port 4500 with non-ESP marker), and compatible versions.

NAT cannot be performed on IPsec packets in ESP tunnel mode because the packets do not contain a port number. As a result, the packets cannot be demultiplexed. To work around this, the FortiGate unit provides a way to protect IPsec packet headers from NAT modifications. When the Nat-traversal option is enabled, outbound encrypted packets are wrapped inside a UDP IP header that contains a port number. This extra encapsulation allows NAT devices to change the port number without modifying the IPsec packet directly.

To provide the extra layer of encapsulation on IPsec packets, the Nat-traversal option must be enabled whenever a NAT device exists between two FortiGate VPN peers or a FortiGate unit and a dialup client such as FortiClient. On the receiving end, the FortiGate unit or FortiClient removes the extra layer of encapsulation before decrypting the packet.

## NAT keepalive frequency

When a NAT device performs network address translation on a flow of packets, the NAT device determines how long the new address will remain valid if the flow of traffic stops (for example, the connected VPN peer may be idle). The device may reclaim and reuse a NAT address when a connection remains idle for too long.

To work around this, when you enable NAT traversal specify how often the FortiGate unit sends periodic keepalive packets through the NAT device in order to ensure that the NAT address mapping does not change during the lifetime of a session. To be effective, the keepalive interval must be smaller than the session lifetime value used by the NAT device.

The keepalive packet is a 138-byte ISAKMP exchange.

## Dead peer detection

Sometimes, due to routing issues or other difficulties, the communication link between a FortiGate unit and a VPN peer or client may go down. Packets could be lost if the connection is left to time out on its own. The FortiGate unit provides a mechanism called Dead Peer Detection, sometimes referred to as gateway detection or ping server, to prevent this situation and reestablish IKE negotiations automatically before a connection times out: the active phase 1 security associations are caught and renegotiated (rekeyed) before the phase 1 encryption key expires.

By default, Dead Peer Detection sends probe messages every five seconds by default (see `dpd-retryinterval` in the [FortiGate CLI Reference](#)). If you are experiencing high network traffic, you can experiment with increasing the ping interval. However longer intervals will require more traffic to detect dead peers which will result in more traffic.

In the web-based manager, the Dead Peer Detection option can be enabled when you define advanced phase 1 options. The `config vpn ipsec phase1` CLI command supports additional options for specifying a retry count and a retry interval.

For more information about these commands and the related `config router gwdetect` CLI command, see the [FortiGate CLI Reference](#).

For example, enter the following CLI commands to configure dead peer detection on the existing IPsec Phase1 configuration called `test` to use 15 second intervals and to wait for 3 missed attempts before declaring the peer dead and taking action.

```
config vpn ipsec phase1
 edit test
 set dpd enable
 set dpd-retryinterval 15
 set dpd-retrycount 3
 next
end
```

## Using XAuth authentication

Extended authentication (XAuth) increases security by requiring the remote dialup client user to authenticate in a separate exchange at the end of phase 1. XAuth draws on existing FortiGate user group definitions and uses established authentication mechanisms such as PAP, CHAP, RADIUS and LDAP to authenticate dialup clients. You can configure a FortiGate unit to function either as an XAuth server or an XAuth client. If the server or client is attempting a connection using XAuth and the other end is not using XAuth, the failed connection attempts that are logged will not specify XAuth as the reason.

## Using the FortiGate unit as an XAuth server

A FortiGate unit can act as an XAuth server for dialup clients. When the phase 1 negotiation completes, the FortiGate unit challenges the user for a user name and password. It then forwards the user's credentials to an external RADIUS or LDAP server for verification.

If the user records on the RADIUS server have suitably configured Framed-IP-Address fields, you can assign client virtual IP addresses by XAuth instead of from a DHCP address range. See [“Assigning VIPs by RADIUS user group” on page 1713](#).

The authentication protocol to use for XAuth depends on the capabilities of the authentication server and the XAuth client:

- Select PAP whenever possible.
- You must select PAP for all implementations of LDAP and some implementations of Microsoft RADIUS.
- Select AUTO when the authentication server supports CHAP but the XAuth client does not. The FortiGate unit will use PAP to communicate with the XAuth client and CHAP to communicate with the authentication server.

Before you begin, create user accounts and user groups to identify the dialup clients that need to access the network behind the FortiGate dialup server. If password protection will be provided through an external RADIUS or LDAP server, you must configure the FortiGate dialup server to forward authentication requests to the authentication server. For information about these topics, see the [FortiGate User Authentication Guide](#).

### To authenticate a dialup user group using XAuth settings

1. At the FortiGate dialup server, go to *VPN > IPsec > Auto Key (IKE)*.
2. In the list, select the *Edit* icon of a phase 1 configuration to edit its parameters for a particular remote gateway.
3. Select *Advanced*.
4. Under *XAuth*, select *Enable as Server*.
5. The *Server Type* setting determines the type of encryption method to use between the XAuth client, the FortiGate unit and the authentication server. Select one of the following options:
  - *PAP*—Password Authentication Protocol.
  - *CHAP*— Challenge-Handshake Authentication Protocol.
  - *AUTO*—Use PAP between the XAuth client and the FortiGate unit, and CHAP between the FortiGate unit and the authentication server.
6. From the *User Group* list, select the user group that needs to access the private network behind the FortiGate unit. The group must be added to the FortiGate configuration before it can be selected here.
7. Select *OK*.

## Using the FortiGate unit as an XAuth client

If the FortiGate unit acts as a dialup client, the remote peer, acting as an XAuth server, might require a user name and password. You can configure the FortiGate unit as an XAuth client, with its own user name and password, which it provides when challenged.

### To configure the FortiGate dialup client as an XAuth client

1. At the FortiGate dialup client, go to *VPN > IPsec > Auto Key (IKE)*.
2. In the list, select a phase 1 configuration and select *Edit*.
3. Select *Advanced*.

4. Under *XAuth*, select *Enable as Client*.
5. In the *Username* field, type the FortiGate PAP, CHAP, RADIUS, or LDAP user name that the FortiGate XAuth server will compare to its records when the FortiGate XAuth client attempts to connect.
6. In the *Password* field, type the password to associate with the user name.
7. Select *OK*.

# Phase 2 parameters

This section describes the phase 2 parameters that are required to establish communication through a VPN.

The following topics are included in this section:

- [Basic phase 2 settings](#)
- [Advanced phase 2 settings](#)
- [Configure the phase 2 parameters](#)

## Basic phase 2 settings

After IPsec VPN phase 1 negotiations complete successfully, phase 2 negotiation begins. Phase 2 parameters define the algorithms that the FortiGate unit can use to encrypt and transfer data for the remainder of the session. The basic phase 2 settings associate IPsec phase 2 parameters with a phase 1 configuration.

When defining phase 2 parameters, you can choose any set of phase 1 parameters to set up a secure connection and authenticate the remote peer.

For more information on phase 2 settings in the web-based manager, see [“Phase 2 configuration” on page 1629](#)

## Advanced phase 2 settings

The following additional advanced phase 2 settings are available to enhance the operation of the tunnel:

- [P2 Proposals](#)
- [Replay detection](#)
- [Perfect forward secrecy \(PFS\)](#)
- [Keylife](#)
- [Quick mode selectors](#)

### P2 Proposals

In phase 2, the VPN peer or client and the FortiGate unit exchange keys again to establish a secure communication channel. The P2 Proposal parameters select the encryption and authentication algorithms needed to generate keys for protecting the implementation details of Security Associations (SAs). The keys are generated automatically using a Diffie-Hellman algorithm.

### Replay detection

IPsec tunnels can be vulnerable to replay attacks. Replay detection enables the FortiGate unit to check all IPsec packets to see if they have been received before. If any encrypted packets arrive out of order, the FortiGate unit discards them.

## Perfect forward secrecy (PFS)

By default, phase 2 keys are derived from the session key created in phase 1. Perfect forward secrecy forces a new Diffie-Hellman exchange when the tunnel starts and whenever the phase 2 keylife expires, causing a new key to be generated each time. This exchange ensures that the keys created in phase 2 are unrelated to the phase 1 keys or any other keys generated automatically in phase 2.

## Keylife

The Keylife setting sets a limit on the length of time that a phase 2 key can be used. The default units are seconds. Alternatively, you can set a limit on the number of kilobytes (KB) of processed data, or both. If you select both, the key expires when either the time has passed or the number of KB have been processed. When the phase 2 key expires, a new key is generated without interrupting service.

## Auto-negotiate

By default, the phase 2 security association (SA) is not negotiated until a peer attempts to send data. The triggering packet and some subsequent packets are dropped until the SA is established. Applications normally resend this data, so there is no loss, but there might be a noticeable delay in response to the user.

Automatically establishing the SA can also be important for a dialup peer. This ensures that the VPN tunnel is available for peers at the server end to initiate traffic to the dialup peer. Otherwise, the VPN tunnel does not exist until the dialup peer initiates traffic.

When enabled, auto-negotiate initiates the phase 2 SA negotiation automatically, repeating every five seconds until the SA is established.

The auto-negotiate feature is available only through the Command Line Interface (CLI). Use the following commands to enable it.

```
config vpn ipsec phase2
 edit <phase2_name>
 set auto-negotiate enable
 end
```

If the tunnel goes down, the auto-negotiate feature will attempt to re-establish it. However, the Autokey Keep Alive feature is a better method to ensure your VPN remains up.

## Autokey Keep Alive

The phase 2 SA has a fixed duration. If there is traffic on the VPN as the SA nears expiry, a new SA is negotiated and the VPN switches to the new SA without interruption. If there is no traffic, the SA expires and the VPN tunnel goes down. A new SA will not be generated until there is traffic.

The Autokey Keep Alive option ensures that a new SA is negotiated even if there is no traffic so that the VPN tunnel stays up.

## DHCP-IPsec

Select this option if the FortiGate unit assigns VIP addresses to FortiClient dialup clients through a DHCP server or relay. This option is available only if the Remote Gateway in the phase 1 configuration is set to Dialup User and it works only on policy-based VPNs.

With the DHCP-IPsec option, the FortiGate dialup server acts as a proxy for FortiClient dialup clients that have VIP addresses on the subnet of the private network behind the FortiGate unit. In this case, the FortiGate dialup server acts as a proxy on the local private network for the FortiClient dialup client. When a host on the network behind the dialup server issues an ARP request that corresponds to the device MAC address of the FortiClient host (when a remote server sends an ARP to the local FortiClient dialup client), the FortiGate unit answers the ARP request on behalf of the FortiClient host and forwards the associated traffic to the FortiClient host through the tunnel.

This feature prevents the VIP address assigned to the FortiClient dialup client from causing possible arp broadcast problems — the normal and VIP addresses can confuse some network switches by two addresses having the same MAC address.

## Quick mode selectors

Quick Mode selectors determine which IP addresses can perform IKE negotiations to establish a tunnel. By only allowing authorized IP addresses access to the VPN tunnel, the network is more secure.

The default settings are as broad as possible: any IP address or configured address object, using any protocol, on any port.



While the drop down menus for specifying an address also show address groups, the use of address groups is not supported.

To make it easy to determine if one of the choices in the drop down menu is an address or an address group the two types of objects have been broken into sections with the address groups at the bottom of the list.

---

When configuring Quick Mode selector *Source Address* and *Destination address*, valid options include IPv4 and IPv6 single addresses, IPv4 subnet, or IPv6 subnet. For more information on IPv6 IPsec VPN, see [“Overview of IPv6 IPsec support” on page 1772](#).

There are some configurations that require specific selectors:

- the VPN peer is a third-party device that uses specific phase2 selectors
- the FortiGate unit connects as a dialup client to another FortiGate unit, in which case you must specify a source IP address, IP address range or subnet

With FortiOS VPNs, your network has multiple layers of security, with quick mode selectors being an important line of defence.

- Routes guide traffic from one IP address to another.
- Phase 1 and phase 2 connection settings ensure there is a valid remote end point for the VPN tunnel that agrees on the encryption and parameters.
- Quick mode selectors allow IKE negotiations only for allowed peers.
- Security policies control which IP addresses can connect to the VPN.
- Security policies also control what protocols are allowed over the VPN along with any bandwidth limiting.

## Configure the phase 2 parameters

If you are creating a hub-and-spoke configuration or an Internet-browsing configuration, you may have already started defining some of the required phase 2 parameters. If so, edit the existing definition to complete the configuration.

## Specifying the phase 2 parameters

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. Select *Create Phase 2*.
3. Enter a *Name* for the phase 2 configuration, and select a *Phase 1* configuration from the drop-down list.
4. Select *Advanced*.
5. Include appropriate entries and select *OK*:

---

<b>P2 Proposal</b>	<p>Select the encryption and authentication algorithms that will be used to change data into encrypted code.</p> <p>Add or delete encryption and authentication algorithms as required. Select a minimum of one and a maximum of three combinations. The remote peer must be configured to use at least one of the proposals that you define.</p> <p>It is invalid to set both <i>Encryption</i> and <i>Authentication</i> to null.</p>
<b>Encryption</b>	<p>Select a symmetric-key algorithms:</p> <p><b>NULL</b> — Do not use an encryption algorithm.</p> <p><b>DES</b> — Digital Encryption Standard, a 64-bit block algorithm that uses a 56-bit key.</p> <p><b>3DES</b> — Triple-DES; plain text is encrypted three times by three keys.</p> <p><b>AES128</b> — A 128-bit block algorithm that uses a 128-bit key.</p> <p><b>AES192</b> — A 128-bit block algorithm that uses a 192-bit key.</p> <p><b>AES256</b> — A 128-bit block algorithm that uses a 256-bit key.</p>
<b>Authentication</b>	<p>You can select either of the following message digests to check the authenticity of messages during an encrypted session:</p> <ul style="list-style-type: none"><li>• <b>NULL</b> — Do not use a message digest.</li><li>• <b>MD5</b> — Message Digest 5.</li><li>• <b>SHA1</b> — Secure Hash Algorithm 1 - a 160-bit message digest.</li></ul> <p>To specify one combination only, set the <i>Encryption</i> and <i>Authentication</i> options of the second combination to NULL. To specify a third combination, use the <i>Add</i> button beside the fields for the second combination.</p>
<b>Enable replay detection</b>	<p>Optionally enable or disable replay detection. Replay attacks occur when an unauthorized party intercepts a series of IPsec packets and replays them back into the tunnel.</p>
<b>Enable perfect forward secrecy (PFS)</b>	<p>Enable or disable PFS. Perfect forward secrecy (PFS) improves security by forcing a new Diffie-Hellman exchange whenever keylife expires.</p>
<b>DH Group</b>	<p>Select one Diffie-Hellman group (1, 2, 5, or 14). The remote peer or dialup client must be configured to use the same group.</p>

---



<b>Keylife</b>	Select the method for determining when the phase 2 key expires: <i>Seconds</i> , <i>KBytes</i> , or <i>Both</i> . If you select <i>Both</i> , the key expires when either the time has passed or the number of KB have been processed. The range is from 120 to 172800 seconds, or from 5120 to 2147483648 KB.
<b>Autokey Keep Alive</b>	Enable the option if you want the tunnel to remain active when no data is being processed.
<b>DHCP-IPsec</b>	<p>Select <i>Enable</i> if the FortiGate unit acts as a dialup server and FortiGate DHCP server or relay will be used to assign VIP addresses to FortiClient dialup clients. The DHCP server or relay parameters must be configured separately.</p> <p>If the FortiGate unit acts as a dialup server and the FortiClient dialup client VIP addresses match the network behind the dialup server, select <i>Enable</i> to cause the FortiGate unit to act as a proxy for the dialup clients.</p> <p>This is available only for phase 2 configurations associated with a dialup phase 1 configuration. It works only on policy-based VPNs.</p>
<b>Quick Mode Selector</b>	<p>Optionally specify the source and destination IP address to be used as selectors for IKE negotiations. If the FortiGate unit is a dialup server, keep the default value 0.0.0.0/0 unless you need to circumvent problems caused by ambiguous IP addresses between one or more of the private networks making up the VPN.</p> <p>Note that IKEv1 does not support the use of multiple addresses in selectors. Instead, use the default 0.0.0.0/0 subnet selector and rely on the firewall policy to limit destination addresses. Only use the Addressing objects if they are carried over from earlier versions of FortiOS.</p> <p>If you are editing an existing phase 2 configuration, the <i>Source address</i> and <i>Destination address</i> fields are unavailable if the tunnel has been configured to use firewall addresses as selectors. This option exists only in the CLI. See the <i>dst-addr-type</i>, <i>dst-name</i>, <i>src-addr-type</i> and <i>src-name</i> keywords for the <code>vpn ipsec phase2</code> command in the <a href="#">FortiGate CLI Reference</a>.</p>
<b>Source address</b>	<p>If the FortiGate unit is a dialup server, type the source IP address that corresponds to the local sender(s) or network behind the local VPN peer (for example, 172.16.5.0/24 or 172.16.5.0/255.255.255.0 for a subnet, or 172.16.5.1/32 or 172.16.5.1/255.255.255.255 for a server or host, or 192.168.10.[80-100] or 192.168.10.80-192.168.10.100 for an address range). A value of 0.0.0.0/0 means all IP addresses behind the local VPN peer.</p> <p>If the FortiGate unit is a dialup client, source address must refer to the private network behind the FortiGate dialup client.</p>
<b>Source port</b>	Type the port number that the local VPN peer uses to transport traffic related to the specified service (protocol number). The range is 0 to 65535. To specify all ports, type 0.

---

<b>Destination address</b>	Type the destination IP address that corresponds to the recipient(s) or network behind the remote VPN peer (for example, 192.168.20.0/24 for a subnet, or 172.16.5.1/32 for a server or host, or 192.168.10.[80-100] for an address range). A value of 0.0.0.0/0 means all IP addresses behind the remote VPN peer.
<b>Destination port</b>	Type the port number that the remote VPN peer uses to transport traffic related to the specified service (protocol number). The range is 0 to 65535. To specify all ports, type 0.
<b>Protocol</b>	Type the IP protocol number of the service. The range is 1 to 255. To specify all services, type 0.

---

# Defining VPN security policies

This section explains how to specify the source and destination IP addresses of traffic transmitted through an IPsec VPN, and how to define appropriate security policies.

The following topics are included in this section:


- [Defining policy addresses](#)
- [Defining VPN security policies](#)

## Defining policy addresses

A VPN tunnel has two end points. These end points may be VPN peers such as two FortiGate gateways. Encrypted packets are transmitted between the end points. At each end of the VPN tunnel, a VPN peer intercepts encrypted packets, decrypts the packets, and forwards the decrypted IP packets to the intended destination.

You need to define firewall addresses for the private networks behind each peer. You will use these addresses as the source or destination address depending on the security policy.

**Figure 253:**Example topology for the following policies



In general:

- In a gateway-to-gateway, hub-and-spoke, dynamic DNS, redundant-tunnel, or transparent configuration, you need to define a policy address for the private IP address of the network behind the remote VPN peer (for example, `192.168.10.0/255.255.255.0` or `192.168.10.0/24`).
- In a peer-to-peer configuration, you need to define a policy address for the private IP address of a server or host behind the remote VPN peer (for example, `172.16.5.1/255.255.255.255` or `172.16.5.1/32` or `172.16.5.1`).

For a FortiGate dialup server in a dialup-client or Internet-browsing configuration:

- If you are not using VIP addresses, or if the FortiGate dialup server assigns VIP addresses to FortiClient dialup clients through FortiGate DHCP relay, select the predefined destination address “all” in the security policy to refer to the dialup clients.
- If you assign VIP addresses to FortiClient dialup clients manually, you need to define a policy address for the VIP address assigned to the dialup client (for example, 10.254.254.1/32), or a subnet address from which the VIP addresses are assigned (for example, 10.254.254.0/24 or 10.254.254.0/255.255.255.0).
- For a FortiGate dialup client in a dialup-client or Internet-browsing configuration, you need to define a policy address for the private IP address of a host, server, or network behind the FortiGate dialup server.

#### To define a security IP address

1. Go to *Firewall Objects > Address > Addresses* and select *Create New*.
2. In the *Name* field, type a descriptive name that represents the network, server(s), or host(s).
3. In *Type*, select *Subnet*.
4. In the *Subnet/IP Range* field, type the corresponding IP address and subnet mask.  
For a subnet you could use the format 172.16.5.0/24 or its equivalent 172.16.5.0/255.255.255.0. For a server or host it would likely be 172.16.5.1/32. Alternately you can use an IP address range such as 192.168.10.[80-100] or 192.168.10.80-192.168.10.100.
5. Select *OK*.

## Defining VPN security policies

Security policies allow IP traffic to pass between interfaces on a FortiGate unit. You can limit communication to particular traffic by specifying source address and destination addresses. Then only traffic from those addresses will be allowed.

Policy-based and route-based VPNs require different security policies.

- A policy-based VPN requires an IPsec security policy. You specify the interface to the private network, the interface to the remote peer and the VPN tunnel. A single policy can enable traffic inbound, outbound, or in both directions.
- A route-based VPN requires an Accept security policy for each direction. As source and destination interfaces, you specify the interface to the private network and the virtual IPsec interface (phase 1 configuration) of the VPN. The IPsec interface is the destination interface for the outbound policy and the source interface for the inbound policy. One security policy must be configured for each direction of each VPN interface.

There are examples of security policies for both policy-based and route-based VPNs throughout this guide. See [“Route-based or policy-based VPN” on page 1697](#).



If the security policy, which grants the VPN Connection is limited to certain services, DHCP must be included, otherwise the client won't be able to retrieve a lease from the FortiGate's (IPSec) DHCP server, because the DHCP Request (coming out of the tunnel) will be blocked.

---

## Defining an IPsec security policy for a policy-based VPN

An IPsec security policy enables the transmission and reception of encrypted packets, specifies the permitted direction of VPN traffic, and selects the VPN tunnel. In most cases, a single policy is needed to control both inbound and outbound IP traffic through a VPN tunnel.

### Allow traffic to be initiated from the remote site

In addition to these operations, security policies specify which IP addresses can initiate a tunnel. By default, traffic from the local private network initiates the tunnel. When the *Allow traffic to be initiated from the remote site* option is selected, traffic from a dialup client or computers on the remote network initiates the tunnel. Both can be enabled at the same time for bi-directional initiation of the tunnel.

### Outbound and inbound NAT

When a FortiGate unit operates in NAT mode, you can also enable inbound or outbound NAT. Outbound NAT may be performed on outbound encrypted packets, or on IP packets before they are sent through the tunnel. Inbound NAT is performed on IP packets emerging from the tunnel. By default, these options are not selected in security policies.

When used in conjunction with the `natip` CLI attribute (see the “config firewall” chapter of the [FortiGate CLI Reference](#)), outbound NAT enables you to change the source addresses of IP packets before they go into the tunnel. This feature is often used to resolve ambiguous routing when two or more of the private networks making up a VPN have the same or overlapping IP addresses.

When inbound NAT is enabled, inbound encrypted packets are intercepted and decrypted, and the source IP addresses of the decrypted packets are translated into the IP address of the FortiGate interface to the local private network before they are routed to the private network. If the computers on the local private network can communicate only with devices on the local private network (that is, the FortiGate interface to the private network is not the default gateway) and the remote client (or remote private network) does not have an IP address in the same network address space as the local private network, enable inbound NAT.

### Source and destination addresses

Most security policies control outbound IP traffic. A VPN outbound policy usually has a source address originating on the private network behind the local FortiGate unit, and a destination address belonging to a dialup VPN client or a network behind the remote VPN peer. The source address that you choose for the security policy identifies from where outbound cleartext IP packets may originate, and also defines the local IP address or addresses that a remote server or client will be allowed to access through the VPN tunnel. The destination address that you choose identifies where IP packets must be forwarded after they are decrypted at the far end of the tunnel, and determines the IP address or addresses that the local network will be able to access at the far end of the tunnel.

### Enabling other policy features

You can fine-tune a policy for services such as HTTP, FTP, and POP3; enable logging, traffic shaping, antivirus protection, web filtering, email filtering, file transfer, and email services throughout the VPN; and optionally allow connections according to a predefined schedule.

As an option, differentiated services (diffserv or DSCP) can be enabled in the security policy through CLI commands. For more information on this feature, see [Traffic Shaping chapter](#) or the “firewall” chapter of the [FortiGate CLI Reference](#).

When a remote server or client attempts to connect to the private network behind a FortiGate gateway, the security policy intercepts the connection attempt and starts the VPN tunnel. The

FortiGate unit uses the remote gateway specified in its phase 1 tunnel configuration to reply to the remote peer. When the remote peer receives a reply, it checks its own security policy, including the tunnel configuration, to determine which communications are permitted. As long as one or more services are allowed through the VPN tunnel, the two peers begin to negotiate the tunnel. To follow this negotiation in the web-based manager, go to *VPN > Monitor > IPsec Monitor*. There you will find a list of the VPN tunnels, their status, and the data flow both incoming and outgoing.

## Before you begin

Before you define the IPsec policy, you must:

- Define the IP source and destination addresses. See “[Defining policy addresses](#)” on page 1659.
- Specify the phase 1 authentication parameters. See “[Auto Key phase 1 parameters](#)” on page 1637.
- Specify the phase 2 parameters. See “[Phase 2 parameters](#)” on page 1653.

### To define an IPsec security policy

1. Go to *Policy > Policy > Policy*.
2. Select *Create New* and select *VPN*.
3. Complete the options:

<b>Local Interface</b>	Select the local interface to the internal (private) network.
<b>Local Protected Subnet</b>	Select the name that corresponds to the local network, server(s), or host(s) from which IP packets may originate.
<b>Outgoing VPN Interface</b>	Select the local interface to the external (public) network.
<b>Remote Protected Subnet</b>	Select the name that corresponds to the remote network, server(s), or host(s) to which IP packets may be delivered.
<b>Schedule</b>	Keep the default setting (always) unless changes are needed to meet specific requirements.
<b>Service</b>	Keep the default setting (ANY) unless changes are needed to meet your specific requirements.
<b>VPN Tunnel</b>	Select <i>Use Existing</i> and select the tunnel from the drop-down list.
<b>Allow traffic to be initiated from the remote site</b>	Select if traffic from the remote network will be allowed to initiate the tunnel.

4. You may enable UTM features, and/or event logging, or select advanced settings to authenticate a user group, or shape traffic. For more information, see the [Firewall](#) chapter of *The Handbook*.
5. Select *OK*.
6. Place the policy in the policy list above any other policies having similar source and destination addresses.

## Defining multiple IPsec policies for the same tunnel

You must define at least one IPsec policy for each VPN tunnel. If the same remote server or client requires access to more than one network behind a local FortiGate unit, the FortiGate unit

must be configured with an IPsec policy for each network. Multiple policies may be required to configure redundant connections to a remote destination or control access to different services at different times.

To ensure a secure connection, the FortiGate unit must evaluate IPSEC policies before ACCEPT and DENY security policies. Because the FortiGate unit reads policies starting at the top of the list, you must move all IPsec policies to the top of the list. When you define multiple IPsec policies for the same tunnel, you must reorder the IPsec policies that apply to the tunnel so that specific constraints can be evaluated before general constraints.



Adding multiple IPsec policies for the same VPN tunnel can cause conflicts if the policies specify similar source and destination addresses but have different settings for the same service. When policies overlap in this manner, the system may apply the wrong IPsec policy or the tunnel may fail.

---

For example, if you create two equivalent IPsec policies for two different tunnels, it does not matter which one comes first in the list of IPsec policies — the system will select the correct policy based on the specified source and destination addresses. If you create two different IPsec policies for the same tunnel (that is, the two policies treat traffic differently depending on the nature of the connection request), you might have to reorder the IPsec policies to ensure that the system selects the correct IPsec policy. Reordering is especially important when the source and destination addresses in both policies are similar (for example, if one policy specifies a subset of the IP addresses in another policy). In this case, place the IPsec policy having the most specific constraints at the top of the list so that it can be evaluated first.

## Defining security policies for a route-based VPN

When you define a route-based VPN, you create a virtual IPsec interface on the physical interface that connects to the remote peer. You create ordinary Accept security policies to enable traffic between the IPsec interface and the interface that connects to the private network. This makes configuration simpler than for policy-based VPNs, which require IPsec security policies.

### To define security policies for a route-based VPN

1. Go to *Policy > Policy > Policy*.
2. Select *Create New* and leave the *Policy Type* as *Firewall*, and the *Policy Subtype* as *Address*.

3. Define an *ACCEPT* security policy to permit communications between the local private network and the private network behind the remote peer. Enter these settings in particular:

---

<b>Incoming Interface</b>	Select the interface that connects to the private network behind this FortiGate unit.
<b>Source Address</b>	Select the address name that you defined for the private network behind this FortiGate unit.
<b>Outgoing Interface</b>	Select the IPsec Interface you configured.
<b>Destination Address</b>	Select the address name that you defined for the private network behind the remote peer.
<b>Action</b>	Select <i>ACCEPT</i> .
<b>Enable NAT</b>	Disable.

---

To permit the remote client to initiate communication, you need to define a security policy for communication in that direction.

4. Select *Create New* and leave the *Policy Type* as *Firewall*, and the *Policy Subtype* as *Address*
5. Enter these settings in particular:

---

<b>Incoming Interface</b>	Select the IPsec Interface you configured.
<b>Source Address</b>	Select the address name that you defined for the private network behind the remote peer.
<b>Outgoing Interface</b>	Select the interface that connects to the private network behind this FortiGate unit.
<b>Destination Address</b>	Select the address name that you defined for the private network behind this FortiGate unit.
<b>Action</b>	Select <i>ACCEPT</i> .
<b>Enable NAT</b>	Disable.

---



# Gateway-to-gateway configurations

This section explains how to set up a basic gateway-to-gateway (site-to-site) IPsec VPN.


The following topics are included in this section:

- [Configuration overview](#)
- [General configuration steps](#)
- [Configuring the two VPN peers](#)
- [How to work with overlapping subnets](#)
- [Testing](#)

## Configuration overview

In a gateway-to-gateway configuration, two FortiGate units create a VPN tunnel between two separate private networks. All traffic between the two networks is encrypted and protected by FortiGate security policies.

**Figure 254:**Example gateway-to-gateway configuration




In some cases, computers on the private network behind one VPN peer may (by co-incidence) have IP addresses that are already used by computers on the network behind the other VPN peer. In this type of situation (ambiguous routing), conflicts may occur in one or both of the FortiGate routing tables and traffic destined for the remote network through the tunnel may not be sent. To resolve issues related to ambiguous routing, see [“How to work with overlapping subnets” on page 1672](#).

In other cases, computers on the private network behind one VPN peer may obtain IP addresses from a local DHCP server. However, unless the local and remote networks use different private network address spaces, unintended ambiguous routing and/or IP-address overlap issues may arise. For a discussion of the related issues, see [“FortiGate dialup-client configurations” on page 1724](#).


You can set up a fully meshed or partially meshed configuration (see [Figure 255](#) and [Figure 256](#)).

**Figure 255:** Fully meshed configuration



In a fully meshed network, all VPN peers are connected to each other, with one hop between peers. This topology is the most fault-tolerant: if one peer goes down, the rest of the network is not affected. This topology is difficult to scale because it requires connections between all peers. In addition, unnecessary communication can occur between peers. Best practices dictates a hub-and-spoke configuration instead (see [“Hub-and-spoke configurations” on page 1679](#)).

**Figure 256:** Partially meshed configuration



A partially meshed network is similar to a fully meshed network, but instead of having tunnels between all peers, tunnels are only configured between peers that communicate with each other regularly.

## General configuration steps

The FortiGate units at both ends of the tunnel must be operating in NAT mode and have static public IP addresses.

When a FortiGate unit receives a connection request from a remote VPN peer, it uses IPsec phase 1 parameters to establish a secure connection and authenticate that VPN peer. Then, if the security policy permits the connection, the FortiGate unit establishes the tunnel using IPsec phase 2 parameters and applies the IPsec security policy. Key management, authentication, and security services are negotiated dynamically through the IKE protocol.

To support these functions, the following general configuration steps must be performed by both FortiGate units:

- Define the phase 1 parameters that the FortiGate unit needs to authenticate the remote peer and establish a secure connection.
- Define the phase 2 parameters that the FortiGate unit needs to create a VPN tunnel with the remote peer.
- Create security policies to control the permitted services and permitted direction of traffic between the IP source and destination addresses.

### Using auto-ipsec

In some cases, it may be easier to use the `auto-ipsec` CLI command to notify and push the IPsec configuration to the branch offices. For more information, refer to the [FortiGate CLI Reference](#).

## Configuring the two VPN peers

Configure the VPN peers as follows. Each step is required, but these are general steps. For more detailed information on each step follow the cross references. See [“Auto Key phase 1 parameters” on page 1637](#). All steps are required. Cross references point to required information that is repeated. No steps are optional.

### Configuring Phase 1 and Phase 2 for both peers

This procedure applies to both peers. Repeat the procedure on each FortiGate unit, using the correct IP address for each. You may wish to vary the Phase 1 names but this is optional. Otherwise all steps are the same for each peer.

The phase 1 configuration defines the parameters that FortiGate\_1 will use to authenticate FortiGate\_2 and establish a secure connection. For the purposes of this example, a preshared key will be used to authenticate FortiGate\_2. The same preshared key must be specified at both FortiGate units.

Before you define the phase 1 parameters, you need to:

- Reserve a name for the remote gateway.
- Obtain the IP address of the public interface to the remote peer.
- Reserve a unique value for the preshared key.

The key must contain at least 6 printable characters and best practices dictate that it only be known by network administrators. For optimum protection against currently known attacks, the key must have a minimum of 16 randomly chosen alphanumeric characters.

At the local FortiGate unit, define the phase 1 configuration needed to establish a secure connection with the remote peer. See [“Phase 1 configuration” on page 1625](#).

### To create phase 1 to establish a secure connection with the remote peer

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. Select *Create Phase 1*.
3. Enter the following information, and select *OK*.

<b>Name</b>	Enter <code>peer_1</code> .  A name to identify the VPN tunnel. This name appears in phase 2 configurations, security policies and the VPN monitor.
<b>Remote Gateway</b>	Select <i>Static IP Address</i> .
<b>IP Address</b>	Enter <code>172.20.0.2</code> when configuring <code>FortiGate_1</code> . Enter <code>172.18.0.2</code> when configuring <code>FortiGate_2</code> . The IP address of the remote peer public interface.
<b>Local Interface</b>	Select <i>wan1</i> .
<b>Enable IPsec Interface Mode</b>	Select <i>Advanced</i> to see this setting.  Enable <i>IPsec Interface Mode</i> to have the FortiGate unit create a virtual IPsec interface for a route-based VPN.  Disable this option to create a policy-based VPN. For more information, see <a href="#">“Comparing policy-based or route-based VPNs” on page 1621</a> .  After selecting <i>OK</i> , you cannot change this setting.

The basic phase 2 settings associate IPsec phase 2 parameters with the phase 1 configuration and specify the remote end point of the VPN tunnel. Before you define the phase 2 parameters, you need to reserve a name for the tunnel. See [“Phase 2 configuration” on page 1629](#).

### To configure phase 2 settings

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. Select *Create Phase 2*.
3. Enter a *Name* of `peer_1_p2`.
4. Select *peer\_1* from the *Phase 1* drop-down menu.

## Creating security policies

Security policies control all IP traffic passing between a source address and a destination address.

An IPsec security policy is needed to allow the transmission of encrypted packets, specify the permitted direction of VPN traffic, and select the VPN tunnel that will be subject to the policy. A single policy is needed to control both inbound and outbound IP traffic through a VPN tunnel.

Before you define security policies, you must first specify the IP source and destination addresses. In a gateway-to-gateway configuration:

- The IP source address corresponds to the private network behind the local FortiGate unit.
- The IP destination address refers to the private network behind the remote VPN peer.

When you are creating security policies, choose one of either route-based or policy-based methods and follow it for both VPN peers. DO NOT configure both route-based and policy-based policies on the same FortiGate unit for the same VPN tunnel.

The configuration of FortiGate\_2 is similar to that of FortiGate\_1. You must:

- Define the phase 1 parameters that FortiGate\_2 needs to authenticate FortiGate\_1 and establish a secure connection.
- Define the phase 2 parameters that FortiGate\_2 needs to create a VPN tunnel with FortiGate\_1.
- Create the security policy and define the scope of permitted services between the IP source and destination addresses.

When creating security policies it is good practice to include a comment describing what the policy does.

When creating security policies you need to be

- [Creating firewall addresses](#)
- [Creating route-based VPN security policies](#)
- [Configuring a default route for VPN interface](#)

or

- [Creating firewall addresses](#)
- [Creating policy-based VPN security policy](#)

## Creating firewall addresses

Define names for the addresses or address ranges of the private networks that the VPN links. These addresses are used in the security policies that permit communication between the networks.

### To define the IP address of the network behind FortiGate\_1

1. Go to *Firewall Objects > Address > Addresses* and select *Create New*.
2. Enter the *Name* of `Finance_network`.
3. Select a *Type* of *Subnet*.
4. Enter the *Subnet* of `10.21.101.0/24`.
5. Select *OK*.

### To specify the address of the network behind FortiGate\_2

1. Go to *Firewall Objects > Address > Addresses* and select *Create New*.
2. Enter the *Name* of `HR_network`.
3. Select a *Type* of *Subnet*.
4. Enter the *Subnet/IP Range* of `10.31.101.0/24`.
5. Select *OK*.

## Creating route-based VPN security policies

Define an ACCEPT security policy to permit communications between the source and destination addresses.

### To create route-based VPN security policies

1. Go to *Policy > Policy > Policy* and select *Create New*
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter the following, and select *OK*.

---

<b>Incoming Interface</b>	Select <i>internal</i> . The interface that connects to the private network behind this FortiGate unit.
<b>Source Address</b>	Select <i>Finance_network</i> when configuring FortiGate_1. Select <i>HR_network</i> when configuring FortiGate_2. The address name for the private network behind this FortiGate unit.
<b>Outgoing Interface</b>	Select <i>peer_1</i> . The VPN Tunnel (IPsec Interface) you configured earlier.
<b>Destination Address</b>	Select <i>HR_network</i> when configuring FortiGate_1. Select <i>Finance_network</i> when configuring FortiGate_2. The address name that you defined for the private network behind the remote peer.
<b>Action</b>	Select <i>ACCEPT</i> .
<b>Enable NAT</b>	Disable.
<b>Comments</b>	Allow Internal to remote VPN network traffic.

---

4. Configure any additional features such as UTM or traffic shaping you may want. (optional).
5. Select *Create New* to create another policy for the other direction.
6. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
7. Enter the following information, and select *OK*.

---

<b>Incoming Interface</b>	Select <i>peer_1</i> . The VPN Tunnel (IPsec Interface) you configured.
<b>Source Address</b>	Select <i>HR_network</i> when configuring FortiGate_1. Select <i>Finance_Network</i> when configuring FortiGate_2. The address name defined for the private network behind the remote peer.
<b>Outgoing Interface</b>	Select <i>internal</i> . The interface that connects to the private network behind this FortiGate unit.
<b>Destination Address</b>	Select <i>Finance_Network</i> when configuring FortiGate_1. Select <i>HR_network</i> when configuring FortiGate_2. The address name defined for the private network behind this FortiGate unit.

---

<b>Action</b>	Select <i>ACCEPT</i> .
<b>Enable NAT</b>	Disable.
<b>Comments</b>	Allow remote VPN network traffic to Internal.

8. Configure any additional features such as UTM or traffic shaping you may want. (optional).

### Configuring a default route for VPN interface

All network traffic must have a static route to direct its traffic to the proper destination. Without a route, traffic will not flow even if the security policies are configured properly. You may need to create a static route entry for both directions of VPN traffic if your security policies allow bi-directional tunnel initiation.

#### To configure the route for a route-based VPN

1. On FortiGate\_2, go to *Router > Static > Static Routes* and select *Create New*.  
For low-end FortiGate units, go to *System > Network > Routing* and select *Create New*.
2. Enter the following information, and then select *OK*:

<b>Destination IP / Mask</b>	10.21.101.0/24
<b>Device</b>	FGT2_to_FGT1_Tunnel
<b>Gateway</b>	Leave as default: 0.0.0.0.
<b>Distance (Advanced)</b>	Leave this at its default.  If there are other routes on this FortiGate unit, you may need to set the distance on this route so the VPN traffic will use it as the default route. However, this normally happens by default because this route is typically a better match than the generic default route.

### Creating policy-based VPN security policy

Define an IPsec security policy to permit communications between the source and destination addresses.

1. Go to *Policy > Policy > Policy*.
2. Select the *Policy Type* as *VPN* and leave the *Policy Subtype* as *IPsec*.
3. Complete the following:

<b>Local Interface</b>	Select <i>internal</i> .  The interface that connects to the private network behind this FortiGate unit.
<b>Local Protected Subnet</b>	Select <i>Finance_network</i> when configuring FortiGate_1. Select <i>HR_network</i> when configuring FortiGate_2.  The address name defined for the private network behind this FortiGate unit.
<b>Outgoing VPN Interface</b>	Select <i>wan1</i> .  The FortiGate unit's public interface.

<b>Remote Protected Subnet</b>	Select <i>HR_network</i> when configuring FortiGate_1. Select <i>Finance_network</i> when configuring FortiGate_2.  The address name that you defined in Step for the private network behind the remote peer.
<b>VPN Tunnel</b>	Select <i>Use Existing</i> and select <i>peer_1</i> from the <i>VPN Tunnel</i> drop-down list.  Select <i>Allow traffic to be initiated from the remote site</i> to enable traffic from the remote network to initiate the tunnel.
<b>Comments</b>	Bidirectional policy-based VPN policy.


Place VPN policies in the policy list above any other policies having similar source and destination addresses.

## How to work with overlapping subnets

A site-to-site VPN configuration sometimes has the problem that the private subnet addresses at each end are the same. You can resolve this problem by remapping the private addresses using virtual IP addresses (VIP).

VIPs allow computers on those overlapping private subnets to each have another set of IP addresses that can be used without confusion. The FortiGate unit maps the VIP addresses to the original addresses. This means if PC1 starts a session with PC2 at 10.31.101.10, FortiGate\_2 directs that session to 10.11.101.10 — the actual IP address of PC2. [Figure 257](#) shows this — Finance network VIP is 10.21.101.0/24 and the HR network is 10.31.101.0/24.

**Figure 257:**Overlapped subnets example





## Solution for route-based VPN

You need to:

- Configure IPsec Phase 1 and Phase 2 as you usually would for a route-based VPN. In this example, the resulting IPsec interface is named `FGT1_to_FGT2`.
- Configure virtual IP (VIP) mapping:
  - the 10.21.101.0/24 network mapped to the 10.11.101.0/24 network on FortiGate\_1
  - the 10.31.101.0/24 network mapped to the 10.11.101.0/24 network on FortiGate\_2
- Configure an outgoing security policy with ordinary source NAT on both FortiGates.
- Configure an incoming security policy with the VIP as the destination on both FortiGates.
- Configure a route to the remote private network over the IPsec interface on both FortiGates.

### To configure VIP mapping on both FortiGates

1. Go to *Firewall Objects > Virtual IPs > Virtual IPs* and select *Create New*.
2. Enter the following information, and select *OK*:

<b>Name</b>	Enter a name, for example, <code>my_vip</code> .
<b>External Interface</b>	Select <code>FGT1_to_FGT2</code> . The IPsec interface.
<b>Type</b>	Static NAT
<b>External IP Address/Range</b>	For the external IP address field enter: <ul style="list-style-type: none"><li>• <code>10.21.101.1</code> when configuring FortiGate_1, or</li><li>• <code>10.31.101.1</code> when configuring FortiGate_2.</li></ul>
<b>Mapped IP Address/Range</b>	For the Mapped IP Address enter <code>10.11.101.1</code> . For the Range enter <code>10.11.101.254</code> .
<b>Port Forwarding</b>	Disable

Repeat this procedure on both FortiGate\_1 and FortiGate\_2.

### To configure the outbound security policy on both FortiGates

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	Select <i>Port 1</i> .
<b>Source Address</b>	Select <i>all</i> .
<b>Outgoing Interface</b>	Select <code>FGT1_to_FGT2</code> . The IPsec interface.
<b>Destination Address</b>	Select <i>all</i> .
<b>Action</b>	Select <i>ACCEPT</i>
<b>Enable NAT</b>	Enable

Repeat this procedure on both FortiGate\_1 and FortiGate\_2.

### To configure the inbound security policy on both FortiGates

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter the following information, and then select *OK*:

<b>Incoming Interface</b>	Select <i>FGT1_to_FGT2</i> .
<b>Source Address</b>	Select <i>all</i> .
<b>Outgoing Interface</b>	Select <i>Port 1</i> . The IPsec interface.
<b>Destination Address</b>	Select <i>my-vip</i> .
<b>Action</b>	Select <i>ACCEPT</i>
<b>Enable NAT</b>	Disable

Repeat this procedure on both *FortiGate\_1* and *FortiGate\_2*.

### To configure the static route for both FortiGates

1. Go to *Router > Static > Static Routes* and select *Create New*.  
For low-end FortiGate units, go to *System > Network > Routing* and select *Create New*.
2. Enter the following information, and then select *OK*:

<b>Destination IP / Mask</b>	Enter <i>10.31.101.0/24</i> when configuring <i>FortiGate_1</i> . Enter <i>10.21.101.0/24</i> when configuring <i>FortiGate_2</i> .
<b>Device</b>	Select <i>FGT1_to_FGT2</i> .
<b>Gateway</b>	Leave as default: <i>0.0.0.0</i> .
<b>Distance (Advanced)</b>	Leave at default. If you have advanced routing on your network, you may have to change this value

## Solution for policy-based VPN

As with the route-based solution, users contact hosts at the other end of the VPN using an alternate subnet address. PC1 communicates with PC2 using IP address 10.31.101.10, and PC2 communicates with PC1 using IP address 10.21.101.10.

In this solution however, outbound NAT is used to translate the source address of packets from the 10.11.101.0/24 network to the alternate subnet address that hosts at the other end of the VPN use to reply. Inbound packets from the remote end have their destination addresses translated back to the 10.11.101.0/24 network.

For example, PC1 uses the destination address 10.31.101.10 to contact PC2. Outbound NAT on *FortiGate\_1* translates the PC1 source address to 10.21.101.10. At the *FortiGate\_2* end of the tunnel, the outbound NAT configuration translates the destination address to the actual PC2 address of 10.11.101.10. Similarly, PC2 replies to PC1 using destination address 10.21.101.10,

with the PC2 source address translated to 10.31.101.10. PC1 and PC2 can communicate over the VPN even though they both have the same IP address.

- You need to:
- Configure IPsec Phase 1 as you usually would for a policy-based VPN.
- Configure IPsec Phase 2 with the `use-natip disable` CLI option.
- Define a firewall address for the local private network, 10.11.101.0/24.
- Define a firewall address for the remote private network:
  - define a firewall address for 10.31.101.0/24 on FortiGate\_1
  - define a firewall address for 10.21.101.0/24 on FortiGate\_2
- Configure an outgoing IPsec security policy with outbound NAT to map 10.11.101.0/24 source addresses:
  - to the 10.21.101.0/24 network on FortiGate\_1
  - to the 10.31.101.0/24 network on FortiGate\_2

### To configure IPsec Phase 2 - CLI

```
config vpn ipsec phase2
 edit "FGT1_FGT2_p2"
 set keepalive enable
 set pfs enable
 set phase1name FGT1_to_FGT2
 set proposal 3des-sha1 3des-md5
 set replay enable
 set use-natip disable
 end
```

In this example, your phase 1 definition is named FGT1\_to\_FGT2. `use-natip` is set to `disable`, so you can specify the source selector using the `src-addr-type`, `src-start-ip` / `src-end-ip` or `src-subnet` keywords. This example leaves these keywords at their default values, which specify the subnet 0.0.0.0/0.

The `pfs` keyword ensures that perfect forward secrecy (PFS) is used. This ensures that each Phase 2 key created is unrelated to any other keys in use.

### To define the local private network firewall address

1. Go to *Firewall Objects > Address > Addresses* and select *Create New*.
2. Enter the following information and select *OK*.

<b>Name</b>	Enter <code>vpn-local</code> . A meaningful name for the local private network.
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	10.11.101.0 255.255.255.0
<b>Interface</b>	Any

### To define the remote private network firewall address

1. Go to *Firewall Objects > Address > Addresses* and select *Create New*.

2. Enter the following information, and select *OK*:

<b>Name</b>	Enter <code>vpn-remote</code> . A meaningful name for the remote private network.
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	10.31.101.0 255.255.255.0 on FortiGate_1. 10.21.101.0 255.255.255.0 on FortiGate_2.
<b>Interface</b>	Any

### To configure the IPsec security policy

In the CLI on FortiGate\_1, enter the commands:

```
config firewall policy
 edit 1
 set srcintf "port1"
 set dstintf "port2"
 set srcaddr "vpn-local"
 set dstaddr "vpn-remote"
 set action ipsec
 set schedule "always"
 set service "ANY"
 set inbound enable
 set outbound enable
 set vpntunnel "FGT1_to_FGT2"
 set natoutbound enable
 set natip 10.31.101.0 255.255.255.0
 end
```

Optionally, you can set everything except `natip` in the web-based manager and then use the CLI to set `natip`.

Enter the same commands on FortiGate\_2, but set `natip` be `10.21.101.0 255.255.255.0`.

## Testing

The best testing is to look at the packets both as the VPN tunnel is negotiated, and when the tunnel is up.

### To determine what the other end of the VPN tunnel is proposing

1. Start a terminal program such as `puTTY` and set it to log all output.  
When necessary refer to the logs to locate information when output is verbose.
2. Logon to the FortiGate unit using a `super_admin` account.
3. Enter the following CLI commands.
4. Display all the possible IKE error types and the number of times they have occurred:  

```
diag vpn ike errors
```

5. Check for existing debug sessions:  

```
diag debug info
```

If a debug session is running, to halt it enter:  

```
diag debug disable
```
6. Confirm your proposal settings:  

```
diag vpn ike config list
```
7. If your proposal settings do not match what you expect, make a change to it and save it to force an update in memory. If that fixes the problem, stop here.
8. List the current vpn filter:  

```
diag vpn ike filter
```
9. If all fields are set to any, there are no filters set and all VPN ike packets will be displayed in the debug output. If your system has only a few VPNs, skip setting the filter.  

If your system has many VPN connections this will result in very verbose output and make it very difficult to locate the correct connection attempt.
10. Set the VPN filter to display only information from the destination IP address for example 10.10.10.10:  

```
diag vpn ike log-filter dst-addr4 10.10.10.10
```

To add more filter options, enter them one per line as above. Other filter options are:

<b>clear</b>	<b>erase the current filter</b>
dst-addr6	the IPv6 destination address range to filter by
dst-port	the destination port range to filter by
interface	interface that IKE connection is negotiated over
list	display the current filter
name	the phase1 name to filter by
negate	negate the specified filter parameter
src-addr4	the IPv4 source address range to filter by
src-addr6	the IPv6 source address range to filter by
src-port	the source port range to filter by
vd	index of virtual domain. 0 matches all

11. Start debugging:  

```
diag debug app ike 255
```

```
diag debug enable
```
12. Have the remote end attempt a VPN connection.  

If the remote end attempts the connection they become the initiator. This situation makes it easier to debug VPN tunnels because then you have the remote information and all of your local information. by initiate the connection, you will not see the other end's information.
13. If possible go to the web-based manager on your FortiGate unit, go to the VPN monitor and try to bring the tunnel up.
14. Stop the debug output:  

```
diag debug disable
```

**15.** Go back through the output to determine what proposal information the initiator is using, and how it is different from your VPN P1 proposal settings.

Things to look for in the debug output of attempted VPN connections are shown below.

**Table 74:** Important terms to look for in VPN debug output

initiator	Starts the VPN attempt, in the above procedure that is the remote end
responder	Answers the initiator's request
local ID	In aggressive mode, this is not encrypted
error no SA proposal chosen	There was no proposal match — there was no encryption-authentication pair in common, usually occurs after a long list of proposal attempts
R U THERE and R U THERE ack	dead peer detection (dpd), also known as dead gateway detection — after three failed attempts to contact the remote end it will be declared dead, no farther attempts will be made to contact it
negotiation result	lists the proposal settings that were agreed on
SA_life_soft and SA_life_hard	negotiating a new key, and the key life
R U THERE	If you see this, it means Phase 1 was successful
tunnel up	the negotiation was successful, the VPN tunnel is operational

# Hub-and-spoke configurations


This section describes how to set up hub-and-spoke IPsec VPNs. The following topics are included in this section:

- [Configuration overview](#)
- [Configure the hub](#)
- [Configure the spokes](#)
- [Dynamic spokes configuration example](#)

## Configuration overview

In a hub-and-spoke configuration, VPN connections radiate from a central FortiGate unit (the hub) to a number of remote peers (the spokes). Traffic can pass between private networks behind the hub and private networks behind the remote peers. Traffic can also pass between remote peer private networks through the hub.

**Figure 258:**Example hub-and-spoke configuration



The actual implementation varies in complexity depending on

- whether the spokes are statically or dynamically addressed
- the addressing scheme of the protected subnets
- how peers are authenticated.

This guide discusses the issues involved in configuring a hub-and-spoke VPN and provides some basic configuration examples.

## Hub-and-spoke infrastructure requirements

- The FortiGate hub must be operating in NAT mode and have a static public IP address.
- Spokes may have static IP addresses, dynamic IP addresses (see [“FortiGate dialup-client configurations”](#) on page 1724), or static domain names and dynamic IP addresses (see [“Dynamic DNS configuration”](#) on page 1695).

## Spoke gateway addressing

The public IP address of the spoke is the VPN remote gateway as seen from the hub. Statically addressed spokes each require a separate VPN phase 1 configuration on the hub. When there are many spokes, this becomes rather cumbersome.

Using dynamic addressing for spokes simplifies the VPN configuration because then the hub requires only a single phase 1 configuration with “dialup user” as the remote gateway. You can use this configuration even if the remote peers have static IP addresses. A remote peer can establish a VPN connection regardless of its IP address if its traffic selectors match and it can authenticate to the hub. See [“Dynamic spokes configuration example”](#) on page 1689 for an example of this configuration.

## Protected networks addressing


The addresses of the protected networks are needed to configure destination selectors and sometimes for security policies and static routes. The larger the number of spokes, the more addresses there are to manage. You can

- assign spoke subnets as part of a larger subnet, usually on a new network  
or
- create address groups that contain all of the needed addresses

## Using aggregated subnets

If you are creating a new network, where subnet IP addresses are not already assigned, you can simplify the VPN configuration by assigning spoke subnets that are part of a large subnet.

**Figure 259:**Aggregated subnets



All spokes use the large subnet address, 10.1.0.0/16 for example, as

- the IPsec destination selector
- the destination of the security policy from the private subnet to the VPN (required for policy-based VPN, optional for route-based VPN)
- the destination of the static route to the VPN (route-based)

Each spoke uses the address of its own protected subnet as the IPsec source selector and as the source address in its VPN security policy. The remote gateway is the public IP address of the hub FortiGate unit.



## Using an address group

If you want to create a hub-and-spoke VPN between existing private networks, the subnet addressing usually does not fit the aggregated subnet model discussed earlier. All of the spokes and the hub will need to include the addresses of all the protected networks in their configuration.

On FortiGate units, you can define a named firewall address for each of the remote protected networks and add these addresses to a firewall address group. For a policy-based VPN, you can then use this address group as the destination of the VPN security policy.

For a route-based VPN, the destination of the VPN security policy can be set to All. You need to specify appropriate routes for each of the remote subnets.

## Authentication

Authentication is by a common preshared key or by certificates. For simplicity, the examples in this chapter assume that all spokes use the same preshared key.

## Configure the hub

At the FortiGate unit that acts as the hub, you need to

- configure the VPN to each spoke
- configure communication between spokes

You configure communication between spokes differently for a policy-based VPN than for a route-based VPN. For a policy-based VPN, you configure a VPN concentrator. For a route-based VPN, you must either define security policies or group the IPsec interfaces into a zone

## Define the hub-spoke VPNs

Perform these steps at the FortiGate unit that will act as the hub. Although this procedure assumes that the spokes are all FortiGate units, a spoke could also be VPN client software, such as FortiClient Endpoint Security.

### To configure the VPN hub

1. At the hub, define the phase 1 configuration for each spoke. See [“Auto Key phase 1 parameters” on page 1637](#). Enter these settings in particular:

---

<b>Name</b>	Enter a name to identify the VPN in phase 2 configurations, security policies and the VPN monitor.
-------------	----------------------------------------------------------------------------------------------------

---

<b>Remote Gateway</b>	<p>The remote gateway is the other end of the VPN tunnel. There are three options:</p> <p><b>Static IP Address</b> — Enter the spoke’s public <i>IP Address</i>. You will need to create a phase 1 configuration for each spoke. Either the hub or the spoke can establish the VPN connection.</p> <p><b>Dialup User</b> — No additional information is needed. The hub accepts connections from peers with appropriate encryption and authentication settings. Only one phase 1 configuration is needed for multiple dialup spokes. Only the spoke can establish the VPN tunnel.</p> <p><b>Dynamic DNS</b> — If the spoke subscribes to a dynamic DNS service, enter the spoke’s <i>Dynamic DNS</i> domain name. Either the hub or the spoke can establish the VPN connection. For more information, see <a href="#">“Dynamic DNS configuration” on page 1695</a>.</p>
<b>Local Interface</b>	Select the FortiGate interface that connects to the remote gateway. This is usually the FortiGate unit’s public interface.
<b>Enable IPsec Interface Mode</b>	<p>You must select Advanced to see this setting. If <i>IPsec Interface Mode</i> is enabled, the FortiGate unit creates a virtual IPsec interface for a route-based VPN. Disable this option if you want to create a policy-based VPN. For more information, see <a href="#">“Comparing policy-based or route-based VPNs” on page 1621</a>.</p> <p>After you select <i>OK</i> to create the phase 1 configuration, you cannot change this setting.</p>

2. Define the phase 2 parameters needed to create a VPN tunnel with each spoke. See [“Phase 2 parameters” on page 1653](#). Enter these settings in particular:

<b>Name</b>	Enter a name to identify this spoke phase 2 configuration.
<b>Phase 1</b>	Select the name of the phase 1 configuration that you defined for this spoke.

## Define the hub-spoke security policies

1. Define a name for the address of the private network behind the hub. For more information, see [“Defining policy addresses” on page 1659](#).
2. Define names for the addresses or address ranges of the private networks behind the spokes. For more information, see [“Defining policy addresses” on page 1659](#).
3. Define the VPN concentrator. See [“To define the VPN concentrator” on page 1684](#).
4. Define security policies to permit communication between the hub and the spokes. For more information, see [“Defining VPN security policies” on page 1660](#).

## Route-based VPN security policies

Define ACCEPT security policies to permit communications between the hub and the spoke. You need one policy for each direction.

### To add policies

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.

3. Enter these settings in particular:

---

<b>Incoming Interface</b>	Select the VPN Tunnel (IPsec Interface) you configured in Step 1.
<b>Source Address</b>	Select the address name you defined in Step 2 for the private network behind the spoke FortiGate unit.
<b>Outgoing Interface</b>	Select the hub's interface to the internal (private) network.
<b>Destination Address</b>	Select the source address that you defined in Step 1.
<b>Action</b>	Select <i>ACCEPT</i> .
<b>Enable NAT</b>	Enable.

---

---

<b>Incoming Interface</b>	Select the VPN Tunnel (IPsec Interface) you configured in Step 1.
<b>Source Address</b>	Select the address name you defined in Step 2 for the private network behind the spoke FortiGate units.
<b>Outgoing Interface</b>	Select the source address that you defined in Step 1.
<b>Destination Address</b>	Select the hub's interface to the internal (private) network.
<b>Action</b>	Select <i>ACCEPT</i> .
<b>Enable NAT</b>	Enable.

---

### Policy-based VPN security policy

Define an IPsec security policy to permit communications between the hub and the spoke.

#### To add policies

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Select the *Policy Type* as *VPN* and leave the *Policy Subtype* as *IPsec*.
3. Enter these settings in particular:

---

<b>Local Interface</b>	Select the hub's interface to the internal (private) network.
<b>Local Protected Subnet</b>	Select the source address that you defined in Step 1.
<b>Outgoing VPN Interface</b>	Select the hub's public network interface.
<b>Remote Protected Subnet</b>	Select the address name you defined in Step 2 for the private network behind the spoke FortiGate unit.
<b>VPN Tunnel</b>	Select <i>Use Existing</i> and select the name of the phase 1 configuration that you created for the spoke in Step 1.  Select <i>Allow traffic to be initiated from the remote site</i> to enable traffic from the remote network to initiate the tunnel.

---

In the policy list, arrange the policies in the following order:

- IPsec policies that control traffic between the hub and the spokes first
- the default security policy last

## Configuring communication between spokes (policy-based VPN)

For a policy-based hub-and-spoke VPN, you define a concentrator to enable communication between the spokes.

### To define the VPN concentrator

1. At the hub, go to *VPN > IPSEC > Concentrator* and select *Create New*.
2. In the *Concentrator Name* field, type a name to identify the concentrator.
3. From the *Available Tunnels* list, select a VPN tunnel and then select the right-pointing arrow.
4. Repeat Step 3 until all of the tunnels associated with the spokes are included in the concentrator.
5. Select *OK*.

## Configuring communication between spokes (route-based VPN)

For a route-based hub-and-spoke VPN, there are several ways you can enable communication between the spokes:

- put all of the IPsec interfaces into a zone and enable intra-zone traffic. This eliminates the need for any security policy for the VPN, but you cannot apply UTM features to scan the traffic for security threats.
- put all of the IPsec interfaces into a zone and create a single zone-to-zone security policy
- create a security policy for each pair of spokes that are allowed to communicate with each other. The number of policies required increases rapidly as the number of spokes increases.

### Using a zone as a concentrator

A simple way to provide communication among all of the spokes is to create a zone and allow intra-zone communication. You cannot apply UTM features using this method.

1. Go to *System > Network > Interfaces*.
2. Select the down-arrow on the *Create New* button and select *Zone*.
3. In the *Zone Name* field, enter a name, such as `Our_VPN_zone`.
4. Clear *Block intra-zone traffic*.
5. In the *Interface Members* list, select the IPsec interfaces that are part of your VPN.
6. Select *OK*.

### Using a zone with a policy as a concentrator

If you put all of the hub IPsec interfaces involved in the VPN into a zone, you can enable communication among all of the spokes and apply UTM features with just one security policy.

### To create a zone for the VPN

1. Go to *System > Network > Interfaces*.
2. Select the down-arrow on the *Create New* button and select *Zone*.
3. In the *Zone Name* field, enter a name, such as `Our_VPN_zone`.
4. Select *Block intra-zone traffic*.

5. In the *Interface Members* list, select the IPsec interfaces that are part of your VPN.
6. Select *OK*.

#### To create a security policy for the zone

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter the settings: and select *OK*.

<b>Incoming Interface</b>	Select the zone you created for your VPN.
<b>Source Address</b>	Select <i>All</i> .
<b>Outgoing Interface</b>	Select the zone you created for your VPN.
<b>Destination Address</b>	Select <i>All</i> .
<b>Action</b>	Select <i>ACCEPT</i> .
<b>Enable NAT</b>	Enable.

#### Using security policies as a concentrator

To enable communication between two spokes, you need to define an *ACCEPT* security policy for them. To allow either spoke to initiate communication, you must create a policy for each direction. This procedure describes a security policy for communication from Spoke 1 to Spoke 2. Others are similar.

1. Define names for the addresses or address ranges of the private networks behind each spoke. For more information, see [“Defining policy addresses” on page 1659](#).
2. Go to *Policy > Policy > Policy* and select *Create New*.
3. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
4. Enter the settings and select *OK*.

<b>Incoming Interface</b>	Select the IPsec interface that connects to Spoke 1.
<b>Source Address</b>	Select the address of the private network behind Spoke 1.
<b>Outgoing Interface</b>	Select the IPsec interface that connects to Spoke 2.
<b>Destination Address</b>	Select the address of the private network behind Spoke 2.
<b>Action</b>	Select <i>ACCEPT</i> .
<b>Enable NAT</b>	Enable.

## Configure the spokes

Although this procedure assumes that the spokes are all FortiGate units, a spoke could also be VPN client software, such as FortiClient Endpoint Security.

Perform these steps at each FortiGate unit that will act as a spoke.

### To create the phase 1 and phase\_2 configurations

1. At the spoke, define the phase 1 parameters that the spoke will use to establish a secure connection with the hub. See [“Auto Key phase 1 parameters” on page 1637](#). Enter these settings:

<b>Remote Gateway</b>	Select <i>Static IP Address</i> .
<b>IP Address</b>	Type the IP address of the interface that connects to the hub.
<b>Enable IPsec Interface Mode</b>	Enable if you are creating a route-based VPN. Clear if you are creating a policy-based VPN

2. Create the phase 2 tunnel definition. See [“Phase 2 parameters” on page 1653](#). Select the set of phase 1 parameters that you defined for the hub. You can select the name of the hub from the *Static IP Address* part of the list.

### Configuring security policies for hub-to-spoke communication

1. Create an address for this spoke. See [“Defining policy addresses” on page 1659](#). Enter the IP address and netmask of the private network behind the spoke.
2. Create an address to represent the hub. See [“Defining policy addresses” on page 1659](#). Enter the IP address and netmask of the private network behind the hub.
3. Define the security policy to enable communication with the hub.

#### Route-based VPN security policy

Define two security policies to permit communications to and from the hub.

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter these settings:

<b>Incoming Interface</b>	Select the virtual IPsec interface you created.
<b>Source Address</b>	Select the hub address you defined in Step 1.
<b>Outgoing Interface</b>	Select the spoke’s interface to the internal (private) network.
<b>Destination Address</b>	Select the spoke addresses you defined in Step 2.
<b>Action</b>	Select <i>ACCEPT</i> .
<b>Enable NAT</b>	Enable

<b>Incoming Interface</b>	Select the spoke’s interface to the internal (private) network.
<b>Source Address</b>	Select the spoke address you defined in Step 1.
<b>Outgoing Interface</b>	Select the virtual IPsec interface you created.
<b>Destination Address</b>	Select the hub destination addresses you defined in Step 2.

<b>Action</b>	Select <i>ACCEPT</i> .
<b>Enable NAT</b>	Enable

### Policy-based VPN security policy

Define an IPsec security policy to permit communications with the hub. See [“Defining VPN security policies” on page 1660](#).

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Select the *Policy Type* as *VPN* and leave the *Policy Subtype* as *IPsec*.
3. Enter these settings in particular:

<b>Local Interface</b>	Select the spoke’s interface to the internal (private) network.
<b>Local Protected Subnet</b>	Select the spoke address you defined in Step 1.
<b>Outgoing VPN Interface</b>	Select the spoke’s interface to the external (public) network.
<b>Remote Protected Subnet</b>	Select the hub address you defined in Step 2.
<b>VPN Tunnel</b>	Select <i>Use Existing</i> and select the name of the phase 1 configuration you defined.  Select <i>Allow traffic to be initiated from the remote site</i> to enable traffic from the remote network to initiate the tunnel.

### Configuring security policies for spoke-to-spoke communication

Each spoke requires security policies to enable communication with the other spokes. Instead of creating separate security policies for each spoke, you can create an address group that contains the addresses of the networks behind the other spokes. The security policy then applies to all of the spokes in the group.

1. Define destination addresses to represent the networks behind each of the other spokes. Add these addresses to an address group.
2. Define the security policy to enable communication between this spoke and the spokes in the address group you created.

#### Policy-based VPN security policy

Define an IPsec security policy to permit communications with the other spokes. See [“Defining VPN security policies” on page 1660](#). Enter these settings in particular:

### Route-based VPN security policy

Define two security policies to permit communications to and from the other spokes.

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter these settings in particular:

<b>Incoming Interface</b>	Select the virtual IPsec interface you created.
<b>Source Address</b>	Select the spoke address group you defined in Step 1.
<b>Outgoing Interface</b>	Select the spoke’s interface to the internal (private) network.

---

**Destination Address** Select this spoke's address name.

---

**Action** Select *ACCEPT*.

---

**Enable NAT** Enable

---

4. Select *Create New*, leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*, and enter these settings:

---

**Incoming Interface** Select the spoke's interface to the internal (private) network.

---

**Source Address** Select this spoke's address name.

---

**Outgoing Interface** Select the virtual IPsec interface you created.

---

**Destination Address** Select the spoke address group you defined in Step 1.

---

**Action** Select *ACCEPT*.

---

**Enable NAT** Enable

---

### Policy-based VPN security policy

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Select the *Policy Type* as *VPN* and leave the *Policy Subtype* as *IPsec*.
3. Enter the following:

---

**Local Interface** Select this spoke's internal (private) network interface.

---

**Local Protected Subnet** Select this spoke's source address.

---

**Outgoing VPN Interface** Select the spoke's interface to the external (public) network.

---

**Remote Protected Subnet** Select the spoke address group you defined in Step 1.

---

**VPN Tunnel** Select *Use Existing* and select the name of the phase 1 configuration you defined.  
Select *Allow traffic to be initiated from the remote site* to enable traffic from the remote network to initiate the tunnel.

---


Place this policy or policies in the policy list above any other policies having similar source and destination addresses.



## Dynamic spokes configuration example

This example demonstrates how to set up a basic route-based hub-and-spoke IPsec VPN that uses preshared keys to authenticate VPN peers.

**Figure 260:**Example hub-and-spoke configuration



In the example configuration, the protected networks 10.1.0.0/24, 10.1.1.0/24 and 10.1.2.0/24 are all part of the larger subnet 10.1.0.0/16. The steps for setting up the example hub-and-spoke configuration create a VPN among Site 1, Site 2, and the HR Network.

The spokes are dialup. Their addresses are not part of the configuration on the hub, so only one spoke definition is required no matter the number of spokes. For simplicity, only two spokes are shown.

### Configure the hub (FortiGate\_1)

The phase 1 configuration defines the parameters that FortiGate\_1 will use to authenticate spokes and establish secure connections.

For the purposes of this example, one preshared key will be used to authenticate all of the spokes. Each key must contain at least 6 printable characters and best practices dictates that it only be known by network administrators. For optimum protection against currently known attacks, each key must consist of a minimum of 16 randomly chosen alphanumeric characters.

## Define the IPsec configuration

### To define the phase 1 parameters

1. At FortiGate\_1, go to *VPN > IPsec > Auto Key (IKE)*.
2. Define the phase 1 parameters that the hub will use to establish a secure connection to the spokes. Select *Create Phase 1*, enter the following information, and select *OK*:

<b>Name</b>	Enter a name (for example, toSpokes).
<b>Remote Gateway</b>	Dialup user
<b>Local Interface</b>	External
<b>Mode</b>	Main
<b>Authentication Method</b>	Preshared Key
<b>Pre-shared Key</b>	Enter the preshared key.
<b>Peer Options</b>	Accept any peer ID

The basic phase 2 settings associate IPsec phase 2 parameters with the phase 1 configuration and specify the remote end points of the VPN tunnels.

### To define the phase 2 parameters

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. Select *Create Phase 2*, enter the following information, and select *OK*:

<b>Name</b>	Enter a name for the phase 2 definition (for example, toSpokes_ph2).
<b>Phase 1</b>	Select the Phase 1 configuration that you defined previously (for example, toSpokes).

## Define the security policies

security policies control all IP traffic passing between a source address and a destination address. For a route-based VPN, the policies are simpler than for a policy-based VPN. Instead of an IPSEC policy, you use an ACCEPT policy with the virtual IPsec interface as the external interface.

Before you define security policies, you must first define firewall addresses to use in those policies. You need addresses for:

- the HR network behind FortiGate\_1
- the aggregate subnet address for the protected networks

### To define the IP address of the HR network behind FortiGate\_1

1. Go to *Firewall Objects > Address > Addresses*.
2. Select *Create New*, enter the following information, and select *OK*:

<b>Name</b>	Enter an address name (for example, HR_Network).
-------------	--------------------------------------------------

<b>Type</b>	Subnet
<b>Subnet/IP Range</b>	Enter the IP address of the HR network behind FortiGate_1 (for example, 10.1.0.0/24).

### To specify the IP address the aggregate protected subnet

1. Go to *Firewall Objects > Address > Addresses*.
2. Select *Create New*, enter the following information, and select *OK*:

<b>Address Name</b>	Enter an address name (for example, Spoke_net).
<b>Type</b>	Subnet
<b>Subnet/IP Range</b>	Enter the IP address of the aggregate protected network, 10.1.0.0/16

### To define the security policy for traffic from the hub to the spokes

1. Go to *Policy > Policy > Policy*. and select *Create New*,
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	Select the interface to the HR network, <i>port 1</i> .
<b>Source Address</b>	Select <i>HR_Network</i> .
<b>Outgoing Interface</b>	Select the virtual IPsec interface that connects to the spokes, <i>toSpokes</i> .
<b>Destination Address</b>	Select <i>Spoke_net</i> .
<b>Action</b>	Select <i>ACCEPT</i> .

Place the policy in the policy list above any other policies having similar source and destination addresses.

## Configure communication between spokes

Spokes communicate with each other through the hub. You need to configure the hub to allow this communication. An easy way to do this is to create a zone containing the virtual IPsec interfaces even if there is only one, and create a zone-to-zone security policy.

### To create a zone for the VPN

1. Go to *System > Network > Interfaces*.
2. Select the down-arrow on the *Create New* button and select *Zone*.
3. In the *Zone Name* field, enter a name, such as *Our\_VPN\_zone*.
4. Select *Block intra-zone traffic*.  
You could enable intra-zone traffic and then you would not need to create a security policy. But, you would not be able to apply UTM features.
5. In *Interface Members*, select the virtual IPsec interface, *toSpokes*.
6. Select *OK*.

### To create a security policy for the zone

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter these settings:

<b>Incoming Interface</b>	Select <i>Our_VPN_zone</i> .
<b>Source Address</b>	Select <i>All</i> .
<b>Outgoing Interface</b>	Select <i>Our_VPN_zone</i> .
<b>Destination Address</b>	Select <i>All</i> .
<b>Action</b>	Select <i>ACCEPT</i> .
<b>Enable NAT</b>	Enable.

4. Select *OK*.

## Configure the spokes

In this example, all spokes have nearly identical configuration, requiring the following:

- phase 1 authentication parameters to initiate a connection with the hub
- phase 2 tunnel creation parameters to establish a VPN tunnel with the hub
- a source address that represents the network behind the spoke. This is the only part of the configuration that is different for each spoke.
- a destination address that represents the aggregate protected network
- a security policy to enable communications between the spoke and the aggregate protected network

### Define the IPsec configuration

At each spoke, create the following configuration.

#### To define the Phase 1 parameters

1. At the spoke, go to *VPN > IPsec > Auto Key (IKE)*.
2. Select *Create Phase 1*, enter the following information, and select *OK*:

<b>Name</b>	Type a name, for example, <i>toHub</i> .
<b>Remote Gateway</b>	Select <i>Static IP Address</i> .
<b>IP Address</b>	Enter <i>172.16.10.1</i> .
<b>Local Interface</b>	Select <i>Port2</i> .
<b>Mode</b>	Main
<b>Authentication Method</b>	Preshared Key
<b>Pre-shared Key</b>	Enter the preshared key. The value must be identical to the preshared key that you specified previously in the <i>FortiGate_1</i> configuration

<b>Peer Options</b>	Select <i>Accept any peer ID</i> .
<b>Enable IPsec Interface Mode</b>	Select <i>Advanced</i> to see this option. Enable the option to create a route-based VPN.

### To define the Phase 2 parameters

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. Select *Create Phase 2*, enter the following information, and select *OK*:

<b>Name</b>	Enter a name for the tunnel, for example, <code>toHub_ph2</code> .
<b>Phase 1</b>	Select the name of the phase 1 configuration that you defined previously, for example, <code>toHub</code> .
<b>Advanced</b>	Select to show the following <i>Quick Mode Selector</i> settings.
<b>Source</b>	Enter the address of the protected network at this spoke. For <code>spoke_1</code> , this is <code>10.1.1.0/24</code> . For <code>spoke_2</code> , this is <code>10.1.2.0/24</code> .
<b>Destination</b>	Enter the aggregate protected subnet address, <code>10.1.0.0/16</code> .

### Define the security policies

You need to define firewall addresses for the spokes and the aggregate protected network and then create a security policy to enable communication between them.

#### To define the IP address of the network behind the spoke

1. Go to *Firewall Objects > Address > Addresses*.
2. Select *Create New*, enter the following information, and select *OK*:

<b>Address Name</b>	Enter an address name, for example <code>LocalNet</code> .
<b>Type</b>	Subnet
<b>Subnet/IP Range</b>	Enter the IP address of the private network behind the spoke. For <code>spoke_1</code> , this is <code>10.1.1.0/24</code> . For <code>spoke_2</code> , this is <code>10.1.2.0/24</code> .

#### To specify the IP address of the aggregate protected network

1. Go to *Firewall Objects > Address > Addresses*.
2. Select *Create New*, enter the following information, and select *OK*:

<b>Address Name</b>	Enter an address name, for example, <code>Spoke_net</code> .
<b>Type</b>	Subnet
<b>Subnet/IP Range</b>	Enter the IP address of the aggregate protected network, <code>10.1.0.0/16</code> .

#### To define the security policy

1. Go to *Policy > Policy > Policy* and select *Create New*.

2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter the following and select *OK*:

---

<b>Incoming Interface</b>	Select the virtual IPsec interface, <code>toHub</code> .
<b>Source Address</b>	Select the aggregate protected network address <code>Spoke_net</code> .
<b>Outgoing Interface</b>	Select the interface to the internal (private) network, <code>port1</code> .
<b>Destination Address</b>	Select the address for this spoke's protected network <code>LocalNet</code> .
<b>Action</b>	Select <i>ACCEPT</i> .

---

4. Select *Create New*.
5. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
6. Enter the following information, and select *OK*:

---

<b>Incoming Interface</b>	Select the interface to the internal private network, <code>port1</code> .
<b>Source Address</b>	Select the address for this spoke's protected network, <code>LocalNet</code> .
<b>Outgoing Interface</b>	Select the virtual IPsec interface, <code>toHub</code> .
<b>Destination Address</b>	Select the aggregate protected network address, <code>Spoke_net</code> .
<b>Action</b>	Select <i>ACCEPT</i> .

---

Place these policies in the policy list above any other policies having similar source and destination addresses.

# Dynamic DNS configuration

This section describes how to configure a site-to-site VPN, in which one FortiGate unit has a static IP address and the other FortiGate unit has a domain name and a dynamic IP address.

The following topics are included in this section:

- [Dynamic DNS over VPN concepts](#)
- [Dynamic DNS topology](#)
- [General configuration steps](#)
- [Configure the dynamically-addressed VPN peer](#)
- [Configure the fixed-address VPN peer](#)
- [Testing](#)

## Dynamic DNS over VPN concepts

A typical computer has a static IP address and one or more DNS servers to resolve fully qualified domain names (FQDN) into IP addresses. A domain name assigned to this computer is resolved by any DNS server having an entry for the domain name and its static IP address. The IP address never changes or changes only rarely so the DNS server can reliably say it has the correct address for that domain all the time.

### Dynamic DNS (DDNS)

It is different when a computer has a dynamic IP address, such as an IP address assigned dynamically by a DHCP server, and a domain name. Computers that want to contact this computer do not know what its current IP address is. To solve this problem there are dynamic DNS servers. These are public servers that store a DNS entry for your computer that includes its current IP address and associated domain name. These entries are kept up to date by your computer sending its current IP address to the dynamic DNS (DDNS) server to ensure its entry is always up to date. When other computers want to contact your domain, their DNS gets your IP address from your DDNS server. To use DDNS servers, you must subscribe to them and usually pay for their services.

When configuring DDNS on your FortiGate unit, go to *System > Network > DNS* and enable *Enable FortiGuard DDNS*. Then select the interface with the dynamic connection, which DDNS server you have an account with, your domain name, and account information. If your DDNS server is not on the list, there is a generic option where you can provide your DDNS server information.

### Routing

When an interface has some form of changing IP address (DDNS, PPPoE, or DHCP assigned address), routing needs special attention. The standard static route cannot handle the changing IP address. The solution is to use the dynamic-gateway command in the CLI. Say for example you already have four static routes, and you have a PPPoE connection over the wan2 interface and you want to use that as your default route.

The route is configured on the dynamic address VPN peer trying to access the static address FortiGate unit.

## To configure dynamic gateway routing - CLI

```
config router static
 edit 5
 set dst 0.0.0.0 0.0.0.0
 set dynamic-gateway enable
 set device wan2
 next
end
```

[handbook chapter](#)

## Dynamic DNS over VPN

IPsec VPN expects an IP address for each end of the VPN tunnel. All configuration and communication with that tunnel depends on the IP addresses as reference points. However, when the interface the tunnel is on has DDNS enabled there is no set IP address. The remote end of the VPN tunnel now needs another way to reference your end of the VPN tunnel. This is accomplished using Local ID.

A FortiGate unit that has a domain name and a dynamic IP address can initiate VPN connections anytime. The remote peer can reply to the local FortiGate unit using the source IP address that was sent in the packet header because it is current. Without doing a DNS lookup first, the remote peer runs the risk of the dynamic IP changing before it attempts to connect. To avoid this, the remote peer must perform a DNS lookup for the domain name of to be sure of the dynamic IP address before initiating the connection.

### Remote Gateway

When configuring the Phase 1 entry for a VPN tunnel, the Remote Gateway determines the addressing method the remote end of the tunnel uses as one of Static IP Address, Dialup User, or Dynamic DNS. There are different fields for each option.

When you select the Dynamic DNS VPN type there is a related field called Dynamic DNS. The Dynamic DNS field is asking for the FQDN of the remote end of the tunnel. It uses this information to look up the IP address of the remote end of the tunnel through the DDNS server associated with that domain name.

### Local ID (peer ID)

The Local ID or peer ID can be used to uniquely identify one end of a VPN tunnel. This enables a more secure connection. Also if you have multiple VPN tunnels negotiating, this ensures the proper remote and local ends connect. When you configure it on your end, it is your Local ID. When the remote end connects to you, they see it as your peer ID.

If you are debugging a VPN connection, the Local ID is part of the VPN negotiations. You can use it to help troubleshoot connection problems.

### To configure your Local ID

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. Select *Create New Phase 1* or edit an existing Phase 1 entry.
3. Select *Advanced*.
4. In the *P1 Proposal* section, enter your Local ID.
5. Select *OK*.

The default configuration is to accept all local IDs (peer IDs). If you have the Local ID set, the remote end of the tunnel must be configured to accept your Local ID.



### To accept a specific Peer ID

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. Select *Create New Phase 1*.
3. Select *Aggressive mode*.
4. For *Peer Options*, select *Accept this peer ID*. This option becomes visible only when *Aggressive mode* is selected.
5. Enter the string the other end of the tunnel used for its Local ID.
6. Configure the rest of the Phase 1 entry as required.
7. Select *OK*.

### Route-based or policy-based VPN

VPN over dynamic DNS can be configured with either route-based or policy-based VPN settings. Both are valid, but have differences in configuration. Choose the best method based on your requirements. For more information on route-based and policy-based, see [“Types of VPNs” on page 1620](#).

Route-based VPN configuration requires two security policies to be configured (one for each direction of traffic) to permit traffic over the VPN virtual interface, and you must also add a static route entry for that VPN interface or the VPN traffic will not reach its destination. See [“Creating branch\\_2 route-based security policies” on page 1702](#) and [“Creating branch\\_1 route-based security policies” on page 1706](#).


Policy-based VPN configuration uses more complex and often more IPsec security policies, but does not require a static route entry. It has the benefit of being able to configure multiple policies for handling multiple protocols in different ways, such as more scanning of less secure protocols or guaranteeing a minimum bandwidth for protocols such as VoIP. See [“Creating branch\\_2 policy-based security policies” on page 1703](#) and [“Creating branch\\_1 policy-based security policies” on page 1707](#).

## Dynamic DNS topology

In this scenario, two branch offices each have a FortiGate unit and are connected in a gateway-to-gateway VPN configuration. One FortiGate unit has a domain name (example.com) with a dynamic IP address. See `branch_2` in [Figure 261](#).

Whenever the `branch_2` unit connects to the Internet (and possibly also at predefined intervals set by the ISP), the ISP may assign a different IP address to the FortiGate unit. The unit has its domain name registered with a dynamic DNS service. The `branch_2` unit checks in with the DDNS server on a regular basis, and that server provides the DNS information for the domain name, updating the IP address from time to time. Remote peers have to locate the `branch_2` FortiGate unit through a DNS lookup each time to ensure the address they get is current and correct.

**Figure 261:**Example dynamic DNS configuration



When a remote peer (such as the `branch_1` FortiGate unit in [Figure 261](#)) initiates a connection to `example.com`, the local DNS server looks up and returns the IP address that matches the domain name `example.com`. The remote peer uses the retrieved IP address to establish a VPN connection with the `branch_2` FortiGate unit.

## Assumptions

- You have administrator access to both FortiGate units.
- Both FortiGate units have interfaces named `wan1` and `internal`. (If not, you can use the alias feature to assign these labels as “nicknames” to other interfaces to follow this example.)
- Both FortiGate units have the most recent firmware installed, have been configured for their networks, and are currently passing normal network traffic.
- The `branch_2` FortiGate unit has its `wan1` interface defined as a dynamic DNS interface with the domain name of `example.com`.
- A basic gateway-to-gateway configuration is in place (see [“Gateway-to-gateway configurations” on page 1665](#)) except one of the FortiGate units has a static domain name and a dynamic IP address instead of a static IP address.
- The FortiGate unit with the domain name is subscribed to one of the supported dynamic DNS services. Contact one of the services to set up an account. For more information and instructions about how to configure the FortiGate unit to push its dynamic IP address to a dynamic DNS server, see the [System Administration handbook chapter](#).

## General configuration steps

When a FortiGate unit receives a connection request from a remote VPN peer, it uses IPsec phase 1 parameters to establish a secure connection and authenticate the VPN peer. Then, if the security policy permits the connection, the FortiGate unit establishes the tunnel using IPsec phase 2 parameters and applies the security policy. Key management, authentication, and security services are negotiated dynamically through the IKE protocol.


To support these functions, the following general configuration steps must be performed:

- Configure the branch\_2 FortiGate unit with the dynamic IP address. This unit uses a Local ID string instead of an IP address to identify itself to the remote peer. See [“Configure the dynamically-addressed VPN peer”](#) on page 1699.
  - [Configuring branch\\_2 VPN tunnel settings](#)
  - [Configuring branch\\_2 security policies](#)
- Configure the fixed-address VPN peer. To initiate a VPN tunnel with the dynamically-addressed peer, this unit must first retrieve the IP address for the domain from the dynamic DNS service. See [“Configure the fixed-address VPN peer”](#) on page 1704.
  - [Configuring branch\\_1 VPN tunnel settings](#)
  - [Configuring branch\\_1 security policies](#)

## Configure the dynamically-addressed VPN peer

It is assumed that this FortiGate unit (branch\_2) has already had its public facing interface, for example the wan1, configured with the proper dynamic DNS configuration.

**Figure 262:**Configure branch\_2, the dynamic address side



Configuring the dynamically-addressed VPN peer includes:

- [Configuring branch\\_2 VPN tunnel settings](#)
- [Configuring branch\\_2 security policies](#)

### Configuring branch\_2 VPN tunnel settings

Define the phase 1 parameters needed to establish a secure connection with the remote peer. See [“Auto Key phase 1 parameters”](#) on page 1637. During this procedure you need to choose if you will be using route-based or policy-based VPNs.

**To configure branch\_2 VPN tunnel settings**

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. Select *Create New Phase 1*.

3. Enter the following information.

---

<b>Name</b>	Enter <code>branch_2</code> , a name to identify the VPN tunnel. This name appears in phase 2 configurations, security policies, and the VPN monitor.
<b>Remote Gateway</b>	Select <i>Static IP Address</i> .  The remote peer this FortiGate is connecting to has a static IP public address.  If the remote interface is PPPoE do not select <i>Retrieve default gateway from server</i> .
<b>IP Address</b>	Enter <code>172.16.20.1</code> . The IP address of the public interface to the remote peer.
Enter <code>172.16.20.1</code>	Select <i>Aggressive</i> .  The IP address of the public interface to the remote peer.

---

4. Select *Advanced* and complete the following:

---

<b>Enable IPsec Interface Mode</b>	Enable for a route-based VPN and when configuring policies, go to <a href="#">“Creating branch_2 route-based security policies” on page 1702</a> .  Disable for a policy-based VPN and when configuring policies, go to <a href="#">“Creating branch_2 policy-based security policies” on page 1703</a> .  If enabled, default settings are used.
<b>Local ID</b>	Enter <code>example.com</code> .  A character string used by the <code>branch_2</code> FortiGate unit to identify itself to the remote peer.  This value must be identical to the value in the <i>Accept this peer ID</i> field of the phase 1 remote gateway configuration on the <code>branch_1</code> remote peer. See <a href="#">“Configuring branch_1 VPN tunnel settings” on page 1704</a> .

---

5. Select *Create Phase 2*.

Define the phase 2 parameters needed to create a VPN tunnel with the remote peer. For details on phase 2, see [“Phase 2 parameters” on page 1653](#).

6. Enter the following information and select *OK*.

---

<b>Name</b>	Enter <code>branch_2_phase2</code> .  A name to identify this phase 2 configuration.
<b>Phase 1</b>	Select <code>branch_2</code> .  The name of the phase 1 configuration that you defined earlier.

---

## Configuring branch\_2 security policies

Define security policies to permit communications between the private networks through the VPN tunnel. Route-based and policy-based VPNs require different security policies. For detailed information about creating security policies, see [“Defining VPN security policies” on page 1660](#).

After defining the two address ranges, select one of [“Creating branch\\_2 route-based security policies” on page 1702](#) or [“Creating branch\\_2 policy-based security policies” on page 1703](#) to configure the appropriate VPN policies.

### Define address ranges for branch\_2 security policies

Define VPN connection names for the address ranges of the private networks. These addresses are used in the security policies that permit communication between the networks. For more information, see [“Defining policy addresses” on page 1659](#).

Define an address name for the IP address and netmask of the private network behind the local FortiGate unit.

#### To define branch\_2 address ranges

1. Go to *Firewall Objects > Address > Addresses*.
2. Select *Create New*.
3. Enter the following information, and select *OK*.

<b>Name</b>	Enter <code>branch_2_internal</code> . Enter a meaningful name.
<b>Type</b>	Select <i>Subnet</i> .
<b>Subnet / IP Range</b>	Enter <code>10.10.10.0/24</code> . Include the netmask or specify a specific range.
<b>Interface</b>	Select <i>internal</i> . The interface that will be handling the traffic from the internal network.

Define an address name for the IP address and netmask of the private network behind the remote peer.

4. Select *Create New*.
5. Enter the following information, and select *OK*.

<b>Name</b>	Enter <code>branch_1_internal</code> . A meaningful name for the private network at the remote end of the VPN tunnel.
<b>Type</b>	Select <i>Subnet</i> .
<b>Subnet / IP Range</b>	Enter <code>192.168.1.0/24</code> . Include the netmask. Optionally you can specify a range
<b>Interface</b>	Select <i>any</i> .  The interface that will be handling the remote VPN traffic on this FortiGate unit. If you are unsure, or multiple interfaces may be handling this traffic use <i>any</i> .

## Creating branch\_2 route-based security policies

Define ACCEPT security policies to permit communication between the branch\_2 and branch\_1 private networks. Once the route-based policy is configured a routing entry must be configured to route traffic over the VPN interface.

Define a policy to permit the branch\_2 local FortiGate unit to initiate a VPN session with the branch\_1 VPN peer.

### To create route-based security policies

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter the following information, and select *OK*.

---

<b>Incoming Interface</b>	Select <i>internal</i> . The interface that connects to the private network behind this FortiGate unit.
<b>Source Address</b>	Select <i>branch_2_internal</i> . Select the address name for the private network behind this FortiGate unit.
<b>Outgoing Interface</b>	Select <i>branch_2</i> . The VPN Tunnel (IPsec Interface).
<b>Destination Address</b>	Select <i>branch_1_internal</i> . The address name the private network behind the remote peer.
<b>Action</b>	Select <i>ACCEPT</i> .
<b>Enable NAT</b>	Disable.
<b>Comments</b>	Route-based: Initiate a branch_2 to branch_1 VPN tunnel.

---

Define a policy to permit the branch\_1 remote VPN peer to initiate VPN sessions.

1. Select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter the following information, and select *OK*.

---

<b>Incoming Interface</b>	Select <i>branch_2</i> . The VPN Tunnel (IPsec Interface).
<b>Source Address</b>	Select <i>branch_1_internal</i> . The address name for the private network behind the remote peer.
<b>Outgoing Interface</b>	Select <i>internal</i> . The interface connecting the private network behind this FortiGate unit.
<b>Destination Address</b>	Select <i>branch_2_internal</i> . The address name for the private network behind this FortiGate unit.
<b>Action</b>	Select <i>ACCEPT</i> .
<b>Enable NAT</b>	Disable.
<b>Comments</b>	Route-based: Initiate a branch_1 to branch_2 internal VPN tunnel.

---

4. Optionally configure any other security policy settings you require such as UTM or traffic shaping for this policy.
5. Place these policies in the policy list above any other policies having similar source and destination addresses. This will ensure VPN traffic is matched against the VPN policies before any other policies.

#### To create routing entry for VPN interface - CLI

```
config router static
 edit 5
 set dst 0.0.0.0 0.0.0.0
 set dynamic-gateway enable
 set device wan1
 next
end
```

This routing entry must be added in the CLI because the dynamic-gateway option is not available in the web-based manager.

### Creating branch\_2 policy-based security policies

Define an IPsec policy to permit VPN sessions between the private networks. Define an IPsec policy to permit the VPN sessions between the local branch\_2 unit and the remote branch\_1 unit.

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Select the *Policy Type* as *VPN* and leave the *Policy Subtype* as *IPsec*.
3. Enter the following information, and select *OK*.


<b>Local Interface</b>	Select <i>internal</i> . The interface connecting the private network behind this FortiGate unit.
<b>Local Protected Subnet</b>	Select <i>branch_2_internal</i> . The address name for the private network behind this local FortiGate unit.
<b>Outgoing VPN Interface</b>	Select <i>wan1</i> . The FortiGate unit's public interface.
<b>Remote Protected Subnet</b>	Select <i>branch_1_internal</i> . The address name for the private network behind branch_1, the remote peer.
<b>VPN Tunnel</b>	Select <i>Use Existing</i> and select <i>branch_2</i> from the drop-down list. The name of the phase 1 tunnel.
	Select <i>Allow traffic to be initiated from the remote site</i> .
<b>Comments</b>	Policy-based: allows traffic in either direction to initiate the VPN tunnel.

4. Optionally configure any other security policy settings you require such as UTM or traffic shaping for this policy.
5. Place these policies in the policy list above any other policies having similar source and destination addresses. This will ensure VPN traffic is matched against the VPN policies before any other policies.

## Configure the fixed-address VPN peer

The fixed-address VPN peer, `branch_1`, needs to retrieve the IP address from the dynamic DNS service to initiate communication with the dynamically-addressed peer, `branch_2`. It also depends on the peer ID (local ID) to initiate the VPN tunnel with `branch_2`.

**Figure 263:** Configure `branch_1`, the fixed address side



Configuring the fixed-address VPN peer includes:

- [Configuring `branch\_1` VPN tunnel settings](#)
- [Configuring `branch\_1` security policies](#)

### Configuring `branch_1` VPN tunnel settings

Define the phase 1 parameters needed to establish a secure connection with the remote peer. For more information, see [“Auto Key phase 1 parameters” on page 1637](#).

#### To configure `branch_1` phase 1 VPN settings

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. Select *Create New Phase 1*.
3. Enter the following information and select *OK*.

<b>Name</b>	Enter <code>branch_1</code> . A name to identify the VPN tunnel. This name appears in phase 2 configurations, security policies and the VPN monitor.
<b>Remote Gateway</b>	Select <i>Dynamic DNS</i> . The remote peer this FortiGate is connecting to has a dynamic IP address.
<b>Dynamic DNS</b>	Type the fully qualified domain name of the remote peer (for example, <code>example.com</code> ).
<b>Interface</b>	Select <i>wan1</i> . The public facing interface on the fixed-address FortiGate unit.
<b>Mode</b>	Select <i>Aggressive</i> .



<b>Peer Options</b>	Select <i>Accept this peer ID</i> , and enter <code>example.com</code> . This option only appears when the mode is set to Aggressive. The identifier of the FortiGate unit with the dynamic address.
<b>Enable IPsec Interface Mode</b>	Enable for a route-based VPN and when configuring policies, go to <a href="#">“Creating branch_1 route-based security policies” on page 1706</a> . Disable for a policy-based VPN and when configuring policies, go to <a href="#">“Creating branch_1 policy-based security policies” on page 1707</a> . If Interface mode is enabled, default settings are used.

4. Define the phase 2 parameters needed to create a VPN tunnel with the remote peer. See [“Phase 2 parameters” on page 1653](#). Enter these settings in particular:

<b>Name</b>	Enter <code>branch_1_p2</code> . A name to identify this phase 2 configuration.
<b>Phase 1</b>	Select <i>branch_1</i> . The name of the phase 1 configuration that you defined for the remote peer. You can select the name of the remote gateway from the Dynamic DNS part of the list.

## Configuring branch\_1 security policies

The `branch_1` FortiGate unit has a fixed IP address and will be connecting to the `branch_2` FortiGate unit that has a dynamic IP address and a domain name of `example.com`. Remember if you are using route-based security policies that you must add a route for the VPN traffic.

### Defining address ranges for branch\_1 security policies

As with `branch_2` previously, `branch_1` needs address ranges defined as well. See [“Defining policy addresses” on page 1659](#).

1. Go to *Firewall Objects > Address > Addresses*.
2. Select *Create New*.
3. Enter the following information, and select *OK*.

<b>Name</b>	Enter <code>branch_2_internal</code> . A meaningful name for the private network behind the <code>branch_2</code> FortiGate unit.
<b>Type</b>	Select <i>Subnet</i> .
<b>Subnet / IP Range</b>	Enter <code>10.10.10.0/24</code> . Include the netmask or specify a specific range.
<b>Interface</b>	Select <i>internal</i> . This is the interface on this FortiGate unit that will be handling with this traffic.

4. Define an address name for the IP address and netmask of the private network behind the remote peer.
5. Select *Create New*.

- Enter the following information, and select *OK*.

<b>Name</b>	Enter <code>branch_1_internal</code> . A meaningful name for the private network behind the <code>branch_1</code> peer.
<b>Type</b>	Select <i>Subnet</i> .
<b>Subnet / IP Range</b>	Enter <code>192.168.1.0/24</code> . Include the netmask or specify a specific range.
<b>Interface</b>	Select <i>any</i> . The interface on this FortiGate unit that will be handling with this traffic. If you are unsure, or multiple interfaces may be handling this traffic use <i>any</i> .

### Creating `branch_1` route-based security policies

Define an *ACCEPT* security policy to permit communications between the source and destination addresses. See [“Defining VPN security policies” on page 1660](#).

- Go to *Policy > Policy > Policy* and select *Create New*.
- Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
- Enter the following information, and select *OK*.

<b>Incoming Interface</b>	Select <i>internal</i> . The interface that connects to the private network behind the <code>branch_1</code> FortiGate unit.
<b>Source Address</b>	Select <i>branch_1_internal</i> . The address name that you defined for the private network behind this FortiGate unit.
<b>Outgoing Interface</b>	Select <i>branch_1</i> . The VPN Tunnel (IPsec Interface) you configured earlier.
<b>Destination Address</b>	Select <i>branch_2_internal</i> . The address name that you defined for the private network behind the <code>branch_2</code> peer.
<b>Action</b>	Select <i>ACCEPT</i> .
<b>Enable NAT</b>	Disable
<b>Comments</b>	Internal -> branch2

To permit the remote client to initiate communication, you need to define a security policy for communication in that direction.

- Select *Create New*.
- Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
- Enter the following information, and select *OK*.

<b>Incoming Interface</b>	Select <i>branch_1</i> . The VPN Tunnel (IPsec Interface) you configured earlier.
<b>Source Address</b>	Select <i>branch_2_internal</i> . The address name that you defined for the private network behind the <code>branch_2</code> remote peer.
<b>Outgoing Interface</b>	Select <i>internal</i> . The interface that connects to the private network behind this FortiGate unit.

<b>Destination Address</b>	Select <i>branch_1_internal</i> . The address name that you defined for the private network behind this FortiGate unit.
<b>Action</b>	Select <i>ACCEPT</i> .
<b>Enable NAT</b>	Disable
<b>Comments</b>	branch_2 -> Internal

## Creating branch\_1 policy-based security policies

A policy-based security policy allows you the flexibility to allow inbound or outbound traffic or both through this single policy.

This policy-based IPsec VPN security policy allows both inbound and outbound traffic

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Select the *Policy Type* as *VPN* and leave the *Policy Subtype* as *IPsec*.
3. Enter the following information, and select *OK*.

<b>Local Interface</b>	Select <i>internal</i> . The interface that connects to the private network behind this FortiGate unit.
<b>Local Protected Subnet</b>	Select <i>branch_1_internal</i> . The address name that you defined for the private network behind this FortiGate unit.
<b>Outgoing VPN Interface</b>	Select <i>wan1</i> . The FortiGate unit's public interface.
<b>Remote Protected Subnet</b>	Select <i>branch_2_internal</i> . The address name that you defined for the private network behind the remote peer.
<b>VPN Tunnel</b>	Select <i>Use Existing</i> and select <i>branch_1</i> from the drop-down list.  Select <i>Allow traffic to be initiated from the remote site</i> to enable traffic from the remote network to initiate the tunnel.

4. Place this security policy in the policy list above any other policies having similar source and destination addresses.

## Testing

Once both ends are configured, you can test the VPN tunnel.

### To test the VPN initiated by branch\_2

1. On *branch\_2*, go to *VPN > Monitor > IPsec Monitor*.  
All IPsec VPN tunnels will be listed on this page, no matter if they are connected or disconnected.
2. Select the tunnel listed for *branch\_2*, and select the status column for that entry.  
The status will say *Bring Up* and remote port, incoming and outgoing data will all be zero. This indicates an inactive tunnel. When you select *Bring Up*, the FortiGate will try to set up a VPN session over this tunnel. If it is successful, *Bring Up* will change to *Active*, and the arrow icon will change to a green up arrow icon.
3. If this does not create a VPN tunnel with increasing values for incoming and outgoing data, you need to start troubleshooting:

### To test the VPN initiated by branch\_1

1. On branch\_1, go to *VPN > Monitor > IPsec Monitor*.
2. Select the tunnel listed for branch\_1, and select the status column.  
The difference between branch\_2 and branch\_1 at this point is that the tunnel entry for branch-1 will not have a remote gateway IP address. It will be resolved when the VPN tunnel is started.
3. If this does not create a VPN tunnel with increasing values for incoming and outgoing data, you need to start troubleshooting.

Some troubleshooting ideas include:

- If there was no entry for the tunnel on the monitor page, check the Auto Key (IKE) page to verify the phase 1 and phase 2 entries exist.
- Check the security policy or policies, and ensure there is an outgoing policy as a minimum.
- Check that you entered a local ID in the phase 1 configuration, and that branch\_1 has the same local ID.
- Ensure the local DNS server has an up-to-date DNS entry for exmaple.com.

For more information on VPN troubleshooting and testing, see [“VPN troubleshooting tips” on page 1826](#).

# FortiClient dialup-client configurations

The FortiClient Endpoint Security application is an IPsec VPN client with antivirus, antispam and firewall capabilities. This section explains how to configure dialup VPN connections between a FortiGate unit and one or more FortiClient Endpoint Security applications.

FortiClient users are usually mobile or remote users who need to connect to a private network behind a FortiGate unit. For example, the users might be employees who connect to the office network while traveling or from their homes.

For greatest ease of use, the FortiClient application can download the VPN settings from the FortiGate unit to configure itself automatically. This section covers both automatic and manual configuration.

The following topics are included in this section:

- [Configuration overview](#)
- [FortiClient-to-FortiGate VPN configuration steps](#)
- [Configure the FortiGate unit](#)
- [Configure the FortiClient Endpoint Security application](#)
- [Adding XAuth authentication](#)
- [FortiClient dialup-client configuration example](#)

## Configuration overview


Dialup users typically obtain dynamic IP addresses from an ISP through Dynamic Host Configuration Protocol (DHCP) or Point-to-Point Protocol over Ethernet (PPPoE). Then, the FortiClient Endpoint Security application initiates a connection to a FortiGate dialup server.

By default the FortiClient dialup client has the same IP address as the host PC on which it runs. If the host connects directly to the Internet, this is a public IP address. If the host is behind a NAT device, such as a router, the IP address is a private IP address. The NAT device must be NAT traversal (NAT-T) compatible to pass encrypted packets (see [“NAT traversal” on page 1649](#)). The FortiClient application also can be configured to use a virtual IP address (VIP). For the duration of the connection, the FortiClient application and the FortiGate unit both use the VIP address as the IP address of the FortiClient dialup client.

For a faster and easier method of configuring a FortiGate - to - FortiClient VPN, see [“One button FortiGate - to - FortiClient Phase1 VPN” on page 1711](#).

The FortiClient application sends its encrypted packets to the VPN remote gateway, which is usually the public interface of the FortiGate unit. It also uses this interface to download VPN settings from the FortiGate unit. See [“Automatic configuration of FortiClient dialup clients” on page 1710](#).

**Figure 264:**Example FortiClient dialup-client configuration



## Peer identification

The FortiClient application can establish an IPsec tunnel with a FortiGate unit configured to act as a dialup server. When the FortiGate unit acts as a dialup server, it does not identify the client using the phase 1 remote gateway address. The IPsec tunnel is established if authentication is successful and the IPsec security policy associated with the tunnel permits access. If configured, the FortiGate unit could also require FortiClient registration, that is, the remote user would be required to have FortiClient installed before connection is completed.

There are several different ways to authenticate dialup clients and restrict access to private networks based on client credentials. For more information, see [“Authenticating remote peers and clients”](#) on page 1642.

## Automatic configuration of FortiClient dialup clients

The FortiClient application can obtain its VPN settings from the FortiGate VPN server. FortiClient users need to know only the FortiGate VPN server IP address and their user name and password on the FortiGate unit.

The FortiGate unit listens for VPN policy requests from clients on TCP port 8900. When the dialup client connects:

- The client initiates a Secure Sockets Layer (SSL) connection to the FortiGate unit.
- The FortiGate unit requests a user name and password from the FortiClient user. Using these credentials, it authenticates the client and determines which VPN policy applies to the client.
- Provided that authentication is successful, the FortiGate unit downloads a VPN policy to the client over the SSL connection. The information includes IPsec phase 1 and phase 2 settings, and the IP addresses of the private networks that the client is authorized to access.
- The client uses the VPN policy settings to establish an IPsec phase 1 connection and phase 2 tunnel with the FortiGate unit.

## One button FortiGate - to - FortiClient Phase1 VPN

On the FortiOS VPN IKE page there is a method to create a Phase1 portion of a VPN tunnel between the FortiGate and FortiClient. Very little information is required for this configuration. No encryption or authentication method is required. This feature is ideal for setting up quick VPN connections with basic settings.

On the Phase 1 screen (*VPN > IPsec > Phase 1*) is the option *Create a FortiClient VPN*. When selected, the FortiGate unit requires a few basic VPN configuration related questions. Once all the information is added, select *OK*. This will create a new dial-up IPsec-interface mode tunnel. Phase 1 and Phase 2 will be added using the default ike settings.

The following Settings will be used when creating a one-button FortiClient VPN Phase1 object:

- Remote Gateway: Dialup User
- Mode: Aggressive
- Enable IPsec Interface Mode
- Default setting for P1 and P2 Proposal
- XAUTH Enable as Server (Auto)
- IKE mode-config will be enabled
- Peer Option set to "Accept any peer ID"
- Rest of the setting use the current defaults (Default value needs to be the same on FCT side)

Once the completed, you need to create a default Phase2 configuration. This only requires a name for the Phase2 object, and select the FortiClient connection Phase1 name.

### How the FortiGate unit determines which settings to apply

The FortiGate unit follows these steps to determine the configuration information to send to the FortiClient application:

1. Check the virtual domain associated with the connection to determine which VPN policies might apply.
2. Select the VPN policy that matches the dialup client's user group and determine which tunnel (phase 1 configuration) is involved.
3. Check all IPsec security policies that use the specified tunnel to determine which private networks the dialup clients may access.
4. Retrieve the rest of the VPN policy information from the existing IPsec phase 1 and phase 2 parameters in the dialup-client configuration.

### Using virtual IP addresses

When the FortiClient host PC is located behind a NAT device, unintended IP address overlap issues may arise between the private networks at the two ends of the tunnel. For example, the client's host might receive a private IP address from a DHCP server on its network that by co-incidence is the same as a private IP address on the network behind the FortiGate unit. A conflict will occur in the host's routing table and the FortiClient Endpoint Security application will be unable to send traffic through the tunnel. Configuring virtual IP (VIP) addresses for FortiClient applications prevents this problem.

Using VIPs ensures that client IP addresses are in a predictable range. You can then define security policies that allow access only to that source address range. If you do not use VIPs, the security policies must allow all source addresses because you cannot predict the IP address for a remote mobile user.

The FortiClient application must not have the same IP address as any host on the private network behind the FortiGate unit or any other connected FortiClient application. You can

ensure this by reserving a range of IP addresses on the private network for FortiClient users. Or, you can assign FortiClient VIPs from an uncommonly used subnet such as 10.254.254.0/24 or 192.168.254.0/24.

You can reserve a VIP address for a particular client according to its device MAC address and type of connection. The DHCP server then always assigns the reserved VIP address to the client. For more information about this feature, see the “dhcp reserved-address” section in the “system” chapter of the [FortiGate CLI Reference](#).



On the host computer, you can find out the VIP address that the FortiClient Endpoint Security application is using. For example, in Windows command prompt, type `ipconfig /all`

On Linux or Mac OS X, type `ifconfig` in a terminal window. The output will also show the IP address that has been assigned to the host Network Interface Card (NIC).

---

It is best to assign VIPs using DHCP over IPsec. The FortiGate dialup server can act as a DHCP server or relay requests to an external DHCP server. You can also configure VIPs manually on FortiClient applications, but it is more difficult to ensure that all clients use unique addresses.



If you assign a VIP on the private network behind the FortiGate unit and enable DHCP-IPsec (a phase 2 advanced option), the FortiGate unit acts as a proxy on the local private network for the FortiClient dialup client. Whenever a host on the network behind the dialup server issues an ARP request for the device MAC address of the FortiClient host, the FortiGate unit answers the ARP request on behalf of the FortiClient host and forwards the associated traffic to the FortiClient host through the tunnel. For more information, see “[DHCP-IPsec](#)” on page 1654

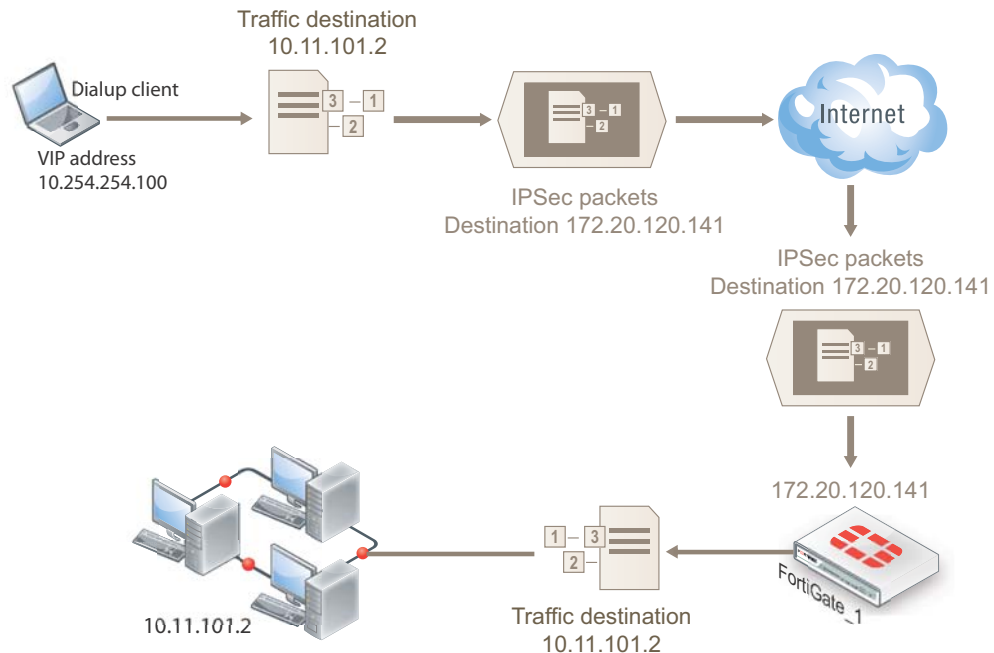
---

FortiGate units fully support [RFC 3456](#). The FortiGate DHCP over IPsec feature can be enabled to allocate VIP addresses to FortiClient dialup clients using a FortiGate DHCP server.

[Figure 265](#) shows an example of a FortiClient-to-FortiGate VPN where the FortiClient application is assigned a VIP on an uncommonly used subnet. The diagram also shows that while the destination for the information in the encrypted packets is the private network behind the FortiGate unit, the destination of the IPsec packets themselves is the public interface of the FortiGate unit that acts as the end of the VPN tunnel.



**Figure 265:** IP address assignments in a FortiClient dialup-client configuration



### Assigning VIPs by RADIUS user group

If you use XAuth authentication, you can assign users the virtual IP address stored in the Framed-IP-Address field of their record on the RADIUS server. (See [RFC 2865](#) and [RFC 2866](#) for more information about RADIUS fields.) To do this:

- Set the DHCP server *IP Assignment Mode* to *User-group defined method*. This is an Advanced setting. See “[To configure a DHCP server on a FortiGate interface](#)” on page 1717.
- Create a new firewall user group and add the RADIUS server to it.
- In your phase 1 settings, configure the FortiGate unit as an XAuth server and select from *User Group* the new user group that you created. For more information, see “[Using the FortiGate unit as an XAuth server](#)” on page 1651.
- Configure the FortiClient application to use XAuth. See “[Adding XAuth authentication](#)” on page 1718.

### FortiClient dialup-client infrastructure requirements

- To support policy-based VPNs, the FortiGate dialup server may operate in either NAT mode or transparent mode. NAT mode is required if you want to create a route-based VPN.
- If the FortiClient dialup clients will be configured to obtain VIP addresses through FortiGate DHCP relay, a DHCP server must be available on the network behind the FortiGate unit and the DHCP server must have a direct route to the FortiGate unit.
- If the FortiGate interface to the private network is not the default gateway, the private network behind the FortiGate unit must be configured to route IP traffic destined for dialup clients back (through an appropriate gateway) to the FortiGate interface to the private network. As an alternative, you can configure the IPsec security policy on the FortiGate unit to perform inbound NAT on IP packets. Inbound NAT translates the source addresses of inbound decrypted packets into the IP address of the FortiGate interface to the local private network.

## FortiClient-to-FortiGate VPN configuration steps

Configuring dialup client capability for FortiClient dialup clients involves the following general configuration steps:

1. If you will be using VIP addresses to identify dialup clients, determine which VIP addresses to use. As a precaution, consider using VIP addresses that are not commonly used.
2. Configure the FortiGate unit to act as a dialup server. See [“Configure the FortiGate unit” on page 1714](#).
3. If the dialup clients will be configured to obtain VIP addresses through DHCP over IPsec, configure the FortiGate unit to act as a DHCP server or to relay DHCP requests to an external DHCP server.
4. Configure the dialup clients. See [“Configure the FortiClient Endpoint Security application” on page 1718](#).



When a FortiGate unit has been configured to accept connections from FortiClient dialup-clients, you can optionally arrange to have an IPsec VPN configuration downloaded to FortiClient dialup clients automatically. For more information, see [“Configuring the FortiGate unit as a VPN policy server” on page 1717](#).

## Configure the FortiGate unit

Configuring the FortiGate unit to establish VPN connections with FortiClient Endpoint Security users involves the following steps:

- Configure the VPN settings
- If the dialup clients use automatic configuration, configure the FortiGate unit as a VPN policy server
- If the dialup clients obtain VIP addresses by DHCP over IPsec, configure an IPsec DHCP server or relay

The procedures in this section cover basic setup of policy-based and route-based VPNs compatible with FortiClient Endpoint Security. A route-based VPN is simpler to configure.

### Configuring FortiGate unit VPN settings

To configure FortiGate unit VPN settings to support FortiClient users, you need to:

- configure the FortiGate Phase 1 VPN settings
  - configure the FortiGate Phase 2 VPN settings
  - add the security policy
1. On the local FortiGate unit, define the phase 1 configuration needed to establish a secure connection with the FortiClient peer. See [“Auto Key phase 1 parameters” on page 1637](#). Enter these settings in particular:

<b>Name</b>	Enter a name to identify the VPN tunnel. This name appears in phase 2 configurations, security policies and the VPN monitor.
<b>Remote Gateway</b>	Select <i>Dialup User</i> .
<b>Local Interface</b>	Select the interface through which clients connect to the FortiGate unit.

<b>Mode</b>	Select <i>Main (ID Protection)</i> .
<b>Authentication Method</b>	Select <i>Pre-shared Key</i> .
<b>Pre-shared Key</b>	Enter the pre-shared key. This must be the same preshared key provided to the FortiClient users.
<b>Peer option</b>	Select <i>Accept any peer ID</i> .
<b>Enable IPsec Interface Mode</b>	You must select <i>Advanced</i> to see this setting. If <i>IPsec Interface Mode</i> is enabled, the FortiGate unit creates a virtual IPsec interface for a route-based VPN.

- Define the phase 2 parameters needed to create a VPN tunnel with the FortiClient peer. See [“Phase 2 parameters” on page 1653](#). Enter these settings in particular:

<b>Name</b>	Enter a name to identify this phase 2 configuration.
<b>Phase 1</b>	Select the name of the phase 1 configuration that you defined.
<b>Advanced</b>	Select to configure the following optional setting.
<b>DHCP-IPsec</b>	Select if you provide virtual IP addresses to clients using DHCP.

- Define names for the addresses or address ranges of the private networks that the VPN links. These addresses are used in the security policies that permit communication between the networks. For more information, see [“Defining policy addresses” on page 1659](#).

Enter these settings in particular:

- Define an address name for the individual address or the subnet address that the dialup users access through the VPN.
- If FortiClient users are assigned VIP addresses, define an address name for the subnet to which these VIPs belong.

- Define security policies to permit communication between the private networks through the VPN tunnel. Route-based and policy-based VPNs require different security policies. For detailed information about creating security policies, see [“Defining VPN security policies” on page 1660](#).

If the security policy, which grants the VPN Connection is limited to certain services, DHCP must be included, otherwise the client won't be able to retrieve a lease from the FortiGate's (IPSec) DHCP server, because the DHCP Request (coming out of the tunnel) will be blocked.

### Route-based VPN security policies

Define an ACCEPT security policy to permit communications between the source and destination addresses.

- Go to *Policy > Policy > Policy* and select *Create New*.
- Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
- Enter these settings in particular:

<b>Incoming Interface</b>	Select the VPN Tunnel (IPsec Interface) you configured in Step 1.
<b>Source Address</b>	Select <i>All</i> .

<b>Outgoing Interface</b>	Select the interface that connects to the private network behind this FortiGate unit.
<b>Destination Address</b>	Select <i>All</i> .
<b>Action</b>	Select <i>ACCEPT</i> .
<b>Enable NAT</b>	Disable.

If you want to allow hosts on the private network to initiate communications with the FortiClient users after the tunnel is established, you need to define a security policy for communication in that direction.

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter these settings in particular:

<b>Incoming Interface</b>	Select the interface that connects to the private network behind this FortiGate unit.
<b>Source Address</b>	Select <i>All</i> .
<b>Outgoing Interface</b>	Select the interface that connects to the private network behind this FortiGate unit.
<b>Destination Address</b>	Select <i>All</i> .
<b>Action</b>	Select <i>ACCEPT</i> .
<b>Enable NAT</b>	Disable.

### Policy-based VPN security policy

Define an IPsec security policy to permit communications between the source and destination addresses.

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Select the *Policy Type* of *VPN* and leave the *Policy Subtype* as *IPsec*.
3. Enter these settings in particular:

<b>Local Interface</b>	Select the interface that connects to the private network behind this FortiGate unit.
<b>Local Protected Subnet</b>	Select the address name that you defined in Step 3 for the private network behind this FortiGate unit.
<b>Outgoing VPN Interface</b>	Select the FortiGate unit's public interface.
<b>Remote Protected Subnet</b>	If FortiClient users are assigned VIPs, select the address name that you defined in Step 3 for the VIP subnet. Otherwise, select <i>All</i> .
<b>VPN Tunnel</b>	Select <i>Use Existing</i> and select the name of the phase 1 configuration that you created in Step 1.  Select <i>Allow traffic to be initiated from the remote site</i> to enable traffic from the remote network to initiate the tunnel.

Place VPN policies in the policy list above any other policies having similar source and destination addresses.

## Configuring the FortiGate unit as a VPN policy server

When a FortiClient application set to automatic configuration connects to the FortiGate unit, the FortiGate unit requests a user name and password. If the user supplies valid credentials, the FortiGate unit downloads the VPN settings to the FortiClient application.

You must do the following to configure the FortiGate unit to work as a VPN policy server for FortiClient automatic configuration:

1. Create user accounts for FortiClient users.
2. Create a user group for FortiClient users and the user accounts that you created in step 1.
3. Connect to the FortiGate unit CLI and configure VPN policy distribution as follows:

```
config vpn ipsec forticlient
 edit <policy_name>
 set phase2name <tunnel_name>
 set usergroupname <group_name>
 set status enable
 end
```

<tunnel\_name> must be the Name you specified in the step 2 of [“Configure the FortiGate unit” on page 1714](#). <group\_name> must be the name of the user group your created for FortiClient users.

## Configuring DHCP services on a FortiGate interface

If the FortiClient dialup clients are configured to obtain a VIP address using DHCP, configure the FortiGate dialup server to either:

- relay DHCP requests to a DHCP server behind the FortiGate unit (see [“To configure DHCP relay on a FortiGate interface”](#) below).
- act as a DHCP server (see [“To configure a DHCP server on a FortiGate interface” on page 1717](#)).

Note that DHCP services are typically configured during the interface creation stage, but you can return to an interface to modify DHCP settings if need be.

### To configure DHCP relay on a FortiGate interface

1. Go to *System > Network > Interfaces* and select the interface that you want to relay DHCP.
2. Under *DHCP Server*, select *Enable* and create a new *DHCP Address Range* and *Netmask*.
3. Open the *Advanced...* menu and select *Relay* for the *Mode* option.
4. For the *Type*, select *IPsec*.
5. Select *OK*.

### To configure a DHCP server on a FortiGate interface

1. Go to *System > Network > Interfaces* and select the interface that you want to act as a DHCP server.
2. Under *DHCP Server*, select *Enable* and create a new *DHCP Address Range* and *Netmask*.
3. For *Default Gateway*, enter the IP address of the default gateway that the DHCP server assigns to DHCP clients.
4. For *DNS Server*, select *Same as System DNS*. If you want to use a different DNS server for VPN clients, select *Specify* and enter an IP address in the available field.

5. Open the *Advanced...* menu and select *Server* for the *Mode* option.
6. For the *Type*, select *IPsec*.
7. Select *OK*.

## Configure the FortiClient Endpoint Security application

The following procedure explains how to configure the FortiClient Endpoint Security application to communicate with a remote FortiGate dialup server using the VIP address that you specify manually. These procedures are based on FortiClient 5.0.

### Configuring FortiClient

This procedure explains how to configure the FortiClient application manually using the default IKE and IPsec settings. For more information, refer to the [FortiClient Administration Guide](#).

This procedure includes instructions for configuring a virtual IP for the FortiClient application, either manually or using DHCP over IPsec.

#### To create a FortiClient VPN configuration

1. Go to *Remote Access* and select the down-arrow for the VPN connection.
2. Select *Add new connection* and complete following information:

<b>VPN Type</b>	Select <i>IPsec VPN</i> .
<b>Connection Name</b>	Enter a descriptive name for the connection.
<b>Remote Gateway</b>	Enter the IP address or the fully qualified domain name (FQDN) of the remote gateway.
<b>Authentication Method</b>	Select <i>Pre-shared Key</i> .
<b>Pre-shared Key</b>	Enter the pre-shared key.
<b>User Name</b>	Enter the user name to connect to the tunnel.

3. Select *OK*.

## Adding XAuth authentication

Extended Authentication (XAuth) increases security by requiring additional user authentication in a separate exchange at the end of the VPN phase 1 negotiation. The FortiGate unit challenges the user for a user name and password. It then forwards the user's credentials to an external RADIUS or LDAP server for verification.

Implementation of XAuth requires configuration at both the FortiGate unit and the FortiClient application. For information about configuring a FortiGate unit as an XAuth server, see [“Using the FortiGate unit as an XAuth server” on page 1651](#). The following procedure explains how to configure the FortiClient application.


Note that XAuth is not compatible with IKE version 2.

For more information on configuring XAuth authentication, see the [FortiClient Administration Guide](#).

## FortiClient dialup-client configuration example

This example demonstrates how to set up a FortiClient dialup-client IPsec VPN that uses preshared keys for authentication purposes. In the example configuration, the DHCP over IPsec feature is enabled in the FortiClient Endpoint Security application so that the FortiClient Endpoint Security application can acquire a VIP address through the FortiGate DHCP server. Both route-based and policy-based solutions are covered.

**Figure 266:**Example FortiClient dialup-client configuration



In the example configuration:

- VIP addresses that are not commonly used (in this case, 10.254.254.0/24) are assigned to the FortiClient dialup clients using a DHCP server.
- The dialup clients have access to the LAN behind FortiGate\_1.
- The other network devices are assigned IP addresses as shown in [Figure 266](#).

### Configuring FortiGate\_1

When a FortiGate unit receives a connection request from a dialup client, it uses IPsec phase 1 parameters to establish a secure connection and authenticate the client. Then, if the security policy permits the connection, the FortiGate unit establishes the tunnel using IPsec phase 2 parameters and applies the IPsec security policy. Key management, authentication, and security services are negotiated dynamically through the IKE protocol.

To support these functions, the following general configuration steps must be performed at the FortiGate unit:

- Define the phase 1 parameters that the FortiGate unit needs to authenticate the dialup clients and establish a secure connection. See [“To define the phase 1 parameters” on page 1720](#).
- Define the phase 2 parameters that the FortiGate unit needs to create a VPN tunnel and enable all dialup clients having VIP addresses on the 10.254.254.0/24 network to connect using the same tunnel definition. See [“To define the phase 2 parameters” on page 1720](#).
- Create security policy to control the permitted services and permitted direction of traffic between the IP source address and the dialup clients. See [“To define the firewall addresses” on page 1721](#).
- Configure the FortiGate unit to service DHCP requests from dialup clients. See [“Configuring the FortiClient Endpoint Security application” on page 1722](#).

### To define the phase 1 parameters

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. Select *Create Phase 1*, enter the following information, and select *OK*:

<b>Name</b>	todialups
<b>Remote Gateway</b>	Dialup User
<b>Local Interface</b>	Port 1
<b>Mode</b>	Main
<b>Authentication Method</b>	Preshared Key
<b>Pre-shared Key</b>	hardtoguess
<b>Peer Options</b>	Accept any peer ID
<b>Advanced</b>	Select
<b>Enable IPsec Interface Mode</b>	Enable for route-based VPN. Disable for policy-based VPN.

### To define the phase 2 parameters

1. Go to *VPN > IPsec > Auto Key (IKE)* and select *Create Phase 2*.
2. Select *Advanced*, enter the following information, and select *OK*:

<b>Name</b>	td_2
<b>Phase 1</b>	todialups
<b>Advanced</b>	DHCP-IPsec



### To define the firewall addresses

1. Go to *Firewall Objects > Address > Addresses*.
2. Select *Create New*, enter the following information, and select *OK*:

<b>Name</b>	internal_net
<b>Type</b>	Subnet
<b>Subnet/IP Range</b>	10.11.101.0/24
<b>Interface</b>	Port 2

3. Select *Create New*, enter the following information, and select *OK*:

<b>Name</b>	dialups
<b>Type</b>	IP Range
<b>Subnet/IP Range</b>	10.254.254.1-10.254.254.10
<b>Interface</b>	Route-based VPN: todialups Policy-based VPN: Any

The security policies for route-based and policy-based VPNs are described in separate sections below.

### To define security policies - route-based VPN

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	todialups
<b>Source Address</b>	dialups
<b>Outgoing Interface</b>	Port 2
<b>Destination Address</b>	internal_net
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Disable

4. Select *Create New*.
5. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
6. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	Port 2
<b>Source Address</b>	internal_net
<b>Outgoing Interface</b>	todialups
<b>Destination Address</b>	dialups

<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Disable

7. Select *Create New*.
8. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
9. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	Port 2
<b>Source Address</b>	internal_net
<b>Outgoing Interface</b>	todialups
<b>Destination Address</b>	all
<b>Service</b>	DHCP
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Disable

10. Place these policies in the policy list above any other policies having similar source and destination addresses.

The policy in step 7 is required for DHCP to function properly for policy-based VPNs. You can omit this policy if you change the *Destination Address Name* to *all* in the step before. Route-based policies are not affected by this.

### To define the security policy - policy-based VPN

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Select the *Policy Type* of *VPN* and leave the *Policy Subtype* as *IPsec*.
3. Enter the following information, and select *OK*:

<b>Local Interface</b>	Port 2
<b>Local Protected Subnet</b>	internal_net
<b>Outgoing VPN Interface</b>	Port 1
<b>Remote Protected Subnet</b>	dialups
<b>VPN Tunnel</b>	Select <i>Use Existing</i> and select <i>todialups</i> from the drop-down list.
<b>Allow traffic to be initiated from the remote site</b>	Enable

4. Place the policy in the policy list above any other policies having similar source and destination addresses.

## Configuring the FortiClient Endpoint Security application

The following procedure explains how to configure the FortiClient Endpoint Security application to connect to FortiGate\_1 and broadcast a DHCP request. The dialup client uses the VIP

address acquired through FortiGate DHCP relay as its IP source address for the duration of the connection.

**To configure FortiClient**

1. Go to *Remote Access* and select the down-arrow for the VPN connection.
2. Select *Add new connection* and complete following information:

<b>VPN Type</b>	Select <i>IPsec VPN</i> .
<b>Connection Name</b>	Headquarters.
<b>Remote Gateway</b>	The port1 IP address.
<b>Authentication Method</b>	Select <i>Pre-shared Key</i> .
<b>Pre-shared Key</b>	hardtoguess
<b>User Name</b>	Enter the user name to connect to the tunnel.

3. Select *OK*.

# FortiGate dialup-client configurations

This section explains how to set up a FortiGate dialup-client IPsec VPN. In a FortiGate dialup-client configuration, a FortiGate unit with a static IP address acts as a dialup server and a FortiGate unit having a dynamic IP address initiates a VPN tunnel with the FortiGate dialup server.


The following topics are included in this section:

- [Configuration overview](#)
- [FortiGate dialup-client configuration steps](#)
- [Configure the server to accept FortiGate dialup-client connections](#)
- [Configure the FortiGate dialup client](#)

## Configuration overview

A dialup client can be a FortiGate unit. The FortiGate dialup client typically obtains a dynamic IP address from an ISP through the Dynamic Host Configuration Protocol (DHCP) or Point-to-Point Protocol over Ethernet (PPPoE) before initiating a connection to a FortiGate dialup server.

**Figure 267:**Example FortiGate dialup-client configuration



In a dialup-client configuration, the FortiGate dialup server does not rely on a phase 1 remote gateway address to establish an IPsec VPN connection with dialup clients. As long as authentication is successful and the IPsec security policy associated with the tunnel permits access, the tunnel is established.

Several different ways to authenticate dialup clients and restrict access to private networks based on client credentials are available. To authenticate FortiGate dialup clients and help to distinguish them from FortiClient dialup clients when multiple clients will be connecting to the VPN through the same tunnel, best practices dictate that you assign a unique identifier (local ID

or peer ID) to each FortiGate dialup client. For more information, see [“Authenticating remote peers and clients” on page 1642](#).



Whenever you add a unique identifier (local ID) to a FortiGate dialup client for identification purposes, you must select Aggressive mode on the FortiGate dialup server and also specify the identifier as a peer ID on the FortiGate dialup server. For more information, see [“Enabling VPN access with user accounts and pre-shared keys” on page 1645](#).

Users behind the FortiGate dialup server cannot initiate the tunnel because the FortiGate dialup client does not have a static IP address. After the tunnel is initiated by users behind the FortiGate dialup client, traffic from the private network behind the FortiGate dialup server can be sent to the private network behind the FortiGate dialup client.

Encrypted packets from the FortiGate dialup client are addressed to the public interface of the dialup server. Encrypted packets from the dialup server are addressed either to the public IP address of the FortiGate dialup client (if the dialup client connects to the Internet directly), or if the FortiGate dialup client is behind a NAT device, encrypted packets from the dialup server are addressed to the public IP address of the NAT device.

If a router with NAT capabilities is in front of the FortiGate dialup client, the router must be NAT-T compatible for encrypted traffic to pass through the NAT device. For more information, see [“NAT traversal” on page 1649](#).

When the FortiGate dialup server decrypts a packet from the FortiGate dialup client, the source address in the IP header may be one of the following values, depending on the configuration of the network at the far end of the tunnel:

- If the FortiGate dialup client connects to the Internet directly, the source address will be the private IP address of a host or server on the network behind the FortiGate dialup client.
- If the FortiGate dialup client is behind a NAT device, the source address will be the public IP address of the NAT device.


In some cases, computers on the private network behind the FortiGate dialup client may (by co-incidence) have IP addresses that are already used by computers on the network behind the FortiGate dialup server. In this type of situation (ambiguous routing), conflicts may occur in one or both of the FortiGate routing tables and traffic destined for the remote network through the tunnel may not be sent.

In many cases, computers on the private network behind the FortiGate dialup client will most likely obtain IP addresses from a local DHCP server behind the FortiGate dialup client. However, unless the local and remote networks use different private network address spaces, unintended ambiguous routing and IP-address overlap issues may arise.

To avoid these issues, you can configure FortiGate DHCP relay on the dialup client instead of using a DHCP server on the network behind the dialup client. The FortiGate dialup client can be configured to relay DHCP requests from the local private network to a DHCP server that resides on the network behind the FortiGate dialup server (see [Figure 268 on page 1726](#)). You configure the FortiGate dialup client to pass traffic from the local private network to the remote network by enabling FortiGate DHCP relay on the FortiGate dialup client interface that is connected to the local private network.

Afterward, when a computer on the network behind the dialup client broadcasts a DHCP request, the dialup client relays the message through the tunnel to the remote DHCP server. The remote DHCP server responds with a private IP address for the computer. To avoid ambiguous routing and network overlap issues, the IP addresses assigned to computers behind the dialup client cannot match the network address space used by the private network behind the FortiGate dialup server.

**Figure 268:** Preventing network overlap in a FortiGate dialup-client configuration



When the DHCP server resides on the private network behind the FortiGate dialup server, the IP destination address specified in the IPsec security policy on the FortiGate dialup client must refer to that network.



You must add a static route to the DHCP server FortiGate unit if it is not directly connected to the private network behind the FortiGate dialup server; its IP address does not match the IP address of the private network. Also, the destination address in the IPsec security policy on the FortiGate dialup client must refer to the DHCP server address. The DHCP server must be configured to assign a range of IP addresses different from the DHCP server's local network, and also different from the private network addresses behind the FortiGate dialup server. See [“Routing” on page 1695](#).

## FortiGate dialup-client infrastructure requirements

The requirements are:

- The FortiGate dialup server must have a static public IP address.
- NAT mode is required if you want to create a route-based VPN.
- The FortiGate dialup server may operate in either NAT mode or transparent mode to support a policy-based VPN.
- Computers on the private network behind the FortiGate dialup client can obtain IP addresses either from a DHCP server behind the FortiGate dialup client, or a DHCP server behind the FortiGate dialup server.
  - If the DHCP server resides on the network behind the dialup client, the DHCP server must be configured to assign IP addresses that do not match the private network behind the FortiGate dialup server.
  - If the DHCP server resides on the network behind the FortiGate dialup server, the DHCP server must be configured to assign IP addresses that do not match the private network behind the FortiGate dialup client.

## FortiGate dialup-client configuration steps

The procedures in this section assume that computers on the private network behind the FortiGate dialup client obtain IP addresses from a local DHCP server. The assigned IP addresses do not match the private network behind the FortiGate dialup server.



In situations where IP-address overlap between the local and remote private networks is likely to occur, FortiGate DHCP relay can be configured on the FortiGate dialup client to relay DHCP requests to a DHCP server behind the FortiGate dialup server. For more information, see [“To configure DHCP relay on a FortiGate interface” on page 1717](#).

Configuring dialup client capability for FortiGate dialup clients involves the following general configuration steps:

- Determine which IP addresses to assign to the private network behind the FortiGate dialup client, and add the IP addresses to the DHCP server behind the FortiGate dialup client. Refer to the software supplier’s documentation to configure the DHCP server.
- Configure the FortiGate dialup server. See [“Configure the server to accept FortiGate dialup-client connections” on page 1727](#).
- Configure the FortiGate dialup client. See [“Configure the FortiGate dialup client” on page 1729](#).

## Configure the server to accept FortiGate dialup-client connections

Before you begin, optionally reserve a unique identifier (peer ID) for the FortiGate dialup client. The dialup client will supply this value to the FortiGate dialup server for authentication purposes during the IPsec phase 1 exchange. In addition, the value will enable you to distinguish FortiGate dialup-client connections from FortiClient dialup-client connections. The same value must be specified on the dialup server and on the dialup client.

1. At the FortiGate dialup server, define the phase 1 parameters needed to authenticate the FortiGate dialup client and establish a secure connection. See [“Auto Key phase 1 parameters” on page 1637](#). Enter these settings in particular:

<b>Name</b>	Enter a name to identify the VPN tunnel. This name appears in phase 2 configurations, security policies and the VPN monitor.
<b>Remote Gateway</b>	Select <i>Dialup User</i> .
<b>Local Interface</b>	Select the interface through which clients connect to the FortiGate unit.
<b>Mode</b>	If you will be assigning an ID to the FortiGate dialup client, select <i>Aggressive</i> .

---

<b>Peer Options</b>	If you will be assigning an ID to the FortiGate dialup client, select <i>Accept this peer ID</i> and type the identifier that you reserved for the FortiGate dialup client into the adjacent field.
---------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

---

<b>Enable IPsec Interface Mode</b>	You must select <i>Advanced</i> to see this setting. If <i>IPsec Interface Mode</i> is enabled, the FortiGate unit creates a virtual IPsec interface for a route-based VPN. Disable this option if you want to create a policy-based VPN.  After you select <i>OK</i> to create the phase 1 configuration, you cannot change this setting.
------------------------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

---

2. Define the phase 2 parameters needed to create a VPN tunnel with the FortiGate dialup client. See [“Phase 2 parameters” on page 1653](#). Enter these settings in particular:

---

<b>Name</b>	Enter a name to identify this phase 2 configuration.
-------------	------------------------------------------------------

---

<b>Phase 1</b>	Select the name of the phase 1 configuration that you defined.
----------------	----------------------------------------------------------------

---

3. Define names for the addresses or address ranges of the private networks that the VPN links. See [“Defining policy addresses” on page 1659](#). Enter these settings in particular:
  - Define an address name for the server, host, or network behind the FortiGate dialup server.
  - Define an address name for the private network behind the FortiGate dialup client.
4. Define the security policies to permit communications between the private networks through the VPN tunnel. Route-based and policy-based VPNs require different security policies. For detailed information about creating security policies, see [“Defining VPN security policies” on page 1660](#).

### Route-based VPN security policy

Define an ACCEPT security policy to permit communications between hosts on the private network behind the FortiGate dialup client and the private network behind this FortiGate dialup server. Because communication cannot be initiated in the opposite direction, there is only one policy.

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter these settings in particular:

---

<b>Incoming Interface</b>	Select the VPN tunnel (IPsec interface) created in Step 1.
---------------------------	------------------------------------------------------------

---

<b>Source Address</b>	Select <i>All</i> .
-----------------------	---------------------

---

<b>Outgoing Interface</b>	Select the interface that connects to the private network behind this FortiGate unit.
---------------------------	---------------------------------------------------------------------------------------

---

<b>Destination Address</b>	Select <i>All</i> .
----------------------------	---------------------

---

<b>Action</b>	Select <i>ACCEPT</i> .
---------------	------------------------

---

<b>Enable NAT</b>	Disable
-------------------	---------

---



## Policy-based VPN security policy

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Select the *Policy Type* of *VPN* and leave the *Policy Subtype* as *IPsec*.
3. Enter these settings in particular:

<b>Local Interface</b>	Select the interface that connects to the private network behind this FortiGate unit.
<b>Local Protected Subnet</b>	Select the address name that you defined in Step 3 for the private network behind this FortiGate unit.
<b>Outgoing VPN Interface</b>	Select the FortiGate unit's public interface.
<b>Remote Protected Subnet</b>	Select the address name that you defined in Step 3.
<b>VPN Tunnel</b>	Select <i>Use Existing</i> and select the name of the phase 1 configuration that you created in Step 1. from the drop-down list.  Select <i>Allow traffic to be initiated from the remote site</i> to enable traffic from the remote network to initiate the tunnel.  Clear <i>Allow outbound</i> to prevent traffic from the local network from initiating the tunnel after the tunnel has been established.

4. To prevent traffic from the local network from initiating the tunnel after the tunnel has been established, you need to disable the outbound VPN traffic in the CLI

```
config firewall policy
 edit <policy_number>
 set outbound disable
 end
```

Place the policy in the policy list above any other policies having similar source and destination addresses.

If configuring a route-based policy, configure a default route for VPN traffic on this interface.

## Configure the FortiGate dialup client

Configure the FortiGate dialup client.

1. At the FortiGate dialup client, define the phase 1 parameters needed to authenticate the dialup server and establish a secure connection. See [“Auto Key phase 1 parameters” on page 1637](#). Enter these settings in particular:

<b>Name</b>	Enter a name to identify the VPN tunnel.
<b>Remote Gateway</b>	Select <i>Static IP Address</i> .
<b>IP Address</b>	Type the IP address of the dialup server's public interface.
<b>Local Interface</b>	Select the interface that connects to the public network.

<b>Mode</b>	The FortiGate dialup client has a dynamic IP address, select <i>Aggressive</i> .
<b>Advanced</b>	Select to view the following options.
<b>Local ID</b>	If you defined a peer ID for the dialup client in the FortiGate dialup server configuration, enter the identifier of the dialup client. The value must be identical to the peer ID that you specified previously in the FortiGate dialup server configuration.
<b>Enable IPsec Interface Mode</b>	If <i>IPsec Interface Mode</i> is enabled, the FortiGate unit creates a virtual IPsec interface for a route-based VPN. Disable this option if you want to create a policy-based VPN.  After you select <i>OK</i> to create the phase 1 configuration, you cannot change this setting.

2. Define the phase 2 parameters needed to create a VPN tunnel with the dialup server. See “Phase 2 parameters” on page 1653. Enter these settings in particular:

<b>Name</b>	Enter a name to identify this phase 2 configuration.
<b>Phase 1</b>	Select the name of the phase 1 configuration that you defined.

3. Define names for the addresses or address ranges of the private networks that the VPN links. See “Defining policy addresses” on page 1659. Enter these settings in particular:
  - Define an address name for the server, host, or network behind the FortiGate dialup server.
  - Define an address name for the private network behind the FortiGate dialup client.
4. Define security policies to permit communication between the private networks through the VPN tunnel. Route-based and policy-based VPNs require different security policies. For detailed information about creating security policies, see “Defining VPN security policies” on page 1660.

### Route-based VPN security policy

Define an ACCEPT security policy to permit communications between hosts on the private network behind this FortiGate dialup client and the private network behind the FortiGate dialup server. Because communication cannot be initiated in the opposite direction, there is only one policy.

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* of *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter these settings in particular:

<b>Incoming Interface</b>	Select the interface that connects to the private network behind this FortiGate unit.
<b>Source Address</b>	Select <i>All</i> .
<b>Outgoing Interface</b>	Select the VPN tunnel (IPsec interface) created in Step 1.
<b>Destination Address</b>	Select <i>All</i> .

<b>Action</b>	Select <i>ACCEPT</i> .
<b>Enable NAT</b>	Disable

### Policy-based VPN security policy

Define an IPsec security policy to permit communications between the source and destination addresses.

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Select the *Policy Type* of *VPN* and leave the *Policy Subtype* as *IPsec*.
3. Enter these settings in particular:

<b>Local Interface</b>	Select the interface that connects to the private network behind this FortiGate unit.
<b>Local Protected Subnet</b>	Select the address name that you defined in Step 3 for the private network behind this FortiGate unit.
<b>Outgoing Interface</b>	Select the FortiGate unit's public interface.
<b>Remote Protected Subnet</b>	Select the address name that you defined in Step 3 for the private network behind the dialup server.
<b>VPN Tunnel</b>	Select <i>Use Existing</i> and select the name of the phase 1 configuration that you created in Step 1 from the drop-down list.  Clear <i>Allow traffic to be initiated from the remote site</i> to prevent traffic from the remote network from initiating the tunnel after the tunnel has been established.

Place the policy in the policy list above any other policies having similar source and destination addresses.

# Supporting IKE Mode config clients

IKE Mode Config is an alternative to DHCP over IPsec. A FortiGate unit can be configured as either an IKE Mode Config server or client. This chapter contains the following sections:

- [Automatic configuration overview](#)
- [IKE Mode Config overview](#)
- [Configuring IKE Mode Config](#)
- [Example: FortiGate unit as IKE Mode Config server](#)
- [Example: FortiGate unit as IKE Mode Config client](#)

## Automatic configuration overview

VPN configuration for remote clients is simpler if it is automated. Several protocols support automatic configuration:

- The Fortinet FortiClient Endpoint Security application can completely configure a VPN connection with a suitably configured FortiGate unit given only the FortiGate unit's address. This protocol is exclusive to Fortinet. For more information, see the "[FortiClient dialup-client configurations](#)" chapter.
- DHCP over IPsec can assign an IP address, Domain, DNS and WINS addresses. The user must first configure IPsec parameters such as gateway address, encryption and authentication algorithms.
- IKE Mode Config can configure host IP address, Domain, DNS and WINS addresses. The user must first configure IPsec parameters such as gateway address, encryption and authentication algorithms. Several network equipment vendors support IKE Mode Config, which is described in the ISAKMP Configuration Method document [draft-dukes-ike-mode-cfg-02.txt](#).

This chapter describes how to configure a FortiGate unit as either an IKE Mode Config server or client.

## IKE Mode Config overview

Dialup VPN clients connect to a FortiGate unit that acts as a VPN server, providing the client the necessary configuration information to establish a VPN tunnel. The configuration information typically includes a virtual IP address, netmask, and DNS server address.

IKE Mode Config is available only for VPNs that are route-based, also known as interface-based. A FortiGate unit can function as either an IKE Configuration Method server or client. IKE Mode Config is configurable only in the CLI.

## Configuring IKE Mode Config

IKE Mode Config is configured with the CLI command `config vpn ipsec phase1-interface`. The `mode-cfg` variable enables IKE Mode Config. The `type` field determines whether you are creating an IKE Mode Config server or a client. Setting `type` to `dynamic` creates a server configuration, otherwise the configuration is a client.

## Configuring an IKE Mode Config client

If the FortiGate unit will connect as a dialup client to a remote gateway that supports IKE Mode Config, the relevant `vpn ipsec phase1-interface` variables are as follows:

Variable	Description
<code>ike-version 1</code>	IKE v1 is the default for FortiGate IPsec VPNs. IKE Mode Config is also compatible with IKE v2 (RFC 4306).
<code>mode-cfg enable</code>	Enable IKE Mode Config.
<code>type {ddns   static}</code>	If you set <code>type</code> to <code>dynamic</code> , an IKE Mode Config server is created.
<code>assign-ip {enable   disable}</code>	Enable to request an IP address from the server.
<code>interface &lt;interface_name&gt;</code>	This is a regular IPsec VPN field. Specify the physical, aggregate, or VLAN interface to which the IPsec tunnel will be bound.
<code>proposal &lt;encryption_combination&gt;</code>	This is a regular IPsec VPN field that determines the encryption and authentication settings that the client will accept. For more information, see <a href="#">“Defining IKE negotiation parameters” on page 1646</a> .
<code>mode-cfg-ip-version {4 6}</code>	Select if the Method client receives an IPv4 or IPv6 IP address. The default is 4. the <code>ip-version</code> setting matches this variable’s value.
<code>ip-version &lt;4   6&gt;</code>	This is a regular IPsec VPN field. By default, IPsec VPNs use IPv4 addressing. You can set <code>ip-version</code> to 6 to create a VPN with IPv6 addressing.

For a complete list of available variables, see the [CLI Reference](#).

## Configuring an IKE Mode Config server

If the FortiGate unit will accept connection requests from dialup clients that support IKE Mode Config, the following `vpn ipsec phase1-interface` settings are required before any other configuration is attempted:

Variable	Description
<code>ike-version 1</code>	IKE v1 is the default for FortiGate IPsec VPNs. IKE Mode Config is also compatible with IKE v2 (RFC 4306).
<code>mode-cfg enable</code>	Enable IKE Mode Config.
<code>type dynamic</code>	Any other setting creates an IKE Mode Config client.
<code>interface &lt;interface_name&gt;</code>	This is a regular IPsec VPN field. Specify the physical, aggregate, or VLAN interface to which the IPsec tunnel will be bound.

Variable	Description
proposal <encryption_combination>	This is a regular IPsec VPN field that determines the encryption and authentication settings that the server will accept. For more information, see <a href="#">“Defining IKE negotiation parameters” on page 1646</a> .
ip-version <4   6>	This is a regular IPsec VPN field. By default, IPsec VPNs use IPv4 addressing. You can set <code>ip-version</code> to 6 to create a VPN with IPv6 addressing.

For a complete list of available variables, see the [CLI Reference](#).

After you have enabled the basic configuration, you can configure:

- IP address assignment for clients
- DNS and WINS server assignment

### IP address assignment

Usually you will want to assign IP addresses to clients. The simplest method is to assign addresses from a specific range, similar to a DHCP server.

If your clients are authenticated by a RADIUS server, you can obtain the user’s IP address assignment from the Framed-IP-Address attribute. The user must be authenticated using XAuth.

#### To assign IP addresses from an address range

If your VPN uses IPv4 addresses,

```
config vpn ipsec phase1-interface
 edit vpn1
 set mode-cfg-ipversion 4
 set assign-ip enable
 set assign-ip-type ip
 set assign-ip-from range
 set ipv4-start-ip <range_start>
 set ipv4-end-ip <range_end>
 set ipv4-netmask <netmask>
 end
```

If your VPN uses IPv6 addresses,

```
config vpn ipsec phase1-interface
 edit vpn1
 set mode-cfg-ipversion 6
 set assign-ip enable
 set assign-ip-type ip
 set assign-ip-from range
 set ipv6-start-ip <range_start>
 set ipv6-end-ip <range_end>
 end
```

### To assign IP addresses from a RADIUS server

The users must be authenticated by a RADIUS server and assigned to the FortiGate user group <grpname>. Since the IP address will not be static, `type` is set to `dynamic`, and `mode-cfg` is enabled. This is IKE Configuration Method so that compatible clients can configure themselves with settings that the FortiGate unit provides.

```
config vpn ipsec phase1-interface
 edit vpn1
 set type dynamic
 set mode-cfg enable
 set assign-ip enable
 set assign-ip-from usrgrp
 set xauthtype auto
 set authusrgrp <grpname>
 end
```

## Example: FortiGate unit as IKE Mode Config server

In this example, the FortiGate unit assigns IKE Mode Config clients addresses in the range of 10.11.101.160 through 10.11.101.180. DNS and WINS server addresses are also provided. The public interface of the FortiGate unit is Port 1.

The `ipv4-split-include` variable specifies a firewall address that represents the networks to which the clients will have access. This destination IP address information is sent to the clients.

Only the CLI fields required for IKE Mode Config are shown here. For detailed information about these variables, see the [FortiGate CLI Reference](#).

```
config vpn ipsec phase1-interface
 edit vpn1
 set ip-version 4
 set type dynamic
 set interface port1
 set proposal 3des-sha1 aes128-sha1
 set mode-cfg enable
 set mode-cfg-ipversion 4
 set assign-ip enable
 set assign-ip-type ip
 set assign-ip-from range
 set ipv4-start-ip 10.11.101.160
 set ipv4-end-ip 10.11.101.180
 set ipv4-netmask 255.255.255.0
 set dns-server1 10.11.101.199
 set dns-server2 66.11.168.195
 set wins-server1 10.11.101.191
 set domain example
 set ipv4-split-include OfficeLAN
 end
```

## Example: FortiGate unit as IKE Mode Config client

In this example, the FortiGate unit connects to a VPN gateway with a static IP address that can be reached through Port 1. Only the port, gateway and proposal information needs to be configured. All other configuration information will come from the IKE Mode Config server.

```
config vpn ipsec phase1-interface
 edit vpn1
 set ip-version 4
 set type static
 set remote-gw <gw_address>
 set interface port 1
 set proposal 3des-sha1 aes128-sha1
 set mode-cfg enable
 set mode-cfg-ipversion 4
 set assign-ip enable
 end
```



# Internet-browsing configuration

This section explains how to support secure web browsing performed by dialup VPN clients, and/or hosts behind a remote VPN peer. Remote users can access the private network behind the local FortiGate unit and browse the Internet securely. All traffic generated remotely is subject to the security policy that controls traffic on the private network behind the local FortiGate unit.

The following topics are included in this section:


- [Configuration overview](#)
- [Creating an Internet browsing security policy](#)
- [Routing all remote traffic through the VPN tunnel](#)

## Configuration overview

A VPN provides secure access to a private network behind the FortiGate unit. You can also enable VPN clients to access the Internet securely. The FortiGate unit inspects and processes all traffic between the VPN clients and hosts on the Internet according to the Internet browsing policy. This is accomplished even though the same FortiGate interface is used for both encrypted VPN client traffic and unencrypted Internet traffic.

In [Figure 269](#), FortiGate\_1 enables secure Internet browsing for FortiClient Endpoint Security users such as Dialup\_1 and users on the Site\_2 network behind FortiGate\_2, which could be a VPN peer or a dialup client.

**Figure 269:**Example Internet-browsing configuration



You can adapt any of the following configurations to provide secure Internet browsing:

- a gateway-to-gateway configuration (see [“Gateway-to-gateway configurations”](#) on page 1665)
- a FortiClient dialup-client configuration (see [“FortiClient dialup-client configurations”](#) on page 1709)
- a FortiGate dialup-client configuration (see [“FortiGate dialup-client configurations”](#) on page 1724)

The procedures in this section assume that one of these configurations is in place, and that it is operating properly.

To create an internet-browsing configuration based on an existing gateway-to-gateway configuration, you must edit the gateway-to-gateway configuration as follows:

- On the FortiGate unit that will provide Internet access, create an Internet browsing security policy. See [“Creating an Internet browsing security policy”](#), below.
- Configure the remote peer or client to route all traffic through the VPN tunnel. You can do this on a FortiGate unit or on a FortiClient Endpoint Security application. See [“Routing all remote traffic through the VPN tunnel”](#) on page 1739.

## Creating an Internet browsing security policy

On the FortiGate unit that acts as a VPN server and will provide secure access to the Internet, you must create an Internet browsing security policy. This policy differs depending on whether your gateway-to-gateway configuration is policy-based or route-based.

### To create an Internet browsing policy - policy-based VPN

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Select the *Policy Type* as *VPN* and leave the *Policy Subtype* as *IPsec*.
3. Enter the following information and then select *OK*:

<b>Local Interface</b>	The interface to which the VPN tunnel is bound.
<b>Local Protected Subnet</b>	All
<b>Outgoing VPN Interface</b>	The interface to which the VPN tunnel is bound.
<b>Remote Protected Subnet</b>	The internal range of address of the remote spoke site.
<b>VPN Tunnel</b>	Select <i>Use Existing</i> and select the tunnel that provides access to the private network behind the FortiGate unit.
<b>Allow traffic to be initiated from the remote site</b>	Enable
<b>Inbound NAT</b>	Enable

4. Enable inbound NAT in the CLI.

```
config firewall policy
 edit <policy_number>
 set natinbound enable
 end
```

### To create an Internet browsing policy - route-based VPN

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter the following information and then select *OK*:

<b>Incoming Interface</b>	The IPsec VPN interface.
<b>Source Address</b>	All

<b>Outgoing Interface</b>	The interface that connects to the Internet. The virtual IPsec interface is configured on this physical interface.
<b>Destination Address</b>	The internal range of address of the remote spoke site.
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Enable

The VPN clients must be configured to route all Internet traffic through the VPN tunnel.

## Routing all remote traffic through the VPN tunnel

To make use of the Internet browsing configuration on the VPN server, the VPN peer or client must route all traffic through the VPN tunnel. Usually, only the traffic destined for the private network behind the FortiGate VPN server is sent through the tunnel.

The remote end of the VPN can be a FortiGate unit that acts as a peer in a gateway-to-gateway configuration, or a FortiClient application that protects an individual client PC.

- To configure a remote peer FortiGate unit for Internet browsing via VPN, see [“Configuring a FortiGate remote peer to support Internet browsing”](#).
- To configure a FortiClient Endpoint Security application for Internet browsing via VPN, see [“Configuring a FortiClient application to support Internet browsing”](#) on page 1740.

These procedures assume that your VPN connection to the protected private network is working and that you have configured the FortiGate VPN server for Internet browsing as described in [“Creating an Internet browsing security policy”](#) on page 1738.

### Configuring a FortiGate remote peer to support Internet browsing

The configuration changes to send all traffic through the VPN differ for policy-based and route-based VPNs.

#### To route all traffic through a policy-based VPN

1. At the FortiGate dialup client, go to *Policy > Policy > Policy*.
2. Select the IPsec security policy and then select *Edit*.
3. From the *Remote Protected Subnet* list, select *all*.
4. Select *OK*.

Packets are routed through the VPN tunnel, not just those destined for the protected private network.

#### To route all traffic through a route-based VPN

1. At the FortiGate dialup client, go to *Router > Static > Static Routes*.
2. On a low-end FortiGate unit, go to *System > Network > Routing*.
3. Select the default route (destination IP 0.0.0.0) and then select *Edit*. If there is no default route, select *Create New*. Enter the following information and select *OK*:

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Device</b>	Select the IPsec virtual interface.
<b>Distance</b>	Leave at default.

All packets are routed through the VPN tunnel, not just packets destined for the protected private network.

## Configuring a FortiClient application to support Internet browsing

By default, the FortiClient application configures the PC so that traffic destined for the remote protected network passes through the VPN tunnel but all other traffic is sent to the default gateway. You need to modify the FortiClient settings so that it configures the PC to route all outbound traffic through the VPN.

### To route all traffic through VPN - FortiClient application

1. At the remote host, start FortiClient.
2. Go to *VPN > Connections*.
3. Select the definition that connects FortiClient to the FortiGate dialup server.
4. Select *Advanced* and then select *Edit*.
5. In the *Edit Connection* dialog box, select *Advanced*.
6. In the *Remote Network* group, select *Add*.
7. In the *IP* and *Subnet Mask* fields, type `0.0.0.0/0.0.0.0` and select *OK*.

The address is added to the *Remote Network* list. The first destination IP address in the list establishes a VPN tunnel. The second destination address (`0.0.0.0/0.0.0.0` in this case) forces all other traffic through the VPN tunnel.

8. Select *OK*.

# Redundant VPN configurations

This section discusses the options for supporting redundant and partially redundant IPsec VPNs, using route-based approaches.

The following topics are included in this section:

- [Configuration overview](#)
- [General configuration steps](#)
- [Configure the VPN peers - route-based VPN](#)
- [Redundant route-based VPN configuration example](#)
- [Partially-redundant route-based VPN example](#)
- [Creating a backup IPsec interface](#)

## Configuration overview

A FortiGate unit with two interfaces connected to the Internet can be configured to support redundant VPNs to the same remote peer. If the primary connection fails, the FortiGate unit can establish a VPN using the other connection.

Redundant tunnels do not support Tunnel Mode. You must use Interface Mode.

A fully-redundant configuration requires redundant connections to the Internet on both peers. [Figure 270 on page 1742](#) shows an example of this. This is useful to create a reliable connection between two FortiGate units with static IP addresses.

When only one peer has redundant connections, the configuration is partially-redundant. For an example of this, see [“Partially-redundant route-based VPN example” on page 1758](#). This is useful to provide reliable service from a FortiGate unit with static IP addresses that accepts connections from dialup IPsec VPN clients.


In a fully-redundant VPN configuration with two interfaces on each peer, four distinct paths are possible for VPN traffic from end to end. Each interface on a peer can communicate with both interfaces on the other peer. This ensures that a VPN will be available as long as each peer has one working connection to the Internet.

You configure a VPN and an entry in the routing table for each of the four paths. All of these VPNs are ready to carry data. You set different routing distances for each route and only the shortest distance route is used. If this route fails, the route with the next shortest distance is used.

The redundant configurations described in this chapter use route-based VPNs, otherwise known as virtual IPsec interfaces. This means that the FortiGate unit must operate in NAT mode. You must use auto-keying. A VPN that is created using manual keys (see [“Manual-key configurations” on page 1770](#)) cannot be included in a redundant-tunnel configuration.

The configuration described here assumes that your redundant VPNs are essentially equal in cost and capability. When the original VPN returns to service, traffic continues to use the replacement VPN until the replacement VPN fails. If your redundant VPN uses more expensive facilities, you want to use it only as a backup while the main VPN is down. For information on how to do this, see [“Creating a backup IPsec interface” on page 1765](#).

**Figure 270:**Example redundant-tunnel configuration



A VPN that is created using manual keys (see [“Manual-key configurations”](#) on page 1770) cannot be included in a redundant-tunnel configuration.

## General configuration steps

A redundant configuration at each VPN peer includes:

- one phase 1 configuration (virtual IPsec interface) for each path between the two peers. In a fully-meshed redundant configuration, each network interface on one peer can communicate with each network interface on the remote peer. If both peers have two public interfaces, this means that each peer has four paths, for example.
- one phase 2 definition for each phase 1 configuration
- one static route for each IPsec interface, with different distance values to prioritize the routes
- two Accept security policies per IPsec interface, one for each direction of traffic
- dead peer detection enabled in each phase 1 definition

The procedures in this section assume that two separate interfaces to the Internet are available on each VPN peer.

## Configure the VPN peers - route-based VPN

VPN peers are configured using Interface Mode for redundant tunnels.

Configure each VPN peer as follows:

1. Ensure that the interfaces used in the VPN have static IP addresses.

2. Create a phase 1 configuration for each of the paths between the peers. Enable IPsec Interface mode so that this creates a virtual IPsec interface. Enable dead peer detection so that one of the other paths is activated if this path fails.

Enter these settings in particular, and any other VPN settings as required:

#### Path 1

<b>Remote Gateway</b>	Select <i>Static IP Address</i> .
<b>IP Address</b>	Type the IP address of the primary interface of the remote peer.
<b>Local Interface</b>	Select the primary public interface of this peer.
<b>Enable IPsec Interface Mode</b>	Enable
<b>Dead Peer Detection</b>	Enable

#### Path 2

<b>Remote Gateway</b>	Select <i>Static IP Address</i> .
<b>IP Address</b>	Type the IP address of the secondary interface of the remote peer.
<b>Local Interface</b>	Select the primary public interface of this peer.
<b>Enable IPsec Interface Mode</b>	Enable
<b>Dead Peer Detection</b>	Enable

#### Path 3

<b>Remote Gateway</b>	Select <i>Static IP Address</i> .
<b>IP Address</b>	Type the IP address of the primary interface of the remote peer.
<b>Local Interface</b>	Select the secondary public interface of this peer.
<b>Enable IPsec Interface Mode</b>	Enable
<b>Dead Peer Detection</b>	Enable

#### Path 4

<b>Remote Gateway</b>	Select <i>Static IP Address</i> .
<b>IP Address</b>	Type the IP address of the secondary interface of the remote peer.
<b>Local Interface</b>	Select the secondary public interface of this peer.
<b>Enable IPsec Interface Mode</b>	Enable
<b>Dead Peer Detection</b>	Enable

For more information, see [“Auto Key phase 1 parameters” on page 1637](#).

3. Create a phase 2 definition for each path. See [“Phase 2 parameters” on page 1653](#). Select the phase 1 configuration (virtual IPsec interface) that you defined for this path. You can select the name from the Static IP Address part of the list.
4. Create a route for each path to the other peer. If there are two ports on each peer, there are four possible paths between the peer devices.

---

**Destination IP/Mask** The IP address and netmask of the private network behind the remote peer.

---

**Device** One of the virtual IPsec interfaces on the local peer.

---

**Distance** For each path, enter a different value to prioritize the paths.

---

5. Define the security policy for the local primary interface. See [“Defining VPN security policies” on page 1660](#). You need to create two policies for each path to enable communication in both directions. Enter these settings in particular:

---

**Incoming Interface** Select the local interface to the internal (private) network.

---

**Source Address** All

---

**Outgoing Interface** Select one of the virtual IPsec interfaces you created in Step 2.

---

**Destination Address** All

---

**Schedule** Always

---

**Service** Any

---

**Action** ACCEPT

---

6. Select *Create New*, leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*, and enter these settings:

---

**Incoming Interface** Select one of the virtual IPsec interfaces you created in Step 2.

---

**Source Address** All

---

**Outgoing Interface** Select the local interface to the internal (private) network.

---

**Destination Address** All

---

**Schedule** Always

---

**Service** Any

---

**Action** ACCEPT

---

7. Place the policy in the policy list above any other policies having similar source and destination addresses.
8. Repeat this procedure at the remote FortiGate unit.




## Redundant route-based VPN configuration example

This example demonstrates a fully redundant site-to-site VPN configuration using route-based VPNs. At each site, the FortiGate unit has two interfaces connected to the Internet through different ISPs. This means that there are four possible paths for communication between the two units. In this example, these paths, listed in descending priority, are:

- FortiGate\_1 WAN 1 to FortiGate\_2 WAN 1
- FortiGate\_1 WAN 1 to FortiGate\_2 WAN 2
- FortiGate\_1 WAN 2 to FortiGate\_2 WAN 1
- FortiGate\_1 WAN 2 to FortiGate\_2 WAN 2

**Figure 271:**Example redundant route-based VPN configuration



For each path, VPN configuration, security policies and routing are defined. By specifying a different routing distance for each path, the paths are prioritized. A VPN tunnel is established on each path, but only the highest priority one is used. If the highest priority path goes down, the traffic is automatically routed over the next highest priority path. You could use dynamic routing, but to keep this example simple, static routing is used.

### Configuring FortiGate\_1

You must

- configure the interfaces involved in the VPN
- define the phase 1 configuration for each of the four possible paths, creating a virtual IPsec interface for each one
- define the phase 2 configuration for each of the four possible paths
- configure routes for the four IPsec interfaces, assigning the appropriate priorities
- configure incoming and outgoing security policies between the internal interface and each of the virtual IPsec interfaces

#### To configure the network interfaces

1. Go to *System > Network > Interfaces*.
2. Select the Internal interface and select *Edit*.

3. Enter the following information and then select *OK*:

<b>Addressing mode</b>	Manual
<b>IP/Netmask</b>	10.21.101.0/255.255.255.0

4. Select the WAN1 interface and select *Edit*, enter the following information and then select *OK*:

<b>Addressing mode</b>	Manual
<b>IP/Netmask</b>	192.168.10.2/255.255.255.0

5. Select the WAN2 interface and select *Edit*, enter the following information and then select *OK*:

<b>Addressing mode</b>	Manual
<b>IP/Netmask</b>	172.16.20.2/255.255.255.0

### To configure the IPsec interfaces (phase 1 configurations)

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. Select *Create Phase 1*, enter the following information, and select *OK*:

<b>Name</b>	Site_1_A
<b>Remote Gateway</b>	Static IP Address
<b>IP Address</b>	192.168.20.2
<b>Local Interface</b>	WAN1
<b>Mode</b>	Main
<b>Authentication Method</b>	Preshared Key
<b>Pre-shared Key</b>	Enter the preshared key.
<b>Peer Options</b>	Accept any peer ID
<b>Advanced</b>	
<b>Enable IPsec Interface Mode</b>	Select
<b>Dead Peer Detection</b>	Select

3. Select *Create Phase 1*, enter the following information, and select *OK*:

<b>Name</b>	Site_1_B
<b>Remote Gateway</b>	Static IP Address
<b>IP Address</b>	172.16.30.2
<b>Local Interface</b>	WAN1

<b>Mode</b>	Main
<b>Authentication Method</b>	Preshared Key
<b>Pre-shared Key</b>	Enter the preshared key.
<b>Peer Options</b>	Accept any peer ID
<b>Advanced</b>	
<b>Enable IPsec Interface Mode</b>	Select
<b>Dead Peer Detection</b>	Select

4. Select *Create Phase 1*, enter the following information, and select *OK*:

<b>Name</b>	Site_1_C
<b>Remote Gateway</b>	Static IP Address
<b>IP Address</b>	192.168.20.2
<b>Local Interface</b>	WAN2
<b>Mode</b>	Main
<b>Authentication Method</b>	Preshared Key
<b>Pre-shared Key</b>	Enter the preshared key.
<b>Peer Options</b>	Accept any peer ID
<b>Advanced</b>	
<b>Enable IPsec Interface Mode</b>	Select
<b>Dead Peer Detection</b>	Select

5. Select *Create Phase 1*, enter the following information, and select *OK*:

<b>Name</b>	Site_1_D
<b>Remote Gateway</b>	Static IP Address
<b>IP Address</b>	172.16.30.2
<b>Local Interface</b>	WAN2
<b>Mode</b>	Main
<b>Authentication Method</b>	Preshared Key
<b>Pre-shared Key</b>	Enter the preshared key.
<b>Peer Options</b>	Accept any peer ID
<b>Advanced</b>	

---

<b>Enable IPsec Interface Mode</b>	Select
------------------------------------	--------

---

<b>Dead Peer Detection</b>	Select
----------------------------	--------

---

### To define the phase 2 configurations for the four VPNs

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. Select *Create Phase 2*, enter the following information and select *OK*:

---

<b>Name</b>	Route_A
<b>Phase 1</b>	Site_1_A

---

3. Select *Create Phase 2*, enter the following information and select *OK*:

---

<b>Name</b>	Route_B
<b>Phase 1</b>	Site_1_B

---

4. Select *Create Phase 2*, enter the following information and select *OK*:

---

<b>Name</b>	Route_C
<b>Phase 1</b>	Site_1_C

---

5. Select *Create Phase 2*, enter the following information and select *OK*:

---

<b>Name</b>	Route_D
<b>Phase 1</b>	Site_1_D

---

### To configure routes

1. Go to *Router > Static > Static Routes*.  
For low-end FortiGate units, go to *System > Network > Routing*.
2. Select *Create New*, enter the following default gateway information and then select *OK*:

---

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Device</b>	WAN1
<b>Gateway</b>	192.168.10.1
<b>Distance (Advanced)</b>	10

---

3. Select *Create New*, enter the following information and then select *OK*:

---

<b>Destination IP/Mask</b>	10.31.101.0/255.255.255.0
<b>Device</b>	Site_1_A
<b>Distance (Advanced)</b>	1

---

4. Select *Create New*, enter the following information and then select *OK*:

---

<b>Destination IP/Mask</b>	10.31.101.0/255.255.255.0
<b>Device</b>	Site_1_B
<b>Distance (Advanced)</b>	2

---

5. Select *Create New*, enter the following information and then select *OK*:

---

<b>Destination IP/Mask</b>	10.31.101.0/255.255.255.0
<b>Device</b>	Site_1_C
<b>Distance (Advanced)</b>	3

---

6. Select *Create New*, enter the following information and then select *OK*:

---

<b>Destination IP/Mask</b>	10.31.101.0/255.255.255.0
<b>Device</b>	Site_1_D
<b>Distance (Advanced)</b>	4

---

#### To configure security policies

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter the following information, and then select *OK*:

---

<b>Incoming Interface</b>	Internal
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Site_1_A
<b>Destination Address</b>	All
<b>Schedule</b>	Always
<b>Service</b>	Any
<b>Action</b>	ACCEPT

---

4. Select *Create New*.
5. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
6. Enter the following information, and select *OK*:

---

<b>Incoming Interface</b>	Site_1_A
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Internal
<b>Destination Address</b>	All

---

<b>Schedule</b>	Always
<b>Service</b>	Any
<b>Action</b>	ACCEPT

7. Select *Create New*.
8. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
9. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	Internal
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Site_1_B
<b>Destination Address</b>	All
<b>Schedule</b>	Always
<b>Service</b>	Any
<b>Action</b>	ACCEPT

10. Select *Create New*.
11. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
12. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	Site_1_B
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Internal
<b>Destination Address</b>	All
<b>Schedule</b>	Always
<b>Service</b>	Any
<b>Action</b>	ACCEPT

13. Select *Create New*.
14. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
15. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	Internal
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Site_1_C
<b>Destination Address</b>	All
<b>Schedule</b>	Always

<b>Service</b>	Any
<b>Action</b>	ACCEPT

16. Select *Create New*.

17. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.

18. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	Site_1_C
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Internal
<b>Destination Address</b>	All
<b>Schedule</b>	Always
<b>Service</b>	Any
<b>Action</b>	ACCEPT

19. Select *Create New*.

20. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.

21. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	Internal
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Site_1_D
<b>Destination Address</b>	All
<b>Schedule</b>	Always
<b>Service</b>	Any
<b>Action</b>	ACCEPT

22. Select *Create New*.

23. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.

24. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	Site_1_D
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Internal
<b>Destination Address</b>	All
<b>Schedule</b>	Always

<b>Service</b>	Any
<b>Action</b>	ACCEPT

## Configuring FortiGate\_2

The configuration for FortiGate\_2 is very similar that of FortiGate\_1. You must

- configure the interfaces involved in the VPN
- define the phase 1 configuration for each of the four possible paths, creating a virtual IPsec interface for each one
- define the phase 2 configuration for each of the four possible paths
- configure routes for the four IPsec interfaces, assigning the appropriate priorities
- configure incoming and outgoing security policies between the internal interface and each of the virtual IPsec interfaces

### To configure the network interfaces

1. Go to *System > Network > Interfaces*.
2. Select the Internal interface and then select *Edit*. Enter the following information and then select *OK*:

<b>Addressing mode</b>	Manual
<b>IP/Netmask</b>	10.31.101.0/255.255.255.0

3. Select the WAN1 interface and then select *Edit*. Enter the following information and then select *OK*:

<b>Addressing mode</b>	Manual
<b>IP/Netmask</b>	192.168.20.2/255.255.255.0

4. Select the WAN2 interface and then select *Edit*. Enter the following information and then select *OK*:

<b>Addressing mode</b>	Manual
<b>IP/Netmask</b>	172.16.30.2/255.255.255.0

### To configure the IPsec interfaces (phase 1 configurations)

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. Select *Create Phase 1*, enter the following information, and select *OK*:

<b>Name</b>	Site_2_A
<b>Remote Gateway</b>	Static IP Address
<b>IP Address</b>	192.168.10.2
<b>Local Interface</b>	WAN1
<b>Mode</b>	Main



<b>Authentication Method</b>	Preshared Key
<b>Pre-shared Key</b>	Enter the preshared key.
<b>Peer Options</b>	Accept any peer ID
<b>Advanced</b>	
<b>Enable IPsec Interface Mode</b>	Select
<b>Dead Peer Detection</b>	Select

3. Select *Create Phase 1*, enter the following information, and select *OK*:

<b>Name</b>	Site_2_B
<b>Remote Gateway</b>	Static IP Address
<b>IP Address</b>	172.16.20.2
<b>Local Interface</b>	WAN1
<b>Mode</b>	Main
<b>Authentication Method</b>	Preshared Key
<b>Pre-shared Key</b>	Enter the preshared key.
<b>Peer Options</b>	Accept any peer ID
<b>Advanced</b>	
<b>Enable IPsec Interface Mode</b>	Select
<b>Dead Peer Detection</b>	Select

4. Select *Create Phase 1*, enter the following information, and select *OK*:

<b>Name</b>	Site_2_C
<b>Remote Gateway</b>	Static IP Address
<b>IP Address</b>	192.168.10.2
<b>Local Interface</b>	WAN1
<b>Mode</b>	Main
<b>Authentication Method</b>	Preshared Key
<b>Pre-shared Key</b>	Enter the preshared key.
<b>Peer Options</b>	Accept any peer ID
<b>Advanced</b>	
<b>Enable IPsec Interface Mode</b>	Select
<b>Dead Peer Detection</b>	Select

5. Select *Create Phase 1*, enter the following information, and select *OK*:

<b>Name</b>	Site_2_D
<b>Remote Gateway</b>	Static IP Address
<b>IP Address</b>	172.16.20.2
<b>Local Interface</b>	WAN1
<b>Mode</b>	Main
<b>Authentication Method</b>	Preshared Key
<b>Pre-shared Key</b>	Enter the preshared key.
<b>Peer Options</b>	Accept any peer ID
<b>Advanced</b>	
<b>Enable IPsec Interface Mode</b>	Select
<b>Dead Peer Detection</b>	Select

**To define the phase 2 configurations for the four VPNs**

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. Select *Create Phase 2*, enter the following information and select *OK*:

<b>Name</b>	Route_A
<b>Phase 1</b>	Site_2_A

3. Select *Create Phase 2*, enter the following information and select *OK*:

<b>Name</b>	Route_B
<b>Phase 1</b>	Site_2_B

4. Select *Create Phase 2*, enter the following information and select *OK*:

<b>Name</b>	Route_C
<b>Phase 1</b>	Site_2_C

5. Select *Create Phase 2*, enter the following information and select *OK*:

<b>Name</b>	Route_D
<b>Phase 1</b>	Site_2_D

**To configure routes**

1. Go to *Router > Static > Static Routes*.  
For low-end FortiGate units, go to *System > Network > Routing*.

2. Select *Create New*, enter the following default gateway information and then select *OK*:

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Device</b>	WAN1
<b>Gateway</b>	192.168.10.1
<b>Distance (Advanced)</b>	10

3. Select *Create New*, enter the following information and then select *OK*:

<b>Destination IP/Mask</b>	10.21.101.0/255.255.255.0
<b>Device</b>	Site_2_A
<b>Distance (Advanced)</b>	1

4. Select *Create New*, enter the following information and then select *OK*:

<b>Destination IP/Mask</b>	10.21.101.0/255.255.255.0
<b>Device</b>	Site_2_B
<b>Distance (Advanced)</b>	2

5. Select *Create New*, enter the following information and then select *OK*:

<b>Destination IP/Mask</b>	10.21.101.0/255.255.255.0
<b>Device</b>	Site_2_C
<b>Distance (Advanced)</b>	3

6. Select *Create New*, enter the following information and then select *OK*:

<b>Destination IP/Mask</b>	10.21.101.0/255.255.255.0
<b>Device</b>	Site_2_D
<b>Distance (Advanced)</b>	4

#### To configure security policies

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	Internal
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Site_2_A
<b>Destination Address</b>	All

<b>Schedule</b>	Always
<b>Service</b>	Any
<b>Action</b>	ACCEPT

4. Select *Create New*.
5. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
6. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	Site_2_A
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Internal
<b>Destination Address</b>	All
<b>Schedule</b>	Always
<b>Service</b>	Any
<b>Action</b>	ACCEPT

7. Select *Create New*.
8. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
9. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	Internal
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Site_2_B
<b>Destination Address</b>	All
<b>Schedule</b>	Always
<b>Service</b>	Any
<b>Action</b>	ACCEPT

10. Select *Create New*.
11. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
12. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	Site_2_B
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Internal
<b>Destination Address Name</b>	All
<b>Schedule</b>	Always

<b>Service</b>	Any
<b>Action</b>	ACCEPT

13. Select *Create New*.

14. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.

15. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	Internal
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Site_2_C
<b>Destination Address</b>	All
<b>Schedule</b>	Always
<b>Service</b>	Any
<b>Action</b>	ACCEPT

16. Select *Create New*.

17. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.

18. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	Site_2_C
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Internal
<b>Destination Address</b>	All
<b>Schedule</b>	Always
<b>Service</b>	Any
<b>Action</b>	ACCEPT

19. Select *Create New*.

20. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.

21. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	Internal
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Site_2_D
<b>Destination Address</b>	All
<b>Schedule</b>	Always

<b>Service</b>	Any
<b>Action</b>	ACCEPT

22. Select *Create New*.

23. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.

24. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	Site_2_D
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Internal
<b>Destination Address</b>	All
<b>Schedule</b>	Always
<b>Service</b>	Any
<b>Action</b>	ACCEPT

## Partially-redundant route-based VPN example

This example demonstrates how to set up a partially redundant IPsec VPN between a local FortiGate unit and a remote VPN peer that receives a dynamic IP address from an ISP before it connects to the FortiGate unit. For more information about FortiGate dialup-client configurations, see [“FortiGate dialup-client configurations” on page 1724](#).

When a FortiGate unit has more than one interface to the Internet (see FortiGate\_1 in [Figure 272](#)), you can configure redundant routes. If the primary connection fails, the FortiGate unit can establish a VPN using the redundant connection.

In this case, FortiGate\_2 has only one connection to the Internet. If the link to the ISP were to go down, the connection to FortiGate\_1 would be lost, and the tunnel would be taken down. The tunnel is said to be partially redundant because FortiGate\_2 does not support a redundant connection.

In the configuration example:

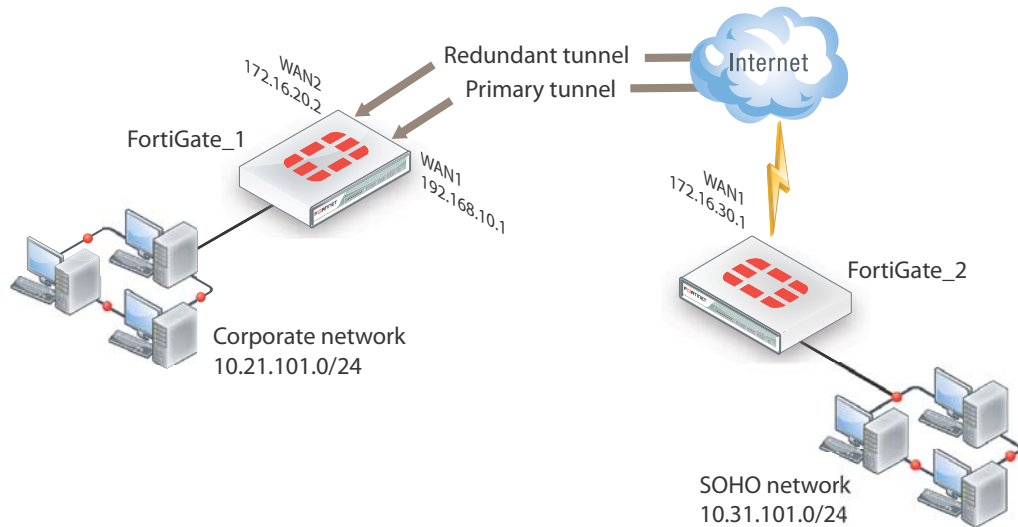
- Both FortiGate units operate in NAT mode.
- Two separate interfaces to the Internet (192.168.10.2 and 172.16.20.2) are available on FortiGate\_1. Each interface has a static public IP address.
- FortiGate\_2 has a single connection to the Internet and obtains a dynamic public IP address (for example, 172.16.30.1) when it connects to the Internet.
- FortiGate\_2 forwards IP packets from the SOHO network (10.31.101.0/24) to the corporate network (10.21.101.0/24) behind FortiGate\_1 through a partially redundant IPsec VPN. Encrypted packets from FortiGate\_2 are addressed to the public interface of FortiGate\_1. Encrypted packets from FortiGate\_1 are addressed to the public IP address of FortiGate\_2.

There are two possible paths for communication between the two units. In this example, these paths, listed in descending priority, are:

- FortiGate\_1 WAN 1 to FortiGate\_2 WAN 1
- FortiGate\_1 WAN 2 to FortiGate\_2 WAN 1

For each path, VPN configuration, security policies and routing are defined. By specifying a different routing distance for each path, the paths are prioritized. A VPN tunnel is established on each path, but only the highest priority one is used. If the highest priority path goes down, the traffic is automatically routed over the next highest priority path. You could use dynamic routing, but to keep this example simple, static routing is used.

**Figure 272:**Example partially redundant route-based configuration



## Configuring FortiGate\_1

You must

- configure the interfaces involved in the VPN
- define the phase 1 configuration for each of the two possible paths, creating a virtual IPsec interface for each one
- define the phase 2 configuration for each of the two possible paths
- configure incoming and outgoing security policies between the internal interface and each of the virtual IPsec interfaces

### To configure the network interfaces

1. Go to *System > Network > Interfaces*.
2. Select the Internal interface and select *Edit*. Enter the following information and select *OK*:

<b>Addressing mode</b>	Manual
<b>IP/Netmask</b>	10.21.101.2/255.255.255.0

3. Select the WAN1 interface and select *Edit*. Enter the following information and select *OK*:

<b>Addressing mode</b>	Manual
<b>IP/Netmask</b>	192.168.10.2/255.255.255.0

4. Select the WAN2 interface and select *Edit*. Enter the following information and select *OK*:

<b>Addressing mode</b>	Manual
<b>IP/Netmask</b>	172.16.20.2/255.255.255.0

**To configure the IPsec interfaces (phase 1 configurations)**

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. Select *Create Phase 1*, enter the following information, and select *OK*:

<b>Name</b>	Site_1_A
<b>Remote Gateway</b>	Dialup User
<b>Local Interface</b>	WAN1
<b>Mode</b>	Main
<b>Authentication Method</b>	Preshared Key
<b>Pre-shared Key</b>	Enter the preshared key.
<b>Peer Options</b>	Accept any peer ID
<b>Advanced</b>	
<b>Enable IPsec Interface Mode</b>	Select
<b>Dead Peer Detection</b>	Select

3. Select *Create Phase 1*, enter the following information, and select *OK*:

<b>Name</b>	Site_1_B
<b>Remote Gateway</b>	Dialup User
<b>Local Interface</b>	WAN2
<b>Mode</b>	Main
<b>Authentication Method</b>	Preshared Key
<b>Pre-shared Key</b>	Enter the preshared key.
<b>Peer Options</b>	Accept any peer ID
<b>Advanced</b>	
<b>Enable IPsec Interface Mode</b>	Select
<b>Dead Peer Detection</b>	Select



### To define the phase 2 configurations for the two VPNs

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. Select *Create Phase 2*, enter the following information and select *OK*:

<b>Name</b>	Route_A
<b>Phase 1</b>	Site_1_A

3. Select *Create Phase 2*, enter the following information and select *OK*:

<b>Name</b>	Route_B
<b>Phase 1</b>	Site_1_B

### To configure routes

1. Go to *Router > Static > Static Routes*.  
For low-end FortiGate units, go to *System > Network > Routing*.
2. Select *Create New*, enter the following default gateway information and select *OK*:

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Device</b>	WAN1
<b>Gateway</b>	192.168.10.1
<b>Distance (Advanced)</b>	10

### To configure security policies

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	Internal
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Site_1_A
<b>Destination Address</b>	All
<b>Schedule</b>	Always
<b>Service</b>	Any
<b>Action</b>	ACCEPT

4. Select *Create New*.
5. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
6. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	Internal
<b>Source Address</b>	All

<b>Outgoing Interface</b>	Site_1_B
<b>Destination Address</b>	All
<b>Schedule</b>	Always
<b>Service</b>	Any
<b>Action</b>	ACCEPT

## Configuring FortiGate\_2

The configuration for FortiGate\_2 is similar to that of FortiGate\_1. You must

- configure the interface involved in the VPN
- define the phase 1 configuration for the primary and redundant paths, creating a virtual IPsec interface for each one
- define the phase 2 configurations for the primary and redundant paths, defining the internal network as the source address so that FortiGate\_1 can automatically configure routing
- configure the routes for the two IPsec interfaces, assigning the appropriate priorities
- configure security policies between the internal interface and each of the virtual IPsec interfaces

### To configure the network interfaces

1. Go to *System > Network > Interfaces*.
2. Select the Internal interface and select *Edit*. Enter the following information and select *OK*:

<b>Addressing mode</b>	Manual
<b>IP/Netmask</b>	10.31.101.2/255.255.255.0

3. Select the WAN1 interface and select *Edit*. Set the *Addressing mode* to *DHCP*.

### To configure the two IPsec interfaces (phase 1 configurations)

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. Select *Create Phase 1*, enter the following information, and select *OK*:

<b>Name</b>	Site_2_A
<b>Remote Gateway</b>	Static IP Address
<b>IP Address</b>	192.168.10.2
<b>Local Interface</b>	WAN1
<b>Mode</b>	Main
<b>Authentication Method</b>	Preshared Key
<b>Pre-shared Key</b>	Enter the preshared key.
<b>Peer Options</b>	Accept any peer ID
<b>Advanced</b>	

---

<b>Enable IPsec Interface Mode</b>	Select
------------------------------------	--------

---

<b>Dead Peer Detection</b>	Select
----------------------------	--------

---

3. Select *Create Phase 1*, enter the following information, and select *OK*:

---

<b>Name</b>	Site_2_B
-------------	----------

---

<b>Remote Gateway</b>	Static IP Address
-----------------------	-------------------

---

<b>IP Address</b>	172.16.20.2
-------------------	-------------

---

<b>Local Interface</b>	WAN1
------------------------	------

---

<b>Mode</b>	Main
-------------	------

---

<b>Authentication Method</b>	Preshared Key
------------------------------	---------------

---

<b>Pre-shared Key</b>	Enter the preshared key.
-----------------------	--------------------------

---

<b>Peer Options</b>	Accept any peer ID
---------------------	--------------------

---

**Advanced**

---

<b>Enable IPsec Interface Mode</b>	Select
------------------------------------	--------

---

<b>Dead Peer Detection</b>	Select
----------------------------	--------

---

**To define the phase 2 configurations for the two VPNs**

1. Go to *VPN > IPsec > Auto Key (IKE)*.
2. Select *Create Phase 2*, enter the following information and select *OK*:

---

<b>Name</b>	Route_A
-------------	---------

---

<b>Phase 1</b>	Site_2_A
----------------	----------

---

**Advanced**

---

<b>Source Address</b>	10.31.101.0/24
-----------------------	----------------

---

3. Select *Create Phase 2*, enter the following information and select *OK*:

---

<b>Name</b>	Route_B
-------------	---------

---

<b>Phase 1</b>	Site_2_B
----------------	----------

---

**Advanced**

---

<b>Source Address</b>	10.31.101.0/24
-----------------------	----------------

---

**To configure routes**

1. Go to *Router > Static > Static Routes*.  
For low-end FortiGate units, go to *System > Network > Routing*.

2. Select *Create New*, enter the following information and then select *OK*:

---

<b>Destination IP/Mask</b>	10.21.101.0/255.255.255.0
----------------------------	---------------------------

---

<b>Device</b>	Site_2_A
---------------	----------

---

<b>Distance (Advanced)</b>	1
----------------------------	---

---

3. Select *Create New*, enter the following information and then select *OK*:

---

<b>Destination IP/Mask</b>	10.21.101.0/255.255.255.0
----------------------------	---------------------------

---

<b>Device</b>	Site_2_B
---------------	----------

---

<b>Distance (Advanced)</b>	2
----------------------------	---

---

### To configure security policies

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter the following information, and select *OK*:

---

<b>Incoming Interface</b>	Internal
---------------------------	----------

---

<b>Source Address</b>	All
-----------------------	-----

---

<b>Outgoing Interface</b>	Site_2_A
---------------------------	----------

---

<b>Destination Address</b>	All
----------------------------	-----

---

<b>Schedule</b>	Always
-----------------	--------

---

<b>Service</b>	Any
----------------	-----

---

<b>Action</b>	ACCEPT
---------------	--------

---

4. Select *Create New*.
5. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
6. Enter the following information, and select *OK*:

---

<b>Incoming Interface</b>	Internal
---------------------------	----------

---

<b>Source Address</b>	All
-----------------------	-----

---

<b>Outgoing Interface</b>	Site_2_B
---------------------------	----------

---

<b>Destination Address</b>	All
----------------------------	-----

---

<b>Schedule</b>	Always
-----------------	--------

---

<b>Service</b>	Any
----------------	-----

---

<b>Action</b>	ACCEPT
---------------	--------

---

## Creating a backup IPsec interface

You can configure a route-based VPN that acts as a backup facility to another VPN. It is used only while your main VPN is out of service. This is desirable when the redundant VPN uses a more expensive facility.

You can configure a backup IPsec interface only in the CLI. The backup feature works only on interfaces with static addresses that have dead peer detection enabled. The `monitor` option creates a backup VPN for the specified phase 1 configuration.

In the following example, `backup_vpn` is a backup for `main_vpn`.

```
config vpn ipsec phase1-interface
 edit main_vpn
 set dpd on
 set interface port1
 set nattraversal enable
 set psksecret "hard-to-guess"
 set remote-gw 192.168.10.8
 set type static
 end
 edit backup_vpn
 set dpd on
 set interface port2
 set monitor main_vpn
 set nattraversal enable
 set psksecret "hard-to-guess"
 set remote-gw 192.168.10.8
 set type static
 end
end
```

# Transparent mode VPNs

This section describes transparent VPN configurations, in which two FortiGate units create a VPN tunnel between two separate private networks transparently.


The following topics are included in this section:

- [Configuration overview](#)
- [Configure the VPN peers](#)

## Configuration overview


In transparent mode, all interfaces of the FortiGate unit except the management interface (which by default is assigned IP address 10.10.10.1/255.255.255.0) are invisible at the network layer. Typically, when a FortiGate unit runs in transparent mode, different network segments are connected to the FortiGate interfaces. [Figure 273](#) shows the management station on the same subnet. The management station can connect to the FortiGate unit directly through the web-based manager.

**Figure 273:**Management station on internal network



An edge router typically provides a public connection to the Internet and one interface of the FortiGate unit is connected to the router. If the FortiGate unit is managed from an external address (see [Figure 274](#) on page 1767), the router must translate (NAT) a routable address to direct management traffic to the FortiGate management interface.

**Figure 274:**Management station on external network



In a transparent VPN configuration, two FortiGate units create a VPN tunnel between two separate private networks transparently. All traffic between the two networks is encrypted and protected by FortiGate security policies.

Both FortiGate units may be running in transparent mode, or one could be running in transparent mode and the other running in NAT mode. If the remote peer is running in NAT mode, it must have a static public IP address.




VPNs between two FortiGate units running in transparent mode do not support inbound/outbound NAT (supported through CLI commands) within the tunnel. In addition, a FortiGate unit running in transparent mode cannot be used in a hub-and-spoke configuration.

Encrypted packets from the remote VPN peer are addressed to the management interface of the local FortiGate unit. If the local FortiGate unit can reach the VPN peer locally, a static route to the VPN peer must be added to the routing table on the local FortiGate unit. If the VPN peer connects through the Internet, encrypted packets from the local FortiGate unit must be routed to the edge router instead. For information about how to add a static route to the FortiGate routing table, see the [Advanced Routing](#) .

In the example configuration shown in [Figure 274](#), Network Address Translation (NAT) is enabled on the router. When an encrypted packet from the remote VPN peer arrives at the router through the Internet, the router performs inbound NAT and forwards the packet to the FortiGate unit. Refer to the software supplier's documentation to configure the router.


If you want to configure a VPN between two FortiGate units running in transparent mode, each unit must have an independent connection to a router that acts as a gateway to the Internet, and both units must be on separate networks that have a different address space. When the two networks linked by the VPN tunnel have different address spaces (see [Figure 275 on page 1768](#)), at least one router must separate the two FortiGate units, unless the packets can be redirected using ICMP (see [Figure 276 on page 1768](#)).

**Figure 275:**Link between two FortiGate units in transparent mode



In [Figure 276](#), interface C behind the router is the default gateway for both FortiGate units. Packets that cannot be delivered on Network\_1 are routed to interface C by default. Similarly, packets that cannot be delivered on Network\_2 are routed to interface C. In this case, the router must be configured to redirect packets destined for Network\_1 to interface A and redirect packets destined for Network\_2 to interface B.


**Figure 276:**ICMP redirecting packets to two FortiGate units in transparent mode



If there are additional routers behind the FortiGate unit (see [Figure 277](#) on page 1769) and the destination IP address of an inbound packet is on a network behind one of those routers, the FortiGate routing table must include routes to those networks. For example, in [Figure 277](#), the FortiGate unit must be configured with static routes to interfaces A and B in order to forward packets to Network\_1 and Network\_2 respectively.



**Figure 277:**Destinations on remote networks behind internal routers



## Transparent VPN infrastructure requirements

- The local FortiGate unit must be operating in transparent mode.
- The management IP address of the local FortiGate unit specifies the local VPN gateway. The management IP address is considered a static IP address for the local VPN peer.
- If the local FortiGate unit is managed through the Internet, or if the VPN peer connects through the Internet, the edge router must be configured to perform inbound NAT and forward management traffic and/or encrypted packets to the FortiGate unit.
- If the remote peer is operating in NAT mode, it must have a static public IP address.

A FortiGate unit operating in transparent mode requires the following basic configuration to operate as a node on the IP network:

- The unit must have sufficient routing information to reach the management station.
- For any traffic to reach external destinations, a default static route to an edge router that forwards packets to the Internet must be present in the FortiGate routing table.
- When all of the destinations are located on the external network, the FortiGate unit may route packets using a single default static route. If the network topology is more complex, one or more static routes in addition to the default static route may be required in the FortiGate routing table.

Only policy-based VPN configurations are possible in transparent mode.

## Before you begin

An IPsec VPN definition links a gateway with a tunnel and an IPsec policy. If your network topology includes more than one virtual domain, you must choose components that were created in the same virtual domain. Therefore, before you define a transparent VPN configuration, choose an appropriate virtual domain in which to create the required interfaces, security policies, and VPN components. For more information, see the [Virtual Domains](#) chapter of *The Handbook*.

## Configure the VPN peers

1. The local VPN peer need to operate in transparent mode.

To determine if your FortiGate unit is in transparent mode, go to *System > Dashboard > Status to the System Information* widget. Select *[change]*. Select transparent for the *Operation Mode*. Two new fields will appear to enter the *Management IP/Netmask*, and the *Default Gateway*.

In transparent mode, the FortiGate unit is invisible to the network. All of its interfaces are on the same subnet and share the same IP address. You only have to configure a management IP address so that you can make configuration changes.

The remote VPN peer may operate in NAT mode or transparent mode.

2. At the local FortiGate unit, define the phase 1 parameters needed to establish a secure connection with the remote peer. See [“Auto Key phase 1 parameters” on page 1637](#). Select *Advanced* and enter these settings in particular:

---

<b>Remote Gateway</b>	Select <i>Static IP Address</i> .
<b>IP Address</b>	Type the IP address of the public interface to the remote peer. If the remote peer is a FortiGate unit running in transparent mode, type the IP address of the remote management interface.
<b>Advanced</b>	Select <i>Nat-traversal</i> , and type a value into the <i>Keepalive Frequency</i> field. These settings protect the headers of encrypted packets from being altered by external NAT devices and ensure that NAT address mappings do not change while the VPN tunnel is open. For more information, see <a href="#">“NAT traversal” on page 1649</a> and <a href="#">“NAT keepalive frequency” on page 1650</a> .

---

3. Define the phase 2 parameters needed to create a VPN tunnel with the remote peer. See [“Phase 2 parameters” on page 1653](#). Select the set of phase 1 parameters that you defined for the remote peer. The name of the remote peer can be selected from the *Static IP Address* list.
4. Define the source and destination addresses of the IP packets that are to be transported through the VPN tunnel. See [“Defining policy addresses” on page 1659](#). Enter these settings in particular:
  - For the originating address (source address), enter the IP address and netmask of the private network behind the local peer network. for the management interface, for example, 10.10.10.0/24. This address needs to be a range to allow traffic from your network through the tunnel. Optionally select *any* for this address.
  - For the remote address (destination address), enter the IP address and netmask of the private network behind the remote peer (for example, 192.168.10.0/24). If the remote peer is a FortiGate unit running in transparent mode, enter the IP address of the remote management interface instead.
5. Define an IPsec security policy to permit communications between the source and destination addresses. See [“Defining VPN security policies” on page 1660](#). Enter these settings in particular:

---

<b>Local Interface</b>	Select the local interface to the internal (private) network.
<b>Local Protected Subnet</b>	Select the source address that you defined in Step 4.

---

<b>Outgoing VPN Interface</b>	Select the interface to the edge router. When you configure the IPsec security policy on a remote peer that operates in NAT mode, you select the public interface to the external (public) network instead.
<b>Remote Protected Subnet</b>	Select the destination address that you defined in Step 4.
<b>VPN Tunnel</b>	Select <i>Use Existing</i> and select the name of the phase 2 tunnel configuration that you created in Step 3 from the drop-down list.  Select <i>Allow traffic to be initiated from the remote site</i> to enable traffic from the remote network to initiate the tunnel.

- Place the policy in the policy list above any other policies having similar source and destination addresses.
- Define another IPsec security policy to permit communications between the source and destination addresses in the opposite direction. This security policy and the previous one form a bi-directional policy pair. See [“Defining VPN security policies” on page 1660](#). Enter these settings in particular:

<b>Local Interface</b>	Select the interface to the edge router. When you configure the IPsec security policy on a remote peer that operates in NAT mode, you select the public interface to the external (public) network instead.
<b>Local Protected Subnet</b>	Select the destination address that you defined in Step 4.
<b>Outgoing VPN Interface</b>	Select the local interface to the internal (private) network.
<b>Remote Protected Subnet</b>	Select the source address that you defined in Step 4.
<b>VPN Tunnel</b>	Select <i>Use Existing</i> and select the name of the phase 2 tunnel configuration that you created in Step 3 from the drop-down list.  Select <i>Allow traffic to be initiated from the remote site</i> to enable traffic from the remote network to initiate the tunnel.

- Repeat this procedure at the remote FortiGate unit to create bidirectional security policies. Use the local interface and address information local to the remote FortiGate unit.

For more information on transparent mode, see the [System Administration Guide handbook chapter](#).

# IPv6 IPsec VPNs

This chapter describes how to configure your FortiGate unit's IPv6 IPsec VPN functionality.



By default IPv6 configurations do not appear on the Web-based Manager. You need to enable the feature first.

## To enable IPv6

1. Go to *System > Admin > Settings*.
2. In the *Display Options on GUI* section, select *IPv6*.
3. Select *Apply*.

The following topics are included in this section:

- [Overview of IPv6 IPsec support](#)
- [Configuring IPv6 IPsec VPNs](#)
- [Site-to-site IPv6 over IPv6 VPN example](#)
- [Site-to-site IPv4 over IPv6 VPN example](#)
- [Site-to-site IPv6 over IPv4 VPN example](#)

## Overview of IPv6 IPsec support

FortiOS supports route-based IPv6 IPsec, but not policy-based. This section describes how IPv6 IPsec support differs from IPv4 IPsec support. FortiOS 4.0 MR3 is IPv6 Ready Logo Program Phase 2 certified.

Where both the gateways and the protected networks use IPv6 addresses, sometimes called IPv6 over IPv6, you can create either an auto-keyed or manually-keyed VPN. You can combine IPv6 and IPv4 addressing in an auto-keyed VPN in the following ways:

<b>IPv4 over IPv6</b>	The VPN gateways have IPv6 addresses. The protected networks have IPv4 addresses. The phase 2 configurations at either end use IPv4 selectors.
<b>IPv6 over IPv4</b>	The VPN gateways have IPv4 addresses. The protected networks use IPv6 addresses. The phase 2 configurations at either end use IPv6 selectors.

Compared with IPv4 IPsec VPN functionality, there are some limitations:

- Except for IPv6 over IPv4, remote gateways with Dynamic DNS are not supported.
- You cannot use RSA certificates in which the common name (cn) is a domain name that resolves to an IPv6 address. This is because FortiOS 3.0 does not support IPv6 DNS.
- DHCP over IPsec is not supported, because FortiOS 3.0 does not support IPv6 DHCP.
- Selectors cannot be firewall address names. Only IP address, address range and subnet are supported.
- Redundant IPv6 tunnels are not supported.

## Certificates

On a VPN with IPv6 phase 1 configuration, you can authenticate using VPN certificates in which the common name (cn) is an IPv6 address. The `cn-type` keyword of the `user peer` command has an option, `ipv6`, to support this.

## Configuring IPv6 IPsec VPNs

Configuration of an IPv6 IPsec VPN follows the same sequence as for an IPv4 route-based VPN: phase 1 settings, phase 2 settings, security policies and routing.



By default IPv6 configurations do not appear on the Web-based Manager. You need to enable the feature first.

### To enable IPv6

1. Go to *System > Admin > Settings*.
2. In the *Display Options on GUI* section, select *IPv6*.
3. Select *Apply*.

## Phase 1 configuration

In the web-based manager, you define the Phase 1 as IPv6 in the Advanced settings. Enable the IPv6 Version check box. You can then enter an IPv6 address for the remote gateway.

In the CLI, you define an IPsec phase 1 configuration as IPv6 by setting `ip-version` to 6. Its default value is 4. Then, the `local-gw` and `remote-gw` keywords are hidden and the corresponding `local-gw6` and `remote-gw6` keywords are available. The values for `local-gw6` and `remote-gw6` must be IPv6 addresses. For example:

```
config vpn ipsec phase1-interface
edit tunnel6
set ip-version 6
set remote-gw6 0:123:4567::1234
set interface port3
set proposal 3des-md5
end
```

## Phase 2 configuration

To create an IPv6 IPsec phase 2 configuration in the web-based manager, you need to define IPv6 selectors in the Advanced settings. Change the default “0.0.0.0/0” address for Source address and Destination address to the IPv6 value “::/0”. If needed, enter specific IPv6 addresses, address ranges or subnet addresses in these fields.

In the CLI, set `src-addr-type` and `dst-addr-type` to `ip6`, `range6` or `subnet6` to specify IPv6 selectors. By default, zero selectors are entered, “`::/0`” for the `subnet6` address type, for example. The simplest IPv6 phase 2 configuration looks like this:

```
config vpn ipsec phase2-interface
 edit tunnel6_p2
 set phase1name tunnel6
 set proposal 3des-md5
 set src-addr-type subnet6
 set dst-addr-type subnet6
 end
```

## Security policies

To complete the VPN configuration, you need a security policy in each direction to permit traffic between the protected network’s port and the IPsec interface. You need IPv6 policies unless the VPN is IPv4 over IPv6.

## Routing

Appropriate routing is needed for both the IPsec packets and the encapsulated traffic within them. You need a route, which could be the default route, to the remote VPN gateway via the appropriate interface. You also need a route to the remote protected network via the IPsec interface.

To create a static route in the web-based manager

1. Go to *Router > Static > Static Routes*.  
On low-end FortiGate units, go to *System > Network > Routing*.
2. Select the drop-down arrow on the *Create New* button and select *IPv6 Route*.
3. Enter the information and select *OK*.

In the CLI, use the `router static6` command. For example, where the remote network is `fec0:0000:0000:0004::/64` and the IPsec interface is `toB`:

```
config router static6
 edit 1
 set device port2
 set dst 0::/0
 next
 edit 2
 set device toB
 set dst fec0:0000:0000:0004::/64
 next
end
```


If the VPN is IPv4 over IPv6, the route to the remote protected network is an IPv4 route. If the VPN is IPv6 over IPv4, the route to the remote VPN gateway is an IPv4 route.

## Site-to-site IPv6 over IPv6 VPN example

In this example, computers on IPv6-addressed private networks communicate securely over public IPv6 infrastructure.

To access IPv6 functionality through the web-based manager, go to *System Admin > Settings* and enable *IPv6* in the section, *Display Options on GUI*.

**Figure 278:**Example IPv6-over-IPv6 VPN topology



## Configure FortiGate A interfaces

Port 2 connects to the public network and port 3 connects to the local network.

```
config system interface
 edit port2
 config ipv6
 set ip6-address fec0::0001:209:0fff:fe83:25f2/64
 end
 next
 edit port3
 config ipv6
 set ip6-address fec0::0000:209:0fff:fe83:25f3/64
 end
 next
end
```

## Configure FortiGate A IPsec settings

The phase 1 configuration creates a virtual IPsec interface on port 2 and sets the remote gateway to the public IP address FortiGate B. This configuration is the same as for an IPv4

route-based VPN, except that `ip-version` is set to 6 and the `remote-gw6` keyword is used to specify an IPv6 remote gateway address.

```
config vpn ipsec phase1-interface
 edit toB
 set ip-version 6
 set interface port2
 set remote-gw6 fec0:0000:0000:0003:209:0fff:fe83:25c7
 set dpd enable
 set psksecret maryhadalittlelamb
 set proposal 3des-md5 3des-sha1
 end
```

By default, phase 2 selectors are set to accept all subnet addresses for source and destination. The default setting for `src-addr-type` and `dst-addr-type` is `subnet`. The IPv6 equivalent is `subnet6`. The default subnet addresses are 0.0.0.0/0 for IPv4, `::/0` for IPv6.

```
config vpn ipsec phase2-interface
 edit toB2
 set phase1name toB
 set proposal 3des-md5 3des-sha1
 set pfs enable
 set replay enable
 set src-addr-type subnet6
 set dst-addr-type subnet6
 end
```

## Configure FortiGate A security policies

Security policies are required to allow traffic between `port3` and the IPsec interface `toB` in each direction. The address `all6` must be defined using the `firewall address6` command as `::/0`.

```
config firewall policy6
 edit 1
 set srcintf port3
 set dstintf toB
 set srcaddr all6
 set dstaddr all6
 set action accept
 set service ANY
 set schedule always
 next
 edit 2
 set srcintf toB
 set dstintf port3
 set srcaddr all6
 set dstaddr all6
 set action accept
 set service ANY
 set schedule always
 end
```



## Configure FortiGate A routing

This simple example requires just two static routes. Traffic to the protected network behind FortiGate B is routed via the virtual IPsec interface toB. A default route sends all IPv6 traffic out on port2.

```
config router static6
 edit 1
 set device port2
 set dst 0::/0
 next
 edit 2
 set device toB
 set dst fec0:0000:0000:0004::/64
end
```

## Configure FortiGate B

The configuration of FortiGate B is very similar to that of FortiGate A. A virtual IPsec interface toA is configured on port2 and its remote gateway is the public IP address of FortiGate A. Security policies enable traffic to pass between the private network and the IPsec interface. Routing ensures traffic for the private network behind FortiGate A goes through the VPN and that all IPv6 packets are routed to the public network.

```
config system interface
 edit port2
 config ipv6
 set ip6-address fec0::0003:209:0fff:fe83:25c7/64
 end
 next
 edit port3
 config ipv6
 set ip6-address fec0::0004:209:0fff:fe83:2569/64
 end
 end
config vpn ipsec phase1-interface
 edit toA
 set ip-version 6
 set interface port2
 set remote-gw6 fec0:0000:0000:0001:209:0fff:fe83:25f2
 set dpd enable
 set psksecret maryhadalittlelamb
 set proposal 3des-md5 3des-sha1
 end
config vpn ipsec phase2-interface
 edit toA2
 set phase1name toA
 set proposal 3des-md5 3des-sha1
 set pfs enable
 set replay enable
 set src-addr-type subnet6
 set dst-addr-type subnet6
```

```


end
config firewall policy6
 edit 1
 set srcintf port3
 set dstintf toA
 set srcaddr all6
 set dstaddr all6
 set action accept
 set service ANY
 set schedule always
 next
 edit 2
 set srcintf toA
 set dstintf port3
 set srcaddr all6
 set dstaddr all6
 set action accept
 set service ANY
 set schedule always
 end
config router static6
 edit 1
 set device port2
 set dst 0::/0
 next
 edit 2
 set device toA
 set dst fec0:0000:0000:0000::/64
 end
end

```

## Site-to-site IPv4 over IPv6 VPN example

In this example, two private networks with IPv4 addressing communicate securely over IPv6 infrastructure.

**Figure 279:**Example IPv4-over-IPv6 VPN topology



## Configure FortiGate A interfaces

Port 2 connects to the IPv6 public network and port 3 connects to the IPv4 LAN.

```
config system interface
 edit port2
 config ipv6
 set ip6-address fec0::0001:209:0fff:fe83:25f2/64
 end
 next
 edit port3
 set 192.168.2.1/24
 end
```

## Configure FortiGate A IPsec settings

The phase 1 configuration is the same as in the IPv6 over IPv6 example.

```
config vpn ipsec phase1-interface
 edit toB
 set ip-version 6
 set interface port2
 set remote-gw6 fec0:0000:0000:0003:209:0fff:fe83:25c7
 set dpd enable
 set psksecret maryhadalittlelamb
 set proposal 3des-md5 3des-sha1
 end
```

The phase 2 configuration is the same as you would use for an IPv4 VPN. By default, phase 2 selectors are set to accept all subnet addresses for source and destination.

```
config vpn ipsec phase2-interface
 edit toB2
 set phase1name toB
 set proposal 3des-md5 3des-sha1
 set pfs enable
 set replay enable
 end
```

## Configure FortiGate A security policies

Security policies are required to allow traffic between port3 and the IPsec interface toB in each direction. These are IPv4 security policies.

```
config firewall policy
 edit 1
 set srcintf port3
 set dstintf toB
 set srcaddr all
 set dstaddr all
 set action accept
 set service ANY
 set schedule always
```

```

next
edit 2
 set srcintf toB
 set dstintf port3
 set srcaddr all
 set dstaddr all
 set action accept
 set service ANY
 set schedule always
end

```

## Configure FortiGate A routing

This simple example requires just two static routes. Traffic to the protected network behind FortiGate B is routed via the virtual IPsec interface toB using an IPv4 static route. A default route sends all IPv6 traffic, including the IPv6 IPsec packets, out on port2.

```

config router static6
 edit 1
 set device port2
 set dst 0::/0
 next
 edit 2
 set device toB
 set dst 192.168.3.0/24
 end
end

```

## Configure FortiGate B

The configuration of FortiGate B is very similar to that of FortiGate A. A virtual IPsec interface toA is configured on port2 and its remote gateway is the public IP address of FortiGate A. The IPsec phase 2 configuration has IPv4 selectors.

IPv4 security policies enable traffic to pass between the private network and the IPsec interface. An IPv4 static route ensures traffic for the private network behind FortiGate A goes through the VPN and an IPv6 static route ensures that all IPv6 packets are routed to the public network.

```

config system interface
 edit port2
 config ipv6
 set ip6-address fec0::0003:fe83:25c7/64
 end
 next
 edit port3
 set 192.168.3.1/24
 end
config vpn ipsec phase1-interface
 edit toA
 set ip-version 6
 set interface port2
 set remote-gw6 fec0:0000:0000:0001:209:0fff:fe83:25f2
 set dpd enable
 end
end

```

```


 set psksecret maryhadalittlelamb
 set proposal 3des-md5 3des-sha1
 end
config vpn ipsec phase2-interface
 edit toA
 set phasename toA
 set proposal 3des-md5 3des-sha1
 set pfs enable
 set replay enable
 end
config firewall policy
 edit 1
 set srcintf port3
 set dstintf toA
 set srcaddr all
 set dstaddr all
 set action accept
 set service ANY
 set schedule always
 next
 edit 2
 set srcintf toA
 set dstintf port3
 set srcaddr all
 set dstaddr all
 set action accept
 set service ANY
 set schedule always
 end
config router static6
 edit 1
 set device port2
 set dst 0::/0
 next
 edit 2
 set device toA
 set dst 192.168.2.0/24
 end

```

## Site-to-site IPv6 over IPv4 VPN example

In this example, IPv6-addressed private networks communicate securely over IPv4 public infrastructure.

**Figure 280:**Example IPv6-over-IPv4 VPN topology



### Configure FortiGate A interfaces

Port 2 connects to the IPv4 public network and port 3 connects to the IPv6 LAN.

```
config system interface
 edit port2
 set 10.0.0.1/24
 next
 edit port3
 config ipv6
 set ip6-address fec0::0001:209:0fff:fe83:25f3/64
 end
 end
```

### Configure FortiGate A IPsec settings

The phase 1 configuration uses IPv4 addressing.

```
config vpn ipsec phase1-interface
 edit toB
 set interface port2
 set remote-gw 10.0.1.1
 set dpd enable
 set psksecret maryhadalittlelamb
 set proposal 3des-md5 3des-sha1
 end
```

The phase 2 configuration uses IPv6 selectors. By default, phase 2 selectors are set to accept all subnet addresses for source and destination. The default setting for `src-addr-type` and `dst-addr-type` is `subnet`. The IPv6 equivalent is `subnet6`. The default subnet addresses are `0.0.0.0/0` for IPv4, `::/0` for IPv6.

```

config vpn ipsec phase2-interface
 edit toB
 set phase1name toB
 set proposal 3des-md5 3des-sha1
 set pfs enable
 set replay enable
 set src-addr-type subnet6
 set dst-addr-type subnet6
 end

```

## Configure FortiGate A security policies

IPv6 security policies are required to allow traffic between port3 and the IPsec interface toB in each direction. Define the address all6 using the firewall address6 command as ::/0.

```

config firewall policy6
 edit 1
 set srcintf port3
 set dstintf toB
 set srcaddr all6
 set dstaddr all6
 set action accept
 set service ANY
 set schedule always
 next
 edit 2
 set srcintf toB
 set dstintf port3
 set srcaddr all6
 set dstaddr all6
 set action accept
 set service ANY
 set schedule always
 end

```

## Configure FortiGate A routing

This simple example requires just two static routes. Traffic to the protected network behind FortiGate B is routed via the virtual IPsec interface toB using an IPv6 static route. A default route sends all IPv4 traffic, including the IPv4 IPsec packets, out on port2.

```

config router static6
 edit 1
 set device toB
 set dst fec0:0000:0000:0004::/64
 end
config router static
 edit 1
 set device port2
 set dst 0.0.0.0/0
 set gateway 10.0.0.254
 end

```

## Configure FortiGate B

The configuration of FortiGate B is very similar to that of FortiGate A. A virtual IPsec interface toA is configured on port2 and its remote gateway is the IPv4 public IP address of FortiGate A. The IPsec phase 2 configuration has IPv6 selectors.

IPv6 security policies enable traffic to pass between the private network and the IPsec interface. An IPv6 static route ensures traffic for the private network behind FortiGate A goes through the VPN and an IPv4 static route ensures that all IPv4 packets are routed to the public network.

```
config system interface
 edit port2
 set 10.0.1.1/24
 next
 edit port3
 config ipv6
 set ip6-address fec0::0004:209:0fff:fe83:2569/64
 end
config vpn ipsec phase1-interface
 edit toA
 set interface port2
 set remote-gw 10.0.0.1
 set dpd enable
 set psksecret maryhadalittlelamb
 set proposal 3des-md5 3des-sha1
 end
config vpn ipsec phase2-interface
 edit toA2
 set phaselname toA
 set proposal 3des-md5 3des-sha1
 set pfs enable
 set replay enable
 set src-addr-type subnet6
 set dst-addr-type subnet6
 end
config firewall policy6
 edit 1
 set srcintf port3
 set dstintf toA
 set srcaddr all6
 set dstaddr all6
 set action accept
 set service ANY
 set schedule always
 next
 edit 2
 set srcintf toA
 set dstintf port3
 set srcaddr all6
 set dstaddr all6
 set action accept
```



```
 set service ANY
 set schedule always
 end
config router static6
 edit 1
 set device toA
 set dst fec0:0000:0000:0000::/64
 end
config router static
 edit 1
 set device port2
 set gateway 10.0.1.254
 end
```

# L2TP and IPsec (Microsoft VPN)

This section describes how to set up a VPN that is compatible with the Microsoft Windows native VPN, which is Layer 2 Tunneling Protocol (L2TP) with IPsec encryption.


The following topics are included in this section:

- [Overview](#)
- [Assumptions](#)
- [Configuring the FortiGate unit](#)
- [Configuring the Windows PC](#)
- [Troubleshooting](#)

## Overview

The topology of a VPN for Microsoft Windows dialup clients is very similar to the topology for FortiClient Endpoint Security clients.

**Figure 281:** Example FortiGate VPN configuration with Microsoft clients



For users, the difference is that instead of installing and using the FortiClient application, they configure a network connection using the software built into the Microsoft Windows operating system. Starting in FortiOS 4.0 MR2, you can configure a FortiGate unit to work with unmodified Microsoft VPN client software.

## Layer 2 Tunneling Protocol (L2TP)

L2TP is a tunneling protocol published in 1999 that is used with VPNs, as the name suggests. Microsoft Windows operating system has a built-in L2TP client starting since Windows 2000. Mac OS X 10.3 system and higher also have a built-in client.

L2TP provides no encryption and used UDP port 1701. IPsec is used to secure L2TP packets. The initiator of the L2TP tunnel is called the L2TP Access Concentrator (LAC).

L2TP and IPsec is supported for native Windows XP, Windows Vista and Mac OSX native VPN clients. However, in Mac OSX (OSX 10.6.3, including patch releases) the L2TP feature does not work properly on the Mac OS side.

## Assumptions

The following assumptions have been made for this example:

- L2TP protocol traffic is allowed through network firewalls (TCP and UDP port 1701)
- User has Microsoft Windows 2000 or higher — a Windows version that supports L2TP

## Configuring the FortiGate unit

To configure the FortiGate unit, you need to:

- configure L2TP users and firewall user group;
- configure the L2TP VPN, including the IP address range it assigns to clients;
- configure an IPsec VPN with encryption and authentication settings that match the Microsoft VPN client;
- configure security policies.

### Configuring L2TP users and firewall user group

Remote users must be authenticated before they can request services and/or access network resources through the VPN. The authentication process can use a password defined on the FortiGate unit or an established external authentication mechanism such as RADIUS or LDAP.

#### Creating user accounts

You need to create user accounts and then add these users to a firewall user group to be used for L2TP authentication. The Microsoft VPN client can automatically send the user's Windows network logon credentials. You might want to use these for their L2TP user name and password.

#### To create a user account - web-based manager

1. Go to *User & Device > User > User Definition* and select *Create New*.
2. Enter the *User Name*.
3. Do one of the following:
  - Select *Password* and enter the user's assigned password.
  - Select *Match user on LDAP server*, *Match user on RADIUS server*, or *Match user on TACACS+ server* and select the authentication server from the list. The authentication server must be already configured on the FortiGate unit.
4. Select *OK*.

### To create a user account - CLI

To create a user account called `user1` with the password `123_user`, enter:

```
config user local
 edit user1
 set type password
 set passwd "123_user"
 set status enable
 end
```

### Creating a user group

When clients connect using the L2TP-over-IPsec VPN, the FortiGate unit checks their credentials against the user group you specify for L2TP authentication. You need to create a firewall user group to use for this purpose.

#### To create a user group - web-based manager

1. Go to *User & Device > User > User Groups*, select *Create New*, and enter the following:

<b>Name</b>	Type or edit the user group name (for example, <code>L2TP_group</code> ).
<b>Type</b>	Select <i>Firewall</i> .
<b>Available Users/Groups</b>	The list of Local users, RADIUS servers, LDAP servers, TACACS+ servers, or PKI users that can be added to the user group. To add a member to this list, select the name and then select the right arrow button.
<b>Members</b>	The list of Local users, RADIUS servers, LDAP servers, TACACS+ servers, or PKI users that belong to the user group. To remove a member, select the name and then select the left arrow button.

2. Select *OK*.

#### To create a user group - CLI

To create the user group `L2TP_group` and add members `User_1`, `User_2`, and `User_3`, enter:

```
config user group
 edit L2TP_group
 set group-type firewall
 set member User_1 User_2 User_3
 end
```

### Configuring L2TP

You can only configure L2TP settings in the CLI. As well as enabling L2TP, you set the range of IP address values that are assigned to L2TP clients and specify the user group that can access the VPN. For example, to allow access to users in the `L2TP_group` and assign them addresses in the range `192.168.0.50` to `192.168.0.59`, enter:

```

config vpn l2tp
 set sip 192.168.0.50
 set eip 192.168.0.59
 set status enable
 set usrgrp "L2TP_group"
end

```

One of the security policies for the L2TP over IPsec VPN uses the client address range, so you need also need to create a firewall address for that range. For example,

```

config firewall address
 edit L2TPclients
 set type iprange
 set start-ip 192.168.0.50
 set end-ip 192.168.0.59
 end

```

Alternatively, you could define this range in the web-based manager.

## Configuring IPsec

The Microsoft VPN client uses IPsec for encryption. The configuration needed on the FortiGate unit is the same as for any other IPsec VPN with the following exceptions.

- Transport mode is used instead of tunnel mode.
- The encryption and authentication proposals must be compatible with the Microsoft client.

L2TP over IPsec is supported on the FortiGate unit using policy-based, not route-based configurations.

### Configuring phase 1 - web-based manager

1. Go to *VPN > IPsec > Auto Key (IKE)* and select *Create Phase 1*.
2. Enter the following information and then select *OK*.

<b>Name</b>	Enter a name for this VPN, dialup_p1 for example.
<b>Remote Gateway</b>	Dialup User
<b>Local Interface</b>	Select the network interface that connects to the Internet. For example, port1.
<b>Mode</b>	Main (ID protection)
<b>Authentication Method</b>	Preshared Key
<b>Pre-shared Key</b>	Enter the preshared key. This key must also be entered in the Microsoft VPN client.
<b>Advanced</b>	Select <i>Advanced</i> to enter the following information.
<b>Enable IPsec Interface Mode</b>	This must <b>not</b> be selected.
<b>P1 Proposal</b>	Enter the following Encryption/Authentication pairs: AES256-MD5, 3DES-SHA1, AES192-SHA1
<b>DH Group</b>	2

<b>NAT Traversal</b>	Enable
<b>Dead Peer Detection</b>	Enable

### Configuring phase 1 - CLI

To create a phase 1 configuration called dialup\_p1 on a FortiGate unit that has port1 connected to the Internet, you would enter:

```
config vpn ipsec phase1
 edit dialup_p1
 set type dynamic
 set interface port1
 set mode main
 set psksecret *****
 set proposal aes256-md5 3des-sha1 aes192-sha1
 set dhgrp 2
 set natTraversal enable
 set dpd enable
 end
```

### Configuring phase 2 - web-based manager

1. Go to *VPN > IPsec > Auto Key (IKE)* and select *Create Phase 2*.
2. Enter the following information and then select *OK*.

<b>Name</b>	Enter a name for this phase 2 configuration.
<b>Phase 1</b>	Select the name of the phase 1 configuration.
<b>Advanced</b>	Select <i>Advanced</i> to enter the following information.
<b>P2 Proposal</b>	Enter the following Encryption/Authentication pairs: AES256-MD5, 3DES-SHA1, AES192-SHA1
<b>Enable replay detection</b>	Enable
<b>Enable perfect forward secrecy (PFS)</b>	Disable
<b>Keylife</b>	3600 seconds

3. Make this a transport-mode VPN. You must use the CLI to do this. If your phase 2 name is dialup\_p2, you would enter:

```
config vpn ipsec phase2
 edit dialup_p2
 set encapsulation transport-mode
 end
```

## Configuring phase 2 - CLI

To configure a phase 2 to work with your phase\_1 configuration, you would enter:

```
config vpn ipsec phase2
 edit dialup_p2
 set phase1name dialup_p1
 set proposal aes256-md5 3des-sha1 aes192-sha1
 set replay enable
 set pfs disable
 set keylifeseconds 3600
 set encapsulation transport-mode
 end
```

## Configuring security policies

The security policies required for L2TP over IPsec VPN are:

- an IPSEC policy, as you would create for any policy-based IPsec VPN
- a regular ACCEPT policy to allow traffic from the L2TP clients to access the protected network

### Configuring the IPSEC security policy - web-based manager

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Select the *Policy Type* as *VPN* and leave the *Policy Subtype* as *IPsec*.
3. Enter the following information and select *OK*:

<b>Local Interface</b>	Select the interface that connects to the private network behind this FortiGate unit.
<b>Local Protected Subnet</b>	All
<b>Outgoing VPN Interface</b>	Select the FortiGate unit's public interface.
<b>Remote Protected Subnet</b>	All
<b>VPN Tunnel</b>	Select <i>Use Existing</i> and select the name of the phase 1 configuration that you created. For example, dialup_p1. See <a href="#">"Configuring IPsec" on page 1789</a> .
<b>Allow traffic to be initiated from the remote site</b>	enable

### Configuring the IPSEC security policy - CLI

If your VPN tunnel (phase 1) is called dialup\_p1, your protected network is on port2, and your public interface is port1, you would enter:

```
config firewall policy
 edit 0
 set srcintf port2
 set dstintf port1
 set srcaddr all
 set dstaddr all
 set action ipsec
 set schedule always
 set service ANY
 set inbound enable
 set vpngroup dialup_p1
 end
```

### Configuring the ACCEPT security policy - web-based manager

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter the following information and select *OK*:

<b>Incoming Interface</b>	Select the FortiGate unit's public interface.
<b>Source Address</b>	Select the firewall address that you defined for the L2TP clients.
<b>Outgoing Interface</b>	Select the interface that connects to the private network behind this FortiGate unit.
<b>Destination Address</b>	All
<b>Action</b>	ACCEPT

### Configuring the ACCEPT security policy - CLI

If your public interface is port1, your protected network is on port2, and L2TPclients is the address range that L2TP clients use, you would enter:

```
config firewall policy
 edit 0
 set srcintf port1
 set dstintf port2
 set srcaddr L2TPclients
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
 end
```



## Configuring the Windows PC

Configuration of the Windows PC for a VPN connection to the FortiGate unit consists of the following:

- In Network Connections, configure a Virtual Private Network connection to the FortiGate unit.
- Ensure that the IPSEC service is running.
- Ensure that IPsec has not been disabled for the VPN client. It may have been disabled to make the Microsoft VPN compatible with an earlier version of FortiOS.

The instructions in this section are based on Windows XP SP3. Other versions of Windows may vary slightly.

### To configure the network connection

1. Open *Network Connections*.  
This is available through the Control Panel.
2. Double-click *New Connection Wizard* and *Select Next*.
3. Select *Connect to the network at my workplace*.
4. Select *Next*.
5. Select *Virtual Private Network connection* and select *Next*.
6. In the *Company Name* field, enter a name for the connection and select *Next*.
7. Select *Do not dial the initial connection* and then select *Next*.
8. Enter the public IP address or FQDN of the FortiGate unit and select *Next*.
9. Optionally, select *Add a shortcut to this connection to my desktop*.
10. Select *Finish*.  
The *Connect* dialog opens on the desktop.
11. Select *Properties* and then select the *Security* tab.
12. Select *IPSec Settings*.
13. Select *Use pre-shared key for authentication*, enter the preshared key that you configured for your VPN, and select *OK*.
14. Select *OK*.

### To check that the IPSEC service is running

1. Open *Administrative Tools* through the Control Panel.
2. Double-click *Services*.
3. Look for IPSEC Services. Confirm that the *Startup Type* is *Automatic* and *Status* is set to *Started*. If needed, double-click *IPSEC Services* to change these settings.

### To check that IPsec has not been disabled

1. Select *Start > Run*.
2. Enter `regedit` and select *OK*.
3. Find the Registry key  
HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\Parameters
4. If there is a `ProhibitIPSec` value, it must be set to 0.

## Troubleshooting

This section describes some checks and tools you can use to resolve issues with L2TP-over-IPsec VPNs.

This section includes:

- [Quick checks](#)
- [Mac OS X and L2TP](#)
- [Setting up logging](#)
- [Using the FortiGate unit debug commands](#)

### Quick checks

The table below is a list of common L2TP over IPsec VPN problems and the possible solutions.

Problem	What to check
IPsec tunnel does not come up.	<p>Check the logs to determine whether the failure is in Phase 1 or Phase 2.</p> <p>Check the settings, including encapsulation setting, which must be transport-mode.</p> <p>Check the user password.</p> <p>Confirm that the user is a member of the user group assigned to L2TP.</p> <p>On the Windows PC, check that the IPsec service is running and has not been disabled. See <a href="#">“Configuring the Windows PC” on page 1793</a>.</p>
Tunnel connects, but there is no communication.	<p>Did you create an ACCEPT security policy from the public network to the protected network for the L2TP clients? See <a href="#">“Configuring security policies” on page 1791</a>.</p>

### Mac OS X and L2TP

FortiOS allows L2TP connections with empty AVP host names and therefore Mac OS X L2TP connections can connect to the FortiGate.

Prior to FortiOS 4.0 MR3, FortiOS refused L2TP connections with empty AVP host names in compliance with RFC 2661 and RFC 3931.

### Setting up logging

L2TP logging must be enabled to record L2TP events. Alert email can be configured to report L2TP errors.

#### To configure FortiGate logging for L2TP over IPsec

1. Go to *Log & Report > Log Config > Log Settings*.
2. Select *Event Log*.
3. Select the *VPN activity event* check box.
4. Select *Apply*.

### To view FortiGate logs

1. Go to *Log & Report > Event Log > VPN*.
2. Select the *Log location* if required.
3. After each attempt to start the L2TP over IPsec VPN, select *Refresh* to view logged events.

## Using the FortiGate unit debug commands

### To view debug output for IKE and L2TP

1. Start an SSH or Telnet session to your FortiGate unit.
2. Enter the following CLI commands

```
diagnose debug application ike -1
diagnose debug application l2tp -1
diagnose debug enable
```
3. Attempt to use the VPN and note the debug output in the SSH or Telnet session.
4. Enter the following command to reset debug settings to default:

```
diagnose debug reset
```

### To use the packet sniffer

1. Start an SSH or Telnet session to your FortiGate unit.
2. Enter the following CLI command

```
diagnose sniffer packet any icmp 4
```
3. Attempt to use the VPN and note the debug output.
4. Enter `Ctrl-C` to end sniffer operation.

### Typical L2TP over IPsec session startup log entries - raw format

```
2010-01-11 16:39:58 log_id=0101037127 type=event subtype=ipsec
pri=notice vd="root" msg="progress IPsec phase 1" action="negotiate"
rem_ip=172.20.120.151 loc_ip=172.20.120.141 rem_port=500 loc_port=500
out_intf="port1" cookies="5f6da1c0e4bbf680/d6a1009eb1dde780" user="N/A"
group="N/A" xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1"
status=success init=remote mode=main dir=outbound stage=1
role=responder result=OK
```

```
2010-01-11 16:39:58 log_id=0101037127 type=event subtype=ipsec
pri=notice vd="root" msg="progress IPsec phase 1" action="negotiate"
rem_ip=172.20.120.151 loc_ip=172.20.120.141 rem_port=500 loc_port=500
out_intf="port1" cookies="5f6da1c0e4bbf680/d6a1009eb1dde780" user="N/A"
group="N/A" xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1"
status=success init=remote mode=main dir=outbound stage=2
role=responder result=OK
```

```
2010-01-11 16:39:58 log_id=0101037127 type=event subtype=ipsec
pri=notice vd="root" msg="progress IPsec phase 1" action="negotiate"
rem_ip=172.20.120.151 loc_ip=172.20.120.141 rem_port=500 loc_port=500
out_intf="port1" cookies="5f6da1c0e4bbf680/d6a1009eb1dde780" user="N/A"
group="N/A" xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1"
status=success init=remote mode=main dir=inbound stage=3 role=responder
result=DONE
```

```
2010-01-11 16:39:58 log_id=0101037127 type=event subtype=ipsec
pri=notice vd="root" msg="progress IPsec phase 1" action="negotiate"
rem_ip=172.20.120.151 loc_ip=172.20.120.141 rem_port=500 loc_port=500
out_intf="port1" cookies="5f6da1c0e4bbf680/d6a1009eb1dde780" user="N/A"
```

group="N/A" xauth\_user="N/A" xauth\_group="N/A" vpn\_tunnel="dialup\_p1\_0"  
status=success init=remote mode=main dir=outbound stage=3  
role=responder result=DONE

2010-01-11 16:39:58 log\_id=0101037129 type=event subtype=ipsec  
pri=notice vd="root" msg="progress IPsec phase 2" action="negotiate"  
rem\_ip=172.20.120.151 loc\_ip=172.20.120.141 rem\_port=500 loc\_port=500  
out\_intf="port1" cookies="5f6da1c0e4bbf680/d6a1009eb1dde780" user="N/A"  
group="N/A" xauth\_user="N/A" xauth\_group="N/A" vpn\_tunnel="dialup\_p1\_0"  
status=success init=remote mode=quick dir=outbound stage=1  
role=responder result=OK

2010-01-11 16:39:58 log\_id=0101037133 type=event subtype=ipsec  
pri=notice vd="root" msg="install IPsec SA" action="install\_sa"  
rem\_ip=172.20.120.151 loc\_ip=172.20.120.141 rem\_port=500 loc\_port=500  
out\_intf="port1" cookies="5f6da1c0e4bbf680/d6a1009eb1dde780" user="N/A"  
group="N/A" xauth\_user="N/A" xauth\_group="N/A" vpn\_tunnel="dialup\_p1\_0"  
role=responder in\_spi=61100fe2 out\_spi=bd70fca1

2010-01-11 16:39:58 log\_id=0101037139 type=event subtype=ipsec  
pri=notice vd="root" msg="IPsec phase 2 status change"  
action="phase2-up" rem\_ip=172.20.120.151 loc\_ip=172.20.120.141  
rem\_port=500 loc\_port=500 out\_intf="port1"  
cookies="5f6da1c0e4bbf680/d6a1009eb1dde780" user="N/A" group="N/A"  
xauth\_user="N/A" xauth\_group="N/A" vpn\_tunnel="dialup\_p1\_0"  
phase2\_name=dialup\_p2

2010-01-11 16:39:58 log\_id=0101037138 type=event subtype=ipsec  
pri=notice vd="root" msg="IPsec connection status change"  
action="tunnel-up" rem\_ip=172.20.120.151 loc\_ip=172.20.120.141  
rem\_port=500 loc\_port=500 out\_intf="port1"  
cookies="5f6da1c0e4bbf680/d6a1009eb1dde780" user="N/A" group="N/A"  
xauth\_user="N/A" xauth\_group="N/A" vpn\_tunnel="dialup\_p1\_0"  
tunnel\_ip=172.20.120.151 tunnel\_id=1552003005 tunnel\_type=ipsec  
duration=0 sent=0 rcvd=0 next\_stat=0 tunnel=dialup\_p1\_0

2010-01-11 16:39:58 log\_id=0101037129 type=event subtype=ipsec  
pri=notice vd="root" msg="progress IPsec phase 2" action="negotiate"  
rem\_ip=172.20.120.151 loc\_ip=172.20.120.141 rem\_port=500 loc\_port=500  
out\_intf="port1" cookies="5f6da1c0e4bbf680/d6a1009eb1dde780" user="N/A"  
group="N/A" xauth\_user="N/A" xauth\_group="N/A" vpn\_tunnel="dialup\_p1\_0"  
status=success init=remote mode=quick dir=inbound stage=2  
role=responder result=DONE

2010-01-11 16:39:58 log\_id=0101037122 type=event subtype=ipsec  
pri=notice vd="root" msg="negotiate IPsec phase 2" action="negotiate"  
rem\_ip=172.20.120.151 loc\_ip=172.20.120.141 rem\_port=500 loc\_port=500  
out\_intf="port1" cookies="5f6da1c0e4bbf680/d6a1009eb1dde780" user="N/A"  
group="N/A" xauth\_user="N/A" xauth\_group="N/A" vpn\_tunnel="dialup\_p1\_0"  
status=success role=responder esp\_transform=ESP\_3DES esp\_auth=HMAC\_SHA1

2010-01-11 16:39:58 log\_id=0103031008 type=event subtype=ppp vd=root  
pri=information action=connect status=success msg="Client  
172.20.120.151 control connection started (id 805), assigned ip  
192.168.0.50"

2010-01-11 16:39:58 log\_id=0103029013 type=event subtype=ppp vd=root  
pri=notice pppd is started

2010-01-11 16:39:58 log\_id=0103029002 type=event subtype=ppp vd=root  
pri=notice user="user1" local=172.20.120.141 remote=172.20.120.151

```
assigned=192.168.0.50 action=auth_success msg="User 'user1' using l2tp
with authentication protocol MSCHAP_V2, succeeded"

2010-01-11 16:39:58 log_id=0103031101 type=event subtype=ppp vd=root
pri=information action=tunnel-up tunnel_id=1645784497 tunnel_type=l2tp
remote_ip=172.20.120.151 tunnel_ip=192.168.0.50 user="user1"
group="L2TPusers" msg="L2TP tunnel established"
```

# GRE over IPsec (Cisco VPN)

This section describes how to configure a FortiGate VPN that is compatible with Cisco-style VPNs that use GRE in an IPsec tunnel.

The following topics are included in this section:


- [Overview](#)
- [Configuring the FortiGate unit](#)
- [Configuring the Cisco router](#)
- [Troubleshooting](#)

## Overview

Cisco products that include VPN support often use Generic Routing Encapsulation (GRE) protocol tunnel over IPsec encryption. This chapter describes how to configure a FortiGate unit to work with this type of Cisco VPN.

Cisco VPNs can use either transport mode or tunnel mode IPsec. Before FortiOS 4.0 MR2, the FortiGate unit was compatible only with tunnel mode IPsec.

**Figure 282:**Example FortiGate to Cisco GRE-over-IPsec VPN



In this example, users on LAN-1 are provided access to LAN-2.

## Configuring the FortiGate unit

There are several steps to the GRE-over-IPsec configuration:

- Enable overlapping subnets. This is needed because the IPsec and GRE tunnels will use the same addresses.
- Configure a route-based IPsec VPN on the external interface.
- Configure a GRE tunnel on the virtual IPsec interface. Set its local gateway and remote gateway addresses to match the local and remote gateways of the IPsec tunnel.
- Configure security policies to allow traffic to pass in both directions between the GRE virtual interface and the IPsec virtual interface.
- Configure security policies to allow traffic to pass in both directions between the protected network interface and the GRE virtual interface.
- Configure a static route to direct traffic destined for the network behind the Cisco router into the GRE-over-IPsec tunnel.

### Enabling overlapping subnets

By default, each FortiGate unit network interface must be on a separate network. The configuration described in this chapter assigns an IPsec tunnel end point and the external interface to the same network. Enable subnet overlap as follows:

```
config system settings
 set allow-subnet-overlap enable
end
```

### Configuring the IPsec VPN

A route-based VPN is required. It must use encryption and authentication algorithms compatible with the Cisco equipment to which it connects. In this chapter, preshared key authentication is shown.

#### To configure the IPsec VPN - web-based manager

1. Define the phase 1 configuration needed to establish a secure connection with the remote Cisco device. Enter these settings in particular:

<b>Name</b>	Enter a name to identify the VPN tunnel, tocsico for example. This is the name of the virtual IPsec interface. It appears in phase 2 configurations, security policies and the VPN monitor.
<b>Remote Gateway</b>	Select <i>Static IP Address</i> .
<b>IP Address</b>	Enter the IP address of the Cisco device public interface. For example, 192.168.5.113.
<b>Local Interface</b>	Select the FortiGate unit's public interface. For example, 172.20.120.141.
<b>Mode</b>	Select <i>Main (ID Protection)</i> .
<b>Authentication Method</b>	Preshared Key
<b>Pre-shared Key</b>	Enter the preshared key. It must match the preshared key on the Cisco device.

<b>Advanced</b>	Select the Advanced button to see the following settings.
<b>Enable IPsec Interface Mode</b>	Enable.
<b>P1 Proposal</b>	3DES-MD5 At least one proposal must match the settings on the Cisco unit.

For more information about these settings, see [“Auto Key phase 1 parameters” on page 1637](#).

- Define the phase 2 parameters needed to create a VPN tunnel with the remote peer. For compatibility with the Cisco router, Quick Mode Selectors must be entered, which includes specifying protocol 47, the GRE protocol. Enter these settings in particular:

<b>Name</b>	Enter a name to identify this phase 2 configuration.
<b>Phase 1</b>	Select the name of the phase 1 configuration that you defined in <a href="#">Step 1</a> .
<b>Advanced</b>	Select <i>Advanced</i> to view the following fields.
<b>P2 Proposal</b>	3DES-MD5 At least one proposal must match the settings on the Cisco unit.
<b>Quick Mode Selector</b>	
<b>Source Address</b>	Enter the GRE local tunnel end IP address. For example 172.20.120.141.
<b>Source Port</b>	0
<b>Destination Address</b>	Enter the GRE remote tunnel end IP address. For example 192.168.5.113.
<b>Destination Port</b>	0
<b>Protocol</b>	47

For more information about these settings, see [“Phase 2 parameters” on page 1653](#).

- If the Cisco device is configured to use transport mode IPsec, you need to use transport mode on the FortiGate VPN. You can configure this only in the CLI. In your phase 2 configuration, set encapsulation to transport-mode as follows:

```
config vpn phase2-interface
 edit to_cisco_p2
 set encapsulation transport-mode
 end
```



## To configure the IPsec VPN - CLI

```
config vpn ipsec phase1-interface
edit tocisco
set interface port1
set proposal 3des-sha1 aes128-sha1
set remote-gw 192.168.5.113
set psksecret xxxxxxxxxxxxxxxxx
end
config vpn ipsec phase2-interface
edit tocisco_p2
set phase1name "tocisco"
set proposal 3des-md5
set encapsulation tunnel-mode // if tunnel mode
set encapsulation transport-mode // if transport mode
set protocol 47
set src-addr-type ip
set dst-start-ip 192.168.5.113
set src-start-ip 172.20.120.141
end
```

## Adding IPsec tunnel end addresses

The Cisco configuration requires an address for its end of the IPsec tunnel. The addresses are set to match the GRE gateway addresses. Use the CLI to set the addresses, like this:

```
config system interface
edit tocisco
set ip 172.20.120.141 255.255.255.255
set remote-ip 192.168.5.113
end
```

## Configuring the GRE tunnel

The GRE tunnel runs between the virtual IPsec public interface on the FortiGate unit and the Cisco router. You must use the CLI to configure a GRE tunnel. In the example, you would enter:

```
config system gre-tunnel
edit gre1
set interface tocisco
set local-gw 172.20.120.141
set remote-gw 192.168.5.113
end
```

interface is the virtual IPsec interface, local-gw is the FortiGate unit public IP address, and remote-gw is the remote Cisco device public IP address

## Adding GRE tunnel end addresses

You will also need to add tunnel end addresses. The Cisco router configuration requires an address for its end of the GRE tunnel. Using the CLI, enter tunnel end addresses that are not used elsewhere on the FortiGate unit, like this:

```
config system interface
 edit gre1
 set ip 10.0.1.1 255.255.255.255
 set remote-ip 10.0.1.2
 end
```

## Configuring security policies

Two sets of security policies are required:

- policies to allow traffic to pass in both directions between the GRE virtual interface and the IPsec virtual interface.
- policies to allow traffic to pass in both directions between the protected network interface and the GRE virtual interface.

### To configure security policies - web-based manager

1. Define an ACCEPT firewall security policy to permit communications between the protected network and the GRE tunnel:

<b>Incoming Interface</b>	Select the interface that connects to the private network behind this FortiGate unit.
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Select the GRE tunnel virtual interface you configured.
<b>Destination Address</b>	All
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Disable

2. To permit the remote client to initiate communication, you need to define a firewall address security policy for communication in that direction:

<b>Incoming Interface</b>	Select the GRE tunnel virtual interface you configured.
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Select the interface that connects to the private network behind this FortiGate unit.
<b>Destination Address</b>	All
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Disable

- Define a pair of ACCEPT firewall address security policies to permit traffic to flow between the GRE virtual interface and the IPsec virtual interface:

<b>Incoming Interface</b>	Select the GRE virtual interface. See <a href="#">“Configuring the GRE tunnel” on page 1801</a> .
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Select the virtual IPsec interface you created. See <a href="#">“Configuring the IPsec VPN” on page 1799</a> .
<b>Destination Address</b>	All
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Disable
<b>Incoming Interface</b>	Select the virtual IPsec interface you created. See <a href="#">“Configuring the IPsec VPN” on page 1799</a> .
<b>Source Address</b>	All
<b>Outgoing Interface</b>	Select the GRE virtual interface. See <a href="#">“Configuring the GRE tunnel” on page 1801</a> .
<b>Destination Address</b>	All
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Disable

### To configure security policies - CLI

```

config firewall policy
 edit 1 // LAN to GRE tunnel
 set srcintf port2
 set dstintf gre1
 set srcaddr all
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
 next
 edit 2 // GRE tunnel to LAN
 set srcintf gre1
 set dstintf port2
 set srcaddr all
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
 next

```

```

edit 3 // GRE tunnel to IPsec interface
 set srcintf "gre1"
 set dstintf "tocisco"
 set srcaddr "all"
 set dstaddr "all"
 set action accept
 set schedule "always"
 set service "ANY"
next
edit 4 // IPsec interface to GRE tunnel
 set srcintf "tocisco"
 set dstintf "gre1"
 set srcaddr "all"
 set dstaddr "all"
 set action accept
 set schedule "always"
 set service "ANY"
end

```

## Configuring routing

Traffic destined for the network behind the Cisco router must be routed to the GRE tunnel. To do this, create a static route

1. Go to *Router > Static > Static Routes* and select *Create New*.  
For low-end FortiGate units, go to *System > Network > Routing* and select *Create New*.
2. Enter the following information and select OK.

<b>Destination IP/Mask</b>	Enter the IP address and netmask for the network behind the Cisco router. For example 10.21.101.0 255.255.255.0.
<b>Device</b>	Select the GRE virtual interface.
<b>Distance (Advanced)</b>	Leave setting at default value.

In the CLI, using the example values, you would enter

```

config router static
 edit 0
 set device gre1
 set dst 10.21.101.0 255.255.255.0
 end

```

## Configuring the Cisco router

Using Cisco IOS, you would configure the Cisco router as follows, using the addresses from the example:

```
config ter
crypto ipsec transform-set myset esp-3des esp-md5-hmac
no mode
exit
no ip access-list extended tunnel
ip access-list extended tunnel
permit gre host 192.168.5.113 host 172.20.120.141
exit
interface Tunnell
ip address 10.0.1.2 255.255.255.0
tunnel source 192.168.5.113
tunnel destination 172.20.120.141
!
ip route 10.11.101.0 255.255.255.0 Tunnell
end
clea crypto sa
clea crypto isakmp
```

For transport mode, change `no mode` to `mode transport`.

This is only the portion of the Cisco router configuration that applies to the GRE-over-IPsec tunnel. For more information, refer to the Cisco documentation.

## Troubleshooting

This section describes some checks and tools you can use to resolve issues with the GRE-over-IPsec VPN.

### Quick checks

Here is a list of common problems and what to verify.

Problem	What to check
No communication with remote network.	Use the <code>execute ping</code> command to ping the Cisco device public interface.  Use the FortiGate VPN Monitor page to see whether the IPsec tunnel is up or can be brought up.

IPsec tunnel does not come up.	<p>Check the logs to determine whether the failure is in Phase 1 or Phase 2.</p> <p>Check that the encryption and authentication settings match those on the Cisco device.</p> <p>Check the encapsulation setting: tunnel-mode or transport-mode. Both devices must use the same mode.</p>
Tunnel connects, but there is no communication.	<p>Check the security policies. See <a href="#">“Configuring security policies” on page 1802</a>.</p> <p>Check routing. See <a href="#">“Configuring routing” on page 1804</a>.</p>

## Setting up logging

### To configure FortiGate logging for IPsec

1. Go to *Log & Report > Log Config > Log Settings*.
2. Select the *Event Logging*.
3. Select *VPN activity event*.
4. Select *Apply*.

### To view FortiGate logs

1. Go to *Log & Report > Event Log > VPN*.
2. Select the log storage type.
3. Select *Refresh* to view any logged events.

## Using diagnostic commands

There are some diagnostic commands that can provide useful information. When using diagnostic commands, it is best practice that you connect to the CLI using a terminal program, such as puTTY, that allows you to save output to a file. This will allow you to review the data later on at your own speed without worry about missed data as the diag output scrolls by.

### To use the packet sniffer

1. Enter the following CLI command:

```
diag sniff packet any icmp 4
```
2. Ping an address on the network behind the FortiGate unit from the network behind the Cisco router.

The output will show packets coming in from the GRE interface going out of the interface that connects to the protected network (LAN) and vice versa. For example:

```
114.124303 gre1 in 10.0.1.2 -> 10.11.101.10: icmp: echo request
114.124367 port2 out 10.0.1.2 -> 10.11.101.10: icmp: echo request
114.124466 port2 in 10.11.101.10 -> 10.0.1.2: icmp: echo reply
114.124476 gre1 out 10.11.101.10 -> 10.0.1.2: icmp: echo reply
```

3. Enter CTRL-C to stop the sniffer.

### To view debug output for IKE

1. Enter the following CLI commands

```
diagnose debug application ike -1
diagnose debug enable
```

2. Attempt to use the VPN or set up the VPN tunnel and note the debug output.
3. Enter CTRL-C to stop the debug output.
4. Enter the following command to reset debug settings to default:  
`diagnose debug reset`

# Protecting OSPF with IPsec

For enhanced security, OSPF dynamic routing can be carried over IPsec VPN links.


The following topics are included in this section:

- [Overview](#)
- [OSPF over IPsec configuration](#)
- [Creating a redundant configuration](#)

## Overview

This chapter shows an example of OSPF routing conducted over an IPsec tunnel between two FortiGate units. The network shown in [Figure 283](#) is a single OSPF area. FortiGate\_1 is an Area border router that advertises a static route to 10.22.10.0/24 in OSPF. FortiGate\_2 advertises its local LAN as an OSPF internal route.

**Figure 283:**OSPF over an IPsec VPN tunnel



The section [“OSPF over IPsec configuration”](#) describes the configuration with only one IPsec VPN tunnel, tunnel\_wan1. Then, the section [“Creating a redundant configuration”](#) on page 1814 describes how you can add a second tunnel to provide a redundant backup path. This is shown in [Figure 283](#) as VPN tunnel “tunnel\_wan2”.

Only the parts of the configuration concerned with creating the IPsec tunnel and integrating it into the OSPF network are described. It is assumed that security policies are already in place to allow traffic to flow between the interfaces on each FortiGate unit.



## OSPF over IPsec configuration

There are several steps to the OSPF-over-IPsec configuration:

- Configure a route-based IPsec VPN on an external interface. It will connect to a corresponding interface on the other FortiGate unit. Define the two tunnel-end addresses.
- Configure a static route to the other FortiGate unit.
- Configure the tunnel network as part of the OSPF network and define the virtual IPsec interface as an OSPF interface.

This section describes the configuration with only one VPN, tunnel\_wan1. The other VPN is added in the section [“Creating a redundant configuration” on page 1814](#).

### Configuring the IPsec VPN

A route-based VPN is required. In this chapter, preshared key authentication is shown. Certificate authentication is also possible. Both FortiGate units need this configuration.

#### To configure Phase 1

- 1 Define the phase 1 configuration needed to establish a secure connection with the other FortiGate unit. For more information, see [“Auto Key phase 1 parameters” on page 1637](#). Enter these settings in particular:

<b>Name</b>	Enter a name to identify the VPN tunnel, tunnel_wan1 for example. This becomes the name of the virtual IPsec interface.
<b>Remote Gateway</b>	Select <i>Static IP Address</i> .
<b>IP Address</b>	Enter the IP address of the other FortiGate unit’s public (Port 2) interface.
<b>Local Interface</b>	Select this FortiGate unit’s public (Port 2) interface.
<b>Mode</b>	Select <i>Main (ID Protection)</i> .
<b>Authentication Method</b>	Preshared Key
<b>Pre-shared Key</b>	Enter the preshared key. It must match the preshared key on the other FortiGate unit.
<b>Advanced</b>	Select <i>Advanced</i> .
<b>Enable IPsec Interface Mode</b>	Enable

#### To assign the tunnel end IP addresses

1. Go to *System > Network > Interfaces*, select the virtual IPsec interface that you just created on Port 2 and select *Edit*.
2. In the *IP* and *Remote IP* fields, enter the following tunnel end addresses:

	<b>FortiGate_1</b>	<b>FortiGate_2</b>
<b>IP</b>	10.1.1.1	10.1.1.2
<b>Remote_IP</b>	10.1.1.2	10.1.1.1

These addresses are from a network that is not used for anything else.

## To configure Phase 2

1. Enter a name to identify this phase 2 configuration, `twan1_p2`, for example.
2. Select the name of the phase 1 configuration that you defined in Step 1, `tunnel_wan1` for example.

## Configuring static routing

You need to define the route for traffic leaving the external interface.

1. Go to *Router > Static > Static Routes*, select *Create New*.
2. For low-end FortiGate units, go to *System > Network > Routing* and select *Create New*.
3. Enter the following information.

<b>Destination IP/Mask</b>	Leave as 0.0.0.0 0.0.0.0.
<b>Device</b>	Select the external interface.
<b>Gateway</b>	Enter the IP address of the next hop router.

## Configuring OSPF

This section does not attempt to explain OSPF router configuration. It focusses on the integration of the IPsec tunnel into the OSPF network. This is accomplished by assigning the tunnel as an OSPF interface, creating an OSPF route to the other FortiGate unit.

This configuration uses loopback interfaces to ease OSPF troubleshooting. The OSPF router ID is set to the loopback interface address. The loopback interface ensures the router is always up. Even though technically the router ID doesn't have to match a valid IP address on the FortiGate unit, having an IP that matches the router ID makes troubleshooting a lot easier.

The two FortiGate units have slightly different configurations. `FortiGate_1` is an AS border router that advertises its static default route. `FortiGate_2` advertises its local LAN as an OSPF internal route.

Setting the router ID for each FortiGate unit to the lowest possible value is useful if you want the FortiGate units to be the designated router (DR) for their respective ASes. This is the router that broadcasts the updates for the AS.

Leaving the IP address on the OSPF interface at 0.0.0.0 indicates that all potential routes will be advertised, and it will not be limited to any specific subnet. For example if this IP address was 10.1.0.0, then only routes that match that subnet will be advertised through this interface in OSPF.

### FortiGate\_1 OSPF configuration

When configuring `FortiGate_1` for OSPF, the loopback interface is created, and then you configure OSPF area networks and interfaces.

With the exception of creating the loopback interface, OSPF for this example can all be configured in either the web-based manager or CLI.

## To create the loopback interface

A loopback interface can be configured in the CLI only. For example, if the interface will have an IP address of 10.0.0.1, you would enter:

```
config system interface
 edit lback1
 set vdom root
 set ip 10.0.0.1 255.255.255.255
 set type loopback
 end
```

The loopback addresses and corresponding router IDs on the two FortiGate units must be different. For example, set the FortiGate 1 loopback to 10.0.0.1 and the FortiGate 2 loopback to 10.0.0.2.

## To configure OSPF area, networks, and interfaces - web-based manager

1. On FortiGate\_1, go to *Router > Dynamic > OSPF*.  
For low end FortiGate units, you first need to enable *Dynamic Routing* by going to *System > Admin > Settings*.
2. Enter the following information to define the router, area, and interface information.

<b>Router ID</b>	Enter 10.0.0.1. Select <i>Apply</i> before entering the remaining information.
<b>Advanced Options</b>	
<b>Redistribute</b>	Select the <i>Connected</i> and <i>Static</i> check boxes. Use their default metric values.
<b>Areas</b>	Select <i>Create New</i> , enter the <i>Area</i> and <i>Type</i> and then select <i>OK</i> .
<b>Area</b>	0.0.0.0
<b>Type</b>	Regular
<b>Interfaces</b>	Enter a name for the OSPF interface, ospf_wan1 for example.
<b>Name</b>	
<b>Interface</b>	Select the virtual IPsec interface, tunnel_wan1.
<b>IP</b>	0.0.0.0

3. For *Networks*, select *Create New*.
4. Enter the *IP/Netmask* of 10.1.1.0/255.255.255.0 and an *Area* of 0.0.0.0.
5. For *Networks*, select *Create New*.
6. Enter the *IP/Netmask* of 10.0.0.1/255.255.255.0 and an *Area* of 0.0.0.0.
7. Select *Apply*.

## To configure OSPF area and interfaces - CLI

Your loopback interface is 10.0.0.1, your tunnel ends are on the 10.1.1.0/24 network, and your virtual IPsec interface is named `tunnel_wan1`. Enter the following CLI commands:

```
config router ospf
 set router-id 10.0.0.1
 config area
 edit 0.0.0.0
 end
 config network
 edit 4
 set prefix 10.1.1.0 255.255.255.0
 next
 edit 2
 set prefix 10.0.0.1 255.255.255.255
 end
 config ospf-interface
 edit ospf_wan1
 set cost 10
 set interface tunnel_wan1
 set network-type point-to-point
 end
 config redistribute connected
 set status enable
 end
 config redistribute static
 set status enable
 end
end
```

## FortiGate\_2 OSPF configuration

When configuring FortiGate\_2 for OSPF, the loopback interface is created, and then you configure OSPF area networks and interfaces.

Configuring FortiGate\_2 differs from FortiGate\_1 in that three interfaces are defined instead of two. The third interface is the local LAN that will be advertised into OSPF.

With the exception of creating the loopback interface, OSPF for this example can all be configured in either the web-based manager or CLI.

### To create the loopback interface

A loopback interface can be configured in the CLI only. For example, if the interface will have an IP address of 10.0.0.2, you would enter:

```
config system interface
 edit lback1
 set vdom root
 set ip 10.0.0.2 255.255.255.255
 set type loopback
 end
```

The loopback addresses on the two FortiGate units must be different. For example, set the FortiGate 1 loopback to 10.0.0.1 and the FortiGate 2 loopback to 10.0.0.2.

## To configure OSPF area and interfaces - web-based manager

1. On FortiGate\_2, go to *Router > Dynamic > OSPF*.  
For low end FortiGate units, you first need to enable *Dynamic Routing* by going to *System > Admin > Settings*.

2. Complete the following.

<b>Router ID</b>	10.0.0.2
<b>Areas</b>	Select <i>Create New</i> , enter the <i>Area</i> and <i>Type</i> and then select <i>OK</i> .
<b>Area</b>	0.0.0.0
<b>Type</b>	Regular
<b>Interfaces</b>	
<b>Name</b>	Enter a name for the OSPF interface, <i>ospf_wan1</i> for example.
<b>Interface</b>	Select the virtual IPsec interface, <i>tunnel_wan1</i> .
<b>IP</b>	0.0.0.0

3. For *Networks*, select *Create New*.
4. Enter the following information for the loopback interface:

<b>IP/Netmask</b>	10.0.0.2/255.255.255.255
<b>Area</b>	0.0.0.0

5. For *Networks*, select *Create New*.
6. Enter the following information for the tunnel interface:

<b>IP/Netmask</b>	10.1.1.0/255.255.255.255
<b>Area</b>	0.0.0.0

7. For *Networks*, select *Create New*.
8. Enter the following information for the local LAN interface:

<b>IP/Netmask</b>	10.31.101.0/255.255.255.255
<b>Area</b>	0.0.0.0

9. Select *Apply*.

## To configure OSPF area and interfaces - CLI

If for example, your loopback interface is 10.0.0.2, your tunnel ends are on the 10.1.1.0/24 network, your local LAN is 10.31.101.0/24, and your virtual IPsec interface is named tunnel\_wan1, you would enter:

```
config router ospf
 set router-id 10.0.0.2
 config area
 edit 0.0.0.0
 end
 config network
 edit 1
 set prefix 10.1.1.0 255.255.255.0
 next
 edit 2
 set prefix 10.31.101.0 255.255.255.0
 next
 edit 2
 set prefix 10.0.0.2 255.255.255.255
 end
 config ospf-interface
 edit ospf_wan1
 set interface tunnel_wan1
 set network-type point-to-point
 end
 end
end
```

## Creating a redundant configuration

You can improve the reliability of the OSPF over IPsec configuration described in the previous section by adding a second IPsec tunnel to use if the default one goes down. Redundancy in this case is not controlled by the IPsec VPN configuration but by the OSPF routing protocol.

To do this you:

- Create a second route-based IPsec tunnel on a different interface and define tunnel end addresses for it.
- Add the tunnel network as part of the OSPF network and define the virtual IPsec interface as an additional OSPF interface.
- Set the OSPF cost for the added OSPF interface to be significantly higher than the cost of the default route.

## Adding the second IPsec tunnel

The configuration is the same as in [“Configuring the IPsec VPN” on page 1809](#), but the interface and addresses will be different. Ideally, the network interface you use is connected to a different Internet service provider for added redundancy.

When adding the second tunnel to the OSPF network, choose another unused subnet for the tunnel ends, 10.1.2.1 and 10.1.2.2 for example.

## Adding the OSPF interface

OSPF uses the metric called cost when determining the best route, with lower costs being preferred. Up to now in this example, only the default cost of 10 has been used. Cost can be set only in the CLI.

The new IPsec tunnel will have its OSPF cost set higher than that of the default tunnel to ensure that it is only used if the first tunnel goes down. The new tunnel could be set to a cost of 200 compared to the default cost is 10. Such a large difference in cost will ensure this new tunnel will only be used as a last resort.

If the new tunnel is called `tunnel_wan2`, you would enter the following on both FortiGate units:

```
config router ospf
 config ospf-interface
 edit ospf_wan2
 set cost 200
 set interface tunnel_wan2
 set network-type point-to-point
 end
 end
end
```

# Hardware offloading and acceleration

FortiGate units incorporate proprietary FortiASIC NPx network processors that can provide accelerated processing for IPsec VPN traffic. This section describes how to configure offloading and acceleration.

The following topics are included in this section:

- [Overview](#)
- [IPsec offloading configuration examples](#)

## Overview

Fortinet's NPx network processors contain features to improve IPsec tunnel performance. For example, network processors can encrypt and decrypt packets, offloading cryptographic work from the FortiGate unit's main processing resources.

On FortiGate units with the appropriate hardware, you can configure offloading of both IPsec sessions and HMAC checking.

## IPsec session offloading requirements

Sessions must be fast path ready. Fast path ready session requirements are:

- Layer 2 type/length must be 0x0800 (IEEE 802.1q VLAN specification is supported); link aggregation between any network interfaces sharing the same network processor(s) may be used (IEEE 802.3ad specification is supported)
- Layer 3 protocol must be IPv4
- Layer 4 protocol must be UDP, TCP or ICMP
- Layer 3 / Layer 4 header or content modification must not require a session helper (for example, SNAT, DNAT, and TTL reduction are supported, but application layer content modification is not supported)
- FortiGate unit security policy must not require antivirus or IPS inspection, although hardware accelerated anomaly checks are acceptable.
- The session must not use an aggregated link or require QoS, including rate limits and bandwidth guarantees (NP1 processor only).
- Ingress and egress network interfaces are both attached to the same network processor(s)
- In Phase I configuration, Local Gateway IP must be specified as an IP address of a network interface attached to a network processor
- In Phase II configuration:
  - encryption algorithm must be DES, 3DES, AES-128, AES-192, AES-256, or null (for NP1 processor, only 3DES is supported)
  - authentication must be MD5, SHA1, or null (for NP1 processor, only MD5 is supported)
  - if replay detection is enabled, encryption and decryption options must be enabled in the CLI (see "[IPsec encryption offloading](#)", below)

If the IPsec session meets the above requirements, the FortiGate unit sends the IPsec security association (SA) and configured processing actions to the network processors.



## Packet offloading requirements

In addition to the session requirements, the packets themselves must meet fast-path requirements:

- Incoming packets must not be fragmented.
- Outgoing packets must be 385 bytes or larger after any fragmentation. This means the configured MTU (Maximum Transmission Unit) for the network processors' interfaces must have an MTU of 385 bytes or larger.

If packet offloading requirements are not met, an individual packet will use the FortiGate unit main processing resources, regardless of whether other packets in the session are offloaded to the specialized network processors.

## IPsec encryption offloading

Network processing unit (NPU) settings configure offloading behavior for IPsec VPNs. Configured behavior applies to all network processors contained by the FortiGate unit itself or any installed AMC modules.

If replay detection is not enabled (IPsec Phase 2 settings), encryption is always offloaded. NPU offloading is supported when the local gateway is a loopback interface.

### To enable offloading of encryption even when replay detection is enabled

```
config system npu
 set enc-offload-antireplay enable
 set offload-ipsec-host enable
end
```

### To enable offloading of decryption even when replay detection is enabled

```
config system npu
 set dec-offload-antireplay enable
end
```

## HMAC check offloading

The Hash-based Message Authentication Code (HMAC) check can also be offloaded to hardware. SHA-256, SHA-384, or SHA-512 cannot be off-loaded to hardware, and must be processed using only software resources.


### To enable HMAC check offloading

```
configure system global
 set ipsec-hmac-offload (enable|disable)
end
```

## IPsec offloading configuration examples

The following examples configure two FortiASIC NPx network processor accelerated VPNs, one route-based, the other policy based. In both cases, the network topology is as shown in [Figure 284](#).

**Figure 284:**Hardware accelerated IPsec VPN topology



## Accelerated route-based VPN configuration

This example uses the accelerated ports on FortiGate-ASM-FB4 modules in each FortiGate unit. These accelerated ports on the modules are paired interfaces that have their own network processor (NPU) to offload work from the FortiGate unit CPU. Beyond this fact, the example is normal VPN example.

Configuring the FortiGate units require the same basic steps:

- Configure VPN Phase 1
- Configure VPN Phase 2
- Create security policies to allow traffic to flow
- Create a static route to allow traffic to flow

When both FortiGates are have the VPN tunnel configured, test to ensure it is working properly.

### To configure FortiGate\_1

1. Go to *VPN > IPsec > Auto Key (IKE)* and select *Create Phase 1*.
2. Configure Phase 1 settings (name *FGT\_1\_IPsec*), plus
  - Select *Advanced*.
  - Select *Enable IPsec Interface Mode*.
  - In *Local Gateway IP*, select *Specify* and enter the VPN IP address 3.3.3.1, which is the IP address of FortiGate\_1's FortiGate-ASM-FB4 module on port 2.
3. Select *OK*.
4. Select *Create Phase 2* and configure Phase 2 settings, including
  - Select *Enable replay detection*.
  - set `enc-offload-antireplay` to enable using the `config system npu` CLI command.
5. Go to *Policy > Policy > Policy*.
6. Configure two firewall address policies (one for each direction) to apply the Phase 1 IPsec configuration you configured in step 2 to traffic leaving from or arriving on FortiGate-ASM-FB4 module port 1.
7. Go to *Router > Static > Static Routes*.  
For low-end FortiGate units, go to *System > Network > Routing*.

8. Configure a static route to route traffic destined for FortiGate\_2's protected network to the virtual IPsec interface, FGT\_1\_IPsec.

**To add the static route from the CLI:**

```
config router static
 edit 2
 set device "FGT_1_IPsec"
 set dst 2.2.2.0 255.255.255.0
 end
```

**To configure FortiGate\_2**

1. Go to *VPN > IPsec > Auto Key (IKE)* and select *Create Phase 1*.
2. Configure Phase 1 settings (name FGT\_2\_IPsec), plus
  - Select *Advanced*.
  - Select *Enable IPsec Interface Mode*.
  - In *Local Gateway IP*, select *Specify* and enter the VPN IP address 3.3.3.2, which is the IP address of FortiGate\_2's FortiGate-ASM-FB4 module on port 2.
3. Select *OK*.
4. Select *Create Phase 2* and configure Phase 2 settings, including
  - Select *Enable replay detection*.
  - set `enc-offload-antireplay` to enable using the `config system npu` CLI command.
5. Go to *Policy > Policy > Policy*.
6. Configure two firewall address policies (one for each direction) to apply the Phase 1 IPsec configuration you configured in step 2 to traffic leaving from or arriving on FortiGate-ASM-FB4 module port 1.
7. Go to *Router > Static > Static Routes*.
8. Configure a static route to route traffic destined for FortiGate\_1's protected network to the virtual IPsec interface, FGT\_2\_IPsec.

**To add the static route from the CLI:**

```
config router static
 edit 2
 set device "FGT_2_IPsec"
 set dst 1.1.1.0 255.255.255.0
 end
```

**To test the VPN**

1. Activate the IPsec tunnel by sending traffic between the two protected networks.
2. To verify tunnel activation, go to *VPN > Monitor > IPsec Monitor*.

## Accelerated policy-based VPN configuration

### To configure FortiGate\_1

1. Go to *VPN > IPsec > Auto Key (IKE)* and select *Create Phase 1*.
2. Configure Phase 1 settings (name FGT\_1\_IPsec), plus
  - Select *Advanced*.
  - Ensure that the *Enable IPsec Interface Mode* check box is not selected.
  - In *Local Gateway IP*, select *Specify* and enter the VPN IP address 3.3.3.1, which is the IP address of FortiGate\_1's FortiGate-ASM-FB4 module on port 2.
3. Select *OK*.
4. Select *Create Phase 2* and configure Phase 2 settings, including
  - Select *Enable replay detection*.
  - `set enc-offload-antireplay to enable` using the `config system npu` CLI command.
5. Go to *Policy > Policy > Policy*.
6. Configure an IPsec VPN policy to apply the Phase 1 IPsec tunnel you configured in step 2 to traffic between FortiGate-ASM-FB4 module ports 1 and 2.
7. Go to *Router > Static > Static Routes*.  
For low-end FortiGate units, go to *System > Network > Routing*.
8. Configure a static route to route traffic destined for FortiGate\_2's protected network to FortiGate\_2's VPN gateway, 3.3.3.2, through the FortiGate-ASM-FB4 module's port 2 (device).

### To add the static route from the CLI:

```
config router static
 edit 0
 set device "AMC-SW1/2"
 set dst 2.2.2.0 255.255.255.0
 set gateway 3.3.3.1
 end
```

### To configure FortiGate\_2

1. Go to *VPN > IPsec > Auto Key (IKE)* and select *Create Phase 1*.
2. Configure Phase 1 settings (name FGT\_2\_IPsec), plus
  - Select *Advanced*.
  - Select *Enable IPsec Interface Mode*.
  - In *Local Gateway IP*, select *Specify* and enter the VPN IP address 3.3.3.2, which is the IP address of FortiGate\_2's FortiGate-ASM-FB4 module on port 2.
3. Select *OK*.
4. Select *Create Phase 2* and configure Phase 2 settings, including
  - Select *Enable replay detection*.
  - `set enc-offload-antireplay to enable` using the `config system npu` CLI command.
5. Go to *Policy > Policy > Policy*.
6. Configure an IPsec VPN policy to apply the Phase 1 IPsec tunnel you configured in step 2 to traffic between FortiGate-ASM-FB4 module ports 1 and 2.

7. Go to *Router > Static > Static Routes*.  
For low-end FortiGate units, go to *System > Network > Routing*.
8. Configure a static route to route traffic destined for FortiGate\_1's protected network to FortiGate\_2's VPN gateway, 3.3.3.2, through the FortiGate-ASM-FB4 module's port 2 (device).

**To add the static route from the CLI:**

```
config router static
 edit 0
 set device "AMC-SW1/2"
 set dst 1.1.1.0 255.255.255.0
 set gateway 3.3.3.2
 end
```

**To test the VPN**

1. Activate the IPsec tunnel by sending traffic between the two protected networks.
2. To verify tunnel activation, go to *VPN > Monitor > IPsec Monitor*.

# Monitoring and troubleshooting

This section provides some general maintenance and monitoring procedures for VPNs.

The following topics are included in this section:

- [Monitoring VPN connections](#)
- [Testing VPN connections](#)
- [Testing VPN connections](#)
- [Logging VPN events](#)
- [VPN troubleshooting tips](#)

## Monitoring VPN connections

You can use the monitor to view activity on IPsec VPN tunnels and to start or stop those tunnels. The display provides a list of addresses, proxy IDs, and timeout information for all active tunnels.

### Monitoring connections to remote peers

The list of tunnels provides information about VPN connections to remote peers that have static IP addresses or domain names. You can use this list to view status and IP addressing information for each tunnel configuration. You can also start and stop individual tunnels from the list.

To view the list of static-IP and dynamic-DNS tunnels go to *VPN > Monitor > IPsec Monitor*.

### Monitoring dialup IPsec connections

The list of dialup tunnels provides information about the status of tunnels that have been established for dialup clients. The list displays the IP addresses of dialup clients and the names of all active tunnels. The number of tunnels shown in the list can change as dialup clients connect and disconnect.

To view the list of dialup tunnels go to *VPN > Monitor > IPsec Monitor*.

If you take down an active tunnel while a dialup client such as FortiClient is still connected, FortiClient will continue to show the tunnel connected and idle. The dialup client must disconnect before another tunnel can be initiated.

The list of dialup tunnels displays the following statistics:

- The Name column displays the name of the tunnel.
- The meaning of the value in the Remote gateway column changes, depending on the configuration of the network at the far end:
  - When a FortiClient dialup client establishes a tunnel, the Remote gateway column displays either the public IP address and UDP port of the remote host device (on which the FortiClient Endpoint Security application is installed), or if a NAT device exists in front

of the remote host, the Remote gateway column displays the public IP address and UDP port of the remote host.

- When a FortiGate dialup client establishes a tunnel, the Remote gateway column displays the public IP address and UDP port of the FortiGate dialup client.
- The Username column displays the peer ID, certificate name, or XAuth user name of the dialup client (if a peer ID, certificate name, or XAuth user name was assigned to the dialup client for authentication purposes).
- The Timeout column displays the time before the next key exchange. The time is calculated by subtracting the time elapsed since the last key exchange from the keylife.
- The Proxy ID Source column displays the IP addresses of the hosts, servers, or private networks behind the FortiGate unit. A network range may be displayed if the source address in the security encryption policy was expressed as a range of IP addresses.
- The meaning of the value in the Proxy ID Destination column changes, depending on the configuration of the network at the far end:
  - When a FortiClient dialup client establishes a tunnel:
    - If VIP addresses are not used and the remote host connects to the Internet directly, the Proxy ID Destination field displays the public IP address of the Network Interface Card (NIC) in the remote host.
    - If VIP addresses are not used and the remote host is behind a NAT device, the Proxy ID Destination field displays the private IP address of the NIC in the remote host.
    - If VIP addresses were configured (manually or through FortiGate DHCP relay), the Proxy ID Destination field displays either the VIP address belonging to a FortiClient dialup client, or a subnet address from which VIP addresses were assigned.
  - When a FortiGate dialup client establishes a tunnel, the Proxy ID Destination field displays the IP address of the remote private network.

## Testing VPN connections

A VPN connection has multiple stages that can be confirmed to ensure the connection is working properly. It is easiest to see if the final stage is successful first since if it is successful the other stages will be working properly. Otherwise, you will need to work back through the stages to see where the problem is located.

When a VPN connection is properly established, traffic will flow from one end to the other as if both ends were physically in the same place. If you can determine the connection is working properly then any problems are likely problems with your applications.

If the connection is not working properly, you can move on to [“Troubleshooting VPN connections” on page 1824](#) to determine the exact problem.

### LAN interface connection

To confirm whether a VPN connection over LAN interfaces has been configured correctly, issue a ping or traceroute command on the network behind the FortiGate unit to test the connection to a computer on the remote network. If the connection is properly configured, a VPN tunnel will be established automatically when the first data packet destined for the remote network is intercepted by the FortiGate unit.

If the ping or traceroute fail, it indicates a connection problem between the two ends of the tunnel. This may or may not indicate problems with the VPN tunnel. You can confirm this by going to *VPN > Monitor > IPsec Monitor* where you will be able to see your connection. A green arrow means the tunnel is up and currently processing traffic. A red arrow means the tunnel is not processing traffic, and this VPN connection has a problem.

If the connection has problems, see [“Troubleshooting VPN connections” on page 1824](#).

## Dialup connection

A dialup VPN connection has additional steps. To confirm that a VPN between a local network and a dialup client has been configured correctly, at the dialup client, issue a ping command to test the connection to the local network. The VPN tunnel initializes when the dialup client attempts to connect.

If the ping or traceroute fail, it indicates a connection problem between the two ends of the tunnel. This may or may not indicate problems with the VPN tunnel, or dialup client. As with the LAN connection, confirm the VPN tunnel is established by checking *VPN > Monitor > IPsec Monitor*.

## Troubleshooting VPN connections

If you have determined that your VPN connection is not working properly through [“Testing VPN connections” on page 1823](#), the next step is to verify that you have a phase2 connection.

If traffic is not passing through the FortiGate unit as you expect, ensure the traffic does not contain IPcomp packets (IP protocol 108, RFC 3173). FortiGate units do not allow IPcomp packets, they compress packet payload, preventing it from being scanned.

Testing phase 1 and 2 connections is a bit more difficult than testing the working VPN. This is because they require diagnose CLI commands. These commands are typically used by Fortinet customer support to discover more information about your FortiGate unit and its current configuration.

Before you start troubleshooting you need to:

- configure FortiGate units on both ends for interface VPN
- record the information in your VPN phase 1 and phase 2 configurations - for our example here the remote IP address is 10.11.101.10 and the names of the phases are Phase1 and Phase2
- install a telnet or SSH client such as putty that allows logging of output
- ensure that the admin interface supports your chosen connection protocol so you can connect to your FortiGate unit admin interface.
- For this example, default values were used unless stated otherwise.

### To get diagnose information for the VPN connection - CLI

1. Log into the CLI as admin with the output being logged to a file.
2. Stop any diagnose debug sessions that are currently running with the CLI command  
`diagnose debug disable`
3. Clear any existing log-filters by running  
`diagnose vpn ike log-filter clear`
4. Set the log-filter to the IP address of the remote computer (10.11.101.10). This filters out all VPN connections except ones to the IP address we are concerned with. The command is  
`diagnose vpn ike log-filter dst-addr4 10.11.101.10.`
5. Set up the commands to output the VPN handshaking. The commands are:  
`diagnose debug app ike 255`  
`diagnose debug enable`



6. Have the remote FortiGate initiate the VPN connection in the web-based manager by going to *VPN > Monitor* and selecting *Bring up*.

This makes the remote FortiGate the initiator and the local FortiGate becomes the responder. Establishing the connection in this manner means the local FortiGate will have its configuration information as well as the information the remote computer sends. Having both sets of information locally makes it easier to troubleshoot your VPN connection.

7. Watch the screen for output, and after roughly 15 seconds enter the following CLI command to stop the output.

```
diagnose debug disable
```

8. If needed, save the log file of this output to a file on your local computer. Saving the output to a file can make it easier to search for a particular phrase, and is useful for comparisons.

### To troubleshoot a phase1 VPN connection

Using the output from [“To get diagnose information for the VPN connection - CLI” on page 1824](#), search for the word `proposal` in the output. It may occur once indicating a successful connection, or it will occur two or more times for an unsuccessful connection — there will be one proposal listed for each end of the tunnel and each possible combination in their settings. For example if 10.11.101.10 selected both DH Group 1 and 5, that would be at least 2 proposals set.

A successful negotiation proposal will look similar to

```
IPsec SA connect 26 10.12.101.10->10.11.101.10:500
config found
created connection: 0x2f55860 26 10.12.101.10->10.11.101.10:500
IPsec SA connect 26 10.12.101.10->10.11.101.10:500 negotiating
no suitable ISAKMP SA, queuing quick-mode request and initiating
 ISAKMP SA negotiation
initiator: main mode is sending 1st message...
cookie 3db6afe559e3df0f/0000000000000000
out [encryption]
sent IKE msg (ident-ilsend): 10.12.101.10:500->10.11.101.10:500,
 len=264, id=3db6afe559e3df0f/0000000000000000
diaike 0: comes 10.12.101.1:500->10.11.101.1:500,ifindex=26....
```

Note the phrase “`initiator: main mode is sending 1st message...`” which shows you the handshake between the ends of the tunnel is in progress. Initiator shows the remote unit is sending the first message.

## Logging VPN events

You can configure the FortiGate unit to log VPN events. For IPsec VPNs, phase 1 and phase 2 authentication and encryption events are logged. For information about how to interpret log messages, see the [FortiGate Log Message Reference](#).

### To log VPN events

1. Go to *Log & Report > Log Config > Log Settings*.
2. Verify that the *VPN activity event* option is selected.
3. Select *Apply*.

### To view event logs

1. Go to *Log & Report > Event Log > VPN*.
2. Select the *Log location*.

## VPN troubleshooting tips

More in-depth VPN troubleshooting can be found in the [Troubleshooting guide handbook chapter](#).

### The VPN proposal is not connecting

One side may attempt to initiate the VPN tunnel unsuccessfully. There are a number of potential reasons for this problem.

### Attempting hardware offloading beyond SHA1

If you are trying to off-load VPN processing to a network processing unit (NPU), remember that only SHA1 authentication is supported. For high levels of authentication such as SHA256, SHA384, and SHA512 hardware offloading is not an option — all VPN processing must be done in software.

### Check Phase 1 proposal settings

Ensure that both sides have at least one Phase 1 proposal in common. Otherwise they will not connect. If there are many proposals in the list, this will slow down the negotiating of Phase 1. If its too slow, the connection may timeout before completing. If this happens, try removing some of the unused proposals.

NPU offloading is supported when the local gateway is a loopback interface.

### Check your routing

If routing is not properly configured with an entry for the remote end of the VPN tunnel, traffic will not flow properly. You may need static routes on both ends of the tunnel. If routing is the problem, the proposal will likely setup properly but no traffic will flow.

### Try enabling XAuth

If one end of an attempted VPN tunnel is using XAuth and the other end is not, the connection attempt will fail. The log messages for the attempted connection will not mention XAuth is the reason, but when connections are failing it is a good idea to ensure both ends have the same XAuth settings. If you do not know the other end's settings enable or disable XAuth on your end to see if that is the problem.

## General troubleshooting tips

Most connection failures are due to a configuration mismatch between the FortiGate unit and the remote peer. In general, begin troubleshooting an IPsec VPN connection failure as follows:

1. Ping the remote network or client to verify whether the connection is up. See [“Testing VPN connections” on page 1823](#).
2. Traceroute the remote network or client. If DNS is working, you can use domain names. Otherwise use IP addresses.
3. Check the routing behind the dialup client. Routing problems may be affecting DHCP. If this appears to be the case, configure a DHCP relay service to enable DHCP requests to be relayed to a DHCP server on or behind the FortiGate server.

4. Verify the configuration of the FortiGate unit and the remote peer. Check the following IPsec parameters:
  - The mode setting for ID protection (main or aggressive) on both VPN peers must be identical.
  - The authentication method (preshared keys or certificates) used by the client must be supported on the FortiGate unit and configured properly.
  - If preshared keys are being used for authentication purposes, both VPN peers must have identical preshared keys.
  - The remote client must have at least one set of phase 1 encryption, authentication, and Diffie-Hellman settings that match corresponding settings on the FortiGate unit.
  - Both VPN peers must have the same NAT traversal setting (enabled or disabled).
  - The remote client must have at least one set of phase 2 encryption and authentication algorithm settings that match the corresponding settings on the FortiGate unit.
  - If you are using manual keys to establish a tunnel, the *Remote SPI* setting on the FortiGate unit must be identical to the *Local SPI* setting on the remote peer, and vice versa.
5. To correct the problem, see the following table.

**Table 75: VPN trouble-shooting tips**

<b>Configuration problem</b>	<b>Correction</b>
<b>Mode settings do not match.</b>	Select complementary mode settings. See <a href="#">“Choosing main mode or aggressive mode”</a> on page 1638.
<b>Peer ID or certificate name of the remote peer or dialup client is not recognized by FortiGate VPN server.</b>	Check Phase 1 configuration. Depending on the Remote Gateway and Authentication Method settings, you have a choice of options to authenticate FortiGate dialup clients or VPN peers by ID or certificate name (see <a href="#">“Authenticating remote peers and clients”</a> on page 1642).  If you are configuring authentication parameters for FortiClient dialup clients, refer to the <a href="#">Authenticating FortiClient Dialup Clients Technical Note</a> .
<b>Preshared keys do not match.</b>	Reenter the preshared key. See <a href="#">“Authenticating remote peers and clients”</a> on page 1642.
<b>Phase 1 or phase 2 key exchange proposals are mismatched.</b>	Make sure that both VPN peers have at least one set of proposals in common for each phase. See <a href="#">“Defining IKE negotiation parameters”</a> on page 1646 and <a href="#">“Configure the phase 2 parameters”</a> on page 1655.
<b>NAT traversal settings are mismatched.</b>	Select or clear both options as required. See <a href="#">“NAT traversal”</a> on page 1649 and <a href="#">“NAT keepalive frequency”</a> on page 1650.
<b>SPI settings for manual key tunnels are mismatched.</b>	Enter complementary SPI settings. See <a href="#">“Manual-key configurations”</a> on page 1770.

## A word about NAT devices

When a device with NAT capabilities is located between two VPN peers or a VPN peer and a dialup client, that device must be NAT traversal (NAT-T) compatible for encrypted traffic to pass through the NAT device. For more information, see [“NAT traversal”](#) on page 1649.

# IPv6 for FortiOS 5.0

The origins of Internet Protocol Version 6 (IPv6) date back to December 1998 with the publication of [RFC 2460](#), which describes IPv6 as the successor to IPv4, the standard communications protocol still in use by the majority of users today. This transition away from IPv4 was a direct response to the foreseeable exhaustion of 32-bit IPv4 addresses, which are virtually all but assigned—all 4.3 billion.

IPv4 uses 32-bit addresses, which means that there is a theoretical address limit of 2 to the power of 32. The IPv6 address scheme is based on a 128-bit address, resulting in a theoretical address limit of 2 to the power of 128.

## Possible addresses:

- IPv4 = Roughly 4.3 billion
- IPv6 = Over 340 undecillion (340 followed by 36 digits)

Assuming a world population of approximately 8 billion people, IPv6 would allow for each individual to have approximately 42,535,295,865,117,200,000,000,000,000 devices with an IP address. That's 42 quintillion devices, so it's unlikely that we will ever need to worry about the availability of IPv6 addresses.

Aside from the difference of possible addresses, there is also the different formatting of the addresses. A computer would view an IPv4 address as a 32-bit string of binary digits made up of 1s and 0s, broken up into 4 octets of 8 digits separated by a period:

```
10101100.00010000.11111110.00000001
```

To make the number more user-friendly, we translate the address into decimal, again 4 octets separated by a period:

```
172.16.254.1
```

A computer would view an IPv6 address as a 128-bit string of binary digits made up of 1s and 0s, broken up into 8 octets of 16 digits separated by a colon:

```
0010000000000001:0000110110111:0000000000000000:00000000
0000010:0000000000000000:0000000000000000:00000000
0000000:000000000100000
```

To make this number a little more user-friendly, we translate it into hexadecimal, again 8 octets separated by a colon, for example:

```
2001:0db8:0000:0002:0000:0000:0000:0020
```

We can further simplify the above address. Because any four-digit group of zeros within an IPv6 address may be reduced to a single zero or altogether omitted, the above address can be reduced to:

```
2001:0db8:0000:0002:0:0:0:20
```

or

```
2001:db8:0:2::20
```

## IPv6 packet structure

Each IPv6 packet consists of a mandatory fixed header and optional extension headers, and carries a payload, which is typically either a datagram and/or Transport Layer information. The

payload could also contain data for the Internet Layer or Link Layer. Unlike IPv4, IPv6 packets aren't fragmented by routers, requiring hosts to implement Maximum Transmission Unit (MTU) Path Discovery for MTUs larger than the smallest MTU (which is 1280 octets).

## Jumbograms and jumbo payloads

In IPv6, packets which exceed the MTU of the underlying network are labelled jumbograms, which consist of a jumbo payload. A jumbogram typically exceeds the IP MTU size limit of 65,535 octets, and provides the jumbo payload option, which can allow up to nearly 4GiB of payload data, as defined in [RFC 2675](#). When the MTU is determined to be too large, the receiving host sends a 'Packet too Big' ICMPv6 type 2 message to the sender.

## Fragmentation and reassembly

As noted, packets that are too large for the MTU require hosts to perform MTU Path Discovery to determine the maximum size of packets to send. Packets that are too large require a 'Fragment' extension header, to divide the payload into segments that are 8 octets in length (except for the last fragment, which is smaller). Packets are reassembled according to the extension header and the fragment offset.

## Benefits of IPv6

Some of the benefits of IPv6 include:

- More efficient routing
- Reduced management requirement
- Stateless auto-reconfiguration of hosts
- Improved methods to change Internet Service Providers
- Better mobility support
- Multi-homing
- Security
- Scoped address: link-local, site-local, and global address space

# IPv6 Features

In order to configure IPv6 features using the web-based manager, IPv6 must be enabled using Feature Select. Go to *System > Config > Features*, enable IPv6, and click *Apply*.

The following IPv6 features are available from the FortiOS 5.0 web-based manager:


- IPv6 policies
- IPv6 Network Address Translation
- ICMPv6
- IPv6 in dynamic routing
- Dual stack routing
- IPv6 tunnelling
- SIP over IPv6
- New Fortinet FortiGate IPv6 MIB fields
- IPv6 Per-IP traffic shaper
- DHCPv6
- IPv6 forwarding—Policies, IPS, Application Control, flow-based antivirus, web filtering, and DLP
- FortiGate interfaces can get IPv6 addresses from an IPv6 DHCP server

## IPv6 policies

IPv6 security policies are created both for an IPv6 network and a transitional network. A transitional network is a network that is transitioning over to IPv6 but must still have access to the Internet or must connect over an IPv4 network.

These policies allow for this specific type of traffic to travel between the IPv6 and IPv4 networks. The IPv6 options for creating these policies is hidden by default. You must enable this feature under *System > Config > Features*.

**Figure 285:** IPv6 policy route



## IPv6 policy routing

IPv6 policy routing functions in the same way as IPv4 policy routing. To add an IPv6 policy route, go to *Router > Static > Policy Routes* and select *Create New > IPv6 Policy Route*.

**Figure 286:** Adding an IPv6 Policy route

You can also use the following command to add IPv6 policy routes:

```
config router policy6
 edit 0
 set input-device <interface>
 set src <ipv6_ip>
 set dst <ipv6_ip>
 set protocol <0-255>
 set gateway <ipv6_ip>
 set output-device <interface>
 set tos <bit_pattern>
 set tos-mask <bit_mask>
 end
```

## IPv6 security policies

IPv6 security policies now support all the features supported by IPv4 security policies. The following new features were added in FortiOS 5.0:

- Policy types and subtypes.
- NAT support including using the destination interface IP address, fixed port, and dynamic IP pools.
- All security features (antivirus, web filtering, application control, IPS, email filtering, DLP, VoIP, and ICAP).
- All traffic shaping options, including: shared traffic shaping, reverse shared traffic shaping, and per-IP traffic shaping.
- All user and device authentication options.

## IPv6 explicit web proxy

With FortiOS 5.0, you can use the explicit web proxy for IPv6 traffic. To do this you need to:

- Enable the Explicit Proxy from the dashboard.
- Enable the IPv6 explicit web proxy from the CLI.
- Enable the explicit web proxy for one or more FortiGate interfaces. These interfaces also need IPv6 addresses.
- Add IPv6 web proxy security policies to allow the explicit web proxy to accept IPv6 traffic.

Use the following steps to set up a FortiGate unit to accept IPv6 traffic for the explicit web proxy at the Internal interface and forward IPv6 explicit proxy traffic out the wan1 interface to the Internet.

1. Go to *System > Dashboard > Status* and turn on *Explicit Proxy* under the *Features > Security Features* widget.

2. Enter the following CLI command to enable the IPv6 explicit web proxy:

```
config web-proxy explicit
 set status enable
 set ipv6-status enable
end
```

3. Go to *System > Network > Interfaces* and edit the *internal* interface, select *Enable Explicit Web Proxy* and select *OK*.

4. Go to *Policy > Policy > IPv6 Policy* and select *Create New* to add an IPv6 explicit web proxy security policy with the following settings shown in [Figure 287](#).

This IPv6 explicit web proxy policy allows traffic from all IPv6 IP addresses to connect through the explicit web proxy and through the wan1 interface to any IPv6 addresses that are accessible from the wan1 interface.



If you have enabled both the IPv4 and the IPv6 explicit web proxy, you can combine IPv4 and IPv6 addresses in a single explicit web proxy policy to allow both IPv4 and IPv6 traffic through the proxy.

**Figure 287:**Example IPv6 Explicit Web Proxy security policy

New Policy	
Policy Type	<input checked="" type="radio"/> Firewall <input type="radio"/> VPN
Policy Subtype	<input checked="" type="radio"/> Address <input type="radio"/> User Identity <input type="radio"/> Device Identity
Incoming Interface	web-proxy
Source Address	Click to add...
Source IPv6 Address	all
Outgoing Interface	port2
Destination Address	Click to add...
Destination IPv6 Address	all
Schedule	always
Service	webproxy
Action	ACCEPT



## Restricting the IP address of the explicit IPv6 web proxy

You can use the following command to restrict access to the IPv6 explicit web proxy using only one IPv6 address. The IPv6 address that you specify must be the IPv6 address of an interface that the explicit HTTP proxy is enabled on. You might want to use this option if the explicit web proxy is enabled on an interface with multiple IPv6 addresses.

For example, to require users to connect to the IPv6 address 2001:db8:0:2::30 to connect to the explicit IPv6 HTTP proxy, use the following command:

```
config web-proxy explicit
 set incoming-ipv6 2001:db8:0:2::30
end
```

## Restricting the outgoing source IP address of the IPv6 explicit web proxy

You can use the following command to restrict the source address of outgoing web proxy packets to a single IPv6 address. The IP address that you specify must be the IPv6 address of an interface that the explicit HTTP proxy is enabled on. You might want to use this option if the explicit HTTP proxy is enabled on an interface with multiple IPv6 addresses.

For example, to restrict the outgoing packet source address to 2001:db8:0:2::50:

```
config http-proxy explicit
 set outgoing-ipv6 2001:db8:0:2::50
end
```

## VIP64

VIP64 policies can be used to configure static NAT virtual IPv6 address for IPv4 addresses. VIP64 can be configured from the CLI using the following commands:

```
config firewall vip64
 edit <zname_str>
 set arp-reply {enable | disable}
 set color <color_int>
 set comment <comment_str>
 set extip <address_ipv6>[-address_ipv6]
 set extport <port_int>
 set id <id_num_str>
 set mappedip [<start_ipv4>-<end_ipv4>]
 set mappedport <port_int>
 set portforward {enable | disable}
 set src-filter <addr_str>
 endIPv6 for FortiOS 5.0
```

**Table 76:** VIP64 CLI Variables and Defaults

Variable	Description	Default
<zname_str>	Enter the name of this virtual IP address.	No default.
arp-reply {enable   disable}	Select to respond to ARP requests for this virtual IP address.	enable
color <color_int>	Enter the number of the color to use for the group icon in the web-based manager.	0

**Table 76:** VIP64 CLI Variables and Defaults

Variable	Description	Default
comment <comment_str>	Enter comments relevant to the configured virtual IP.	No default.
extip <address_ipv6>[-address_ipv6]	<p>Enter the IP address or address range on the external interface that you want to map to an address or address range on the destination network.</p> <p>If <code>mappedip</code> is an IP address range, the FortiGate unit uses <code>extip</code> as the first IP address in the external IP address range, and calculates the last IP address required to create an equal number of external and mapped IP addresses for one-to-one mapping.</p> <p>To configure a dynamic virtual IP that accepts connections destined for any IP address, set <code>extip</code> to <code>::</code>.</p>	::
extport <port_int>	<p>Enter the external port number that you want to map to a port number on the destination network.</p> <p>This option only appears if <code>portforward</code> is enabled.</p> <p>If <code>portforward</code> is enabled and you want to configure a static NAT virtual IP that maps a range of external port numbers to a range of destination port numbers, set <code>extport</code> to the first port number in the range. Then set <code>mappedport</code> to the start and end of the destination port range. The FortiGate unit automatically calculates the end of the <code>extport</code> port number range.</p>	0
id <id_num_str>	Enter a unique identification number for the configured virtual IP. Not checked for uniqueness. Range 0 - 65535.	No default.

**Table 76:** VIP64 CLI Variables and Defaults

Variable	Description	Default
mappedip [<start_ipv4>-<end_ipv4>]	<p>Enter the IP address or IP address range on the destination network to which the external IP address is mapped.</p> <p>If <code>mappedip</code> is an IP address range, the FortiGate unit uses <code>extip</code> as the first IP address in the external IP address range, and calculates the last IP address required to create an equal number of external and mapped IP addresses for one-to-one mapping.</p> <p>If <code>mappedip</code> is an IP address range, the FortiGate unit uses <code>extip</code> as a single IP address to create a one-to-many mapping.</p>	0.0.0.0
mappedport <port_int>	<p>Enter the port number on the destination network to which the external port number is mapped.</p> <p>You can also enter a port number range to forward packets to multiple ports on the destination network.</p> <p>For a static NAT virtual IP, if you add a map to port range the FortiGate unit calculates the external port number range.</p>	0
portforward {enable   disable}	Select to enable port forwarding. You must also specify the port forwarding mappings by configuring <code>extport</code> and <code>mappedport</code> .	disable
src-filter <addr_str>	Enter a source address filter. Each address must be in the form of an IPv4 subnet (x:x:x:x:x:x/n). Separate addresses with spaces.	null

## VIP46

VIP46 policies can be used to configure static NAT virtual IPv4 address for IPv6 addresses. VIP46 can be configured from the CLI using the following commands (see the table below for variable details):

```
config firewall vip46
 edit <name_str>
 set arp-reply {enable | disable}
 set color <color_int>
 set comment <comment_str>
 set extip <address_ipv4>[-address_ipv4]
 set extport <port_int>
 set id <id_num_str>
 set mappedip [<start_ipv6>-<end_ipv6>]
 set mappedport <port_int>
 set portforward {enable | disable}
 set src-filter <add_str>
 end
```

**Table 77:VIP46 CLI Variables and Defaults**

Variable	Description	Default
<name_str>	Enter the name of this virtual IP address.	No default.
arp-reply {enable   disable}	Select to respond to ARP requests for this virtual IP address.	enable
color <color_int>	Enter the number of the color to use for the group icon in the web-based manager.	0
comment <comment_str>	Enter comments relevant to the configured virtual IP.	No default.
extip <address_ipv4>[-address_ipv4]	Enter the IP address or address range on the external interface that you want to map to an address or address range on the destination network.  If mappedip is an IP address range, the FortiGate unit uses extip as the first IP address in the external IP address range, and calculates the last IP address required to create an equal number of external and mapped IP addresses for one-to-one mapping.  To configure a dynamic virtual IP that accepts connections destined for any IP address, set extip to 0.0.0.0.	0.0.0.0

**Table 77:** VIP46 CLI Variables and Defaults

Variable	Description	Default
extport <port_int>	<p>Enter the external port number that you want to map to a port number on the destination network.</p> <p>This option only appears if <code>portforward</code> is enabled.</p> <p>If <code>portforward</code> is enabled and you want to configure a static NAT virtual IP that maps a range of external port numbers to a range of destination port numbers, set <code>extport</code> to the first port number in the range. Then set <code>mappedport</code> to the start and end of the destination port range. The FortiGate unit automatically calculates the end of the <code>extport</code> port number range.</p>	0
id <id_num_str>	Enter a unique identification number for the configured virtual IP. Not checked for uniqueness. Range 0 - 65535.	No default.
mappedip [<start_ipv6>-<end_ipv6>]	<p>Enter the IP address or IP address range on the destination network to which the external IP address is mapped.</p> <p>If <code>mappedip</code> is an IP address range, the FortiGate unit uses <code>extip</code> as the first IP address in the external IP address range, and calculates the last IP address required to create an equal number of external and mapped IP addresses for one-to-one mapping.</p> <p>If <code>mappedip</code> is an IP address range, the FortiGate unit uses <code>extip</code> as a single IP address to create a one-to-many mapping.</p>	::
mappedport <port_int>	<p>Enter the port number on the destination network to which the external port number is mapped.</p> <p>You can also enter a port number range to forward packets to multiple ports on the destination network.</p> <p>For a static NAT virtual IP, if you add a map to port range the FortiGate unit calculates the external port number range.</p>	0
portforward {enable   disable}	Select to enable port forwarding. You must also specify the port forwarding mappings by configuring <code>extport</code> and <code>mappedport</code> .	disable
src-filter <addr_str>	Enter a source address filter. Each address must be in the form of an IPv4 subnet (x.x.x.x/n). Separate addresses with spaces.	null

## IPv6 Network Address Translation

NAT66, NAT64, and DNS64 are now supported for IPv6. These options provide IPv6 NAT and DNS capabilities with IPv6-IPv4 tunnelling or dual stack configurations. The commands are available only in the CLI.


Fortinet supports all features described in [RFC 6146](#). However, for DNS64 there is no support for handling Domain Name System Security Extensions (DNSSEC). DNSSEC is for securing types of information that are provided by the DNS as used on an IP network or networks. You can find more information about DNS64 in [RFC 6147](#).

### NAT64 and DNS64 (DNS proxy)

NAT64 is used to translate IPv6 addresses to IPv4 addresses so that a client on an IPv6 network can communicate transparently with a server on an IPv4 network.

NAT64 is usually implemented in combination with the DNS proxy called DNS64. DNS64 synthesizes AAAA records from A records and is used to synthesize IPv6 addresses for hosts that only have IPv4 addresses. 'DNS proxy' and 'DNS64' are interchangeable terms.

**Figure 288:**Example NAT64 configuration



With a NAT64 and DNS64 configuration in place on a FortiGate unit, clients on an IPv6 network can transparently connect to addresses on an IPv4 network. NAT64 and DNS64 perform the IPv4 to IPv6 transition, allowing clients that have already switched to IPv6 addresses to continue communicating with servers that still use IPv4 addresses.

To enable NAT64 and DNS64, use the following CLI command:

```
config system nat64
 set status enable
end
```

### NAT64 policies

You can configure security policies for NAT64 using the web-based manager. For these options to appear, the feature must be enabled using *Feature Select*. You can then configure the policies under *Policy > Policy > NAT64 Policy*.

NAT64 policies and can also be configured from the CLI using the following command:

```
config firewall policy64
```

In the following section, you will configure a NAT64 policy that allows connections from an internal IPv6 network to an external IPv4 network.

### To configure NAT64 to allow a host on the IPv6 network to connect to the Internet server

In this example, the Internal IPv6 network address is 2001:db8:1::/48 and the external IPv4 network address is 172.20.120.0/24. NAT64 is configured to allow a user on the internal network to connect to the server at IPv4 address 172.20.120.12. In this configuration, sessions exiting the wan1 interface must have their source address changed to an IPv4 address in the range 172.20.120.200 to 172.20.120.210.

1. Enter the following command to enable NAT64.

```
config system nat64
 set status enable
end
```



Enabling NAT64 with the `config system nat64` command means that all IPv6 traffic received by the current VDOM can be subject to NAT64 if the source and destination address matches an NAT64 security policy.

By default, the setting `always-synthesize-aaaa-record` is not enabled. With this setting disabled, the DNS proxy (DNS64) will attempt to find an AAAA records for queries to domain names and therefore resolve the host names to IPv6 addresses. If the DNS proxy cannot find an AAAA record, it synthesizes one by adding the NAT64 prefix to the A record.

By using the `nat64-prefix` option of the `config system nat64` command to change the default nat64 prefix from the well-known prefix of `64:ff9b::/96` and setting `always-synthesize-aaaa-record` to *enable*, the DNS proxy does not check for AAAA records but rather synthesizes AAAA records.

As an alternative to the above entry, there is the optional configuration that would allow the resolution of CNAME queries.

```
config system nat64
 set status enable
 set nat64-prefix 64:ff9b::/96
 set always-synthesize-aaaa-record enable
end
```

2. Enter the following command to add an IPv6 firewall address for the internal network:

```
config firewall address6
 edit internal-net6
 set ip6 2001:db8::/48
 end
```

3. Enter the following command to add an IPv4 firewall address for the external network:

```
config firewall address
 edit external-net4
 set subnet 172.20.120.0/24
 set associated-interface wan1
 end
```

4. Enter the following command to add an IP pool containing the IPv4 address that the should become the source address of the packets exiting the wan1 interface:

```
config firewall ippool
 edit exit-pool4
 set startip 172.20.120.200
 set endip 172.20.120.210
 end
```

5. Enter the following command to add a NAT64 policy that allows connections from the internal IPv6 network to the external IPv4 network:

```
config firewall policy64
 edit 0
 set srcintf internal
 set srcaddr internal-net6
 set dstintf wan1
 set dstaddr external-net4
 set action accept
 set schedule always
 set service ANY
 set logtraffic enable
 set ippool enable
 set poolname exit-pool4
 end
```



The `srcaddr` can be any IPv6 firewall address and the `dstaddr` can be any IPv4 firewall address.

Other NAT64 policy options include `fixedport`, which can be used to prevent NAT64 from changing the destination port. You can also configure traffic shaping for NAT64 policies.

#### How a host on the internal IPv6 network communicates with `example.server.com` that only has IPv4 address on the Internet

1. The host on the internal network does a DNS lookup for `example.server.com` by sending a DNS query for an AAAA record for `example.server.com`.
2. The DNS query is intercepted by the FortiGate DNS proxy.
3. The DNS proxy attempts to resolve the query with a DNS server on the Internet and discovers that there are no AAAA records for `example.server.com`.



The previous step is skipped if `always-synthesize-aaaa-record` is enabled.

4. The DNS proxy performs an A-record query for `example.server.com` and gets back an RRSets containing a single A record with the IPv4 address `172.20.120.12`.
5. The DNS proxy then synthesizes an AAAA record. The IPv6 address in the AAAA record begins with the configured NAT64 prefix in the upper 96 bits and the received IPv4 address in the lower 32 bits. By default, the resulting IPv6 address is `64:ff9b::172.20.120.12`.



6. The host on the internal network receives the synthetic AAAA record and sends a packet to the destination address 64:ff9b::172.20.120.12.
7. The packet is routed to the FortiGate internal interface where it is accepted by the NAT64 security policy.
8. The FortiGate unit translates the destination address of the packets from IPv6 address 64:ff9b::172.20.120.12 to IPv4 address 172.20.120.12 and translates the source address of the packets to 172.20.120.200 (or another address in the IP pool range) and forwards the packets out the wan1 interface to the Internet.

## NAT66

NAT66 is used for translating an IPv6 source or destination address to a different IPv6 source or destination address. NAT66 is not as common or as important as IPv4 NAT, as many IPv6 addresses do not need NAT66 as much as IPv4 NAT. However, NAT66 can be useful for a number of reasons. For example, you may have changed the IP addresses of some devices on your network but want traffic to still appear to be coming from their old addresses. You can use NAT66 to translate the source addresses of packets from the devices to their old source addresses.

In FortiOS 5.0, NAT66 options can be added to an IPv6 security policy from the CLI.

Configuring NAT66 is very similar to configuring NAT in an IPv4 security policy. For example, use the following command to add an IPv6 security policy that translates the source address of IPv6 packets to the address of the destination interface (similar to IPv4 source NAT):

```
config firewall policy6
 edit 0
 set srcintf internal
 set dstintf wan1
 set srcaddr internal_net
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
 set nat enable
 end
```

Its also can be useful to translate one IPv6 source address to another address that is not the same as the address of the exiting interface. You can do this using IP pools. For example, enter the following command to add an IPv6 IP pool containing one IPv6 IP address:

```
configure firewall ippool6
 edit example_6_pool
 set startip 2001:db8::
 set endip 2001:db8::
 end
```

Enter the following command to add an IPv6 firewall address that contains a single IPv6 IP address.

```
configure firewall address6
 edit device_address
 set ip6 2001:db8::132/128
 end
```

Enter the following command to add an IPv6 security policy that accepts packets from a device with IP address 2001:db8::132 and translates the source address to 2001:db8:::

```
config firewall policy6
 edit 0
 set srcintf internal
 set dstintf wan1
 set srcaddr device_address
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
 set nat enable
 set ippool enable
 set poolname example_6_pool
 end
```

### NAT66 destination address translation

NAT66 can also be used to translate destination addresses. This is done in an IPv6 policy by using IPv6 virtual IPs. For example, enter the following command to add an IPv6 virtual IP that maps the destination address 2001:db8::dd to 2001:db8::ee

```
configure firewall vip6
 edit example-vip6
 set extip 2001:db8::dd
 set mappedip 2001:db8::ee
 end
```

Enter the following command to add an IPv6 security policy that accepts packets with a destination address 2001:db8::dd and translates that destination address to 2001:db8::ee

```
config firewall policy6
 edit 0
 set srcintf internal
 set dstintf wan1
 set srcaddr all
 set dstaddr example-vip6
 set action accept
 set schedule always
 set service ANY
 end
```

### NAT64 and NAT66 session failover

The FortiGate Clustering Protocol (FGCP) supports IPv6, NAT64, and NAT66 session failover. If session pickup is enabled, these sessions are synchronized between cluster members and, after an HA failover, the sessions will resume with only minimal interruption.

## NAT46

NAT46 is used to translate IPv4 addresses to IPv6 addresses so that a client on an IPv4 network can communicate transparently with a server on an IPv6 network.

To enable NAT46, use the following CLI command:

```
config system nat46
 set status enable
end
```

### NAT46 policies

Security policies for NAT46 can be configured from the web-based manager. For these options to appear in the web-based manager, this feature must be enabled using *Feature Select*. You can then configure the policies under *Policy > Policy > NAT46 Policy*.

NAT46 policies and can also be configured from the CLI using the following command:

```
config firewall policy46
```

## ICMPv6

Internet Control Message Protocol version 6 (ICMPv6) is the new implementation of the Internet Control Message Protocol (ICMP) that is part of Internet Protocol version 6 (IPv6). The ICMPv6 protocol is defined in [RFC 4443](#).

ICMPv6 is a multipurpose protocol. It performs such things as:

- error reporting in packet processing
- diagnostic functions
- Neighbor Discovery process
- IPv6 multicast membership reporting

It also designed as a framework to use extensions for use with future implementations and changes.

Examples of extensions that have already been written for ICMPv6:

- Neighbor Discovery Protocol (NDP) - a node discovery protocol in IPv6 which replaces and enhances functions of ARP.
- Secure Neighbor Discovery Protocol (SEND) - an extension of NDP with extra security.
- Multicast Router Discovery (MRD) - allows discovery of multicast routers.

ICMPv6 messages use IPv6 packets for transportation and can include IPv6 extension headers. ICMPv6 includes some of the functionality that in IPv4 was distributed among protocols such as ICMPv4, ARP (Address Resolution Protocol), and IGMP (Internet Group Membership Protocol version 3).

ICMPv6 has simplified the communication process by eliminating obsolete messages.

ICMPv6 messages are subdivided into two classes: error messages and information messages.

Error Messages are divided into four categories:

- Destination Unreachable
- Time Exceeded
- Packet Too Big
- Parameter Problems

Information messages are divided into three groups:

- Diagnostic messages
- Neighbor Discovery messages
- Messages for the management of multicast groups.

## ICMPv6 Types and Codes

ICMPv6 has a number of messages that are identified by the “Type” field. Some of these types have assigned “Code” fields as well. The table below shows the different types of ICMP Types with their associated codes if there are any.

Type codes 0 – 127 are error messages and type codes 128 – 255 are for information messages.

**Table 78:**ICMPv6 Types and Codes

Type Number	Type Name	Code
0	Reserved	0 - no route to destination 1 - communication with destination administratively prohibited 2 - beyond scope of source address 3 - address unreachable 4 - port unreachable 5 - source address failed ingress/egress policy 6 - reject route to destination 7 - Error in Source Routing Header
1	Destination Unreachable	
2	Packet Too Big	
3	Time Exceeded	0 - hop limit exceeded in transit 1 - fragment reassembly time exceeded
4	Parameter Problem	0 - erroneous header field encountered 1 - unrecognized Next Header type encountered 2 - unrecognized IPv6 option encountered
100	Private Experimentation	
101	Private Experimentation	
102 - 126	Unassigned	
127	Reserved for expansion if ICMPv6 error messages	
128	Echo Request	

**Table 78:**ICMPv6 Types and Codes (continued)

Type Number	Type Name	Code
129	Echo Replay	
130	Multicast Listener Query	
131	Multicast Listener Report	
132	Multicast Listener Done	
133	Router Solicitation	
134	Router Advertisement	
135	Neighbor Solicitation	
136	Neighbor Advertisement	
137	Redirect Message	
138	Router Renumbering	0 - Router Renumbering Command 1 - Router Renumbering Result 255 - Sequence Number Reset
139	ICMP Node Information Query	0 - The Data field contains an IPv6 address which is the Subject of this Query. 1 - The Data field contains a name which is the Subject of this Query, or is empty, as in the case of a NOOP. 2 - The Data field contains an IPv4 address which is the Subject of this Query. 140 ICMP Node Information Response 0 - A successful reply. The Reply Data field may or may not be empty. 1 - The Responder refuses to supply the answer. The Reply Data field will be empty. 2 - The Qtype of the Query is unknown to the Responder. The Reply Data field will be empty.
140	ICMP Node Information Response	0 - A successful reply. The Reply Data field may or may not be empty. 1 - The Responder refuses to supply the answer. The Reply Data field will be empty. 2 - The Qtype of the Query is unknown to the Responder. The Reply Data field will be empty.

**Table 78:**ICMPv6 Types and Codes (continued)

Type Number	Type Name	Code
141	Inverse Neighbor Discovery Solicitation Message	
142	Inverse Neighbor Discovery Advertisement Message	
143	Version 2 Multicast Listener Report	
144	Home Agent Address Discovery Request Message	
145	Home Agent Address Discovery Reply Message	
146	Mobile Prefix Solicitation	
147	Mobile Prefix Advertisement	
148	Certification Path Solicitation Message	
149	Certification Path Advertisement Message	
150	ICMP messages utilized by experimental mobility protocols such as Seamoby	
151	Multicast Router Advertisement	
152	Multicast Router Solicitation	
153	Multicast Router Termination	
154	FMIPv6 Messages	
155	RPL Control Message	
156	ILNPv6 Locator Update Message	
157	Duplicate Address Request	

**Table 78:**ICMPv6 Types and Codes (continued)

Type Number	Type Name	Code
158	Duplicate Address Confirmation	
159 – 199	Unassigned	
200	Private experimentation	
201	Private experimentation	
255	Reserved for expansion of ICMPv6 informational messages	

## IPv6 in dynamic routing

Unless otherwise stated, routing protocols apply to IPv4 addressing. This is the standard address format used. However, IPv6 is becoming more popular and new versions of the dynamic routing protocols have been introduced.

As with most advanced routing features on your FortiGate unit, IPv6 settings for dynamic routing protocols must be enabled before they will be visible in the GUI. To enable IPv6 configuration in the GUI, enable it in *System > Config > Features*. Alternatively, you can directly configure IPv6 for RIP, BGP, or OSPF protocols using CLI commands.

## Dual stack routing

Dual stack routing implements dual IP layers in hosts and routers, supporting both IPv6 and IPv4. A dual stack architecture supports both IPv4 and IPv6 traffic and routes the appropriate traffic as required to any device on the network. Administrators can update network components and applications to IPv6 on their own schedule, and even maintain some IPv4 support indefinitely if that is necessary. Devices that are on this type of network, and connect to the Internet, can query Internet DNS servers for both IPv4 and IPv6 addresses. If the Internet site supports IPv6, the device can easily connect using the IPv6 address. If the Internet site does not support IPv6, then the device can connect using the IPv4 addresses.

In FortiOS, dual stack architecture it is not comprised merely of basic addressing functions that operate in both versions of IP. The other features of the appliance, such as UTM and routing, can also use both IP stacks.

If an organization with a mixed network uses an Internet service provider that does not support IPv6, they can use an IPv6 tunnel broker to connect to IPv6 addresses that are on the Internet. FortiOS supports IPv6 tunnelling over IPv4 networks to tunnel brokers. The tunnel broker extracts the IPv6 packets from the tunnel and routes them to their destinations.

## IPv6 tunnelling

IPv6 Tunnelling is the act of tunnelling IPv6 packets from an IPv6 network through an IPv4 network to another IPv6 network. This is different than Network Address Translation (NAT) because, once the packet reaches its final destination, the true originating address of the sender will still be readable. The IPv6 packets are encapsulated within packets with IPv4 headers, which carry their IPv6 payload through the IPv4 network.

The key to IPv6 tunnelling is the ability of the two devices, whether they are a host or a network device, to be dual stack compatible in order to work with both IPv4 and IPv6 at the same time. In the process, the entry node of the tunnel portion of the path will create an encapsulating IPv4 header and transmit the encapsulated packet. The exit node at the end of the tunnel receives the encapsulated packet. The IPv4 header is removed, the IPv6 header is updated, and the IPv6 packet is processed.

There are two types of tunnels in IPv6:

**Automatic tunnels:** Automatic tunnels are configured by using IPv4 address information embedded in an IPv6 address – the IPv6 address of the destination host includes information about which IPv4 address the packet should be tunneled to.

**Configured tunnels:** Configured tunnels must be configured manually. These tunnels are used when using IPv6 addresses that do not have any embedded IPv4 information. The IPv6 and IPv4 addresses of the endpoints of the tunnel must be specified.

### Tunnel configuration

There are a few ways in which the tunnelling can be performed depending on which segment of the path between the end points of the session the encapsulation takes place.

**Host to Host:** Dual Stack capable hosts that are interconnected by an IPv4 infrastructure can tunnel IPv6 packets between themselves. In this case, the tunnel spans the entire path taken by the IPv6 packets.

**Network Device to Host:** Dual Stack capable network devices can tunnel IPv6 packets to their final destination IPv6 or IPv4 host. This tunnel spans only the last segment of the path taken by the IPv6 packets.

Regardless of whether the tunnel starts at a host or a network device, the node that does the encapsulation needs to maintain soft state information, such as the maximum transmission unit (MTU), about each tunnel in order to process the IPv6 packets.

Use the following command to tunnel IPv6 traffic over an IPv4 network. The IPv6 interface is configured under `config system interface`. The command to do the



reverse is `config system ipv6-tunnel`. These commands are not available in Transparent mode.

```
config system sit-tunnel
 edit <tunnel name>
 set destination <tunnel _address>
 set interface <name>
 set ip6 <address_ipv6>
 set source <address_ipv4>
 end
```

Variable	Description	Default
edit <tunnel_name>	Enter a name for the IPv6 tunnel.	No default.
destination <tunnel_address>	The destination IPv4 address for this tunnel.	0.0.0.0
interface <name>	The interface used to send and receive traffic for this tunnel.	No default.
ip6 <address_ipv6>	The IPv6 address for this tunnel.	No default.
source <address_ipv4>	The source IPv4 address for this tunnel.	0.0.0.0

## Tunnelling IPv6 through IPsec VPN

A variation on tunnelling IPv6 through IPv4 is to use an IPsec VPN tunnel between two FortiGate devices. FortiOS supports IPv6 over IPsec. In this sort of scenario, two networks using IPv6 behind FortiGate units are separated by the Internet, which uses IPv4. An IPsec VPN tunnel is created between the FortiGate units and a tunnel is created over the IPv4-based Internet, but the traffic in the tunnel is IPv6. This has the additional advantage of securing the traffic.


For configuration information, see [IPv6 IPsec VPN](#).

## SIP over IPv6

FortiOS supports Sessions Initiate Protocol (SIP) over IPv6. The SIP application-level gateway (ALG) can process SIP messages that use IPv6 addresses in the headers, bodies, and in the transport stack. The SIP ALG cannot modify the IPv6 addresses in the SIP headers so FortiGate units cannot perform SIP or RTP NAT over IPv6 and also cannot translate between IPv6 and IPv4 addresses.

In the scenario shown below, a SIP phone connects to the Internet through a FortiGate unit operating. The phone and the SIP and RTP servers all have IPv6 addresses.

The FortiGate unit has IPv6 security policies that accept SIP sessions. The SIP ALG understands IPv6 addresses and can forward IPv6 sessions to their destinations. Using SIP application control features the SIP ALG can also apply rate limiting and other settings to SIP sessions.



To enable SIP support for IPv6 add an IPv6 security policy that accepts SIP packets and includes a VoIP profile.

## New Fortinet FortiGate IPv6 MIB fields

The following IPv6 MIB fields have been added to the Fortinet FortiGate MIB. These MIB entries can be used to display IPv6 session and policy statistics.

- IPv6 Session Counters:
  - `fgSysSes6Count`
  - `fgSysSes6Rate1`
  - `fgSysSes6Rate10`
  - `fgSysSes6Rate30`
  - `fgSysSes6Rate60`
- IPv6 Policy Statistics:
  - `fgFwPol6StatsTable`
  - `fgFwPol6StatsEntry`
  - `FgFwPol6StatsEntry`
  - `fgFwPol6ID`
  - `fgFwPol6PktCount`
  - `fgFwPol6ByteCount`
- IPv6 Session Statistics:
  - `fgIp6SessStatsTable`
  - `fgIp6SessStatsEntry`
  - `FgIp6SessStatsEntry`
  - `fgIp6SessNumber`

The `fgSysSesCount` and `fgSysSesRateX` MIBs report statistics for IPv4 plus IPv6 sessions combined. This behavior was not changed.

## New OIDs

The following OIDs have been added:

FORTINET-FORTIGATE-MIB:fortinet.fnFortiGateMib.fgSystem.fgSystemInfo

.fgSysSes6Count 1.3.6.1.4.1.12356.101.4.1.15

.fgSysSesRate1 1.3.6.1.4.1.12356.101.4.1.16

.fgSysSesRate10 1.3.6.1.4.1.12356.101.4.1.17

.fgSysSesRate30 1.3.6.1.4.1.12356.101.4.1.18

.fgSysSesRate60 1.3.6.1.4.1.12356.101.4.1.19

FORTINET-FORTIGATE-MIB:fortinet.fnFortiGateMib.fgFirewall.fgFwPolicies.fgFwPolTables

.fgFwPol6StatsTable.fgFwPol6StatsEntry.fgFwPol6ID 1.3.6.1.4.1.12356.101.5.1.2.2.1.1

.fgFwPol6StatsTable.fgFwPol6StatsEntry.fgFwPol6PktCount 1.3.6.1.4.1.12356.101.5.1.2.2.1.2

.fgFwPol6StatsTable.fgFwPol6StatsEntry.fgFwPol6ByteCount 1.3.6.1.4.1.12356.101.5.1.2.2.1.3

FORTINET-FORTIGATE-MIB:fortinet.fnFortiGateMib.fgInetProto.fgInetProtoTables

.fgIp6SessStatsTable.fgIp6SessStatsEntry.fgIp6SessNumber 1.3.6.1.4.1.12356.101.11.2.3.1.1

## EXAMPLE SNMP get/walk output

```
// Session6 stats excerpt from sysinfo:
snmpwalk -v2c -cpublic 192.168.1.111 1.3.6.1.4.1.12356.101.4
FORTINET-FORTIGATE-MIB::fgSysSes6Count.0 = Gauge32: 203
FORTINET-FORTIGATE-MIB::fgSysSes6Rate1.0 = Gauge32: 10 Sessions
 Per Second
FORTINET-FORTIGATE-MIB::fgSysSes6Rate10.0 = Gauge32: 2 Sessions
 Per Second
FORTINET-FORTIGATE-MIB::fgSysSes6Rate30.0 = Gauge32: 1 Sessions
 Per Second
FORTINET-FORTIGATE-MIB::fgSysSes6Rate60.0 = Gauge32: 0 Sessions
 Per Second

// FwPolicy6 table:
snmpwalk -v2c -cpublic 192.168.1.111 1.3.6.1.4.1.12356.101.5.1.2.2
FORTINET-FORTIGATE-MIB::fgFwPol6ID.1.3 = INTEGER: 3
FORTINET-FORTIGATE-MIB::fgFwPol6ID.1.4 = INTEGER: 4
FORTINET-FORTIGATE-MIB::fgFwPol6PktCount.1.3 = Counter64: 4329
FORTINET-FORTIGATE-MIB::fgFwPol6PktCount.1.4 = Counter64: 0
FORTINET-FORTIGATE-MIB::fgFwPol6ByteCount.1.3 = Counter64: 317776
FORTINET-FORTIGATE-MIB::fgFwPol6ByteCount.1.4 = Counter64: 0

// IP6SessNumber:
snmpwalk -v2c -cpublic 192.168.1.111
 1.3.6.1.4.1.12356.101.11.2.3.1
FORTINET-FORTIGATE-MIB::fgIp6SessNumber.1 = Counter32: 89
```

## IPv6 Per-IP traffic shaper

You can add any Per-IP traffic shaper to an IPv6 security policy using the following command:

```
config firewall policy6
 edit 0
 set per-ip-shaper "new-perip-shaper"
 end
```

## DHCPv6

You can use DHCP with IPv6 using the CLI. Use the CLI command

```
config system dhcp6.
```

For more information on the configuration options, see the [CLI Reference](#).

## DHCPv6 relay

You can use the following command to configure a FortiGate interface to relay DHCPv6 queries and responses from one network to a network with a DHCPv6 server and back.

The command enables DHCPv6 relay and includes adding the IPv6 address of the DHCP server that the FortiGate unit relays DHCPv6 requests to:

```
config system interface
 edit internal
 config ipv6
 set dhcp6-relay-service enable
 set dhcp6-relay-type regular
 set dhcp6-relay-ip 2001:db8:0:2::30
 end
 end
```

## IPv6 forwarding—Policies, IPS, Application Control, flow-based antivirus, web filtering, and DLP

FortiOS 5.0 fully supports flow-based inspection of IPv6 traffic. This includes full support for IPS, application control, virus scanning, and web filtering.

To add flow-based inspection to IPv6 traffic go to *Policy > Policy > IPv6 Policy* and select *Create New* to add an IPv6 Security Policy. Configure the policy to accept the traffic to be scanned. Select *Security Profiles* and select the profiles to apply to the traffic.

## FortiGate interfaces can get IPv6 addresses from an IPv6 DHCP server

From the CLI, you can configure any FortiGate interface to get an IPv6 address from an IPv6 DHCP server. For example, to configure the wan2 interface to get an IPv6 address from an IPv6 DHCP server enter the following command:

```
config system interface
 edit wan2
 config ipv6
 set ip6-mode dhcp
 end
 end
```



# IPv6 Configuration

This section contains configuration information for IPv6 on FortiOS. Attempts are made to include scenarios in each section to better assist with the configuration and to orient the information toward a particular task.

You will find information on the following:

- [IPv6 address groups](#)
- [IPv6 firewall addresses](#)
- [ICMPv6](#)
- [IPv6 IPsec VPN](#)
- [BGP and IPv6](#)
- [RIPng — RIP and IPv6](#)
- [IPv6 IPS](#)
- [Blocking IPv6 packets by extension headers](#)
- [IPv6 Denial of Service policies](#)
- [Configure hosts in an SNMP v1/2c community to send queries or receive traps](#)
- [IPv6 PIM sparse mode multicast routing](#)

## IPv6 address groups

### To create IPv6 address groups from existing IPv6 addresses - web-based manager

Your company has 3 internal servers with IPv6 addresses that it would like to group together for the purposes of a number of policies.

The preconfigured addresses to use will consist of:

- Web\_Server-1
- Web\_Server-2
- Web\_Server-3

Go to *Firewall Objects > Address > Groups* and select *Create New > IPv6 Address Group*.

Fill out the fields with the following information

---

<b>Group Name</b>	Web_Server_Cluster
<b>Members</b>	Web_Server-1
	Web_Server-2
	Web_Server-3

---

Select *OK*.

### To create IPv6 address groups from existing IPv6 addresses - CLI

```
config firewall addrgrp6
 edit Web_Server_Cluster
 set member Web_Server-1 Web_Server-2 Web_Server-3
 end
```

### To verify that the addresses were added correctly:

Go to *Firewall Objects > Address > Groups*. Check that the addresses have been added to the address list and that they are correct.

From the CLI, enter the following commands:

```
config firewall addrgrp6
 edit <the name of the address that you wish to verify>
 Show full-configuration
```

## IPv6 firewall addresses

### Scenario: Mail Server

You need to create an IPv6 address for the Mail Server on Port1 of your internal network.

- These server is on the network off of port1.
- The IP address is 2001:db8:0:2::20/64
- There should be a tag for this address being for a server

Go to *Firewall Objects > Address > Addresses* and select *Create New > IPv6 Address*.

Fill out the fields with the following information

<b>Address Name</b>	Mail_Server
<b>IPv6 Address</b>	2001:db8:0:2::20/64
<b>Add Tags</b>	Server

Select *OK*.

Enter the following CLI command:

```
config firewall address6
 edit Mail_Server
 set type ipmask
 set subnet 2001:db8:0:2::20/64
 set associated-interface port1
 set tags Server
 end
```

### Scenario: First Floor Network

You need to create an IPv4 address for the subnet of the internal network off of Port1.

- These computers connect to port1.
- The network uses the IPv6 addresses: fdde:5a7d:f40b:2e9d:xxxx:xxxx:xxxx:xxxx
- There should be a reference to this being the network for the 1st floor of the building.



Go to *Firewall Objects > Address > Addresses* and select *Create New > IPv6 Address*.

Fill out the fields with the following information

Field Name	Field Value
Address Name	Internal_Subnet_1
Comments	Network for 1st Floor
Type	Subnet / IP Range
Subnet / IP Range	2001:db8:0:2::/64

Select *OK*.

Enter the following CLI command:

```
config firewall address6
 edit Internal_Subnet_1
 set comment "Network for 1st Floor"
 set subnet 2001:db8:0:2::/64
 end
```

**To verify that the addresses were added correctly:**

Go to *Firewall Objects > Address > Addresses*. Check that the addresses have been added to the address list and that they are correct.

Enter the following CLI command:

```
config firewall address6
 edit <the name of the address that you wish to verify>
 Show full-configuration
```

## ICMPv6

The IT Manager is doing some diagnostics and would like to temporarily block the successful replies of ICMP Node information Responses between 2 IPv6 networks.

The ICMP type for ICMP Node informations responses is 140. The codes for a successful response is 0.

**Web-based Manager Instructions**

Go to *Firewall Objects > Service > Services* and select *Create New > Custom Service*.

Fill out the fields with the following information

Field Name	Field Value
Name	diagnostic-test1
Service Type	Firewall
Show in Service List	Enabled
Category	Uncategorized
Protocol Type	ICMP6

<b>Type</b>	140
<b>Code</b>	0

Select *OK*.

Enter the following CLI command:

```
config firewall service custom
edit diagnostic-test1
set protocol ICMP6
set icmptype 140
set icmpcode 0
set visibility enable
end
```

**To verify that the category was added correctly:**

Go to *Firewall Objects > Service > Services*. Check that the services have been added to the services list and that they are correct.

Enter the following CLI command:

```
config firewall service custom
edit <the name of the service that you wish to verify>
show full-configuration
```

## IPv6 IPsec VPN

This chapter describes how to configure your FortiGate unit's IPv6 IPsec VPN functionality. The following topics are included in this section:

- [Overview of IPv6 IPsec support](#)
- [Configuring IPv6 IPsec VPNs](#)
- [Site-to-site IPv6 over IPv6 VPN example](#)
- [Site-to-site IPv4 over IPv6 VPN example](#)
- [Site-to-site IPv6 over IPv4 VPN example](#)

### Overview of IPv6 IPsec support

FortiOS supports route-based IPv6 IPsec, but not policy-based. This section describes how IPv6 IPsec support differs from IPv4 IPsec support.

Where both the gateways and the protected networks use IPv6 addresses, sometimes called IPv6 over IPv6, you can create either an auto-keyed or manually-keyed VPN. You can also combine IPv6 and IPv4 addressing in an auto-keyed VPN in the following ways:

<b>IPv4 over IPv6</b>	The VPN gateways have IPv6 addresses.  The protected networks have IPv4 addresses. The phase 2 configurations at either end use IPv4 selectors.
<b>IPv6 over IPv4</b>	The VPN gateways have IPv4 addresses.  The protected networks use IPv6 addresses. The phase 2 configurations at either end use IPv6 selectors.

Compared with IPv4 IPsec VPN functionality, there are some limitations:

- Except for IPv6 over IPv4, remote gateways with Dynamic DNS are not supported.
- Selectors cannot be firewall address names. Only IP address, address range and subnet are supported.
- Redundant IPv6 tunnels are not supported.

## Certificates

On a VPN with IPv6 phase 1 configuration, you can authenticate using VPN certificates in which the common name (cn) is an IPv6 address. The `cn-type` keyword of the `user peer` command has an option, `ipv6`, to support this.

## Configuring IPv6 IPsec VPNs

Configuration of an IPv6 IPsec VPN follows the same sequence as for an IPv4 route-based VPN: phase 1 settings, phase 2 settings, security policies, and routing.

### Phase 1 configuration

In the web-based manager, you define the Phase 1 as IPv6 in the Advanced settings. Enable the IPv6 Version check box. You can then enter an IPv6 address for the remote gateway.

In the CLI, you define an IPsec phase 1 configuration as IPv6 by setting `ip-version` to 6. Its default value is 4. Then, the `local-gw` and `remote-gw` keywords are hidden and the corresponding `local-gw6` and `remote-gw6` keywords are available. The values for `local-gw6` and `remote-gw6` must be IPv6 addresses. For example:

```
config vpn ipsec phase1-interface
 edit tunnel6
 set ip-version 6
 set remote-gw6 0:123:4567::1234
 set interface port3
 set proposal 3des-md5
 end
```

### Phase 2 configuration

To create an IPv6 IPsec phase 2 configuration in the web-based manager, you need to define IPv6 selectors in the Advanced settings. Change the default “0.0.0.0/0” address for Source address and Destination address to the IPv6 value “::/0”. If needed, enter specific IPv6 addresses, address ranges or subnet addresses in these fields.

In the CLI, set `src-addr-type` and `dst-addr-type` to `ip6`, `range6` or `subnet6` to specify IPv6 selectors. By default, zero selectors are entered, “::/0” for the `subnet6` address type, for example. The simplest IPv6 phase 2 configuration looks like this:

```
config vpn ipsec phase2-interface
 edit tunnel6_p2
 set phase1name tunnel6
 set proposal 3des-md5
 set src-addr-type subnet6
 set dst-addr-type subnet6
 end
```

## Security policies

To complete the VPN configuration, you need a security policy in each direction to permit traffic between the protected network's port and the IPsec interface. You need IPv6 policies unless the VPN is IPv4 over IPv6.

## Routing

Appropriate routing is needed for both the IPsec packets and the encapsulated traffic within them. You need a route, which could be the default route, to the remote VPN gateway via the appropriate interface. You also need a route to the remote protected network via the IPsec interface.

### To create a static route in the web-based manager:

1. Go to *Router > Static > Static Routes*.  
On some desktop FortiGate models, go to *System > Network > Routing*.
2. Select the drop-down arrow on the *Create New* button and select *IPv6 Route*.
3. Enter the information and select *OK*.

In the CLI, use the `router static6` command. For example, where the remote network is `fec0:0000:0000:0004::/64` and the IPsec interface is `toB`:

```
config router static6
 edit 1
 set device port2
 set dst 0::/0
 next
 edit 2
 set device toB
 set dst fec0:0000:0000:0004::/64
 next
end
```


If the VPN is IPv4 over IPv6, the route to the remote protected network is an IPv4 route. If the VPN is IPv6 over IPv4, the route to the remote VPN gateway is an IPv4 route.

## Site-to-site IPv6 over IPv6 VPN example

In this example, computers on IPv6-addressed private networks communicate securely over public IPv6 infrastructure.

To access IPv6 functionality through the web-based manager, go to *System Admin > Settings* and enable *IPv6* in the section, *Display Options on GUI*.

**Figure 285:**Example IPv6-over-IPv6 VPN topology



### Configure FortiGate A interfaces

Port 2 connects to the public network and port 3 connects to the local network.

```
config system interface
 edit port2
 config ipv6
 set ip6-address fec0::0001:209:0fff:fe83:25f2/64
 end
 next
 edit port3
 config ipv6
 set ip6-address fec0::0000:209:0fff:fe83:25f3/64
 end
 next
end
```

### Configure FortiGate A IPsec settings

The phase 1 configuration creates a virtual IPsec interface on port 2 and sets the remote gateway to the public IP address FortiGate B. This configuration is the same as for an IPv4 route-based VPN, except that `ip-version` is set to 6 and the `remote-gw6` keyword is used to specify an IPv6 remote gateway address.

```
config vpn ipsec phase1-interface
 edit toB
 set ip-version 6
 set interface port2
 set remote-gw6 fec0:0000:0000:0003:209:0fff:fe83:25c7
 set dpd enable
 set psksecret maryhadalittlelamb
 set proposal 3des-md5 3des-sha1
 end
```

By default, phase 2 selectors are set to accept all subnet addresses for source and destination. The default setting for `src-addr-type` and `dst-addr-type` is `subnet`. The IPv6 equivalent is `subnet6`. The default subnet addresses are `0.0.0.0/0` for IPv4, `::/0` for IPv6.

```
config vpn ipsec phase2-interface
 edit toB
 set phase1name toB
 set proposal 3des-md5 3des-sha1
 set pfs enable
 set replay enable
 set src-addr-type subnet6
 set dst-addr-type subnet6
 end
```

## Configure FortiGate A security policies

Security policies are required to allow traffic between `port3` and the IPsec interface `toB` in each direction. The address `all6` must be defined using the `firewall address6` command as `::/0`.

```
config firewall policy6
 edit 1
 set srcintf port3
 set dstintf toB
 set srcaddr all6
 set dstaddr all6
 set action accept
 set service ANY
 set schedule always
 next
 edit 2
 set srcintf toB
 set dstintf port3
 set srcaddr all6
 set dstaddr all6
 set action accept
 set service ANY
 set schedule always
 end
```

## Configure FortiGate A routing

This simple example requires just two static routes. Traffic to the protected network behind FortiGate B is routed via the virtual IPsec interface toB. A default route sends all IPv6 traffic out on port2.

```
config router static6
 edit 1
 set device port2
 set dst 0::/0
 next
 edit 2
 set device toB
 set dst fec0:0000:0000:0004::/64
end
```

## Configure FortiGate B

The configuration of FortiGate B is very similar to that of FortiGate A. A virtual IPsec interface toA is configured on port2 and its remote gateway is the public IP address of FortiGate A. Security policies enable traffic to pass between the private network and the IPsec interface. Routing ensures traffic for the private network behind FortiGate A goes through the VPN and that all IPv6 packets are routed to the public network.

```
config system interface
 edit port2
 config ipv6
 set ip6-address fec0::0003:209:0fff:fe83:25c7/64
 end
 next
 edit port3
 config ipv6
 set ip6-address fec0::0004:209:0fff:fe83:2569/64
 end
 end
config vpn ipsec phase1-interface
 edit toA
 set ip-version 6
 set interface port2
 set remote-gw6 fec0:0000:0000:0001:209:0fff:fe83:25f2
 set dpd enable
 set psksecret maryhadalittlelamb
 set proposal 3des-md5 3des-sha1
 end
config vpn ipsec phase2-interface
 edit toA2
 set phase1name toA
 set proposal 3des-md5 3des-sha1
 set pfs enable
 set replay enable
 set src-addr-type subnet6
 set dst-addr-type subnet6
```

```


end
config firewall policy6
 edit 1
 set srcintf port3
 set dstintf toA
 set srcaddr all6
 set dstaddr all6
 set action accept
 set service ANY
 set schedule always
 next
 edit 2
 set srcintf toA
 set dstintf port3
 set srcaddr all6
 set dstaddr all6
 set action accept
 set service ANY
 set schedule always
 end
config router static6
 edit 1
 set device port2
 set dst 0::/0
 next
 edit 2
 set device toA
 set dst fec0:0000:0000:0000::/64
 end
end

```

## Site-to-site IPv4 over IPv6 VPN example

In this example, two private networks with IPv4 addressing communicate securely over IPv6 infrastructure.

**Figure 286:**Example IPv4-over-IPv6 VPN topology





## Configure FortiGate A interfaces

Port 2 connects to the IPv6 public network and port 3 connects to the IPv4 LAN.

```
config system interface
 edit port2
 config ipv6
 set ip6-address fec0::0001:209:0fff:fe83:25f2/64
 end
 next
 edit port3
 set 192.168.2.1/24
 end
```

## Configure FortiGate A IPsec settings

The phase 1 configuration is the same as in the IPv6 over IPv6 example.

```
config vpn ipsec phase1-interface
 edit toB
 set ip-version 6
 set interface port2
 set remote-gw6 fec0:0000:0000:0003:209:0fff:fe83:25c7
 set dpd enable
 set psksecret maryhadalittlelamb
 set proposal 3des-md5 3des-sha1
 end
```

The phase 2 configuration is the same as you would use for an IPv4 VPN. By default, phase 2 selectors are set to accept all subnet addresses for source and destination.

```
config vpn ipsec phase2-interface
 edit toB2
 set phase1name toB
 set proposal 3des-md5 3des-sha1
 set pfs enable
 set replay enable
 end
```

## Configure FortiGate A security policies

Security policies are required to allow traffic between port3 and the IPsec interface toB in each direction. These are IPv4 security policies.

```
config firewall policy
 edit 1
 set srcintf port3
 set dstintf toB
 set srcaddr all
 set dstaddr all
 set action accept
 set service ANY
 set schedule always
 next
```

```

edit 2
 set srcintf toB
 set dstintf port3
 set srcaddr all
 set dstaddr all
 set action accept
 set service ANY
 set schedule always
end

```

## Configure FortiGate A routing

This simple example requires just two static routes. Traffic to the protected network behind FortiGate B is routed via the virtual IPsec interface toB using an IPv4 static route. A default route sends all IPv6 traffic, including the IPv6 IPsec packets, out on port2.

```

config router static6
 edit 1
 set device port2
 set dst 0::/0
 next
 edit 2
 set device toB
 set dst 192.168.3.0/24
 end
end

```

## Configure FortiGate B

The configuration of FortiGate B is very similar to that of FortiGate A. A virtual IPsec interface toA is configured on port2 and its remote gateway is the public IP address of FortiGate A. The IPsec phase 2 configuration has IPv4 selectors.

IPv4 security policies enable traffic to pass between the private network and the IPsec interface. An IPv4 static route ensures traffic for the private network behind FortiGate A goes through the VPN and an IPv6 static route ensures that all IPv6 packets are routed to the public network.

```

config system interface
 edit port2
 config ipv6
 set ip6-address fec0::0003:fe83:25c7/64
 end
 next
 edit port3
 set 192.168.3.1/24
 end
config vpn ipsec phase1-interface
 edit toA
 set ip-version 6
 set interface port2
 set remote-gw6 fec0:0000:0000:0001:209:0fff:fe83:25f2
 set dpd enable
 set psksecret maryhadalittlelamb
 set proposal 3des-md5 3des-sha1
 end
end


```

```
end
config vpn ipsec phase2-interface
edit toA2
 set phase1name toA
 set proposal 3des-md5 3des-sha1
 set pfs enable
 set replay enable
end
config firewall policy
edit 1
 set srcintf port3
 set dstintf toA
 set srcaddr all
 set dstaddr all
 set action accept
 set service ANY
 set schedule always
next
edit 2
 set srcintf toA
 set dstintf port3
 set srcaddr all
 set dstaddr all
 set action accept
 set service ANY
 set schedule always
end
config router static6
edit 1
 set device port2
 set dst 0::/0
next
edit 2
 set device toA
 set dst 192.168.2.0/24
end
```

## Site-to-site IPv6 over IPv4 VPN example

In this example, IPv6-addressed private networks communicate securely over IPv4 public infrastructure.

**Figure 287:**Example IPv6-over-IPv4 VPN topology



### Configure FortiGate A interfaces

Port 2 connects to the IPv4 public network and port 3 connects to the IPv6 LAN.

```
config system interface
 edit port2
 set 10.0.0.1/24
 next
 edit port3
 config ipv6
 set ip6-address fec0::0001:209:0fff:fe83:25f3/64
 end
 end
```

### Configure FortiGate A IPsec settings

The phase 1 configuration uses IPv4 addressing.

```
config vpn ipsec phase1-interface
 edit toB
 set interface port2
 set remote-gw 10.0.1.1
 set dpd enable
 set psksecret maryhadalittlelamb
 set proposal 3des-md5 3des-sha1
 end
```

The phase 2 configuration uses IPv6 selectors. By default, phase 2 selectors are set to accept all subnet addresses for source and destination. The default setting for `src-addr-type` and `dst-addr-type` is `subnet`. The IPv6 equivalent is `subnet6`. The default subnet addresses are `0.0.0.0/0` for IPv4, `::/0` for IPv6.

```
config vpn ipsec phase2-interface
 edit toB2
 set phasename toB
 set proposal 3des-md5 3des-sha1
```

```
 set pfs enable
 set replay enable
 set src-addr-type subnet6
 set dst-addr-type subnet6
end
```

## Configure FortiGate A security policies

IPv6 security policies are required to allow traffic between port3 and the IPsec interface toB in each direction. Define the address all6 using the firewall address6 command as ::/0.

```
config firewall policy6
 edit 1
 set srcintf port3
 set dstintf toB
 set srcaddr all6
 set dstaddr all6
 set action accept
 set service ANY
 set schedule always
 next
 edit 2
 set srcintf toB
 set dstintf port3
 set srcaddr all6
 set dstaddr all6
 set action accept
 set service ANY
 set schedule always
end
```

## Configure FortiGate A routing

This simple example requires just two static routes. Traffic to the protected network behind FortiGate B is routed via the virtual IPsec interface toB using an IPv6 static route. A default route sends all IPv4 traffic, including the IPv4 IPsec packets, out on port2.

```
config router static6
 edit 1
 set device toB
 set dst fec0:0000:0000:0004::/64
 end
config router static
 edit 1
 set device port2
 set dst 0.0.0.0/0
 set gateway 10.0.0.254
 end
```

## Configure FortiGate B

The configuration of FortiGate B is very similar to that of FortiGate A. A virtual IPsec interface toA is configured on port2 and its remote gateway is the IPv4 public IP address of FortiGate A. The IPsec phase 2 configuration has IPv6 selectors.

IPv6 security policies enable traffic to pass between the private network and the IPsec interface. An IPv6 static route ensures traffic for the private network behind FortiGate A goes through the VPN and an IPv4 static route ensures that all IPv4 packets are routed to the public network.

```
config system interface
 edit port2
 set 10.0.1.1/24
 next
 edit port3
 config ipv6
 set ip6-address fec0::0004:209:0fff:fe83:2569/64
 end
config vpn ipsec phase1-interface
 edit toA
 set interface port2
 set remote-gw 10.0.0.1
 set dpd enable
 set psksecret maryhadalittlelamb
 set proposal 3des-md5 3des-sha1
 end
config vpn ipsec phase2-interface
 edit toA2
 set phasename toA
 set proposal 3des-md5 3des-sha1
 set pfs enable
 set replay enable
 set src-addr-type subnet6
 set dst-addr-type subnet6
 end
config firewall policy6
 edit 1
 set srcintf port3
 set dstintf toA
 set srcaddr all6
 set dstaddr all6
 set action accept
 set service ANY
 set schedule always
 next
 edit 2
 set srcintf toA
 set dstintf port3
 set srcaddr all6
 set dstaddr all6
 set action accept
 set service ANY
 set schedule always
 end
config router static6
 edit 1
```

```

 set device toA
 set dst fec0:0000:0000:0000::/64
end
config router static
edit 1
 set device port2
 set gateway 10.0.1.254
end

```

## BGP and IPv6

FortiGate units support IPv6 over BGP using the same `config router bgp` command as IPv4, but different subcommands.

The main CLI keywords have IPv6 equivalents that are identified by the “6” on the end of the keyword, such as with `config network6` or `set allowas-in6`.

IPv6 BGP commands include:

```

config router bgp
 set activate6 {enable | disable}
 set allowas-in6 <max_num_AS_integer>
 set allowas-in-enable6 {enable | disable}
 set as-override6 {enable | disable}
 set attribute-unchanged6 [as-path] [med] [next-hop]
 set capability-default-originate6 {enable | disable}
 set capability-graceful-restart6 {enable | disable}
 set default-originate-route-map6 <routemap_str>
 set distribute-list-in6 <access-list-name_str>
 set distribute-list-out6 <access-list-name_str>
 set filter-list-in6 <aspath-list-name_str>
 set filter-list-out6 <aspath-list-name_str>
 set maximum-prefix6 <prefix_integer>
 set maximum-prefix-threshold6 <percentage_integer>
 set maximum-prefix-warning-only6 {enable | disable}
 set next-hop-self6 {enable | disable}
 set prefix-list-in6 <prefix-list-name_str>
 set prefix-list-out6 <prefix-list-name_str>
 set remove-private-as6 {enable | disable}
 set route-map-in6 <routemap-name_str>
 set route-map-out6 <routemap-name_str>
 set route-reflector-client6 {enable | disable}
 set route-server-client6 {enable | disable}
 set send-community6 {both | disable | extended | standard}
 set soft-reconfiguration6 {enable | disable}
 set unsuppress-map6 <route-map-name_str>
 config network6
 config redistribute6
end

```

## RIPng – RIP and IPv6

RIP next generation, or RIPng, is the version of RIP that supports IPv6.

This is an example of a typical small network configuration using RIPng routing.

Your internal R&D network is working on a project for a large international telecom company that uses IPv6. For this reason, you have to run IPv6 on your internal network and you have decided to use only IPv6 addresses.

Your network has two FortiGate units running the RIPng dynamic routing protocol. Both FortiGate units are connected to the ISP router and the internal network. This configuration provides some redundancy for the R&D internal network enabling it to reach the internet at all times.

This section includes the following topics:

- [Network layout and assumptions](#)
- [Configuring the FortiGate units system information](#)
- [Configuring RIPng on FortiGate units](#)
- [Configuring other network devices](#)
- [Testing the configuration](#)
- [Debugging IPv6 on RIPng](#)

### Network layout and assumptions

#### Basic network layout

Your internal R&D network is working on a project for a large international telecom company that uses IPv6. For this reason, you have to run IPv6 on your internal network and you have decided to use only IPv6 addresses.

Your network has two FortiGate units running the RIPng dynamic routing protocol. Both FortiGate units are connected to the ISP router and the internal network. This configuration provides some redundancy for the R&D internal network enabling it to reach the internet at all times.

All internal computers use RIP routing, so no static routing is required. And all internal computers use IPv6 addresses.

Where possible in this example, the default values will be used or the most general settings. This is intended to provide an easier configuration that will require less troubleshooting.


In this example the routers, networks, interfaces used, and IP addresses are as follows.

**Table 76:** Rip example network topology

Network	Router	Interface & Alias	IPv6 address
R&D	Router1	port1 (internal)	2002:A0B:6565:0:0:0:0:0
		port2 (ISP)	2002:AC14:7865:0:0:0:0:0
	Router2	port1 (internal)	2002:A0B:6566:0:0:0:0:0
		port2 (ISP)	2002:AC14:7866:0:0:0:0:0



**Figure 288:**Network topology for the IPV6 RIPng example



### Assumptions

The following assumptions have been made concerning this example.

- All FortiGate units have 5.0 firmware, and are running factory default settings.
- All CLI and web-based manager navigation assumes the unit is running in NAT/Route operating mode, with VDOMs disabled.
- All FortiGate units have interfaces labelled port1 and port2 as required.
- All firewalls have been configured for each FortiGate unit to allow the required traffic to flow across interfaces.
- All network devices are support IPv6 and are running RIPng.

### Configuring the FortiGate units system information

Each FortiGate unit needs IPv6 enabled, a new hostname, and interfaces configured.

#### To configure system information on Router1 - web-based manager

1. Go to *System > Dashboard > Status*.
2. For *Host name*, select *Change*.
3. Enter "Router1".
4. Enable *IPv6* on the GUI.
5. Go to *System > Network > Interfaces*.
6. Edit port1 (internal) interface.
7. Set the following information, and select *OK*.

<b>Alias</b>	internal
<b>IP/Netmask</b>	2002:A0B:6565::/0

<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Internal RnD network
<b>Administrative Status</b>	Up

8. Edit port2 (ISP) interface.
9. Set the following information, and select *OK*.

<b>Alias</b>	ISP
<b>IP/Netmask</b>	2002:AC14:7865::/0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	ISP and internet
<b>Administrative Status</b>	Up

### To configure system information on Router1 - CLI

```

config system global
 set hostname Router1
 set gui-ipv6 enable
end
config system interface
 edit port1
 set alias internal
 set allowaccess https ping ssh
 set description "Internal RnD network"
 config ipv6
 set ip6-address 2002:a0b:6565::/0
 end
 next
 edit port2
 set alias ISP
 set allowaccess https ping ssh
 set description "ISP and internet"
 config ipv6
 set ip6-address 2002:AC14:7865::
 end
 end
end

```

### To configure system information on Router2 - web-based manager

1. Go to *System > Dashboard > Status*.
2. For *Host name*, select *Change*.
3. Enter "Router2".
4. Enable *IPv6* on the GUI.
5. Go to *System > Network > Interfaces*.
6. Edit port1 (internal) interface.

7. Set the following information, and select *OK*.

<b>Alias</b>	internal
<b>IP/Netmask</b>	2002:A0B:6566::/0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	Internal RnD network
<b>Administrative Status</b>	Up

8. Edit port2 (ISP) interface.
9. Set the following information, and select *OK*.

<b>Alias</b>	ISP
<b>IP/Netmask</b>	2002:AC14:7866::/0
<b>Administrative Access</b>	HTTPS SSH PING
<b>Description</b>	ISP and internet
<b>Administrative Status</b>	Up

#### To configure system information on Router2 - CLI

```
config system global
 set hostname Router2
 set gui-ipv6 enable
end
config system interface
 edit port1
 set alias internal
 set allowaccess https ping ssh
 set description "Internal RnD network"
 config ipv6
 set ip6-address 2002:a0b:6566::/0
 end
 next
 edit port2
 set alias ISP
 set allowaccess https ping ssh
 set description "ISP and internet"
 config ipv6
 set ip6-address 2002:AC14:7866::
 end
 end
end
```

## Configuring RIPng on FortiGate units

Now that the interfaces are configured, you can configure RIPng on the FortiGate units.

There are only two networks and two interfaces to include — the internal network, and the ISP network. There is no redistribution, and no authentication. In RIPng there is no specific command to include a subnet in the RIP broadcasts. There is also no information required for the interfaces beyond including their name.

As this is a CLI only configuration, configure the ISP router and the other FortiGate unit as neighbors. This was not part of the previous example as this feature is not offered in the web-based manager. Declaring neighbors in the configuration like this will reduce the discovery traffic when the routers start up.

Since RIPng is not supported in the web-based manager, this section will only be entered in the CLI.

### To configure RIPng on Router1 - CLI

```
config router ripng
 config interface
 edit port1
 next
 edit port2
 end
 config neighbor
 edit 1
 set interface port1
 set ipv6 2002:a0b:6566::/0
 next
 edit 2
 set interface port2
 set ipv6 2002:AC14:7805::/0
 end
```

### To configure RIPng on Router2 - CLI

```
config router ripng
 config interface
 edit port1
 next
 edit port2
 end
 config neighbor
 edit 1
 set interface port1
 set ipv6 2002:a0b:6565::/0
 next
 edit 2
 set interface port2
 set ipv6 2002:AC14:7805::/0
 end
```

## Configuring other network devices

The other devices on the internal network all support IPv6, and are running RIPng where applicable. They only need to know the internal interface network addresses of the FortiGate units.

The ISP routers need to know the FortiGate unit information such as IPv6 addresses.

## Testing the configuration

In addition to normal testing of your network configuration, you must also test the IPv6 part of this example.

For troubleshooting problems with your network, see the [Troubleshooting](#) chapter.

### Testing the IPv6 RIPng information

There are some commands to use when checking that your RIPng information is correct on your network. These are useful to check on your RIPng FortiGate units on your network. Comparing the output between devices will help you understand your network better, and also track down any problems.

```
diagnose ipv6 address list
```

View the local scope IPv6 addresses used as next-hops by RIPng on the FortiGate unit.

```
diagnose ipv6 route list
```

View ipv6 addresses that are installed in the routing table.

```
get router info6 routing-table
```

View the routing table. This information is almost the same as the previous command (diagnose ipv6 route list) however it is presented in an easier to read format.

```
get router info6 rip interface external
```

View brief output on the RIP information for the interface listed. The information includes if the interface is up or down, what routing protocol is being used, and whether passive interface or split horizon are enabled.

```
get router info6 neighbor-cache list
```

View the IPv6/MAC address mapping. This also displays the interface index and name associated with the address.

## Debugging IPv6 on RIPng

The debug commands are very useful to see what is happening on the network at the packet level. There are a few changes to debugging the packet flow when debugging IPv6.

The following CLI commands specify both IPv6 and RIP, so only RIPng packets will be reported. The output from these commands will show you the RIPng traffic on your FortiGate unit including RECV, SEND, and UPDATE actions.

The addresses are in IPv6 format.

```
diagnose debug enable
diagnose ipv6 router rip level info
diagnose ipv6 router rip all enable
```

These three commands will:

- turn on debugging in general
- set the debug level to information, a verbose reporting level
- turn on all rip router settings

Part of the information displayed from the debugging is the metric (hop count). If the metric is 16, then that destination is unreachable since the maximum hop count is 15.

In general, you should see an update announcement, followed by the routing table being sent out, and a received reply in response.

## IPv6 IPS

IPv6 IPS signature scan can be enabled by interface policy. The user can create an normal IPS sensor and assign it to the IPv6 interface policy.

```
config firewall interface-policy6
 edit 1
 set interface "port1"
 set srcaddr6 "all"
 set dstaddr6 "all"
 set service6 "ANY"
 set ips-sensor-status enable
 set ips-sensor "all_default"
 next
end
```

## Blocking IPv6 packets by extension headers

FortiOS can now block IPv6 packets based on the extension headers, using the CLI syntax

```
config firewall ipv6-eh-filter.
```

The following commands are now available:

```
set hop-opt {disable | enable}: Block packets with Hop-by-Hop Options header.
set dest-opt {disable | enable}: Block packets with Destination Options header.
set hdopt-type <integer>: Block specific Hop-by-Hop and/or Destination Option types
(maximum 7 types, each between 0 and 255).
set routing {disable | enable}: Block packets with Routing header.
set routing-type <integar>: Block specific Routing header types (maximum 7 types,
each between 0 and 255).
set fragment {disable | enable}: Block packets with Fragment header.
set auth {disable | enable}: Block packets with Authentication header.
set no-next {disable | enable}: Block packets with No Next header.
```

## IPv6 Denial of Service policies

Denial of Service (DoS) policies can now be configured by going to *Policy > Policy > IPv6 Dos Policy*.

## Configure hosts in an SNMP v1/2c community to send queries or receive traps

When you add a host to an SNMP v1/2c community you can now decide whether the FortiGate unit will accept queries from the host or whether the FortiGate unit will send traps to the host.

You can also configure the host for both traps and queries. You can add up to 16 IPv4 hosts and up to 16 IPv6 hosts.

- An IPv4 host that can send queries to the FortiGate unit
- An IPv6 host that the FortiGate unit will send traps to

Use the following command to add two hosts to an SNMP community:

```
config system snmp community
 config hosts
 edit 1
 set interface port1
 set ip 172.20.120.1
 set host-type query
 end
 config hosts6
 edit 1
 set interface port6
 set ip 2001:db8:0:2::30
 set host-type trap end
```

## IPv6 PIM sparse mode multicast routing

FortiOS supports PIM sparse mode multicast routing for IPv6 multicast (multicast6) traffic and is compliant with [RFC 4601](#). You can use the following command to configure IPv6 PIM sparse multicast routing.

```
config router multicast6
 set multicast-routing {enable | disable}
 config interface
 edit <interface-name>
 set hello-interval <1-65535 seconds>
 set hello-holdtime <1-65535 seconds>
 end
 config pim-sm-global
 config rp-address
 edit <index>
 set ipv6-address <ipv6-address>
 end
 end
```

The following diagnose commands for IPv6 PIM sparse mode are also available:

```
diagnose ipv6 multicast status
diagnose ipv6 multicast vif
diagnose ipv6 multicast mroute
```

# Chapter 12 Load Balancing for FortiOS

## 5.0

FortiOS server load balancing includes the features you would expect of any server load balancing solution. Traffic can be distributed across multiple backend servers based on multiple methods including static (failover), round robin, weighted to account for different sized servers, or based on the health and performance of the server including round trip time, number of connections. The load balancer supports HTTP, HTTPS, IMAPS, POP3S, SMTPS, SSL or generic TCP/UDP or IP protocols. Session persistence is supported based on the SSL session ID or based on an injected HTTP cookie.

### Before you begin

Before you begin to configure load balancing, take a moment to note the following:

- To be able to configure load balancing from the web-based manager you should begin by going to the System Information dashboard widget and enabling *Load Balance*.

### How this chapter is organized

This document contains detailed information about how to configure firewall server load balancing to load balance various types of traffic to multiple backend servers. This document describes all server load balancing configuration options and contains detailed configuration examples.

This FortiOS Handbook chapter contains the following sections:

[Configuring load balancing](#) describes FortiGate firewall load balancing.

[Load balancing configuration examples](#) describes includes basic and advanced load balancing configurations.



# Configuring load balancing

This section describes how to use the FortiGate firewall load balancing configuration to load balance traffic to multiple backend servers.


This section describes:

- [Load balancing overview](#)
- [Basic load balancing configuration example](#)
- [HTTP and HTTPS load balancing, multiplexing, and persistence](#)
- [SSL/TLS load balancing](#)
- [IP, TCP, and UDP load balancing](#)

## Load balancing overview

You can configure FortiOS load balancing to intercept incoming traffic with a virtual server and share it among one or more backend real servers. By doing so, the FortiGate unit enables multiple real servers to respond as if they were a single device or virtual server. This in turn means that more simultaneous requests can be handled.

**Figure 289:**Load balancing configuration



Traffic can be balanced across multiple backend real servers based on a selection of load balancing methods including static (failover), round robin, weighted to account for different sized servers, or based on the health and performance of the server including round trip time, number of connections. The load balancer can balance layer 7 HTTP, HTTPS, SSL, generic layer

4 TCP, UDP and generic layer 3 IP protocols. Session persistence is supported based on injected HTTP/HTTPS cookies or the SSL session ID.

You can bind up to 8 real servers can to one virtual server. The real server topology is transparent to end users, and the users interact with the system as if it were only a single server with the IP address and port number of the virtual server. The real servers may be interconnected by high-speed LAN or by geographically dispersed WAN. The FortiGate unit schedules requests to the real servers and makes parallel services of the virtual server to appear to involve a single IP address.

There are additional benefits to load balancing. First, because the load is distributed across multiple servers, the service being provided can be highly available. If one of the servers breaks down, the load can still be handled by the other servers. Secondly, this increases scalability. If the load increases substantially, more servers can be added behind the FortiGate unit in order to cope with the increased load.

## Load balancing, UTM, authentication, and other FortiOS features

Flow-based and proxy-based UTM features such as virus scanning, IPS, DLP, application control, and web filtering can be applied to sessions that are to be load balanced. This includes SSL offloading and multiplexing. Applying these UTM features to load balancing traffic may reduce load balancing performance.

Authentication and dynamic profiles are not supported for load balancing sessions. Usually FortiGate load balancing is used to allow public access to services on servers protected by a FortiGate unit. Authentication is not generally not required for this kind of configuration.

Features such web proxying, web caching, and WAN optimization also do not work with load balanced sessions. However, most other features that can be applied by a security policy are supported.

## Configuring load balancing virtual servers

A virtual server is a specialized firewall virtual IP that performs server load balancing. From the web-based manager you add load balancing virtual server by going to *Firewall Objects > Load Balance > Virtual Server*.

---

<b>Name</b>	Enter the name for the virtual server.
<b>Color</b>	Select <i>Change</i> beside the icon to change the color of the icon. When you select <i>Change</i> , a color palette window appears; select a color from the palette window.

---

<b>Type</b>	<p>Select the protocol to be load balanced by the virtual server. If you select a general protocol such as <i>IP</i>, <i>TCP</i>, or <i>UDP</i> the virtual server load balances all <i>IP</i>, <i>TCP</i>, or <i>UDP</i> sessions. If you select specific protocols such as <i>HTTP</i>, <i>HTTPS</i>, or <i>SSL</i> you can apply additional server load balancing features such as <i>Persistence</i> and <i>HTTP Multiplexing</i>.</p> <ul style="list-style-type: none"> <li>• Select <i>HTTP</i> to load balance only <i>HTTP</i> sessions with destination port number that matches the <i>Virtual Server Port</i> setting. Change <i>Virtual Server Port</i> to match the destination port of the sessions to be load balanced (usually port 80 for <i>HTTP</i> sessions). You can also select <i>HTTP Multiplex</i>. You can also set <i>Persistence</i> to <i>HTTP Cookie</i> to select cookie-based persistence.</li> <li>• Select <i>HTTPS</i> to load balance only <i>HTTPS</i> sessions with destination port number that matches the <i>Virtual Server Port</i> setting. Change <i>Virtual Server Port</i> to match the destination port of the sessions to be load balanced (usually port 443 for <i>HTTPS</i> sessions). You can also select <i>Multiplex HTTP requests/responses</i>. You can also set <i>Persistence</i> to <i>HTTP Cookie</i> to select cookie-based persistence. You can also set <i>Persistence</i> to <i>SSL Session ID</i>.</li> <li>• Select <i>IMAPS</i> to load balance only <i>IMAPS</i> sessions with destination port number that matches the <i>Virtual Server Port</i> setting. Change <i>Virtual Server Port</i> to match the destination port of the sessions to be load balanced (usually port 993 for <i>IMAPS</i> sessions). You can also set <i>Persistence</i> to <i>SSL Session ID</i>.</li> <li>• Select <i>POP3S</i> to load balance only <i>POP3S</i> sessions with destination port number that matches the <i>Virtual Server Port</i> setting. Change <i>Virtual Server Port</i> to match the destination port of the sessions to be load balanced (usually port 995 for <i>POP3S</i> sessions). You can also set <i>Persistence</i> to <i>SSL Session ID</i>.</li> <li>• Select <i>SMTPS</i> to load balance only <i>SMTPS</i> sessions with destination port number that matches the <i>Virtual Server Port</i> setting. Change <i>Virtual Server Port</i> to match the destination port of the sessions to be load balanced (usually port 465 for <i>SMTPS</i> sessions). You can also set <i>Persistence</i> to <i>SSL Session ID</i>.</li> <li>• Select <i>SSL</i> to load balance only <i>SSL</i> sessions with destination port number that matches the <i>Virtual Server Port</i> setting. Change <i>Virtual Server Port</i> to match the destination port of the sessions to be load balanced.</li> <li>• Select <i>TCP</i> to load balance only <i>TCP</i> sessions with destination port number that matches the <i>Virtual Server Port</i> setting. Change <i>Virtual Server Port</i> to match the destination port of the sessions to be load balanced.</li> <li>• Select <i>UDP</i> to load balance only <i>UDP</i> sessions with destination port number that matches the <i>Virtual Server Port</i> setting. Change <i>Virtual Server Port</i> to match the destination port of the sessions to be load balanced.</li> <li>• Select <i>IP</i> to load balance all sessions accepted by the security policy that contains this virtual server.</li> </ul>
<b>Interface</b>	<p>Select the virtual server external interface from the list. The external interface is connected to the source network and receives the packets to be forwarded to the destination network.</p>
<b>Virtual Server IP</b>	<p>The IP address of the virtual server. This is an IP address on the external interface that you want to map to an address on the destination network.</p>

<b>Virtual Server Port</b>	Enter the external port number that you want to map to a port number on the destination network. Sessions with this destination port are load balanced by this virtual server.
<b>Load Balance Method</b>	Select the load balancing method used by the virtual server. See <a href="#">“Load balancing methods” on page 1885</a> .
<b>Persistence</b>	Configure persistence to make sure that a user is connected to the same server every time they make a request that is part of the same session. Session persistence is supported for HTTP and SSL sessions. See <a href="#">“Session persistence” on page 1886</a> . For HTTP and HTTPS sessions, see <a href="#">“HTTP and HTTPS persistence” on page 1898</a> .
<b>HTTP Multiplexing</b>	Select to use the FortiGate unit to multiplex multiple client connections into a few connections between the FortiGate unit and the real server. See <a href="#">“HTTP and HTTPS multiplexing” on page 1898</a> .
<b>Preserve Client IP</b>	Select to preserve the IP address of the client in the <code>X-Forwarded-For</code> HTTP header. This can be useful if you want log messages on the real servers to the client’s original IP address. If this option is not selected, the header will contain the IP address of the FortiGate unit.  This option appears only if <i>HTTP</i> or <i>HTTPS</i> are selected for <i>Type</i> , and is available only if <i>HTTP Multiplexing</i> is selected.
<b>SSL Offloading</b>	Select to accelerate clients’ SSL connections to the server by using the Fortinet FortiGate unit to perform SSL operations, then select which segments of the connection will receive SSL offloading. See <a href="#">“SSL offloading” on page 1903</a>
<b>Certificate</b>	Select the certificate to use with <i>SSL Offloading</i> . The certificate key size must be 1024 or 2048 bits. 4096-bit keys are not supported.  This option appears only if <i>HTTPS</i> or <i>SSL</i> are selected for <i>Type</i> , and is available only if <i>SSL Offloading</i> is selected.
<b>Health Check</b>	Select which health check monitor configuration will be used to determine a server’s connectivity status. See <a href="#">“Health check monitoring” on page 1888</a> .

From the CLI you configure a virtual server by added a firewall virtual IP and setting the virtual IP type to server load balance:

```
config firewall vip
 edit Vserver-HTTP-1
 set type server-load-balance
 ...
```

A virtual server includes a virtual server IP address bound to an interface. The virtual server IP address is the destination address incoming packets to be load balanced and the virtual server is bound to the interface that receives the packets to be load balanced.

For example, if you want to load balance incoming HTTP traffic from the Internet to a group of web servers on a DMZ network, the virtual server IP address is the known Internet IP address of the web servers and the virtual server binds this IP address to the FortiGate interface connected to the Internet.

When you bind the virtual server’s external IP address to a FortiGate unit interface, by default, the network interface responds to ARP requests for the bound IP address. Virtual servers use proxy ARP, as defined in [RFC 1027](#), so that the FortiGate unit can respond to ARP requests on a network for a real server that is actually installed on another network. In some cases you may

not want the network interface sending ARP replies. You can use the `arp-reply` option disable sending ARP replies:

```
config firewall vip
 edit Vserver-HTTP-1
 set type server-load-balance
 set arp-reply disable
 ...
```

The load balancing virtual server configuration also includes the virtual server port. This is the TCP port on the bound interface that the virtual server listens for traffic to be load balanced on. The virtual server can listen on any port.

## Load balancing methods

The load balancing method defines how sessions are load balanced to real servers. A number of load balancing methods are available as listed in [Table 77](#).

All load balancing methods will not send traffic to real servers that are down or not responding. However, the FortiGate unit can only determine if a real server is not responding by using a health check monitor. You should always add at least one health check monitor to a virtual server or to individual real servers, or load balancing methods may attempt to distribute sessions to real servers that are not functioning.

**Table 77:** Load balancing methods

Method	Description
<b>Source IP Hash</b>	The traffic load is statically spread evenly across all real servers. However, sessions are not assigned according to how busy individual real servers are. This load balancing method provides some persistence because all sessions from the same source address always go to the same real server. However, the distribution is stateless, so if a real server is added or removed (or goes up or down) the distribution is changed and persistence could be lost.
<b>Round Robin</b>	Directs new requests to the next real server, and treats all real servers as equals regardless of response time or number of connections. Dead real servers or non responsive real servers are avoided.
<b>Weighted</b>	Real servers with a higher weight value receive a larger percentage of connections. Set the real server weight when adding a real server.
<b>First Alive</b>	Always directs sessions to the first alive real server. This load balancing schedule provides real server failover protection by sending all sessions to the first alive real server and if that real server fails, sending all sessions to the next alive real server. Sessions are not distributed to all real servers so all sessions are processed by the “first” real server only.  First refers to the order of the real servers in the virtual server configuration. For example, if you add real servers A, B and C in that order, then all sessions always go to A as long as it is alive. If A goes down then sessions go to B and if B goes down sessions go to C. If A comes back up sessions go back to A. Real servers are ordered in the virtual server configuration in the order in which you add them, with the most recently added real server last. If you want to change the order you must delete and re-add real servers in the required order.
<b>Least RTT</b>	Directs sessions to the real server with the least round trip time. The round trip time is determined by a Ping health check monitor and is defaulted to 0 if no Ping health check monitors are added to the virtual server.

**Table 77:** Load balancing methods

Method	Description
<b>Least Session</b>	Directs requests to the real server that has the least number of current connections. This method works best in environments where the real servers or other equipment you are load balancing all have similar capabilities. This load balancing method uses the FortiGate session table to track the number of sessions being processed by each real server. The FortiGate unit cannot detect the number of sessions actually being processed by a real server.
<b>HTTP Host</b>	Load balances HTTP host connections across multiple real servers using the host's HTTP header to guide the connection to the correct real server.

## Session persistence

Use persistence to make sure that a user is connected to the same real server every time they make an HTTP, HTTPS, or SSL request that is part of the same user session. For example, if you are load balancing HTTP and HTTPS sessions to a collection of eCommerce web servers, when a user is making a purchase they will be starting multiple sessions as they navigate the eCommerce site. In most cases all of the sessions started by this user during on eCommerce session should be processed by the same real server. Typically, the HTTP protocol keeps track of these related sessions using cookies. HTTP cookie persistence makes sure that all sessions that are part of the same user session are processed by the same real server

When you configure persistence, the FortiGate unit load balances a new session to a real server according to the load balance method. If the session has an HTTP cookie or an SSL session ID, the FortiGate unit sends all subsequent sessions with the same HTTP cookie or SSL session ID to the same real server. For more information about HTTP and HTTPS persistence, see [“HTTP and HTTPS persistence” on page 1898](#).

## Real servers

Add real servers to a load balancing virtual server to provide the information the virtual server requires to be able to send sessions to the server. A real server configuration includes the IP address of the real server and port number that the real server receives sessions on. The FortiGate unit sends sessions to the real server's IP address using the destination port number in the real server configuration.

When configuring a real server you can also specify the weight (used if the load balance method is set to weighted) and you can limit the maximum number of open connections between the FortiGate unit and the real server. If the maximum number of connections is reached for the real server, the FortiGate unit will automatically switch all further connection requests other real servers until the connection number drops below the specified limit. Setting Maximum Connections to 0 means that the FortiGate unit does not limit the number of connections to the real server.

### Real server active, standby, and disabled modes

By default the real server mode setting is active indicating that the real server is available to receive connections. If the real server is removed from the network (for example, for routine maintenance or because of a hardware or software failure) you can change the mode to standby or disabled. In disabled mode the FortiGate unit no longer sends sessions to the real server.

If a real server is in standby mode the FortiGate also does not send sessions to it unless other real servers added to the same virtual server become unavailable. For example:

- A virtual server that includes two real servers one in active mode and one in standby mode. If the real server in active mode fails, the real server in standby mode is changed to active mode and all sessions are sent to this real server.
- A virtual server includes three real servers, two in active mode and one in standby mode, if one of the real servers in active mode fails, the real server in standby mode is changed to active mode and sessions are load balanced between it and still operating real server. If both real servers in active mode fail, all sessions are sent to the real server in standby mode.

### Adding real servers

To add a real server from the web-based manager go to *Firewall Objects > Load Balance > Real Server*.

<b>Virtual Server</b>	Select the virtual server that will send sessions to this real server.
<b>IP Address</b>	Enter the IP address of the real server.
<b>Port</b>	Enter the port number on the destination network to which the external port number is mapped.
<b>Weight</b>	Enter the weight value of the real server. The higher the weight value, the higher the percentage of connections the server will handle. A range of 1-255 can be used. This option is available only if the associated virtual server's load balance method is <i>Weighted</i> .
<b>Max Connections</b>	Enter the limit on the number of active connections directed to a real server. A range of 1-99999 can be used. If the maximum number of connections is reached for the real server, the FortiGate unit will automatically switch all further connection requests to another server until the connection number drops below the specified limit.  Setting <i>Maximum Connections</i> to 0 means that the FortiGate unit does not limit the number of connections to the real server.
<b>HTTP Host</b>	Enter the HTTP header for load balancing across multiple real servers. This feature is used for load balancing HTTP host connections across multiple real servers using the host's HTTP header to guide the connection to the correct real server, providing better load balancing for those specific connections.
<b>Mode</b>	Select a mode for the real server.

To add a real server from the CLI you configure a virtual server and add real servers to it. For example, to add three real servers to a virtual server that load balances UDP sessions on port 8190 using weighted load balancing. For each real server the port is not changed. The default real server port is 0 resulting in the traffic being sent the real server with destination port 8190.

Each real sever is given a different weight. Servers with higher weights have a max-connections limit to prevent too many sessions from being sent to them.

```
config firewall vip
 edit Vserver-UDP-1
 set type server-load-balance
 set server-type udp
 set ldb-method weighted
 set extip 172.20.120.30
 set extintf wan1
 set extport 8190
 set monitor ping-mon-1
 config realservers
 edit 1
 set ip 10.31.101.30
 set weight 100
 set max-connections 10000
 next
 edit 2
 set ip 10.31.101.40
 set weight 100
 set max-connections 10000
 next
 edit 3
 set ip 10.31.101.50
 set weight 10
 end
 end
```

## Health check monitoring

From the FortiGate web-based manager you can go to *Firewall Objects > Load Balance > Health Check* and configure health check monitoring so that the FortiGate unit can verify that real servers are able respond to network connection attempts. If a real server responds to connection attempts the load balancer continues to send sessions to it. If a real server stops responding to connection attempts the load balancer assumes that the server is down and does not send sessions to it. The health check monitor configuration determines how the load balancer tests the real servers. You can use a single health check monitor for multiple load balancing configurations.

You can configure TCP, HTTP and Ping health check monitors. Usually you would want the health check monitor to use the same protocol for checking the health of the server as the traffic being load balanced to it. For example, for an HTTP load balancing configuration you would normally use an HTTP health check monitor.

For the TCP and HTTP health check monitors you can specify the destination port to use to connect to the real servers. If you set the port to 0, the health check monitor uses the port defined in the real server. This allows you to use the same health check monitor for multiple real servers using different ports. You can also configure the interval, timeout and retry. A health check occurs every number of seconds indicated by the interval. If a reply is not received within the timeout period the health check is repeated. If no response is received after the number of configured retries, the virtual server is considered unresponsive, and load balancing will disabling traffic to that real server. The health check monitor will continue to contact the real



server and if successful, the load balancer can resume sending sessions to the recovered real server.

For HTTP health check monitors, you can add URL that the FortiGate unit connects to when sending a get request to check the health of a HTTP server. The URL should match an actual URL for the real HTTP servers. The URL is optional.

The URL would not usually include an IP address or domain name. Instead it should start with a “/” and be followed by the address of an actual web page on the real server. For example, if the IP address of the real server is 10.31.101.30, the URL “/test\_page.htm” causes the FortiGate unit to send an HTTP get request to “http://10.31.101.30/test\_page.htm”.

For HTTP health check monitors, you can also add a matched content phrase that a real HTTP server should include in response to the get request sent by the FortiGate unit using the content of the URL option. If the URL returns a web page, the matched content should exactly match some of the text on the web page. You can use the URL and Matched Content options to verify that an HTTP server is actually operating correctly by responding to get requests with expected web pages. Matched content is only required if you add a URL.

For example, you can set matched content to “server test page” if the real HTTP server page defined by the URL option contains the phrase “server test page”. When the FortiGate unit receives the web page in response to the URL get request, the system searches the content of the web page for the matched content phrase.

---

<b>Name</b>	Enter the name of the health check monitor configuration.
<b>Type</b>	Select the protocol used to perform the health check. <ul style="list-style-type: none"><li>• TCP</li><li>• HTTP</li><li>• PING</li></ul>
<b>Port</b>	Enter the port number used to perform the health check. If you set the <i>Port</i> to 0, the health check monitor uses the port defined in the real server. This way you can use a single health check monitor for different real servers.  This option does not appear if the <i>Type</i> is <i>PING</i> .
<b>Interval</b>	Enter the number of seconds between each server health check.
<b>URL</b>	For HTTP health check monitors, add a URL that the FortiGate unit uses when sending a get request to check the health of a HTTP server. The URL should match an actual URL for the real HTTP servers. The URL is optional.  The URL would not usually include an IP address or domain name. Instead it should start with a “/” and be followed by the address of an actual web page on the real server. For example, if the IP address of the real server is 10.10.10.1, the <i>URL</i> “/test_page.htm” causes the FortiGate unit to send an HTTP get request to “http://10.10.10.1/test_page.htm”.  This option appears only if <i>Type</i> is <i>HTTP</i> .

---

<b>Matched Content</b>	<p>For HTTP health check monitors, add a phrase that a real HTTP server should include in response to the get request sent by the FortiGate unit using the content of the <i>URL</i> option. If the <i>URL</i> returns a web page, the <i>Matched Content</i> should exactly match some of the text on the web page. You can use the <i>URL</i> and <i>Matched Content</i> options to verify that an HTTP server is actually operating correctly by responding to get requests with expected web pages. Matched content is only required if you add a URL.</p> <p>For example, you can set <i>Matched Content</i> to “server test page” if the real HTTP server page defined by the URL option contains the phrase “server test page”. When the FortiGate unit receives the web page in response to the URL get request, the system searches the content of the web page for the <i>Matched Content</i> phrase.</p> <p>This option appears only if <i>Type</i> is <i>HTTP</i>.</p>
<b>Timeout</b>	Enter the number of seconds which must pass after the server health check to indicate a failed health check.
<b>Retry</b>	Enter the number of times, if any, a failed health check will be retried before the server is determined to be inaccessible.

### Virtual IP, load balance virtual server and load balance real server limitations

The following limitations apply when adding virtual IPs, Load balancing virtual servers, and load balancing real servers. Load balancing virtual servers are actually server load balancing virtual IPs. You can add server load balance virtual IPs from the CLI.

- Virtual IP *External IP Address/Range* entries or ranges cannot overlap with each other or with load balancing virtual server *Virtual Server IP* entries.
- A virtual IP *Mapped IP Address/Range* cannot be 0.0.0.0 or 255.255.255.255.
- A real server *IP* cannot be 0.0.0.0 or 255.255.255.255.
- If a static NAT virtual IP *External IP Address/Range* is 0.0.0.0, the *Mapped IP Address/Range* must be a single IP address.
- If a load balance virtual IP *External IP Address/Range* is 0.0.0.0, the *Mapped IP Address/Range* can be an address range.
- When port forwarding, the count of mapped port numbers and external port numbers must be the same. The web-based manager does this automatically but the CLI does not.
- Virtual IP and virtual server names must be different from firewall address or address group names.

## Monitoring load balancing

From the web-based manager you can go to *Firewall Objects > Monitor > Load Balance Monitor* to monitor the status of configured virtual servers and real server and start or stop the real servers. You can also use the `get test ipldb` command from the CLI to display similar information.

For each real server the monitor displays health status (up or down), active sessions, round trip time and the amount of bytes of data processed. From the monitor page you can also stop sending new sessions to any real server. When you select to stop sending sessions the FortiGate unit performs a graceful stop by continuing to send data for sessions that were established or persistent before you selected stop. However, no new sessions are started.

<b>Virtual Server</b>	The IP addresses of the existing virtual servers.
<b>Real Server</b>	The IP addresses of the existing real servers.
<b>Health Status</b>	Displays the health status according to the health check results for each real server. A green arrow means the server is up. A red arrow means the server is down.
<b>Mode</b>	The mode of the health check monitor. Can be active, standby, or disabled.
<b>Monitor Events</b>	Display each real server's up and down times.
<b>Active Sessions</b>	Display each real server's active sessions.
<b>RTT (ms)</b>	Displays the Round Trip Time (RTT) of each real server. By default, the RTT is "<1". This value will change only when ping monitoring is enabled on a real server.
<b>Bytes Processed</b>	Displays the traffic processed by each real server.
<b>Graceful Stop/Start</b>	Select to start or stop real servers. When stopping a server, the FortiGate unit will not accept new sessions but will wait for the active sessions to finish.

## Load balancing get command

The following get command is available to display testing and debug information for the FortiGate virtual server process:

```
get test vs <test-level_int>
```

Where <test-level\_int> can be:

- 3 to display the virtual server process id.
- 8 to display the virtual server log configuration.
- 30 to display the virtual server configuration statistics.
- 99 to restart the virtual server process.

## Load balancing diagnose commands

You can also use the following diagnose commands to view status information for load balancing virtual servers and real servers:

```
diagnose firewall vip realserver {down | flush | healthcheck | list | up}
diagnose firewall vip virtual-server {filter | log | real-server | session | stats}
```

For example, the following command lists and displays status information for all real servers:

```
diagnose firewall vip virtual-server real-server

vd root/0 vs vs/2 addr 10.31.101.30:80 status 1/1
conn: max 0 active 0 attempts 0 success 0 drop 0 fail 0

vd root/0 vs vs/2 addr 10.31.101.20:80 status 1/1
conn: max 0 active 0 attempts 0 success 0 drop 0 fail 0
```

Many of the diagnostic commands involve retrieving information about one or more virtual servers. To control which servers are queried you can define a filter:

```
diagnose firewall vip virtual-server filter <filter_str>
```

Where <filter\_str> can be:

- clear erase the current filter
- dst the destination address range to filter by
- dst-port the destination port range to filter by
- list display the current filter
- name the vip name to filter by
- negate negate the specified filter parameter
- src the source address range to filter by
- src-port the source port range to filter by
- vd index of virtual domain. -1 matches all

The default filter is empty so no filtering is done.

## Logging Diagnostics

The logging diagnostics provide information about two separate features:

```
diagnose firewall vip virtual-server log {console | filter}
```

Where

console {disable | enable} enables or disables displaying the event log messages generated by virtual server traffic on the console to simplify debugging.

filter sets a filter for the virtual server debug log

The filter option controls what entries the virtual server daemon will log to the console if diagnose debug application vs level is non-zero. The filtering can be done on source, destination, virtual-server name, virtual domain, and so on:

```
diagnose firewall vip virtual-server log filter <filter_str>
```

where <filter\_str> can be

- clear erase the current filter
- dst the destination address range to filter by
- dst-port the destination port range to filter by
- list display the current filter
- name the virtual-server name to filter by
- negate negate the specified filter parameter

`src` the source address range to filter by  
`src-port` the source port range to filter by  
`vd` index of virtual domain. -1 matches all  
The default filter is empty so no filtering is done.

## Real server diagnostics

Enter the following command to list all the real servers:

```
diag firewall vip virtual-server real-server list
```

In the following example there is only one virtual server called `slb` and it has two real-servers:

```
diag firewall vip virtual-server server
vd root/0 vs slb/2 addr 172.16.67.191:80 status 1/1
 conn: max 10 active 0 attempts 0 success 0 drop 0 fail 0
 http: available 0 total 0
```

```
vd root/0 vs slb/2 addr 172.16.67.192:80 status 1/1
 conn: max 10 active 1 attempts 4 success 4 drop 0 fail 0
 http: available 1 total 1
```

The `status` indicates the administrative and operational status of the real-server.

`max` indicates that the real-server will only allow 10 concurrent connections.

`active` is the number of current connections to the server `attempts` is the total number of connections attempted `success` is the total number of connections that were successful.

`drop` is the total number of connections that were dropped because the active count hit max.

`fail` is the total number of connections that failed to complete due to some internal problem (for example, lack of memory).


If the virtual server has HTTP multiplexing enabled then the HTTP section indicates how many established connections to the real-server are available to service a HTTP request and also the total number of connections.

## Basic load balancing configuration example

This section describes the steps required to configure the load balancing configuration shown in [Figure 290](#). In this configuration a FortiGate-51B unit is load balancing HTTP traffic from the Internet to three HTTP servers on the Internal network. HTTP sessions are accepted at the `wan1` interface with destination IP address 172.20.120.121 on TCP port 8080 and forwarded from the internal interface to the web servers. When forwarded the destination address of the sessions is translated to the IP address of one of the web servers.

The load balancing configuration also includes session persistence using HTTP cookies, round-robin load balancing, and TCP health monitoring for the real servers. Ping health monitoring consists of the FortiGate unit using ICMP ping to make sure the web servers can respond to network traffic.

**Figure 290:**Virtual server and real servers setup



**To configure the example load balancing configuration - general configuration steps**

- 1 Add a load balance ping health check monitor  
A ping health check monitor causes the FortiGate unit to ping the real servers every 10 seconds. If one of the servers does not respond within 2 seconds, the FortiGate unit will retry the ping 3 times before assuming that the HTTP server is not responding.
- 2 Add a load balance virtual server.
- 3 Add the three load balance real servers. Include the virtual server in each real server configuration.
- 4 Add a security policy that includes the load balance virtual server as the destination address.

**To configure the example load balancing configuration - web-based manager**

- 1 Go to go to *Firewall Objects > Load Balance > Health Check* and add the following health check monitor.

<b>Name</b>	Ping-mon-1
<b>Type</b>	Ping
<b>Interval</b>	10 seconds
<b>Timeout</b>	2 seconds
<b>Retry</b>	3

- Go to *Firewall Objects > Load Balance > Virtual Server* and add virtual server that accepts the traffic to be load balanced.

<b>Name</b>	Vserver-HTTP-1
<b>Type</b>	HTTP
<b>Interface</b>	wan1
<b>Virtual Server IP</b>	172.20.120.121
<b>Virtual Server Port</b>	8080
<b>Load Balance Method</b>	Round Robin
<b>Persistence</b>	HTTP Cookie
<b>HTTP Multiplexing</b>	Do not select
<b>Health Check</b>	Move Ping-mon-1 to the Selected list.

- Go to go to *Firewall Objects > Load Balance > Real Server* and add the real servers.

<b>Virtual Server</b>	Vserver-HTTP-1
<b>IP Address</b>	10.31.101.30
<b>Port</b>	80
<b>Weight</b>	n/a
<b>Max Connections</b>	0
<b>Mode</b>	Active
<b>Virtual Server</b>	Vserver-HTTP-1
<b>IP Address</b>	10.31.101.40
<b>Port</b>	80
<b>Weight</b>	n/a
<b>Max Connections</b>	0
<b>Mode</b>	Active
<b>Virtual Server</b>	Vserver-HTTP-1
<b>IP Address</b>	10.31.101.50
<b>Port</b>	80
<b>Weight</b>	n/a

<b>Max Connections</b>	0
<b>Mode</b>	Active

- Go to *Policy > Policy > Policy* and add a wan1 to internal security policy that includes the virtual server. This policy also applies an Antivirus profile to the load balanced sessions.

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	wan1
<b>Source Address</b>	all
<b>Outgoing Interface</b>	internal
<b>Destination Address</b>	Vserver-HTTP-1
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Select this option and select <i>Use Destination Interface Address</i> .
<b>Use Standard UTM Profiles</b>	Select
<b>Antivirus</b>	Turn ON and select an Antivirus profile.
<b>UTM Proxy Options</b>	Select a profile.

- Select OK.

#### To configure the example load balancing configuration- CLI

- Use the following command to add a Ping health check monitor.

```
config firewall ldb-monitor
 edit ping-mon-1
 set type ping
 set interval 10
 set timeout 2
 set retry 3
 end
```



- 2 Use the following command to add the virtual server that accepts HTTP sessions on port 8080 at the wan1 interface and load balances the traffic to three real servers.

```
config firewall vip
 edit Vserver-HTTP-1
 set type server-load-balance
 set server-type http
 set ldb-method round-robin
 set extip 172.20.120.30
 set extintf wan1
 set extport 8080
 set persistence http-cookie
 set monitor tcp-mon-1
 config realservers
 edit 1
 set ip 10.31.101.30
 set port 80
 next
 edit 2
 set ip 10.31.101.40
 set port 80
 end
 end
```

- 3 Use the following command to add a security policy that includes the load balance virtual server as the destination address.

```
config firewall policy
 edit 0
 set srcintf wan1
 set srcaddr all
 set dstintf internal
 set dstaddr Vserver-HTTP-1
 set action accept
 set schedule always
 set service ALL
 set nat enable
 set utm-status enable
 set profile-protocol-options default
 set av-profile scan
 end
```

## HTTP and HTTPS load balancing, multiplexing, and persistence

In a firewall load balancing virtual server configuration, you can select HTTP to load balance only HTTP sessions. The virtual server will load balance HTTP sessions received at the virtual server interface with destination IP address that matches the configured virtual server IP and destination port number that matches the configured virtual server port. The default virtual server port for HTTP load balancing is 80, but you can change this to any port number. Similarly for HTTPS load balancing, set the virtual server type to HTTPS and then select the interface, virtual server IP, and virtual server port that matches the HTTPS traffic to be load balanced. Usually HTTPS traffic uses port 443.

You can also configure load balancing to offload SSL processing for HTTPS and SSL traffic. See [“SSL offloading” on page 1903](#) for more information.

## HTTP and HTTPS multiplexing

For both HTTP and HTTPS load balancing you can multiplex HTTP requests and responses over a single TCP connection. HTTP multiplexing is a performance saving feature of HTTP/1.1 compliant web servers that provides the ability to pipeline many unrelated HTTP or HTTPS requests on the same connection. This allows a single HTTPD process on the server to interleave and serve multiple requests. The result is fewer idle sessions on the web server so server resources are used more efficiently. HTTP multiplexing can take multiple separate inbound sessions and multiplex them over the same internal session. This may reduce the load on the backend server and increase the overall performance.

HTTP multiplexing may improve performance in some cases. For example, if users web browsers are only compatible with HTTP 1.0. HTTP multiplexing can also improve performance between a web server and the FortiGate unit if the FortiGate unit is performing SSL acceleration. However, in most cases HTTP multiplexing should only be used if enabling it leads to a measurable improvement in performance.

To enable HTTP multiplexing from the web-based manager, select multiplex HTTP requests/responses over a single TCP connection. To enable HTTP multiplexing from the CLI enable the `http-multiplex` option.

## Preserving the client IP address

Select preserve client IP from the web-based manager or enable the `http-ip-header` option from the CLI to preserve the IP address of the client in the `X-Forwarded-For` HTTP header. This can be useful in an HTTP multiplexing configuration if you want log messages on the real servers to the client's original IP address. If this option is not selected, the header will contain the IP address of the FortiGate unit.

## HTTP and HTTPS persistence

Configure load balancing persistence for HTTP or HTTPS to make sure that a user is connected to the same server every time they make a request that is part of the same session. HTTP cookie persistence uses injected cookies to enable persistence.

When you configure persistence, the FortiGate unit load balances a new session to a real server according to the *Load Balance Method*. If the session has an HTTP cookie or an SSL session ID, the FortiGate unit sends all subsequent sessions with the same HTTP cookie or SSL session ID to the same real server.

The following example shows how to enable cookie persistence and set the cookie domain to `.example.org`.

```
config firewall vip
 edit HTTP_Load_Balance
 set type server-load-balance
 set server-type http
 set extport 8080
 set extintf port2
 set extip 192.168.20.20
 set persistence http-cookie
 set http-cookie-domain .example.org
 config realservers
 edit 1
 set ip 10.10.10.1
 set port 80
 next
 edit 2
 set ip 10.10.10.2
 set port 80
 next
 edit 3
 set ip 10.10.10.3
 set port 80
 end
end
```

## How HTTP cookie persistence options work

The following options are available for the `config firewall vip` command when `type` is set to `server-load-balance`, `server-type` is set to `http` or `https` and `persistence` is set to `http-cookie`:

```
http-cookie-domain-from-host
http-cookie-domain
http-cookie-path
http-cookie-generation
http-cookie-age
http-cookie-share
https-cookie-share
```

When HTTP cookie persistence is enabled the FortiGate unit inserts a header of the following form into each HTTP response unless the corresponding HTTP request already contains a `FGTServer` cookie:

```
Set-Cookie: FGTServer=E7D01637C4B08E89A6714213A9D85D9C7E4D8158;
Version=1; Max-Age=3600
```

The value of the `FGTServer` cookie encodes the server that traffic should be directed to. The value is encoded so as to not leak information about the internal network.

Enable `http-cookie-domain-from-host` to extract the cookie domain from the `host` header in the HTTP request. For example, to restrict the cookie to `.server.com`, enter:

The generated cookies could have the following form if the *Host:* header contains *exhost.com*:

```
Set-Cookie: FGTSerVer=E7D01637C4B08E89A6714213A9D85D9C7E4D8158;
Version=1; Domain=.exhost.com; Max-Age=3600
```

For more information, see [“HTTP host-based load balancing” on page 1901](#).

Use `http-cookie-domain` to restrict the domain that the cookie should apply to. For example, to restrict the cookie to `.server.com`, enter:

```
set http-cookie-domain .server.com
```

Now all generated cookies will have the following form:

```
Set-Cookie: FGTSerVer=E7D01637C4B08E89A6714213A9D85D9C7E4D8158;
Version=1; Domain=.server.com; Max-Age=3600
```

Use `http-cookie-path` to limit the cookies to a particular path. For example, to limit cookies to the path `/sales`, enter:

```
set http-cookie-path /sales
```

Now all generated cookies will have the following form:

```
Set-Cookie: FGTSerVer=E7D01637C4B08E89A6714213A9D85D9C7E4D8158;
Version=1; Domain=.server.com; Path=/sales; Max-Age=3600
```

Use `http-cookie-age` to change how long the browser caches the cookie. You can enter an age in minutes or set the age to 0 to make the browser keep the cookie indefinitely:

```
set http-cookie-age 0
```

Now all generated cookies will have the following form:

```
Set-Cookie: FGTSerVer=E7D01637C4B08E89A6714213A9D85D9C7E4D8158;
Version=1; Domain=.server.com; Path=/sales
```

Use `http-cookie-generation` to invalidate all cookies that have already been generated. The exact value of the generation is not important, only that it is different from any generation that has already been used for cookies in this domain. The simplest approach is to increment the generation by one each time invalidation is required. Since the default is 0, enter the following to invalidate all existing cookies:

```
set http-cookie-generation 1
```

Use `http-cookie-share {disable | same-ip}` to control the sharing of cookies across virtual servers in the same virtual domain. The default setting `same-ip` means that any `FGTSerVer` cookie generated by one virtual server can be used by another virtual server in the same virtual domain. For example, if you have an application that starts on HTTP and then changes to HTTPS and you want to make sure that the same server is used for the HTTP and HTTPS traffic then you can create two virtual servers, one for port 80 (for HTTP) and one for port 443 (for HTTPS). As long as you add the same real servers to both of these virtual servers (and as long as both virtual servers have the same number of real servers with the same IP addresses), then cookies generated by accessing the HTTP server are reused when the application changes to the HTTPS server.

If for any reason you do not want this sharing to occur then select `disable` to make sure that a cookie generated for a virtual server cannot be used by other virtual servers.

Use `https-cookie-secure` to enable or disable using secure cookies. Secure cookies are disabled by default because secure cookies can interfere with cookie sharing across HTTP and HTTPS virtual servers. If enabled, then the `Secure` tag is added to the cookie inserted by the FortiGate unit:

```
Set-Cookie: FGTSerVer=E7D01637C4B08E89A6714213A9D85D9C7E4D8158;
Version=1; Max-Age=3600; Secure
```

## HTTP host-based load balancing

When configuring HTTP or HTTPS load balancing you can select HTTP host load balancing to load balance HTTP host connections across multiple real servers using the host's HTTP header to guide the connection to the correct real server. HTTP 1.1 includes the concept of a virtual server which allows a HTTP or HTTPS server with a single external IP address to serve requests for multiple DNS domains by using the mandatory `Host:` header in a HTTP request to indicate which DNS domain the request is destined for.

FortiOS can load-balance HTTP and HTTPS connections among multiple real servers using the `Host:` header to guide the connection to the correct real server. The host load balancing method allows a real server to specify a `http-host` attribute which is the domain name of the traffic for that real server. Each real server can only specify a single domain name. The same domain name can appear in more than one real server but only the first one that is up will be used, any others are purely for redundancy. If the `Host:` header contains a domain that does not match any `http-host` entry then the connection will be dropped. A real server with no `http-host` can be matched by any `Host:` domain.

For example, consider a FortiGate unit that is load-balancing traffic to three real servers. Traffic for `www.example1.com` should go to `192.168.2.1`, traffic for `www.example2.com` should go to `192.168.2.2` and traffic to any other domain should go to `192.168.2.3`. To enable this configuration you would add a virtual server and set the load balance method to HTTP host. Then you would add three real servers and set the HTTP host of the real server with IP address `192.168.2.1` to `www.example1.com`, the HTTP host of the real server with IP address `192.168.2.2` to `www.example2.com` and you would not specify an HTTP host for the third real server.

The configuration of a virtual IP to achieve this result would be:

```
config firewall vip
 edit "http-host-ldb"
 set type server-load-balance
 set extip 172.16.67.195
 set extintf "lan"
 set server-type http
 set ldb-method http-host
 set extport 80
 config realservers
 edit 1
 set http-host "www.example1.com"
 set ip 192.168.2.1
 set port 80
 next
 edit 2
 set http-host "www.example2.com"
 set ip 192.168.2.2
 set port 80
 next
 edit 3
 set ip 192.168.2.3
 set port 80
 next
 end
 end
end
```

## Host load balancing and HTTP cookie persistence

In an HTTP host-based load balancing configuration with HTTP cookie persistence enabled you can optionally configure cookie persistence to use the domain set in the host header as the cookie domain. You can do this by enabling the `http-cookie-domain-from-host` option, for example:

```
config firewall vip
 edit "http-host-ldb"
 set type server-load-balance
 set extip 172.16.67.195
 set extintf "lan"
 set server-type http
 set ldb-method http-host
 set extport 80
 set persistence http-cookie
 set http-cookie-domain-from-host enable
 config realservers
 edit 1
 set http-host "www.example1.com"
 set ip 192.168.2.1
 set port 80
 next
 edit 2
 set http-host "www.example2.com"
 set ip 192.168.2.2
 set port 80
 next
 edit 3
 set ip 192.168.2.3
 set port 80
 next
 end
end
```

## SSL/TLS load balancing

In a firewall load balancing virtual server configuration, you can select SSL to load balance only SSL and TLS sessions. The virtual server will load balance SSL and TLS sessions received at the virtual server interface with destination IP address that matches the configured virtual server IP and destination port number that matches the configured virtual server port. Change this port to match the destination port of the sessions to be load balanced.

For SSL load balancing you can also set persistence to SSL session ID. Persistence is achieved by the FortiGate unit sending all sessions with the same SSL session ID to the same real server. When you configure persistence, the FortiGate unit load balances a new session to a real server according to the *Load Balance Method*. If the session has an SSL session ID, the FortiGate unit sends all subsequent sessions with the same SSL session ID to the same real server.

## SSL offloading


Use SSL offloading to accelerate clients' SSL or HTTPS connections to real servers by using the FortiGate unit to perform SSL operations (offloading them from the real servers using the FortiGate unit's SSL acceleration hardware). FortiGate units can offload SSL 3.0 and TLS 1.0. SSL offloading is available on FortiGate units that support SSL acceleration.

To configure SSL offloading from the web-based manager go to *Firewall Objects > Load Balance > Virtual Server*. Add a virtual server and set the type to HTTPS or SSL and select the SSL offloading type (Client <-> FortiGate or Client <-> FortiGate <->Server).

Select Client <-> FortiGate to apply hardware accelerated SSL processing only to the part of the connection between the client and the FortiGate unit. This mode is called half mode SSL offloading. The segment between the FortiGate unit and the server will use clear text communications. This results in best performance, but cannot be used in failover configurations where the failover path does not have an SSL accelerator.

Select Client <-> FortiGate <->Server to apply hardware accelerated SSL processing to both parts of the connection: the segment between client and the FortiGate unit, and the segment between the FortiGate unit and the server. This mode is called full mode SSL offloading. The segment between the FortiGate unit and the server will use encrypted communications, but the handshakes will be abbreviated. This results in performance which is less than the other option, but still improved over communications without SSL acceleration, and can be used in failover configurations where the failover path does not have an SSL accelerator. If the server is already configured to use SSL, this also enables SSL acceleration without requiring changes to the server's configuration.

**Figure 291:SSL Offloading modes**



Configuring SSL offloading also requires selecting a certificate to use for the SSL offloading sessions. The certificate key size must be 1024 or 2048 bits. 4096-bit keys are not supported.



The following CLI command shows an example half mode HTTPS SSL offloading configuration. In the example the `ssl-mode` option sets the SSL offload mode to `half` (which is the default mode).

```
config firewall vip
 edit Vserver-ssl-offload
 set type server-load-balance
 set server-type https
 set ldb-method round-robin
 set extip 172.20.120.30
 set extintf wan1
 set extport 443
 set persistence ssl-session-id
 set ssl-mode half
 set ssl-certificate my-cert
 set monitor t cp-mon-1
 config realservers
 edit 1
 set ip 10.31.101.30
 set port 443
 next
 edit 2
 set ip 10.31.101.40
 set port 443
 end
 end
```

### Additional SSL load balancing options

The following SSL load balancing and SSL offloading options are only available from the CLI:

```
ssl-client-session-state-max <sessionstates_int>
```

Enter the maximum number of SSL session states to keep for the segment of the SSL connection between the client and the FortiGate unit.

```
ssl-client-session-state-timeout <timeout_int>
```

Enter the number of minutes to keep the SSL session states for the segment of the SSL connection between the client and the FortiGate unit.

```
ssl-client-session-state-type {both | client | disable | time}
```

Select which method the FortiGate unit should use when deciding to expire SSL sessions for the segment of the SSL connection between the client and the FortiGate unit.

- `both`: Select to expire SSL session states when either `ssl-client-session-state-max` or `ssl-client-session-state-timeout` is exceeded, regardless of which occurs first.
- `count`: Select to expire SSL session states when `ssl-client-session-state-max` is exceeded.
- `disable`: Select to keep no SSL session states.
- `time`: Select to expire SSL session states when `ssl-client-session-state-timeout` is exceeded.

```
ssl-dh-bits <bits_int>
```

Enter the number of bits of the prime number used in the Diffie-Hellman exchange for RSA encryption of the SSL connection. Larger prime numbers are associated with greater cryptographic strength.

```
ssl-http-location-conversion {enable | disable}
```

Select to replace `http` with `https` in the reply's `Location` HTTP header field. For example, in the reply, `Location: http://example.com/` would be converted to `Location: https://example.com/`

```
ssl-http-match-host {enable | disable}
```

Select to apply `Location` conversion to the reply's HTTP header only if the host name portion of `Location` matches the request's `Host` field, or, if the `Host` field does not exist, the host name portion of the request's URI. If disabled, conversion occurs regardless of whether the host names in the request and the reply match.

For example, if host matching is enabled, and a request contains `Host: example.com` and the reply contains `Location: http://example.cc/`, the `Location` field does not match the host of the original request and the reply's `Location` field remains unchanged. If the reply contains `Location: http://example.com/`, however, then the FortiGate unit detects the matching host name and converts the reply field to `Location: https://example.com/`.

This option appears only if `ssl-http-location-conversion` is enable.

```
ssl-max-version {ssl-3.0 | tls-1.0}
```

Enter the maximum version of SSL/TLS to accept in negotiation.

```
ssl-min-version {ssl-3.0 | tls-1.0}
```

Enter the minimum version of SSL/TLS to accept in negotiation.

```
ssl-send-empty-frags {enable | disable}
```

Select to precede the record with empty fragments to thwart attacks on CBC IV. You might disable this option if SSL acceleration will be used with an old or buggy SSL implementation which cannot properly handle empty fragments.

```
ssl-server-session-state-max <sessionstates_int>
```

Enter the maximum number of SSL session states to keep for the segment of the SSL connection between the server and the FortiGate unit.

```
ssl-server-session-state-timeout <timeout_int>
```

Enter the number of minutes to keep the SSL session states for the segment of the SSL connection between the server and the FortiGate unit. This option appears only if `ssl-mode` is `full`.

```
ssl-server-session-state-type {both | count | disable | time}
```

Select which method the FortiGate unit should use when deciding to expire SSL sessions for the segment of the SSL connection between the server and the FortiGate unit. This option appears only if `ssl-mode` is `full`.

- `both`: Select to expire SSL session states when either `ssl-server-session-state-max` or `ssl-server-session-state-timeout` is exceeded, regardless of which occurs first.
- `count`: Select to expire SSL session states when `ssl-server-session-state-max` is exceeded.
- `disable`: Select to keep no SSL session states.
- `time`: Select to expire SSL session states when `ssl-server-session-state-timeout` is exceeded.

## SSL offloading support or Internet Explorer 6

In some cases the Internet Explorer 6 web browser may be able to access real servers. To resolve this issue, disable the `ssl-send-empty-frags` option:

```
config firewall vip
 edit vip_name
 set ssl-send-empty-frags disable
 end
```

You can disable this option if SSL acceleration will be used with an old or buggy SSL implementation that cannot properly handle empty fragments.

## Disabling SSL/TLS re-negotiation

The vulnerability [CVE-2009-3555](#) affects all SSL/TLS servers that support re-negotiation. FortiOS when configured for SSL/TLS offloading is operating as a SSL/TLS server. The IETF is working on a TLS protocol change that will fix the problem identified by CVE-2009-3555 while still supporting re-negotiation. Until that protocol change is available, you can use the `ssl-client-renegotiation` option to disable support for SSL/TLS re-negotiation. The default value of this option is `allow`, which allows an SSL client to renegotiate. You can change the setting to `deny` to abort any attempts by an SSL client to renegotiate. If you select `deny` as soon as a `ClientHello` message indicating a re-negotiation is received from the client FortiOS terminates the TCP connection.

Since SSL offloading does not support requesting client certificates the only circumstance in which a re-negotiation is required is when more than  $2^{32}$  bytes of data are exchanged over a single handshake. If you are sure that this volume of traffic will not occur then you can disable re-negotiation and avoid any possibility of the attack described in CVE-2009-3555.

The re-negotiation behavior can be tested using OpenSSL. The OpenSSL `s_client` application has the feature that the user can request that it do renegotiation by typing "R". For example, the following shows a successful re-negotiation against a FortiGate unit configured with a VIP for 192.168.2.100:443:

```
$ openssl s_client -connect 192.168.2.100:443
CONNECTED(00000003)
depth=1 /C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
Authority/CN=support/emailAddress=support@fortinet.com
verify error:num=19:self signed certificate in certificate chain
verify return:0

Certificate chain
 0
s:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Fortigate/CN=FW80CM390
9604325/emailAddress=support@fortinet.com
 i:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
 Authority/CN=support/emailAddress=support@fortinet.com
 1 s:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
 Authority/CN=support/emailAddress=support@fortinet.com
 i:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
 Authority/CN=support/emailAddress=support@fortinet.com

Server certificate
-----BEGIN CERTIFICATE-----
---certificate not shown---
-----END CERTIFICATE-----
```

```

subject=/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Fortigate/CN=FW8
OCM3909604325/emailAddress=support@fortinet.com
 issuer=/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
 Authority/CN=support/emailAddress=support@fortinet.com

No client certificate CA names sent

SSL handshake has read 2370 bytes and written 316 bytes

New, TLSv1/SSLv3, Cipher is DHE-RSA-AES256-SHA
Server public key is 1024 bit
Compression: NONE
Expansion: NONE
SSL-Session:
 Protocol : TLSv1
 Cipher : DHE-RSA-AES256-SHA
 Session-ID:
 02781E1E368DCCE97A95396FAA82E8F740F5BBA96CF022F6FEC3597B0CC88095
 Session-ID-ctx:
 Master-Key:

A6BBBD8477A2422D56E57C1792A4EA9C86F37D731E67D0A66E5CDB2B5C76650780C0E7
F01CFF851EC4466186F4C48397
 Key-Arg : None
 Start Time: 1264453027
 Timeout : 300 (sec)
 Verify return code: 19 (self signed certificate in certificate
 chain)

GET /main.c HTTP/1.0
R
RENEGOTIATING
depth=1 /C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
Authority/CN=support/emailAddress=support@fortinet.com
verify error:num=19:self signed certificate in certificate chain
verify return:0
HTTP/1.0 200 ok
Content-type: text/plain

/*
 * Copyright (C) 2004-2007 Fortinet
 */

#include <stdio.h>
#include "vsd_ui.h"

int main(int argc, char **argv)
{
 return vsd_ui_main(argc, argv);
}
closed
$

```

The following is the same test, but this time with the VIP configuration changed to ssl-client-renegotiation deny:

```

$ openssl s_client -connect 192.168.2.100:443
CONNECTED(00000003)
depth=1 /C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
Authority/CN=support/emailAddress=support@fortinet.com
verify error:num=19:self signed certificate in certificate chain
verify return:0

Certificate chain
 0
s:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Fortigate/CN=FW80CM390
9604325/emailAddress=support@fortinet.com
 i:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
 Authority/CN=support/emailAddress=support@fortinet.com
 1 s:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
 Authority/CN=support/emailAddress=support@fortinet.com
 i:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
 Authority/CN=support/emailAddress=support@fortinet.com

Server certificate
-----BEGIN CERTIFICATE-----
---certificate not shown---
-----END CERTIFICATE-----

subject=/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Fortigate/CN=FW8
0CM3909604325/emailAddress=support@fortinet.com
 issuer=/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
 Authority/CN=support/emailAddress=support@fortinet.com

No client certificate CA names sent

SSL handshake has read 2370 bytes and written 316 bytes

New, TLSv1/SSLv3, Cipher is DHE-RSA-AES256-SHA
Server public key is 1024 bit
Compression: NONE
Expansion: NONE
SSL-Session:
 Protocol : TLSv1
 Cipher : DHE-RSA-AES256-SHA
 Session-ID:
 8253331D266DDE38E4D8A04AFCA9CBDED5B1134932CE1718EED6469C1FBC7474
 Session-ID-ctx:
 Master-Key:

ED05A3EF168AF2D06A486362FE91F1D6CAA55CEFC38A3C36FB8BD74236BF2657D4701B
6C1456CEB5BB5EFAA7619EF12D
 Key-Arg : None
 Start Time: 1264452957
 Timeout : 300 (sec)
 Verify return code: 19 (self signed certificate in certificate
 chain)

GET /main.c HTTP/1.0
R
RENEGOTIATING

```

```
19916:error:1409E0E5:SSL routines:SSL3_WRITE_BYTES:ssl handshake
failure:s3_pkt.c:530:
```

Use the following command to check the SSL stats to see that the renegotiations blocked counter is now 1:

```
firewall vip virtual-server stats ssl
ssl
 client
 connections total 0 active 0 max 0
 handshakes total 4 active 0 max 0 completed 4 abbreviated 0
 session states total 4 active 4 max 4
 cipher-suite failures 0
 embryonics total 0 active 0 max 0 terminated 0
 renegotiations blocked 1
 server
 connections total 0 active 0 max 0
 handshakes total 3 active 0 max 0 completed 2 abbreviated 1
 session states total 1 active 1 max 1
 cipher-suite failures 0
 internal error 0
 bad handshake length 0
 bad change cipher spec length 0
 pubkey too big 0
 persistence
 find 0 found 0 clash 0 addr 0 error 0
```

If the virtual server debug log is examined (diag debug appl vs -1) then at the point the re-negotiation is blocked there is a log:

```
vs ssl 12 handshake recv ClientHello
vs ssl 12 handshake recv 1
(0100005403014b5e056c7f573a563bebe0258c3254bbaff7046a461164f34f94f4f3d
019c41800002600390038003500160013000a00330032002f000500040015001200090
0140011000800060003020100000400230000)
vs ssl 12 client renegotiation attempted rejected, abort
vs ssl 12 closing 0 up
vs src 12 close 0 in
vs src 12 error closing
vs dst 14 error closing
vs dst 14 closed
vs ssl 14 close
vs sock 14 free
vs src 12 closed
vs ssl 12 close
vs sock 12 free
```

## IP, TCP, and UDP load balancing

You can load balance all IP, TCP or UDP sessions accepted by the security policy that includes a load balancing virtual server with the type set to IP, TCP, or UDP. Traffic with destination IP and port that matches the virtual server IP and port is load balanced. For these protocol-level load balancing virtual servers you can select a load balance method and add real servers and health checking. However, you can't configure persistence, HTTP multiplexing and SSL offloading.

# Load balancing configuration examples

This chapter includes the following examples:


- Example: HTTP load balancing to three real web servers
- Example: Basic IP load balancing configuration
- Example: Adding a server load balance port forwarding virtual IP
- Example: Weighted load balancing configuration
- Example: HTTP and HTTPS persistence configuration

## Example: HTTP load balancing to three real web servers

In this example, the virtual web server IP address 192.168.37.4 on the Internet, is mapped to three real web servers connected to the FortiGate unit dmz1 interface. The real servers have IP addresses 10.10.123.42, 10.10.123.43, and 10.10.123.44. The virtual server uses the *First Alive* load balancing method. The configuration also includes an HTTP health check monitor that includes a URL used by the FortiGate unit for get requests to monitor the health of the real servers.

Connections to the virtual web server at IP address 192.168.37.4 from the Internet are translated and load balanced to the real servers by the FortiGate unit. First alive load balancing directs all sessions to the first real server. The computers on the Internet are unaware of this translation and load balancing and see a single virtual server at IP address 192.168.37.4 rather than the three real servers behind the FortiGate unit.

**Figure 292:**Virtual server configuration example



## Web-based manager configuration

Use the following procedures to configure this load balancing setup from the web-based manager.

### To add an HTTP health check monitor

In this example, the HTTP health check monitor includes the *URL* “/index.html” and the *Matched Phrase* “Fortinet products”.

1. Go to *Firewall Objects > Load Balance > Health Check*.
2. Select *Create New*.
3. Add an HTTP health check monitor that sends get requests to `http://<real_server_IP_address>/index.html` and searches the returned web page for the phrase “Fortinet products”.

<b>Name</b>	HTTP_health_chk_1
<b>Type</b>	HTTP
<b>Port</b>	80
<b>URL</b>	/index.html
<b>Matched Content</b>	Fortinet products
<b>Interval</b>	10 seconds
<b>Timeout</b>	2 seconds
<b>Retry</b>	3

4. Select *OK*.

### To add the HTTP virtual server

1. Go to *Firewall Objects > Load Balance > Virtual Server*.
2. Select *Create New*.
3. Add an HTTP virtual server that allows users on the Internet to connect to the real servers on the internal network. In this example, the FortiGate wan1 interface is connected to the Internet.

<b>Name</b>	Load_Bal_VS1
<b>Type</b>	HTTP
<b>Interface</b>	wan1



<b>Virtual Server IP</b>	192.168.37.4  The public IP address of the web server.  The virtual server IP address is usually a static IP address obtained from your ISP for your web server. This address must be a unique IP address that is not used by another host and cannot be the same as the IP address of the external interface the virtual IP will be using. However, the external IP address must be routed to the selected interface. The virtual IP address and the external IP address can be on different subnets. When you add the virtual IP, the external interface responds to ARP requests for the external IP address.
<b>Virtual Server Port</b>	80
<b>Load Balance Method</b>	First Alive
<b>Persistence</b>	HTTP cookie
<b>HTTP Multiplexing</b>	Select.  The FortiGate unit multiplexes multiple client into a few connections between the FortiGate unit and each real HTTP server. This can improve performance by reducing server overhead associated with establishing multiple connections.
<b>Preserve Client IP</b>	Select  The FortiGate unit preserves the IP address of the client in the X-Forwarded-For HTTP header.
<b>Health Check</b>	Move the HTTP_health_chk_1 health check monitor to the <i>Selected</i> list.

4. Select *OK*.

**To add the real servers and associate them with the virtual server**

1. Go to *Firewall Objects > Load Balance > Real Server*.
2. Select *Create New*.

3. Configure three real servers that include the virtual server Load\_Bal\_VS1. Each real server must include the IP address of a real server on the internal network.

Configuration for the first real server.

---

<b>Virtual Server</b>	Load_Bal_VS1
<b>IP Address</b>	10.10.10.42
<b>Port</b>	80
<b>Weight</b>	Cannot be configured because the virtual server does not include weighted load balancing.
<b>Maximum Connections</b>	0

Setting *Maximum Connections* to 0 means the FortiGate unit does not limit the number of connections to the real server. Since the virtual server uses *First Alive* load balancing you may want to limit the number of connections to each real server to limit the traffic received by each server. In this example, the *Maximum Connections* is initially set to 0 but can be adjusted later if the real servers are getting too much traffic.

---

Configuration for the second real server.

---

<b>Virtual Server</b>	Load_Bal_VS1
<b>IP Address</b>	10.10.10.43
<b>Port</b>	80
<b>Weight</b>	Cannot be configured because the virtual server does not include weighted load balancing.
<b>Maximum Connections</b>	0

Setting *Maximum Connections* to 0 means the FortiGate unit does not limit the number of connections to the real server. Since the virtual server uses *First Alive* load balancing you may want to limit the number of connections to each real server to limit the traffic received by each server. In this example, the *Maximum Connections* is initially set to 0 but can be adjusted later if the real servers are getting too much traffic.

---

Configuration for the third real server.

---

<b>Virtual Server</b>	Load_Bal_VS1
<b>IP Address</b>	10.10.10.44
<b>Port</b>	80

---

---

<b>Weight</b>	Cannot be configured because the virtual server does not include weighted load balancing.
---------------	-------------------------------------------------------------------------------------------

---

**Maximum Connections 0**

Setting *Maximum Connections* to 0 means the FortiGate unit does not limit the number of connections to the real server. Since the virtual server uses *First Alive* load balancing you may want to limit the number of connections to each real server to limit the traffic received by each server. In this example, the *Maximum Connections* is initially set to 0 but can be adjusted later if the real servers are getting too much traffic.

---

**To add the virtual server to a security policy**

Add a wan1 to dmz1 security policy that uses the virtual server so that when users on the Internet attempt to connect to the web server's IP address, packets pass through the FortiGate unit from the wan1 interface to the dmz1 interface. The virtual IP translates the destination address of these packets from the virtual server IP address to the real server IP addresses.

1. Go to *Policy > Policy > Policy*.
2. Select *Create New*.
3. Configure the security policy:

---

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	wan1
<b>Source Address</b>	all (or a more specific address)
<b>Outgoing Interface</b>	dmz1
<b>Destination Address</b>	Load_Bal_VS1
<b>Schedule</b>	always
<b>Service</b>	HTTP
<b>Action</b>	ACCEPT
<b>Log Allowed Traffic</b>	Select to log virtual server traffic
<b>Enable NAT</b>	Select this option and select <i>Use Destination Interface Address</i> .

---

4. Select other security policy options as required.
5. Select *OK*.

## CLI configuration

Use the following procedure to configure this load balancing setup from the CLI.

## To configure HTTP load balancing

1. Use the following command to add an HTTP health check monitor that sends get requests to `http://<real_server_IP_address>/index.html` and searches the returned web page for the phrase "Fortinet products".

```
config firewall ldb-monitor
 edit HTTP_health_chk_1
 set type http
 set port 80
 set http-get /index.html
 set http-match "Fortinet products"
 set interval 10
 set timeout 2
 set retry 3
 end
```

2. Use the following command to add an HTTP virtual server that allows users on the Internet to connect to the real servers on the internal network. In this example, the FortiGate wan1 interface is connected to the Internet.

```
config firewall vip
 edit Load-Bal_VS1
 set type server-load-balance
 set server-type http
 set ldb-method first-alive
 set http-multiplex enable
 set http-ip-header enable
 set extip 192.168.37.4
 set extintf wan1
 set extport 80
 set persistence http-cookie
 set monitor HTTP_health_chk_1
 config realservers
 edit 1
 set ip 10.10.10.42
 set port 80
 next
 edit 2
 set ip 10.10.10.43
 set port 80
 next
 edit 3
 set ip 10.10.10.44
 set port 80
 end
 end
```

3. Use the following command to add a security policy that includes the load balance virtual server as the destination address.

```
config firewall policy
 edit 0
 set srcintf wan1
 set srcaddr all
 set dstintf dmz1
 set dstaddr Load-Bal_VS1
 set action accept
 set schedule always
 set service ALL
 set nat enable
 end
```

Configure other security policy settings as required.

## Example: Basic IP load balancing configuration

This example shows how to add a server load balancing virtual IP that load balances all traffic among 3 real servers. In the example the Internet is connected to `port2` and the virtual IP address of the virtual server is `192.168.20.20`. The load balancing method is `weighted`. The IP addresses of the real servers are `10.10.10.1`, `10.10.10.2`, and `10.10.10.3`. The weights for the real servers are 1, 2, and 3. The default weight is 1 and does not have to be changed for the first real server.

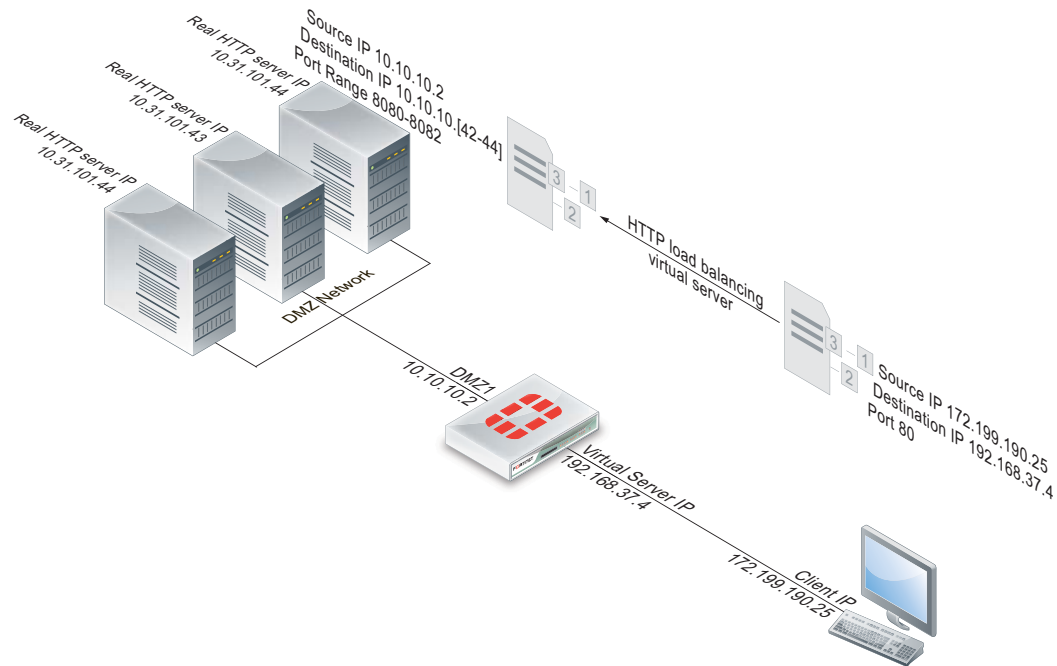
```
config firewall vip
 edit All_Load_Balance
 set type server-load-balance
 set server-type ip
 set extintf port2
 set extip 192.168.20.20
 set ldb-method weighted
 config realservers
 edit 1
 set ip 10.10.10.1
 next
 edit 2
 set ip 10.10.10.2
 set weight 2
 next
 edit 3
 set ip 10.10.10.3
 set weight 3
 end
 end
```

## Example: Adding a server load balance port forwarding virtual IP

This example is the same as the example described in [“Example: HTTP load balancing to three real web servers” on page 1911](#) except that each real server accepts HTTP connections on a

different port number. The first real server accepts connections on port 8080, the second on port 8081, and the third on 8082.

**Figure 293:**Server load balance virtual IP port forwarding



To complete this configuration, all of the steps would be the same as in “[Example: HTTP load balancing to three real web servers](#)” on page 1911 except for configuring the real servers.

**To add the real servers and associate them with the virtual server**

Use the following steps to configure the FortiGate unit to port forward HTTP packets to the three real servers on ports 8080, 8081, and 8082.

1. Go to *Firewall Objects > Load Balance > Real Server*.
2. Select *Create New*.
3. Configure three real servers that include the virtual server Load\_Bal\_VS1. Each real server must include the IP address of a real server on the internal network and have a different port number.

Configuration for the first real server.

<b>Virtual Server</b>	Load_Bal_VS1
<b>IP</b>	10.10.10.42
<b>Port</b>	8080
<b>Weight</b>	Cannot be configured because the virtual server does not include weighted load balancing.
<b>Maximum Connections</b>	0

Configuration for the second real server.

<b>Virtual Server</b>	Load_Bal_VS1
<b>IP</b>	10.10.10.43
<b>Port</b>	8081
<b>Weight</b>	Cannot be configured because the virtual server does not include weighted load balancing.
<b>Maximum Connections</b>	0

Configuration for the third real server.

<b>Virtual Server</b>	Load_Bal_VS1
<b>IP</b>	10.10.10.44
<b>Port</b>	8082
<b>Weight</b>	Cannot be configured because the virtual server does not include weighted load balancing.
<b>Maximum Connections</b>	0

## Example: Weighted load balancing configuration

This example shows how to using firewall load balancing to load balances all traffic among 3 real servers. In the example the Internet is connected to `port2` and the virtual IP address of the virtual server is 192.168.20.20. The load balancing method is `weighted`. The IP addresses of the real servers are 10.10.10.1, 10.10.10.2, and 10.10.10.3. The weights for the real servers are 1, 2, and 3.

This configuration does not include an health check monitor.

### Web-based manager configuration

Use the following procedures to configure this load balancing setup from the web-based manager.

#### To add the HTTP virtual server

1. Go to *Firewall Objects > Load Balance > Virtual Server*.
2. Select *Create New*.

3. Add an IP virtual server that allows users on the Internet to connect to the real servers on the internal network. In this example, the FortiGate port2 interface is connected to the Internet.

<b>Name</b>	HTTP_weghted_LB
<b>Type</b>	IP
<b>Interface</b>	port2
<b>Virtual Server IP</b>	192.168.20.20
<b>Load Balance Method</b>	Weighted

All other virtual server settings are not required or cannot be changed.

4. Select *OK*.

**To add the real servers and associate them with the virtual server**

1. Go to *Firewall Objects > Load Balance > Real Server*.
2. Select *Create New*.
3. Configure three real servers that include the virtual server *All\_Load\_Balance*. Because the *Load Balancing Method* is *Weighted*, each real server includes a weight. Servers with a greater weight receive a greater proportion of forwarded connections,  
Configuration for the first real server.

<b>Virtual Server</b>	HTTP_weghted_LB
<b>IP Address</b>	10.10.10.1
<b>Port</b>	Cannot be configured because the virtual server is an IP server.
<b>Weight</b>	1
<b>Maximum Connections</b>	0  Setting <i>Maximum Connections</i> to 0 means the FortiGate unit does not limit the number of connections to the real server. Since the virtual server uses <i>First Alive</i> load balancing you may want to limit the number of connections to each real server to limit the traffic received by each server. In this example, the <i>Maximum Connections</i> is initially set to 0 but can be adjusted later if the real servers are getting too much traffic.



Configuration for the second real server.

<b>Virtual Server</b>	HTTP_weghted_LB
<b>IP Address</b>	10.10.10.2
<b>Port</b>	Cannot be configured because the virtual server is an IP server.
<b>Weight</b>	2
<b>Maximum Connections</b>	0
	Setting <i>Maximum Connections</i> to 0 means the FortiGate unit does not limit the number of connections to the real server. Since the virtual server uses <i>First Alive</i> load balancing you may want to limit the number of connections to each real server to limit the traffic received by each server. In this example, the <i>Maximum Connections</i> is initially set to 0 but can be adjusted later if the real servers are getting too much traffic.

Configuration for the third real server.

<b>Virtual Server</b>	HTTP_weghted_LB
<b>IP Address</b>	10.10.10.3
<b>Port</b>	Cannot be configured because the virtual server is an IP server.
<b>Weight</b>	3
<b>Maximum Connections</b>	0
	Setting <i>Maximum Connections</i> to 0 means the FortiGate unit does not limit the number of connections to the real server. Since the virtual server uses <i>First Alive</i> load balancing you may want to limit the number of connections to each real server to limit the traffic received by each server. In this example, the <i>Maximum Connections</i> is initially set to 0 but can be adjusted later if the real servers are getting too much traffic.

### To add the virtual server to a security policy

Add a prot2 to port1 security policy that uses the virtual server so that when users on the Internet attempt to connect to the web server's IP address, packets pass through the FortiGate unit from the wan1 interface to the dmz1 interface. The virtual IP translates the destination address of these packets from the virtual server IP address to the real server IP addresses.

1. Go to *Policy > Policy > Policy*.
2. Select *Create New*.
3. Configure the security policy:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	port2
<b>Source Address</b>	all (or a more specific address)

<b>Outgoing Interface</b>	port1
<b>Destination Address</b>	HTTP_weghted_LB
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Select this option and select <i>Use Destination Interface Address</i> .

4. Select other security policy options as required.
5. Select *OK*.

## CLI configuration

Load balancing is configured from the CLI using the `config firewall vip` command and by setting `type` to `server-load-balance`. The default weight is 1 and does not have to be changed for the first real server.

Use the following command to add the virtual server and the three weighted real servers.

```
config firewall vip
 edit HTTP_weghted_LB
 set type server-load-balance
 set server-type ip
 set extintf port2
 set extip 192.168.20.20
 set ldb-method weighted
 config realservers
 edit 1
 set ip 10.10.10.1
 next
 edit 2
 set ip 10.10.10.2
 set weight 2
 next
 edit 3
 set ip 10.10.10.3
 set weight 3
 end
 end
```

## Example: HTTP and HTTPS persistence configuration

This example shows how to add a virtual server named *Http\_Load\_Balance* that load balances HTTP traffic using port 80 and a second virtual server named *Https\_Load\_Balance* that load balances HTTPS traffic using port 443. The Internet is connected to port2 and the virtual IP address of the virtual server is 192.168.20.20. Both server load balancing virtual IPs load balance sessions to the same three real servers with IP addresses 10.10.10.2, 10.10.10.2, and 10.10.10.3. The real servers provide HTTP and HTTPS services.

For both virtual servers, persistence is set to *HTTP Cookie* to enable HTTP cookie persistence.

**To add the HTTP and HTTPS virtual servers**

1. Go to *Firewall Objects > Load Balance > Virtual Server*.
2. Add the HTTP virtual server that includes HTTP Cookie persistence.

<b>Name</b>	HTTP_Load_Balance
<b>Type</b>	HTTP
<b>Interface</b>	port2
<b>Virtual Server IP</b>	192.168.20.20
<b>Virtual Server Port</b>	80  In this example the virtual server uses port 8080 for HTTP sessions instead of port 80.
<b>Load Balance Method</b>	Static
<b>Persistence</b>	HTTP cookie

3. Select *OK*.
4. Select *Create New*.
5. Add the HTTPS virtual server that also includes HTTP Cookie persistence.

<b>Name</b>	HTTPS_Load_Balance
<b>Type</b>	HTTPS
<b>Interface</b>	port2
<b>Virtual Server IP</b>	192.168.20.20
<b>Virtual Server Port</b>	443
<b>Load Balance Method</b>	Static
<b>Persistence</b>	HTTP cookie

6. Select *OK*.

**To add the real servers and associate them with the virtual servers**

1. Go to *Firewall Objects > Load Balance > Real Server*.
2. Select *Create New*.

3. Configure three real servers for HTTP that include the virtual server HTTP\_Load\_Balance. Configuration for the first HTTP real server.

---

<b>Virtual Server</b>	HTTP_Load_Balance
<b>IP Address</b>	10.10.10.1
<b>Port</b>	80
<b>Weight</b>	Cannot be configured because the virtual server does not include weighted load balancing.
<b>Maximum Connections</b>	0

---

Configuration for the second HTTP real server.

---

<b>Virtual Server</b>	HTTP_Load_Balance
<b>IP Address</b>	10.10.10.2
<b>Port</b>	80
<b>Weight</b>	Cannot be configured because the virtual server does not include weighted load balancing.
<b>Maximum Connections</b>	0

---

Configuration for the third HTTP real server.

---

<b>Virtual Server</b>	HTTP_Load_Balance
<b>IP Address</b>	10.10.10.3
<b>Port</b>	80
<b>Weight</b>	Cannot be configured because the virtual server does not include weighted load balancing.
<b>Maximum Connections</b>	0

---

- Configure three real servers for HTTPS that include the virtual server HTTPS\_Load\_Balance. Configuration for the first HTTPS real server.

<b>Virtual Server</b>	HTTP_Load_Balance
<b>IP Address</b>	10.10.10.1
<b>Port</b>	443
<b>Weight</b>	Cannot be configured because the virtual server does not include weighted load balancing.
<b>Maximum Connections</b>	0

Configuration for the second HTTPS real server.

<b>Virtual Server</b>	HTTP_Load_Balance
<b>IP Address</b>	10.10.10.2
<b>Port</b>	443
<b>Weight</b>	Cannot be configured because the virtual server does not include weighted load balancing.
<b>Maximum Connections</b>	0

Configuration for the third HTTPS real server.

<b>Virtual Server</b>	HTTPS_Load_Balance
<b>IP Address</b>	10.10.10.3
<b>Port</b>	443
<b>Weight</b>	Cannot be configured because the virtual server does not include weighted load balancing.
<b>Maximum Connections</b>	0

#### To add the virtual servers to security policies

Add a port2 to port1 security policy that uses the virtual server so that when users on the Internet attempt to connect to the web server's IP address, packets pass through the FortiGate unit from the wan1 interface to the dmz1 interface. The virtual IP translates the destination address of these packets from the virtual server IP address to the real server IP addresses.

- Go to *Policy > Policy > Policy*.
- Select *Create New*.
- Configure the HTTP security policy:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	port2

<b>Source Address</b>	all
<b>Outgoing Interface</b>	port1
<b>Destination Address</b>	HTTP_Load_Balance
<b>Schedule</b>	always
<b>Service</b>	HTTP
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Select this option and select <i>Use Destination Interface Address</i> .

4. Select other security policy options as required.
5. Select *OK*.
6. Select *Create New*.
7. Configure the HTTP security policy:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	port2
<b>Source Address</b>	all
<b>Outgoing Interface</b>	port1
<b>Destination Address</b>	HTTPS_Load_Balance
<b>Schedule</b>	always
<b>Service</b>	HTTPS
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Select this option and select <i>Use Destination Interface Address</i> .

8. Select other security policy options as required.
9. Select *OK*.

## CLI configuration: adding persistence for a specific domain

Load balancing is configured from the CLI using the `config firewall vip` command and by setting `type` to `server-load-balance`.

For the CLI configuration, both virtual servers include setting `http-cookie-domain` to `.example.org` because HTTP cookie persistence is just required for the `example.org` domain.

First, the configuration for the HTTP virtual IP:

```
config firewall vip
 edit HTTP_Load_Balance
 set type server-load-balance
 set server-type http
 set extport 8080
 set extintf port2
 set extip 192.168.20.20
 set persistence http-cookie
 set http-cookie-domain .example.org
 config realservers
 edit 1
 set ip 10.10.10.1
 next
 edit 2
 set ip 10.10.10.2
 next
 edit 3
 set ip 10.10.10.3
 end
 end
 end
```

Second, the configuration for the HTTPS virtual IP. In this configuration you don't have to set extport to 443 because extport is automatically set to 443 when server-type is set to https.

```
config firewall vip
 edit HTTPS_Load_Balance
 set type server-load-balance
 set server-type https
 set extport 443
 set extintf port2
 set extip 192.168.20.20
 set persistence http-cookie
 set http-cookie-domain .example.org
 config realservers
 edit 1
 set ip 10.10.10.1
 next
 edit 2
 set ip 10.10.10.2
 next
 edit 3
 set ip 10.10.10.3
 end
 end
 end
```

# Chapter 13 Logging and Reporting



This FortiOS Handbook chapter contains the following sections:

[Logging and reporting overview](#) provides general information about logging. We recommend that you begin with this chapter as it contains information for both beginners and advanced users as well. It contains an explanation of log messages, files, and devices, and an overview of the Reporting functions.

[Logging and reporting for small networks](#) provides an overview of setting up a small network for logging, with a look at a possible setup with a backup solution and a customized report.

[Logging and reporting for large networks](#) provides an overview of setting up a larger, enterprise-level network, with configuration of multiple FortiGate units, multiple FortiAnalyzer units as a backup solution, and a sample procedure for creating a more intensive and broad report to suit the larger network.

[Advanced logging](#) provides a series of separate tutorials for possible tasks and procedures an advanced user may want to undertake with their FortiGate-powered network. It contains explanations of advanced backup, logging, and report solutions.

[Troubleshooting and logging](#) provides a short overview of how log messages can be used to identify and solve problems within the network, how to identify and solve logging database issues, and how to solve connection issues between FortiGate and FortiAnalyzer units.

[Appendix: FortiGate report charts](#) provides a detailed list of the variety of charts available for use in FortiGate reporting.

# Logging and reporting overview

Logging and reporting in FortiOS can help you in determining what is happening on your network, as well as informing you of certain network activity, such as detection of a virus or IPsec VPN tunnel errors. Logging and reporting go hand in hand, and can become a valuable tool for information as well as helping to show others the activity that is happening on the network.

This section explains logging and reporting features that are available in FortiOS, and how they can be used to help you manage or troubleshoot issues. This includes how the FortiGate unit records logs, what a log message is, and what the log database is.

The following topics are included in this section:

- [What is logging?](#)
- [Log messages](#)
- [Log files and types](#)
- [Log database and datasets](#)
- [Notifications about network activity](#)
- [Log devices](#)
- [Reports](#)
- [Best Practices: Log management](#)



In FortiOS 5.0, logs have changed and as a result, most log data is preserved and will be updated to the new log format when 5.0 is first installed. However, this log update may take a while, as this includes updating the report charts and datasets. A warning will appear to let you know that this process is still in progress even after the upgrade has finished.

---

## What is logging?

Logging records the traffic passing through the FortiGate unit to your network and what action the FortiGate unit took during its scanning process of the traffic. This recorded information is called a log message.

After a log message is recorded, it is stored within a log file which is then stored on a log device. A log device is a central storage location for log messages. The FortiGate unit supports several log devices, such as the FortiCloud service, or a FortiAnalyzer unit. A FortiGate unit's system memory and local disk can also be configured to store logs, and because of this, are also considered log devices.



You must subscribe to FortiCloud before you will be able to configure the FortiGate unit to send logs to a FortiCloud server.

---

When the recorded activity needs to be read in a more human way, the FortiGate unit can generate a Report. A report gathers all the log information that is needed for the report, and presents it in a graphical format, with customizable design and automatically generated charts. Reports can be used to present a graphical representation of what is going on in the network.

Reports can also be generated on a FortiAnalyzer unit; if you want to generate reports on a FortiAnalyzer, see the [FortiAnalyzer Setup and Administration Guide](#) to help you create and generate those reports.

## How the FortiGate unit records log messages

The FortiGate unit records log messages in a specific order, storing them on a log device. The order of how the FortiGate unit records log messages is as follows:

1. Incoming traffic is scanned.
2. During the scanning process, the FortiGate unit performs necessary actions, and simultaneously records the actions and results.
3. Log messages are sent to the log device.

### Example: How the FortiGate unit records a DLP event

1. The FortiGate unit receives incoming traffic and scans for any matches associated within its firewall policies containing a DLP sensor.
2. A match is found; the DLP sensor, `dlp_sensor`, had a rule within it called All-HTTP with the action Exempt applied to the rule. The sensor also has Enable Logging selected, which indicates to the FortiGate unit that the activity should be recorded and placed in the DLP log file.
3. The FortiGate unit exempts the match, and places the recorded activity (the log message) within the DLP log file.
4. According to the log settings that were configured, logs are stored on the FortiGate unit's local hard drive. The FortiGate unit places the DLP log file on the local hard drive.

## FortiOS features available for logging

Logs record FortiGate activity, providing detailed information about what is happening on your network. This recorded activity is found in log files, which are stored on a log device. However, logging FortiGate activity requires configuring certain settings so that the FortiGate unit can record the activity. These settings are often referred to as log settings, and are found in most Security Features, such as profiles, but also the event log settings.

Log settings provide the information that the FortiGate unit needs so that it knows what activities to record. This topic explains what activity each log file records, as well as additional information about the log file, which will help you determine what FortiGate activity the FortiGate unit should record.

## Traffic

Traffic logs record the traffic that is flowing through your FortiGate unit. Since traffic needs firewall policies to properly flow through the unit, this type of logging is also referred to as firewall policy logging. Firewall policies control all traffic that attempts to pass through the FortiGate unit, between FortiGate interfaces, zones and VLAN sub-interfaces.

Logging traffic works in the following way:

- firewall policy has logging enabled on it (Log Allowed Traffic or Log Violation Traffic)
- packet comes into an inbound interface
- a possible log packet is sent regarding a match in the firewall policy, such as URL filter
- traffic log packet is sent, per firewall policy
- packet passes and is sent out an interface

Traffic log messages are stored in the traffic log file. Traffic logs can be stored any log device, even system memory.

All Security Feature-related logs are now tracked within the Traffic logs, as of FortiOS 5.0, so all forward traffic can be searched in one place, such as if you are looking to see all activity from a particular address, security feature or traffic.

The *Security Feature Log* section has been removed from the default interface, but if your device is registered with FortiCloud, *Security Log* may still appear in the web interface, and will list the security feature traffic separately, as FortiCloud tracks it separately from traffic. If you would like to be able to view the security feature logs both within and isolated from the Traffic logs, registering with FortiCloud is necessary.

If you have enabled and configured WAN Optimization, you can enable logging of this activity in the CLI using the `config wanopt setting` command. These logs contain information about WAN Optimization activity and are found in the traffic log file. When configuring logging of this activity, you must also enable logging within the security policy itself, so that the activity is properly recorded.

## Other Traffic

The traffic log also records interface traffic logging, which is referred to as other traffic. Other traffic is enabled only in the CLI. When enabled, the FortiGate unit records traffic activity on interfaces as well as firewall policies. Logging other traffic puts a significant system load on the FortiGate unit and should be used only when necessary.

Logging other traffic works in the following way:

- firewall policy has logging enabled on it (Log Allowed Traffic or Log Violation Traffic) and other-traffic
- packet comes into an interface
- interface log packet is sent to the traffic log that is enabled on that particular interface
- possible log packet is sent regarding a match in the firewall policy, such as URL filter
- interface log packet is sent to the traffic log if enabled on that particular interface
- packet passes and is sent out an interface
- interface log packet is sent to traffic (if enabled) on that particular interface

## Event

The event log records administration management as well as FortiGate system activity, such as when a configuration has changed, admin login, or high availability (HA) events occur. Event logs are an important log file to record because they record FortiGate system activity, which provides valuable information about how your FortiGate unit is performing.

Event logs help you in the following ways:

- keeping track of configuration setting changes
- IPsec negotiation, SSL VPN and tunnel activity
- quarantine events, such as banned users
- system performance
- HA events and alerts
- firewall authentication events
- wireless events on models with WiFi capabilities
- activities concerning modem and internet protocols L2TP, PPP and PPPoE
- VIP activities
- AMC disk's bypass mode
- VoIP activities that include SIP and SCCP protocols.

The FortiGate unit records event logs only when events are enabled.

## Traffic Shaping

Traffic shaping, per-IP traffic shaping and reverse direction traffic shaping settings can be applied to a firewall policy, appearing within the traffic log messages.

By enabling this feature, you can see what traffic shaping, per-IP traffic shaping and reverse direction traffic shaping settings are being used.

## Data Leak Prevention

Data Leak Prevention logs, or DLP logs, provide valuable information about the sensitive data trying to get through to your network as well as any unwanted data trying to get into your network. The DLP rules within a DLP sensor can log the following traffic types:

- email (SMTP, POP3 or IMAP; if SSL content SMTPS, POP3S, and IMAPS)
- HTTP
- HTTPS
- FTP
- NNTP
- IM

A DLP sensor must have log settings enabled for each DLP rule and compound rule, as well as applied to a firewall policy so that the FortiGate unit records this type of activity. A DLP sensor can also contain archiving options, which these logs are then archived to the log device.

## NAC Quarantine

Within the DLP sensor, there is an option for enabling NAC Quarantine. The NAC Quarantine option allows the FortiGate unit to record details of DLP operation that involve the ban and quarantine actions, and sends these to the event log file. The NAC Quarantine option must also be enabled within the Event Log settings. When enabling NAC quarantine within a DLP Sensor, you must enable this in the CLI because it is a CLI-only command.

## Media Access Control (MAC) Address

MAC address logs provide information about MAC addresses that the FortiGate unit sees on the network as well as those removed from the network. These log messages are stored in the

event log (as subtype network; you can view these log messages in Log & Report > Event > Event Log) and are, by default, disabled in the CLI. You can enable logging MAC addresses using the following command syntax:

```
config log setting
 set neighbor-event enable
end
```

When enabled, a new log message is recorded every time a MAC address entry is added to the ARP table, and also when a MAC address is removed as well. A MAC address log message is also recorded when MAC addresses are connected to the local switch, or from a FortiAP or FortiSwitch unit.

## Application control

Application control logs provide detailed information about the traffic that internet applications such as Skype are generating. The application control feature controls the flow of traffic from a specific application, and the FortiGate unit examines this traffic for signatures that the application generates.

The log messages that are recorded provide information such as the type of application being used (such as P2P software), and what type of action the FortiGate unit took. These log messages can also help you to determine the top ten applications that are being used on your network. This feature is called application control monitoring and you can view the information from a widget on the Executive Summary page.

The application control list that is used must have *Enabled Logging* selected within the list, as well as logging enabled within each application entry. Each application entry can also have packet logging enabled. Packet logging for application control records the packet when an application type is identified, similar to IPS packet logging.

Logging of application control activity can only be recorded when an application control list is applied to a firewall policy, regardless of whether or not logging is enabled within the application control list.

## Antivirus

Antivirus logs are recorded when, during the antivirus scanning process, the FortiGate unit finds a match within the antivirus profile, which includes the presence of a virus or grayware signature. Antivirus logs provide a way to understand what viruses are trying to get in, as well as additional information about the virus itself, without having to go to the FortiGuard Center and do a search for the detected virus. The link is provided within the log message itself.

These logs provide valuable information such as:

- the name of the detected virus
- the name of the oversized file or infected file
- the action the FortiGate unit took, for example, a file was blocked
- URL link to the FortiGuard Center which gives detailed information about the virus itself

The antivirus profile must have log settings enabled within it so that the FortiGate unit can record this activity, as well as having the antivirus profile applied to a firewall policy.

## Web Filter

Web filter logs record HTTP traffic activity. These log messages provide valuable and detailed information about this particular traffic activity on your network. Web filtering activity is important to log because it can inform you about:

- what types of web sites employees are accessing
- users attempting to access banned web sites and how often this occurs
- network congestion due to employees accessing the Internet at the same time
- web-based threats resulting from users visiting non-business-related web sites

Web Filter logs are an effective tool to help you determine if you need to update your web filtering settings within a web filter profile due to unforeseen threats or network congestion. These logs also inform you about web filtering quotas that have been configured for filtering HTTP traffic.

You must configure logging settings within the web filter profile and apply the filter to a firewall policy so that the FortiGate unit can record the activity.

## IPS (attack)

IPS logs, also referred to as attack logs, record attacks that occurred against your network. Attack logs contain detailed information about whether the FortiGate unit protected the network using anomaly-based defense settings or signature-based defense settings, as well as what the attack was.

The IPS or attack log file is especially useful because the log messages that are recorded contain a link to the FortiGuard Center, where you can find more information about the attack. This is similar to antivirus logs, where a link to the FortiGuard Center is provided as well that informs you of the virus that was detected by the FortiGate unit.

An IPS sensor with log settings enabled must be applied to a firewall policy so that the FortiGate unit can record the activity.

## Packet logs

When you enable packet logging within an IPS signature override or filter, the FortiGate unit examines network packets, and if a match is found, saves them to the attack log. Packet logging is designed to be used as a diagnostic tool that can focus on a narrow scope of diagnostics, rather than a log that informs you of what is occurring on your network.

You should use caution when enabling packet logging, especially within IPS filters. Filter configuration that contains thousands of signatures could potentially cause a flood of saved packets, which would take up a lot of storage space on the log device. It would also take a great deal of time to sort through all the log messages, as well as consume considerable system resources to process.

You can archive packets, but you must enable this option on the Log Settings page. If your log configuration includes multiple FortiAnalyzer units, packet logs are only sent to the primary (first) FortiAnalyzer unit. Sending packet logs to the other FortiAnalyzer units is not supported.

## Email filter

Email filter logs, also referred to as spam filter logs, records information regarding the content within email messages. For example, within an email filter profile, a match is found that finds the email message to be considered spam.

Email filter logs are recorded when the FortiGate unit finds a match within the email filter profile and logging settings are enabled within the profile.

## Archives (DLP)

Recording DLP logs for network use is called DLP archiving. The DLP engine examines email, FTP, IM, NNTP, and web traffic. Archived logs are usually saved for historical use and can be accessed at any time. IPS packet logs can also be archived, within the Log Settings page.

You can start with the two default DLP sensors that have been configured specifically for archiving log data, Content\_Archive and Content\_Summary. They are available in *Security Profiles > Data Leak Prevention > Sensors*. Content\_Archive provides full content archiving, while Content\_Summary provides summary archiving. For more information about how to configure DLP sensors, see the Security Features chapter of the FortiOS Handbook.

You must enable the archiving to record log archives. Logs are not archived unless enabled, regardless of whether or not the DLP sensor for archiving is applied to the firewall policy.

## Network scan

Network scan logs are recorded when a scheduled scan of the network occurs. These log messages provide detailed information about the network's vulnerabilities regarding software, as well as the discovery of any further vulnerabilities.

A scheduled scan must be configured and logging enabled within the Event Log settings, for the FortiGate unit to record these log messages.

## Log messages

Log messages are recorded by the FortiGate unit, giving you detailed information about the network activity. Each log message has a unique number that helps identify it, as well as containing fields; these fields, often called log fields, organize the information so that it can be easily extracted for reports.

These log fields are organized in such a way that they form two groups: the first group, made up of the log fields that come first, is called the log header. The log header contains general information, such as the unique log identification and date and time that indicates when the activity was recorded. The log body is the second group, and contains all the other information about the activity. There are no two log message bodies that are alike, however, there may be fields common to most log bodies, such as the `srcintf` or `identidix` log fields.

The log header also contains information about the log priority level which is indicated in the `level` field. The priority level indicates the immediacy and the possible repercussions of the logged action. For example, if the field contains 'alert', you need to take immediate action with regards to what occurred. There are six log priority levels.

The log severity level is the level at and above which the FortiGate unit records logs. The log severity level is defined by you when configuring the logging location. The FortiGate unit will log all messages at and above the priority level you select. For example, if you select Error, the unit will log only Error, Critical, Alert, and Emergency level messages.

**Table 78:** Log priority levels

Levels	Description
0 - Emergency	The system has become unstable.
1 - Alert	Immediate action is required.
2 - Critical	Functionality is affected.



**Table 78:** Log priority levels

Levels	Description
3 - Error	An error condition exists and functionality could be affected.
4 - Warning	Functionality could be affected.
5 - Notification	Information about normal events.
6 - Information	General information about system operations.

The Debug priority level, not shown in [Table 78](#), is rarely used. It is the lowest log priority level and usually contains some firmware status information that is useful when the FortiGate unit is not functioning properly.

**Table 79:** Example log header fields

Log header	
date=(2010-08-03)	The year, month and day of when the event occurred in yyyy-mm-dd format.
time=(12:55:06)	The hour, minute and second of when the event occurred in the format hh:mm:ss.
log_id=(2457752353)	A five or ten-digit unique identification number. The number represents that log message and is unique to that log message. This ten-digit number helps to identify the log message.
type=(dlp)	The section of system where the event occurred.
subtype=(dlp)	The subtype category of the log message. See <a href="#">Table 78 on page 1936</a> .
level=(notice)	The priority level of the event. See <a href="#">Table 78 on page 1936</a> .
vd=(root)	The name of the virtual domain where the action/event occurred in. If no virtual domains exist, this field always contains root.

**Table 80:** Example log body fields

Log body	
policyid=(1)	The ID number of the firewall policy that applies to the session or packet. Any policy that is automatically added by the FortiGate will have an index number of zero.
identidx=(0)	The identity-based policy identification number. This field displays zero if the firewall policy does not use an identity-based policy; otherwise, it displays the number of the identity-based policy entry that the traffic matched. This number is not globally unique, it is only locally unique within a given firewall policy.
sessionid=(311)	The serial number of the firewall session of which the event happened.

**Table 80:** Example log body fields

Log body	
srcip=(10.10.10.1)	The source IP address.
srcport=(1190)	The source port number.
srcintf=(internal)	The source interface name.
dstip=(192.168.1.122)	The destination IP address.
dstport=(80)	The destination port number.
dstintf=(wan1)	The destination interface name.
service=(https)	The IP network service that applies to the session or packet. The services displayed correspond to the services configured in the firewall policy.
status=(detected)	The action the FortiGate unit took.
hostname=(example.com)	The home page of the web site.
url=(/image/trees_pine_forest/)	The URL address of the web page that the user was viewing.
msg=(data leak detected(Data Leak Prevention Rule matched)	Explains the FortiGate activity that was recorded. In this example, the data leak that was detected matched the rule, All-HTTP, in the DLP sensor.
rulename=(All-HTTP)	The name of the DLP rule within the DLP sensor.
action=(log-only)	The action that was specified within the rule. In some rules within sensors, you can specify content archiving. If no action type is specified, this field display log-only.
severity=(1)	The level of severity for that specific rule.

Logs from other devices, such as the FortiAnalyzer unit and Syslog server, contain a slightly different log header. For example, when viewing FortiGate log messages on the FortiAnalyzer unit, the log header contains the following log fields when viewed in the Raw format:

```
itime=1302788921 date=20110401 time=09:04:23 devname=FG50BH3G09601792
device_id=FG50BH3G09601792 log_id=0100022901 type=event
subtype=system level=notice vd=root
```

The log body contains the rest of the information of the log message, and this information is unique to the log message itself.

For detailed information on all log messages, see the [FortiGate Log Message Reference](#). For more information about how to understand a log message, see the [FortiGate Cookbook](#).

## Explanation of a debug log message

Debug log messages are only generated if the log severity level is set to Debug. The Debug severity level is the lowest log severity level and is rarely used. This severity level usually contains some firmware status information that is useful when the FortiGate unit is not functioning properly. Debug log messages are generated by all types of FortiGate features.

The following is an example of a debug log message:

```
date=2010-01-25 time=17:25:54 logid=9300000000 type=utm
 subtype=webfilter eventtype=urlfilter level=debug msg="found in
 cache"
```

**Table 81:** Example of a Debug log message

Debug log	
date=(2010-01-25)	The year, month and day of when the event occurred in the format yyyy-mm-dd.
time=(17:25:54)	The hour, minute and second of when the event occurred in the format hh:mm:ss.
logid=(9300000000)	A ten-digit unique identification number. The number represents that log message and is unique to that log message. This ten-digit number helps to identify the log message.
type=(utm)	The section of system where the event occurred.
subtype=(webfilter)	The subtype of the log message, the UTM profile in a UTM log message.
eventtype=(urlfilter)	The event type of the log message. This represents a policy applied to the FortiGate feature in the firewall policy.
level=(debug)	The priority level of the event. There are six priority levels to specify.
msg=("found in cache")	Explains the activity or event that the FortiGate unit recorded.
msg=("found in cache")	Explains the activity or event that the FortiGate unit recorded.

## Viewing log messages and archives

Depending on the log device, you may be able to view logs within the web-based manager or CLI on the FortiGate unit. If you have configured a FortiAnalyzer unit, local hard disk, or system memory, you can view log messages from within the web-based manager or CLI. If you have configured either a Syslog or WebTrends server, you will not be able to view log messages from the web-based manager or CLI. There is also no support for viewing log messages stored on a FortiCloud server, from the FortiGate unit's web-based manager or CLI.

You do not have to view log messages from only the web-based manager. You can view log messages from the CLI as well, using the `execute log display` command. This command allows you to see specific log messages that you already configured within the `execute log filter` command. The `execute log filter` command configures what log messages you will see, how many log messages you can view at one time (a maximum of 1000 lines of log messages), and the type of log messages you can view. For more information about viewing log messages in the CLI, see ["Viewing logs from the CLI" on page 1971](#).

There are two log viewing options in FortiOS: Format and Raw. The Raw format displays logs as they appear within the log file. You can view log messages in the Raw format using the CLI or a text editor, such as Notepad. Format is in a more human-readable format, and you can easily filter information when viewing log messages this way. The Format view is what you see when viewing logs in the web-based manager.

When you download the log messages from within the log message page (for example, *Log & Report > Traffic Log > Forward Traffic*), you are downloading log messages in the Raw format.

## Viewing log messages in detail

From any log page, you can view detailed information about the log message in the log viewer table, located (by default) at the bottom of the page. Each page contains this log viewer table. The Log Viewer Table can contain the Archive tab, which allows you to see the archived version of the log message. The Archive tab only displays the archived log's details if archiving is enabled and logs are being archived by the FortiGate unit, but archived logs will also be recorded when using a FortiAnalyzer unit or the FortiCloud service.

When you are viewing traffic log messages, some of the categories (such as 'Application Name') have entries that can be selected to open a dialog box containing FortiGuard information about the entry. From within the dialog box, you can select the Reference link and go directly to the corresponding FortiGuard page, which contains additional information.

Viewing logs in Raw format allows you to view all log fields at once, as well as have a log file available regardless of whether you are archiving logs or not. You download the log file by selecting *Download Raw Log*. The log file is named in the following format: `<log_type><log_location><log_date/time>.<log_number>.log`. For example, `SystemEventLog-disk-2012-09-19T12_13_46.933949.log`, which is an event log. The time period is the day and month of when the log was downloaded, not the time period of the log messages within the file itself.

## Quarantine

Within the Log & Report menu, you can view detailed information about each quarantined file. The information can either be sorted or filtered, depending on what you want to view.

You must enable quarantine settings within an antivirus profile and the destination must be configured in the CLI using the `config antivirus quarantine` command. The destination can be either a FortiAnalyzer unit or local disk.

Sort the files by file name, date, service, status, duplicate count (DC), or time to live (TTL). Filter the list to view only quarantined files with a specific status or from a specific service.

On *Log & Report > Quarantine*, the file quarantine list displays the following information about each quarantined file.

### Quarantine page

Lists all files that are considered quarantined by the unit. On this page you can filter information so that only specific files are displayed on the page.

GUI Item	Description
Source	Either <i>FortiAnalyzer</i> or <i>Local disk</i> , depending where you configure to quarantined files to be stored.
Sort by	Sort the list. Choose from: <i>Status</i> , <i>Service</i> , <i>File Name</i> , <i>Date</i> , <i>TTL</i> , or <i>Duplicate Count</i> . Select <i>Apply</i> to complete the sort.

<b>Filter</b>	<p>Filter the list. Choose either <i>Status</i> (infected, blocked, or heuristics) or <i>Service</i> (IMAP, POP3, SMTP, FTP, HTTP, MM1, MM3, MM4, MM7, IM, or NNTP). Select <i>Apply</i> to complete the filtering. Heuristics mode is configurable through the CLI only.</p> <p>If your unit supports SSL content scanning and inspection Service can also be IMAPS, POP3S, SMTPS, or HTTPS. For more information, see the Security Features chapter of the FortiOS Handbook.</p>
<b>Apply</b>	Select to apply the sorting and filtering selections to the list of quarantined files.
<b>Delete</b>	Select to delete the selected files.
<b>Page Controls</b>	Use the controls to page through the list.
<b>Remove All Entries</b>	<p>Removes all quarantined files from the local hard disk.</p> <p>This icon only appears when the files are quarantined to the hard disk.</p>
<b>File Name</b>	<p>The file name of the quarantined file. When a file is quarantined, all spaces are removed from the file name, and a 32-bit checksum is performed on the file. The checksum appears in the replacement message but not in the quarantined file. The file is stored on the Fortinet hard disk with the following naming convention:</p> <p>&lt;32bit_CRC&gt;.&lt;processed_filename&gt;</p> <p>For example, a file named Over Size.exe is stored as 3fc155d2.oversize.exe.</p>
<b>Date</b>	The date and time the file was quarantined, in the format dd/mm/yyyy hh:mm. This value indicates the time that the first file was quarantined if duplicates are quarantined.
<b>Service</b>	The service from which the file was quarantined (HTTP, FTP, IMAP, POP3, SMTP, MM1, MM3, MM4, MM7, IM, NNTP, IMAPS, POP3S, SMTPS, or HTTPS).
<b>Status</b>	The reason the file was quarantined: <i>infected</i> , <i>heuristics</i> , or <i>blocked</i> .
<b>Status Description</b>	Specific information related to the status, for example, "File is infected with "W32/Klez.h"" or "File was stopped by file block pattern."
<b>DC</b>	Duplicate count. A count of how many duplicates of the same file were quarantined. A rapidly increasing number can indicate a virus outbreak.
<b>TTL</b>	<p>Time to live in the format hh:mm. When the TTL elapses, the Fortinet unit labels the file as EXP under the TTL heading. In the case of duplicate files, each duplicate found refreshes the TTL.</p> <p>The TTL information is not available if the files are quarantined on a FortiAnalyzer unit.</p>
<b>Upload status</b>	<p>Y indicates the file has been uploaded to Fortinet for analysis, N indicates the file has not been uploaded.</p> <p>This option is available only if the Fortinet unit has a local hard disk.</p>

- Download** Select to download the corresponding file in its original format.  
This option is available only if the Fortinet unit has a local hard disk.
- Submit** Select to upload a suspicious file to Fortinet for analysis.  
This option is available only if the Fortinet unit has a local hard disk.

## Customizing the display of log messages on the web-based manager

Customizing log messages on the web-based manager allows you to remove or add columns from the page and filter the information that appears. For example, you can view only log messages that appeared on December 4, between the hours of 8:00 and 8:30 am.

1. Select the submenu in **Log & Report** in which you want to customize the display of log messages, such as *Log & Report > Traffic Log > Forward Traffic*.
2. Right click on the title bar at the top of any column, and select *Column Settings*. This will change what columns will display on the page.
3. In the section at the top, uncheck a column title such as **Date/Time** to remove it from the interface. Check other columns to add them to the interface. When you are finished, click outside the menu and the page will refresh with the new column settings in place.
4. Choose a column you'd like to filter, and select the funnel icon next to the title of the column. For example, select the funnel in the Src (Source) column. In the text field, enter the source IP address 1.1.1.1 and then select the check box beside **NOT**.  
This filters out the all log messages that have the 1.1.1.1 source IP address in the source IP log field, such as the ones generated when running log tests in the CLI.
5. Select **OK** to save the customize settings, and then view the log messages on the page.  
Log messages that originate from the 1.1.1.1 source address will no longer appear in the list.

## How to download log messages and view them from on a computer

After recording some activity, you can download log messages to view them from a computer. This is can be very useful when in a remote location, or if you want to view log messages at your convenience, or to view packet logs or traffic logs.

1. In Log & Report, select the submenu that you want to download log messages from.  
For example, *Log & Report > Traffic Log > Forward Traffic*.
2. Select the **Download Raw Log** option and save the log file to your computer.  
The log file will be downloaded like any other file. Log file names contain their log type and date in the name, so it is recommended to create a folder in which to archive your log messages, as they can be sorted easily.
3. Open a text editor such as Notepad, open the log file, and then scroll to view all the log messages.  
You can easily search or scroll through the logs to see the information that is available.

## Log files and types

As the log messages are being recorded, log messages are also being put into different log files. The log file contains the log messages that belong to that log type, for example, traffic log messages are put in the traffic log file.

When downloading the log file from within *Log & Report*, the file name indicates the log type and the device on which it is stored, as well as the date, time, and a unique id for that log.

This name is in the format <logtype> - <logdevice> - <date> T <time> . <id>.log.

For example, AntiVirusLog-disk-2012-09-13T11\_07\_57.922495.log.

Below, each of the different log files are explained. Traffic and Event logs come in multiple types, but all contain the base type such as 'Event' in the filename.

**Table 82: Log Types based on network traffic**

Log Type	Description
Traffic	The traffic logs records all traffic to and through the FortiGate interface. Different categories monitor different kinds of traffic, whether it be external, internal, or multicast.
Event	The event logs record management and activity events within the device in particular areas: System, Router, VPN, User, WAN, and WiFi. For example, when an administrator logs in or logs out of the web-based manager, it is logged both in System and in User events.
Antivirus	The antivirus log records virus incidents in Web, FTP, and email traffic.
Web Filter	The web filter log records HTTP FortiGate log rating errors including web content blocking actions that the FortiGate unit performs.
Intrusion	The intrusion log records attacks that are detected and prevented by the FortiGate unit.
Email Filter	The email filter log records blocking of email address patterns and content in SMTP, IMAP, and POP3 traffic.
Vulnerability Scan	The Vulnerability Scan (Netscan) log records vulnerabilities found during the scanning of the network.
Data Leak Prevention	The Data Leak Prevention log records log data that is considered sensitive and that should not be made public. This log also records data that a company does not want entering their network.
VoIP	The VoIP log records VoIP traffic and messages. It only appears if VoIP is enabled on the Administrator Settings page.

## Log database and datasets

The log database, also known as the SQL log database, is used to store logs on the FortiGate unit itself. The log database uses Structured Query Language (SQL), specifically it uses SQLite which is an embedded Relational Database Management System (RDBMS).



If you have disabled SQL logging and have factory defaults on the FortiGate unit, and then you upgrade the firmware, the upgrade will automatically disable SQL logging. When this occurs, you must re-enable SQL logging manually.

The FortiGate unit creates a database table for each log type, when log data is recorded. If the FortiGate unit is not recording log data, it does not create log tables for that device.

The command syntax, `get report database schema`, allows you to view all the tables, column names and types that are available to use when creating SQL statements for datasets.

If you want to view the size of the database, as well as the log database table entries, use the `get report sql status` command. This command displays the amount of free space that is available as well as the first and last log database entry time and date.

The output of the `get log sql status` command contains information similar to the following:

```
Database size: 294912
Free size in database: 0
Database Page Size: 8192
Entry number:
 Event: 49
 Traffic: 370
 Attack: 2
 AntiVirus: 4
 WebFilter: 254
 AntiSpam: 2
 Netscan: 18
 Total: 699
First entry time: 2012-09-10 11:41:02
Last entry time: 2012-09-13 02:59:59
```

The log database is not only used to store logs, but also used to extract the information for reports. Reports are built from datasets, which are SQL statements that tell the FortiGate unit how to extract the information from the database. You can create your own datasets; however, SQL knowledge is required. Default datasets are available for reports.

## How to view datasets

If you want to view the list of all default datasets, use the following command syntax in the CLI (the question mark on the end is to display the list of existing entries):

```
diagnose report dataset ?
```

If you want to view the format of a created dataset, use the following command syntax in the CLI:

```
config report dataset
 show report dataset <dataset_name>
```

## How to create datasets (advanced)

Creating a dataset requires SQL language knowledge. The following is a short high-level summary of how to create a dataset. It assumes that you have SQL language knowledge, and therefore is for advanced users only.

1. Log in to the CLI and use the `config report dataset` command to start creating the dataset.

You can name the dataset any name; however, it should be a name that describes what information is contained in the dataset. For example, `wanopt.traffic.bandwidth.24h` dataset name indicates that the dataset contains only WAN Optimization traffic and the amount of bandwidth that has been used for the past 24 hours.

2. Set the various variables within the dataset as you need for your report. Refer to the [FortiGate CLI Reference](#) to see the complete list of commands and variables you will need.



The following table shows a possible method for customizing a dataset to output a list of the top ten bandwidth consuming applications within an hour.

<pre>config report dataset  edit appctrl.Count.Bandwidth.Top10.Apps  set query "select (timestamp-timestamp%3600) as hourstamp"  (Case WHEN app!=\'N/A\' and app!=\'\' then app ELSE service END) as appname  as bandwidth from traffic_log where ###timestamp_to_oid(traffic_log)###  and (appname in (select (CASE WHEN app!=\' N/A\' and app!=\'\' then app ELSE service END) as appname from traffic_log where ###timestamp_to_oid(traffic_log)### group by appname order by sum(sent+rcvd) desc limit 10)) group by hourstamp, appname order by hourstamp desc"</pre>	<p>This calculates an "hourstamp" to indicate the bandwidth per hour.</p> <p>Uses the application name, appname, or if it is undefined, uses service instead.</p> <p>This states to retrieve the information from the traffic log file.</p> <p>From this information, the FortiGate unit selects the application name from within the specific traffic log message and groups these names in order by the top ten, as well as by the "hourstamp" in descending order.</p>
----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

3. Create a corresponding chart using the `config report chart` command, making sure to apply the dataset that you just configured to the chart.

The custom dataset is now available for use in a report.

If you want to modify an existing dataset that you've created, the process is the same. Enter `config report dataset` in the CLI followed by the name of the existing dataset.

To see the list of created datasets, enter the following in the CLI:

```
config report dataset
show
```

## Notifications about network activity

Alert email messages provide notification about activities or events logged. These email messages also provide notification about log severities that are recorded, such as a critical or emergency.

You can send alert email messages to up to three email addresses. Alert messages are also logged and can be viewed from the Event Log menu, in the System Event log file.

You can use the alert email feature to monitor logs for log messages, and to send email notification about a specific activity or event logged. For example, if you require notification about administrators logging in and out, you can configure an alert email that is sent whenever an administrator logs in and out. You can also base alert email messages on the severity levels of the logs. The FortiGate unit does not currently support SSL/TLS connections for SMTP servers, so you must choose an SMTP server that does not need SSL/TLS when configuring the SMTP server settings.

Before configuring alert email, you must configure at least one DNS server if you are configuring with an Fully Qualified Domain Server (FQDN). The FortiGate unit uses the SMTP server name to

connect to the mail server, and must look up this name on your DNS server. You can also specify an IP address.



The default minimum log severity level is Alert. If the FortiGate unit collects more than one log message before an interval is reached, the FortiGate unit combines the messages and sends out one alert email.

## How to configure email notifications

The following explains how to configure an alert email notification for IPsec tunnel errors, firewall authentication failure, configuration changes and FortiGuard license expiry.

1. In **System > Config > Messaging Services**, configure the SMTP server.

The SMTP server settings allow the FortiGate unit to know exactly where the email will be sent from, as well as who to send it to. The SMTP server must be a server that does not support SSL/TLS connections; if the SMTP server does, the alert email configuration will not work. The FortiGate unit does not currently support SSL/TLS connections for SMTP servers.

2. In **Log & Report > Log Config > Alert E-mail**, enter the source email in the Email From field, and up to three target addresses in the Email To fields.
3. Below the email entry, you can configure the email responses. By default, the **Send alert email for the following** is enabled. Select the check boxes beside **IPsec tunnel errors**, **Configuration changes** and **Firewall authentication failure**.

These alerts will be sent to the email address specified when the trigger occurs. For example, a user attempts to connect to the branch office of the company but cannot; the FortiGate unit detects an IPsec tunnel error, records the event, and then sends the notice to the email address specified in the SMTP server settings.

4. Select **FortiGuard license expiry time:** and then enter 10 so that the email notification will be sent ten days prior to the FortiGuard license expiration.

You can choose up to 100 days prior to when the license will expire. The default time is 15 days. By using this alert email notification, you can easily know when to send an re-registration request long before the expiry.

## Log devices

The FortiGate unit supports a variety of log devices, including the FortiCloud service and FortiAnalyzer units. This provides greater flexibility not only when choosing a log device, but also when your logging requirements need updating.

When you have developed a plan that meets your logging needs and requirements, you need to select the log device that is appropriate for that plan. A log device must be able to store all the logs you need, and if you require archiving those logs, you must consider what log devices support this option.

During this process of deciding what log device meets your needs and requirements, you must also figure out how to provide a backup solution in the event the log device that the FortiGate unit is sending logs to has become unavailable. A backup solution should be an important part of your log setup because it helps you to maintain all logs and prevents lost logs, or logs that are not sent to the log device. For example, a daily backup of log files to the FortiAnalyzer unit occurs at 5 pm.

Log devices provide a central location for storing logs recorded by the FortiGate unit. The following are log devices that the FortiGate unit supports:

- FortiGate system memory
- Hard disk or AMC
- SQL database (for FortiGate units that have a hard disk)
- FortiAnalyzer unit
- FortiCloud service
- Syslog server
- NetIQ WebTrends server

These log devices, except for the FortiGate system memory and local hard disk, can also be used as a backup solution. For example, you can configure logging to the FortiGate unit's local disk, but also configure logging to a FortiCloud server and archive logs to both the FortiCloud server and a FortiAnalyzer unit, and disk logging does not need to be enabled for these alternate log solutions to continue logging.



If you are formatting a disk that contains more than just logs, all information on the disk will be lost.

---

## FortiGate unit's system memory and hard disk

The FortiGate unit's system memory and hard disk can store all log types, including log archives and traffic logs. Traffic logs and log archives are larger files, and need a lot of room when being logged by the FortiGate unit.

When the system memory is full, the FortiGate unit overwrites the oldest messages, and all log messages stored in memory are cleared when the FortiGate unit restarts. By default, logging to memory is enabled. This means that most of the time you will only need to modify the default settings to your network logging requirements. Real-time logging occurs whenever memory logging is enabled, and is enabled by default. Real-time logging means that the activity is being recorded as it happens.

All FortiGate units 100D and larger are capable of disk logging, but it is disabled by default, as it is not recommended. For flash memory-based units, constant rewrites to flash drives can reduce the lifetime and efficiency of the memory. For hard-disk units, it can affect performance under heavy strain. Therefore, disk logging must be manually enabled in the CLI under `config log disk setting` to appear in the interface at all.



Models without a hard disk are not recommended for disk logging. For all units, disk logging must be enabled in the CLI. For some low-end and older models, disk logging is unavailable. Check a product's Feature Matrix for more information. In either case, Fortinet recommends using either a FortiAnalyzer unit or the FortiCloud service.

---

When logging to the FortiGate unit's hard disk or memory, you can configure logging to a FortiAnalyzer unit, which will log independently of the hard disk.

If you are registered with the FortiCloud service, your unit will log both locally and to the service by default. In order to configure the rate and time of uploads to the service, you must register a contract account for the FortiCloud service, which will also grant you additional space.

## FortiAnalyzer unit

The FortiAnalyzer unit can log all FortiGate features, which includes log archives. You can also configure the FortiGate unit to upload logs to the FortiAnalyzer unit at a scheduled time.

Encryption of the logs is supported by default and logs are sent using IPsec or SSL VPN. When the FortiAnalyzer and FortiGate units have SSL encryption, both must choose a setting for the enc-algorithm command (CLI) for encryption to take place. By default, this is enabled and the default setting is a SSL communication with high and medium encryption algorithms. The setting that you choose must be the same for both.

FortiGate units can support logging to multiple FortiAnalyzer units. This logging solution is a backup redundancy solution, since logs are sent to all three units and whenever one of the FortiAnalyzer units fails, the others still carry on storing logs.

If you are using evaluation software FortiGate and FortiAnalyzer-VM images, you will only be able to use low-level encryption.

The FortiGate unit can also connect to a FortiAnalyzer unit using Automatic Discovery. Automatic Discovery is a method of establishing a connection to a FortiAnalyzer unit by using the FortiGate unit to find a FortiAnalyzer unit on the network. The Fortinet Discovery Protocol (FDP) is used to locate the FortiAnalyzer unit. Both the FortiGate and FortiAnalyzer units must be on the same subnet to use FDP, and they must also be able to connect using UDP.

When you enable automatic discovery in the CLI, the FortiGate unit uses HELLO packets to locate any FortiAnalyzer units that are available on the network within the same subnet. When the FortiGate unit discovers a FortiAnalyzer unit, the FortiGate unit automatically enables logging to the FortiAnalyzer unit and begins sending log data.

## Syslog server

The Syslog server is a remote computer running syslog software. Syslog is a standard for forwarding log messages in an IP network, and can be used when considering a log backup solution for your network logging requirements.

FortiGate units support the reliable syslog feature, which is based on RFC 3195. Reliable syslog logging uses TCP, which ensures that connections are set up, including that packets are transmitted.

There are several profiles available for reliable syslog, but only the RAW profile is currently supported on the FortiGate units. The RAW profile is designed to provide a high-performance, low-impact footprint using essentially the same format as the existing UDP-based syslog service.

When enabling the reliable syslog (available only in the CLI), TCP is used. The feature is disabled by default, and when enabled, the FortiGate unit automatically changes the port number to TCP 601. This is based on RFC 3195. The default port for syslog is port 514.

## WebTrends server

A WebTrends server is a remote computer, similar to a Syslog server, running NetIQ WebTrends firewall reporting server. FortiGate log formats comply with WebTrends Enhanced Log Format (WELF) and are compatible with NetIQ WebTrends Security Reporting Center and Firewall Suite 4.1.

You can configure a WebTrends server only in the CLI.

## How to choose a log device for your network topology

When planning the log requirements, you must also consider your network's topology and whether archiving is required, such as if there is a legal requirement to keep a historical record of network activity. The following explains what steps to take when choosing a log device for your specific network topology.

1. What is the scope of your network topology?

If it is a SOHO/SMB network, then logging to the FortiGate unit's local hard disk or the default FortiCloud service would be efficient. If the network topology is a large enterprise, you will need FortiAnalyzer units, a FortiCloud contract, Syslog servers, or any combination.

2. Is archiving required?

If the network activity that is being logged needs to be archived, then, depending on your network topology, you would choose a FortiAnalyzer unit. FortiAnalyzer units store archives in the same way that FortiGate units do, but are able to store large amounts of logs and archives.

3. When troubleshooting issues, you may want to log features such as traffic, that would not be logged otherwise; how much storage space will you need?

Logs can be configured to roll, which is similar to zipping a file; this will lower the space requirements needed to contain them. You can also download logs from the FortiGate unit and save them on a server or on a computer to view and access later, to prevent them from piling up and being overwritten. If you're only temporarily logging a larger amount of traffic, you could upload the logs to an FTP server or cloud service to ensure that these logs are not lost even after you are finished correcting the issue.


4. Should I invest in a log device that can grow as my network grows?

All networks grow, so investing in a device that can grow with your network and that can be expanded is a good investment. For example, if you currently have a SOHO/SMB topology, but see growth already starting, a FortiAnalyzer unit would be best. A FortiAnalyzer unit provides ample storage space, as well as you can add two more FortiAnalyzer units to access additional storage and create a redundancy log backup solution.

## How to create a backup solution for logging

The following helps to explain how to create a log backup solution for a small network topology. This example has one FortiAnalyzer unit and a subscription to the FortiCloud Service.

**Figure 294:**Example of an integrated FortiAnalyzer unit and Syslog servers in a network



1. Log in to the CLI and modify what features will be logged to the FortiAnalyzer unit as well as the settings to the default log device, the FortiGate unit's hard drive.  
By default, the FortiGate unit logs to either the system memory or hard drive, whichever is available on the FortiGate unit. Low-end FortiGate units usually have the system memory enabled for logging by default.
2. In the CLI, use the `config log fortianalyzer setting` command to configure logging to the FortiAnalyzer unit.  
You can only configure log settings for the FortiAnalyzer unit in the CLI. Configuring to upload logs to a FortiAnalyzer unit can be configured in both the CLI and web-based manager.
3. In the CLI, configure the settings for the Syslog server; also enable reliable syslog as well.  
Reliable syslog ensures that logs are sent to the syslog server. When you enable this setting, the default port becomes port 601.

## Reports

Reports provide a way to analyze log data without manually going through a large amount of logs to get to the information you need. This section explains how to configure a FortiOS report and how to modify the existing default FortiOS Security Features report. The FortiOS default Security Features report is a report that gathers security feature activity information and compiles it into a report. This section also explains how to view these reports.

Reports provide a clear, concise overview of what is happening on your network based on log data, without manually going through large amounts of logs. Reports can be configured on a FortiGate unit or a FortiAnalyzer unit. However, in this document only FortiOS reports are explained. FortiOS reports are the reports that are generated on the FortiGate unit. FortiAnalyzer reports are configured on a FortiAnalyzer unit and for information about those reports, see the [FortiAnalyzer Administration Guide](#).

Disk or memory logging must be enabled for reporting to be enabled. Local Reporting can then be enabled in *Log & Report > Log Setting*, in order to view and edit reports.

## What are FortiOS reports?

FortiOS reports are configured from logs stored on the FortiGate unit's hard drive. These reports, generated by the FortiGate unit itself, provide a central location for both configuring and generating reports. A default FortiOS report, called the FortiGate Security Feature Daily Activity Report, is available for you to modify to your requirements. The default report provides a way to quickly and easily set up your own report from within the web-based manager. The default FortiOS report is a report that compiles security feature activity from various security-related logs, such as virus and attack logs.

FortiOS reports consist of multiple parts, regardless of whether its the default FortiOS report or a report that you have configured from scratch, and these parts are configured separately and added to the layout. These parts of a FortiOS report are:

- charts (including datasets within the charts themselves)
- themes (including styles which are within the themes themselves)
- images
- layout

## The parts of a FortiOS report

Charts are used to display the log information in a clear and concise way using graphs and tables. Charts contain datasets, which are SQLite queries that help the FortiGate unit to add specific log information into the chart using the log information that is stored in the SQLite database on the local hard disk. If you want to configure a chart, you must configure the dataset first. Datasets are required for each chart, and if there is no dataset included in a chart, the chart will not be saved.

Themes provide a one-step style application for report layouts. Themes contain various styles, including styles for the table of contents, headings, headers and footers, as well as the margins of the report's pages. Themes are applied to layouts. The styles that are applied to themes are configured separately in the CLI.

You can easily upload your company or organization's logo to use within a report. By uploading your company or organization's logo and applying it to a report, you provide a personalized report that is recognizable as your company or organization's report. The image must be in JPEG, JPG or PNG format.

Layouts provide a way to incorporate the charts, images, and themes that are configured to create a formatted report. A layout is used as a template by the FortiGate unit to compile and then generate the report. The layout is also coded in the CLI.

## What you can do with the default FortiOS report

You can reset the reports you have configured, as well as the default FortiOS report you modified, to default settings. When you reset reports to default settings, any configured reports that you created from scratch are lost. The `execute report-config reset` command resets the reports to default settings. If you are going to reset the reports to their default settings, you should back up the current configuration file before doing so, in the event you want to revert back to the reports you previously created and/or modified.

The default FortiOS report can be modified so that it meets your requirements for a report. This default report is located in *Log & Report > Report > Local*. Select *Customize* to edit it.

The FortiOS default report contains several pages, which appear as stacked boxes in the editing interface. Each page contains one or multiple charts (depending on the configuration of that page in the interface), and each page in the finished report will contain information about the FortiGate unit at the top of each section.

You can select *Run Now* on the *Local* page to immediately create a report with the current layout and design. More complex reports may take longer to generate. After generating a report, you can view it by selecting it from the list below *Run Now*. Historical reports will be marked as 'Scheduled' if created automatically, or 'On Demand' if created by selecting *Run Now*.

## How to modify the default FortiOS report

The following is a sample modification of the default FortiOS report, which includes adding an image.

1. In *Log & Report > Report > Config*, modify the page by adding a new Chart, which will appear on its own page in the final report.
2. Add an information Text field below the chart.  
You should always save the changes you make by selecting *Save*; otherwise, the changes you just made will be lost.
3. Modify the header to add the company's image.  
The company's image will appear in all headers throughout the report. If you select *Save now*, it will appear on all the report's pages.
4. Add other charts to the list so they will appear within the report.  
Charts marked as 'FortiGate Security Feature Security Analysis Report' are autogenerated and take up an entire page or multiple pages on their own. All other charts take up half a page, so two consecutive charts will appear on the same page in the report.
5. Modify the report settings so that the report is generated every Monday at 6 pm, and is emailed to specific employees in the company.  
Reports can be sent to others after the report has been generated, if Messaging Servers are configured.
6. Test the report's modified settings, by selecting *Run Now* in the Config page; after it is generated, go to *Log & Report > Report > Local* and view the report.  
You can tell that it has been generated because the Bandwidth Usage page's charts will be populated, and the text added below each chart appears as well.

## How to create a FortiOS report

This creates a FortiOS report from within the CLI. This is not modifying a default report, rather, it is creating one from scratch. This can take quite a while since there is a lot of commands to use. This is a time-consuming activity, so it is recommended that you allow yourself plenty of time to configure a FortiOS report. To see some sample configuration, view "[Customizing FortiOS reports with CLI](#)" on page 1976.

1. Log in to the CLI and, using the `config report style` command, configure the styles for the footer and header, page, cover page, charts, table of contents, and page headings.  
It is recommended that you outline what you will be doing first, before configuring a report from scratch. When configuring new charts, you must also configure datasets. Charts will not be saved if a dataset is not associated with the chart. When configuring datasets, you must know SQL language.
2. Configure the theme using the `config report theme` command.  
A theme applies the styles that you configured previously in step 1.
3. Configure the new datasets that you want to have for the new charts using the `config report dataset` command.  
A dataset contains the information the FortiGate unit needs to gather the information for a chart. The log database is an SQL database, and that is why the datasets must be written in SQL; specifically, the log database uses SQLite so you must make sure that the information



is correct in a dataset. If the information is not correct, the FortiGate unit will not gather the proper information for that chart.

4. Configure the new charts that you want and make sure to apply the correct dataset with the chart using the `config report chart` command.
5. Configure the report layout using the `config report layout` command.  
The layout puts all the pieces of the report together. The layout provides a template for the FortiGate unit to know how the pieces should be put together and where they need to be on the pages of the report. The layout also allows for specifying when the report will be generated, and what style will be applied.
6. Test the report by going back to the report layout and setting the time so that the report immediately generates; it will then appear in the list of reports in the web-based manager.

## Best Practices: Log management

When the FortiGate unit records FortiGate activity, valuable information is collected that provides insight into how to better protect network traffic against attacks, including misuse and abuse. There is a lot to consider before enabling logging on a FortiGate unit, such as what FortiGate activities to enable and which log device is best suited for your network's logging needs. A plan can help you in deciding the FortiGate activities to log, a log device, as well as a backup solution in the event the log device fails.

This plan should provide you with an outline, similar to the following:

- what FortiGate activities you want and/or need logged (for example, security features)
- the logging device best suited for your network structure
- if you want or require archiving of log files
- ensuring logs are not lost in the event a failure occurs.

After the plan is implemented, you need to manage the logs and be prepared to expand on your log setup when the current logging requirements are outgrown. Good log management practices help you with these tasks.

Log management practices help you to improve and manage logging requirements. Logging is an ever-expanding tool that can seem to be a daunting task to manage. The following management practices will help you when issues arise, or your logging setup needs to be expanded.

1. Revisit your plan on a yearly basis to verify that your logging needs are being met by your current log setup. For example, your company or organization may require archival logging, but not at the beginning of your network's lifespan. Archival logs are stored on a FortiGate unit's local hard drive, a FortiAnalyzer unit, or a FortiCloud server, in increasing order of size.
2. Configure an alert message that will notify you of activities that are important to be aware about. For example: if a branch office does not have a FortiGate administrator, you will need to know at all times that the IPSec VPN tunnel is still up and running. An alert email notification message can be configured to send only if IPSec tunnel errors occur.
3. If your organization or company uses peer-to-peer programs such as Skype or other instant messaging software, use the IM usage dashboard widget or the Executive Summary's report widget (Top 10 Application Bandwidth Usage Per Hour Summary) to help you monitor the usage of these types of instant messaging software. These widgets can help you in determining how these applications are being used, including if there is any misuse and abuse. Their information is taken from application log messages; however, application log messages should be viewed as well since they contain the most detailed information.
4. Ensure that your backup solution is up-to-date. If you have recently expanded your log setup, you should also review your backup solution. The backup solution provides a way to

ensure that all logs are not lost in the event that the log device fails or issues arise with the log device itself.

# Logging and reporting for small networks

This section explains how to configure the FortiGate unit for logging and reporting in a small office or SOHO/SMB network. To properly configure this type of network, you will be modifying the default log settings, as well as the default FortiOS report.

The following procedures are examples and can be used to help you when configuring your own network's log topology. Since some of these settings must be modified or enabled or disabled in the CLI, it is recommended to review the FortiGate CLI Reference for any additional information about the commands used herein, as well as any that you would need to use in your own network's log topology.

The following topics are included in this section:

- [Modifying default log device settings](#)
- [Configuring the backup solution](#)
- [Modifying the default FortiOS report](#)

## Modifying default log device settings

The default log device settings must be modified so that system performance is not compromised. The FortiGate unit, by default, has all logging of FortiGate features enabled, except for traffic logging. The default logging location will be either the FortiGate unit's system memory or hard disk, depending on the model. Units with a flash disk are not recommended for disk logging.

### Modifying the FortiGate unit's system memory default settings

When the FortiGate unit's default log device is its system memory, the following is modified for a small network topology. The following is an example of how to modify these default settings.

#### To modify the default system memory settings

1. Log in to the CLI.
2. Enter the following command syntax to modify the logging settings:

```
config log memory setting
 set ips-archive disable
 set status enable
end
```

3. The following example command syntax modifies which FortiGate features that are enabled for logging:

```
config log memory filter
 set attack enable
 set forward-traffic enable
 set local-traffic enable
 set netscan enable
 set email-log-imap disable
 set multicast-traffic enable
 set scanerror enable
 set app-ctrl enable
end
```

## Modifying the FortiGate unit's hard disk default settings

When the FortiGate unit's default log device is its hard disk, you need to modify those settings to your network's logging needs so that you can effectively log what you want logged. The following is an example of how to modify these default settings.

### To modify the default hard disk settings

1. Log in to the CLI.
2. Enter the following command syntax to modify the logging settings:

```
config log disk setting
 set ips-archive disable
 set status enable
 set max-log-file-size 1000
 set storage FLASH
 set log-quota 100
 set report-quota 100
 set sql-max-size 10000
 set sql-max-size-action overwrite
 set sql-oldest-entry 1024
 set rows-per-transaction 100
 set ms-per-transaction 6000
 config sql-logging
 set netscan disable
 set spam disable
end
```

3. In the CLI, enter the following to disable certain event log messages that you do not want logged:

```
config log disk filter
 set discovery disable
 set email-log-imap disable
end
```

## Testing sending logs to the log device

After modifying both the settings and the FortiGate features for logging, you can test that the modified settings are working properly. This test is done in the CLI.

## To test sending logs to the log device

1. In the CLI, enter the following command syntax:

```
diag log test
```

When you enter the command, the following appears:

```
generating a system event message with level - warning
generating an infected virus message with level - warning
generating a blocked virus message with level - warning
generating a URL block message with level - warning
generating a DLP message with level - warning
generating an attack detection message with level - warning
generating an application control IM message with level - information
generating an antispam message with level - notification
generating an allowed traffic message with level - notice
generating a multicast traffic message with level - notice
generating a ipv6 traffic message with level - notice
generating a wanopt traffic log message with level - notification
generating a HA event message with level - warning
generating netscan log messages with level - notice
generating a VOIP event message with level - information
generating a DNS event message with level - information
generating authentication event messages
generating a Forticlient message with level - information
generating a NAC QUARANTINE message with level - information
generating a URL block message with level - warning
```

2. In the web-based manager, go to *Log & Report > Event Log > User*, and view the logs to see the recently generated test log messages.

You will be able to tell the test log messages from real log messages because they do not have “real” information; for example, the test log messages for the vulnerability scan contain the destination IP address of 1.1.1.1 or 2.2.2.2. If you have disabled certain logs that are logged to the SQL log database, the following will appear in red at the top of the Log Table: ‘Warning: SQL Logging is not enabled.’

## Configuring the backup solution

A backup solution provides a way to ensure logs are not lost. The following backup solution explains logging to a FortiCloud server and uploading logs to a FortiAnalyzer unit. With this backup solution, there can be three simultaneous storage locations for logs, the first being the FortiGate unit itself, the FortiAnalyzer unit and then the FortiCloud server.

### Configuring logging to a FortiCloud server

The FortiCloud server can be used as a redundant backup, or your primary logging solution. The following assumes that this service has already been registered, and a subscription has been purchased for expanded space. The following is an example of how to these settings are configured for a network’s log configuration. You need to have access to both the CLI and the web-based manager when configuring uploading of logs. The upload time and interval settings can only be done in the CLI; however, for uploading of logs to a FortiAnalyzer unit, you can modify the upload time and interval settings from the web-based manager after you have configured it in the CLI.

### To configure logging to the FortiCloud server

1. Go to *System > Dashboard > Status* and click *Login* next to *FortiCloud* in the License Information widget.
2. Enter your username and password, and click OK.
3. The logs will automatically be uploaded to FortiCloud as long as you are logged in.
4. To configure the upload time and interval, go to *Log & Report > Log Config > Log Settings*.
5. Under the Logging and Archiving header, you can select your desired upload time.

If your logs do not appear to be reaching the FortiCloud server, go to *System > Config > FortiGuard* and see what the status of *FortiGuard Availability* is.

In this video, you'll be learning how to use FortiCloud to view log data and reports.

With FortiCloud you can easily store and access FortiGate logs and reports that can give you valuable insight into the health and security of your network.

## Configuring uploading logs to the FortiAnalyzer unit

The logs will be uploaded to the FortiAnalyzer unit at a scheduled time. The following is an example of how to upload logs to a FortiAnalyzer unit.

### To upload logs to a FortiAnalyzer unit

1. Go to *Log & Report > Log Config > Log Settings*.
2. In the *Logging and Archiving* section, select the check box beside *Send Logs to FortiAnalyzer/FortiManager*.
3. Select *FortiAnalyzer (Daily at 00:00)*.
4. Enter the FortiAnalyzer unit's IP address in the *IP Address* field.
5. To configure the daily upload time, log in to the CLI.
6. Enter the following to configure when the upload occurs, and the time when the unit uploads the logs:

```
config log fortianalyzer setting
 set upload-interval {daily | weekly | monthly}
 set upload-time <hh:mm>
end
```

7. To change the upload time, in the web-based manager, select *Change* beside the upload time period, and then make the changes in the Upload Schedule window. Select *OK*.

## Testing uploading logs to a FortiAnalyzer unit

You should test that the FortiGate unit can upload logs to the FortiAnalyzer unit, so that the settings are configured properly.

### To test the FortiAnalyzer upload settings

1. Go to *Log & Report > Log Config > Log Settings*.
2. In the *Logging and Archiving* section, under *Send Logs to FortiAnalyzer/FortiManager*, change the time to the current time by selecting *Change*.  
For example, the current time is 11:10 am, so *Change* now has the time 11:10.
3. Select *OK*.

The logs will be immediately sent to the FortiAnalyzer unit, and will be available to view from within the FortiAnalyzer's interface.

## Modifying the default FortiOS report

The default FortiOS report is provided to help you quickly and easily configure and generate a report. The following is an example of how to modify the default FortiOS report.

### To modify the default FortiOS report

1. In the web-based manager, go to *Log & Report > Report > Local*.
2. Select *Customize* to open the Report Editor.
3. Change the default Fortinet image to the new image: select the Fortinet image and right-click so that *Delete* icon appears, and then select *Delete*; drag the Image icon to the box where the Fortinet image was previous; choose or upload a new image and then select *OK*.
4. Return to *Log & Report > Report > Local*.
5. Under *Report Options*, set the Generate report schedule to *Daily* and set a Time for the report to be compiled every day.
6. Enable *Email Generated Reports*. You may have to configure an SMTP server to send the reports before this option can be enabled. The SMTP configuration can be found in *System > Config > Messaging Servers*.
7. Select *Apply* to save the changes.
8. Select *Run Now* to generate a new On Demand report based on your changes.
9. Select the report from the Historical Reports list to view it.

Running On Demand reports can be a good way to compare report modifications as you configure.

# Logging and reporting for large networks

This section explains how to configure the FortiGate unit for logging and reporting in a larger network, such as an enterprise network. To set up this type of network, you are modifying the default log settings, and you are also modifying the default report.

The following procedures are examples and can be used to help you when configuring your own network's log topology. Since some of these settings must be modified or enabled or disabled in the CLI, it is recommended to review the FortiGate CLI Reference for any additional information about the commands used herein, as well as any that you would need to use in your own network's log topology.

The following topics are included in this section:

- [Modifying default log device settings](#)
- [Configuring the backup solution](#)
- [Modifying the default FortiOS report](#)

## Modifying default log device settings

The default log device settings must be modified so that system performance is not compromised. The FortiGate unit, by default, has all logging of FortiGate features enabled and well as logging to either the FortiGate unit's system memory or hard disk, depending on the model.

### Modifying multiple FortiGate units' system memory default settings

When the FortiGate unit's default log device is its system memory, you can modify it to fit your log network topology. In this topic, the following is an example of how you can modify these default settings.

#### To modify the default system memory settings

1. Log in to the CLI.
2. Enter the following command syntax to modify the logging settings:

```
config log memory setting
 set ips-archive disable
 set status enable
end
```



3. Enter the following command syntax to modify the FortiGate features that are enabled for logging:

```
config log memory filter
 set attack enable
 set forward-traffic enable
 set local-traffic enable
 set netscan enable
 set email-log-imap enable
 set multicast-traffic enable
 set scanerror enable
 set app-ctrl enable
end
```

4. Repeat steps 2 and 3 for the other FortiGate units.
5. Test the modified settings using the procedure [“To test sending logs to the log device” on page 1962](#).

## Modifying multiple FortiGate units' hard disk default log settings

You will have to modify each FortiGate unit's hard disk default log settings. The following is an example of how to modify these default settings.

### To modify the default hard disk settings

1. Log in to the CLI.
2. Enter the following command syntax to modify the logging settings:

```
config log disk setting
 set ips-archive disable
 set status enable
 set max-log-file-size 1000
 set storage Internal
 set log-quota 100
 set report-quota 100
 set sql-max-size 10000
 set sql-max-size-action overwrite
 set sql-oldest-entry 1024
 set rows-per-transaction 100
 set ms-per-transaction 6000
 config sql-logging
 set netscan disable
 set spam disable
 end
end
```

3. In the CLI, enter the following to disable certain event log messages that you do not want logged:

```
config log disk filter
 set discovery disable
 set email-log-imap disable
end
```

4. Repeat the steps 2 to 4 for the other FortiGate units.
5. Test the modified settings using the procedure [“To test sending logs to the log device” on page 1962](#).

## Testing the modified log settings

After modifying both the settings and the FortiGate features for logging, you can test that the modified settings are working properly. This test is done in the CLI.

### To test sending logs to the log device

1. In the CLI, enter the following command syntax:

```
diag log test
```

When you enter the command, the following appears:

```
generating a system event message with level - warning
generating an infected virus message with level - warning
generating a blocked virus message with level - warning
generating a URL block message with level - warning
generating a DLP message with level - warning
generating an attack detection message with level - warning
generating an application control IM message with level - information
generating an antispam message with level - notification
generating an allowed traffic message with level - notice
generating a multicast traffic message with level - notice
generating a ipv6 traffic message with level - notice
generating a wanopt traffic log message with level - notification
generating a HA event message with level - warning
generating netscan log messages with level - notice
generating a VOIP event message with level - information
generating a DNS event message with level - information
generating authentication event messages
generating a Forticlient message with level - information
generating a NAC QUARANTINE message with level - information
generating a URL block message with level - warning
```

2. In the web-based manager, go to *Log & Report > Event Log > User*, and view the logs to see the recently generated test log messages.

You will be able to tell the test log messages from real log messages because they do not have “real” information; for example, the test log messages for the vulnerability scan contain the destination IP address of 1.1.1.1 or 2.2.2.2. If you have disabled certain logs that are logged to the SQL log database, the following will appear in red at the top of the Log Table:

```
Warning: SQL Logging is not enabled
```

## Configuring the backup solution

Even though you are logging to multiple FortiAnalyzer units, this is more of a redundancy solution rather than a complete backup solution in this example. The multiple FortiAnalyzer units act similar to a HA cluster, since if one FortiAnalyzer unit fails, the others continue storing the logs they receive. In a backup solution, the logs are backed up to another secure location if something happens to the log device. A good alternate or redundant option is the FortiCloud service, which can provide secure online logging and management for multiple devices.

## Configuring logging to multiple FortiAnalyzer units

The following example shows how to configure logging to multiple FortiAnalyzer units. Configuring multiple FortiAnalyzer units is quick and easy; however, you can only configure up to three FortiAnalyzer units per FortiGate unit.

### To configure multiple FortiAnalyzer units

1. In the CLI, enter the following command syntax to configure the first FortiAnalyzer unit:

```
config log fortianalyzer setting
 set status enable
 set server 172.20.120.22
 set max-buffer-size 1000
 set buffer-max-send 2000
 set address-mode static
 set conn-timeout 100
 set monitor-keepalive-period 120
 set monitor-failure-retry-period 2000
end
```

2. Disable the features that you do not want logged, using the following example command syntax. You can view the CLI Reference to see what commands are available.

```
config log fortianalyzer filter
 set traffic (enable | disable)
 ...
end
```

3. Enter the following commands for the second FortiAnalyzer unit:

```
config log fortianalyzer2 setting
 set status enable
 set server 172.20.120.23
 set max-buffer-size 1000
 set buffer-max-send 2000
 set address-mode static
 set conn-timeout 100
 set monitor-keepalive-period 120
 set monitor-failure-retry-period 2000
end
```

4. Disable the features that you do not want logged, using the following example command syntax.

```
config log fortianalyzer filter
 set web (enable | disable)
 ...
end
```

5. Enter the following commands for the last FortiAnalyzer unit:

```
config log fortianalyzer3 setting
 set status enable
 set server 172.20.120.23
 set max-buffer-size 1000
 set buffer-max-send 2000
 set address-mode static
 set conn-timeout 100
 set monitor-keepalive-period 120
 set monitor-failure-retry-period 2000
end
```

6. Disable the features that you do not want logged, using the following example command syntax.

```
config log fortianalyzer filter
 set web-filter (enable | disable)
 ...
end
```

7. Test the configuration by using the procedure, [“Testing the modified log settings” on page 1962](#).
8. On the other FortiGate units, configure steps 1 through 6, ensuring that logs are being sent to the FortiAnalyzer units.

## Configuring logging to the FortiCloud server

The following explains how to send logs to a FortiCloud server. The service provides 1Gb of free space for every connected device, but you may need a larger amount of space to manage real-time logging for multiple devices. You can subscribe to the paid FortiCloud service to have 200Gb space or more available for logging.

1. Go to *System > Dashboard > Status* and click *Login* next to *FortiCloud* in the License Information widget.
2. Enter your username and password, and click OK.
3. The logs will automatically be uploaded to FortiCloud as long as you are logged in.
4. To configure the upload time and interval, go to *Log & Report > Log Config > Log Settings*.
5. Under the Logging and Archiving header, you can select your desired upload time, or select real-time logging.

## Modifying the default FortiOS report

The default FortiOS report is provided to help you quickly and easily configure and generate a report. Below is a sample configuration with multiple examples of significant customizations that you can make to tailor reports for larger networks.

## Creating datasets

You need to create a new dataset for gathering information about HA, admin activity and configuration changes.



Creating datasets requires SQL knowledge.

---

### To create the datasets

1. Log in to the CLI.

2. Enter the following command syntax:

```
config report dataset
 edit ha
 set query "select subtype_ha count(*) as totalnum from event_log
 where timestamp >= F_TIMESTAMP ('now', 'hour', '-23') and
 group by subtype_ha order by totalnum desc"
 next
```

3. Create a dataset for the admin activity, that includes log ins and log outs from the three FortiGate administrators.

```
set query "select subtype_admin count(*) as totalnum from
 event_log where timestamp >= F_TIMESTAMP ('now', 'hour',
 '-23') and group by subtype_admin order by totalnum desc"
next
```

4. Create a dataset for the configuration changes that the administrators did for the past 24 hours.

```
set query "select subtype_config count(*) as totalnum from
 event_log where timestamp >= F_TIMESTAMP ('now', 'hour',
 '-23') and group by subtype_config order by totalnum desc"
next
end
```

## Creating charts for the datasets

1. Log in to the CLI.

2. Enter the following to create a new chart:

```
config report chart
 edit ha.24h
 set type table
 set period last24h
 set dataset ha
 set category event
 set favorite no
 set style auto
 set title "24 Hour HA Admin Activity"
 end
```

## Uploading the corporate images

You need to upload the corporate images so that they appear on the report's pages, as well as on the cover page. Uploading images is only available in the web-based manager.

### To upload corporate images

1. Go to *Log & Report > Report > Local*.
2. Select the Image icon and drag it to a place on the page.
3. The Graphic Chooser window appears.
4. Select Upload and then locate the image that you want to upload and upload the image. The images are automatically uploaded and saved.
5. Repeat step 4 until the other corporate images are uploaded.
6. Select Cancel to close the Graphic Chooser window and return to the page.

The images can then be placed as you like by reopening the Graphic Chooser as in step 2.

## Adding a new report cover and page

You need to add a new cover for the report, as well as a new page that will display the HA activity, admin activity and configuration changes.

### To add and customize a new report cover

1. Go to *Log & Report > Report > Local*.
2. Select *Customize*.
3. In *Sections*, select the current default report section, and enter `Report Cover` in the field that appears; then press Enter to save the change.
4. Remove all content from the Report Cover section, and select the image icon and drag it into the main portion of the cover page; select a cover page image and then select *OK*.
5. Select the font size you want, and drag the text icon into the area beneath the image to add a title or explanation for the cover page.
6. Select *Save* to save the new report cover.

### To add and customize a new page

1. Go to *Log & Report > Report > Local*.
2. Select *Customize*.
3. Select *Sections*, and select *Create New* to add a new section to the report. Name it `Report Content`, and press Enter, and OK to close the menu.
4. At the bottom of the editing window is the Section selection, where each Section is represented by a box. Select the second box.
5. Edit the content for the report as you like.

For a simpler report structure, make use of the 'FortiGate UTM Security Analysis Report' charts, which automatically format themselves and fill in all necessary information.

For more complex reports, add headings, default and custom charts, and explanatory text.

6. Select *Save* to save the new report content.

The report will automatically combine all sections. You can use headers and text to more clearly separate parts of the report, and all properly configured charts have titles built-in.

# Advanced logging

This section explains how to configure other log features within your existing log configuration. You may want to include other log features after initially configuring the log topology because the network has either outgrown the initial configuration, or you want to add additional features that will help your network's logging requirements.

The following topics are included in this section:

- [Configuring logging to multiple Syslog servers](#)
- [Using Automatic Discovery to connect to a FortiAnalyzer unit](#)
- [Activating a FortiCloud account for logging purposes](#)
- [Viewing log storage space](#)
- [Customizing and filtering log messages](#)
- [Viewing logs from the CLI](#)
- [Configuring NAC quarantine logging](#)
- [Logging local-in policies](#)
- [Tracking specific search phrases in reports](#)
- [Creating a dataset containing attack name instead of attack ID](#)
- [Reverting modified report settings to default settings](#)
- [Customizing FortiOS reports with CLI](#)

## Configuring logging to multiple Syslog servers

When configuring multiple Syslog servers (or one Syslog server), you can configure reliable delivery of log messages from the Syslog server. Configuring of reliable delivery is available only in the CLI.

If VDOMs are enabled, each VDOM will use the default FortiAnalyzer/Syslog server, but an individual override can be enabled in the CLI, allowing you to specify a different FortiAnalyzer/Syslog server for that VDOM.

### To enable logging to multiple Syslog servers

1. Log in to the CLI.

Enter the following commands:

```
config log syslogd setting
 set csv {disable | enable}
 set facility <facility_name>
 set port <port_integer>
 set reliable {disable | enable}
 set server <ip_address>
 set status {disable | enable}
end
```

2. Enter the following commands to configure the second Syslog server:

```
config log syslogd2 setting
 set csv {disable | enable}
 set facility <facility_name>
 set port <port_integer>
 set reliable {disable | enable}
 set server <ip_address>
 set status {disable | enable}
end
```

3. Enter the following commands to configure the third Syslog server:

```
config log syslogd3 setting
 set csv {disable | enable}
 set facility <facility_name>
 set port <port_integer>
 set reliable {disable | enable}
 set server <ip_address>
 set status {disable | enable}
end
```

Most FortiGate features are, by default, enabled for logging. You can disable individual FortiGate features you do not want the Syslog server to record, as in this example:

```
config log syslogd filter
 set traffic {enable | disable}
 set web {enable | disable}
 set url-filter {enable | disable}
end
```

#### **To enable/disable override settings per-VDOM**

```
config log fortianalyzer override-filter
 set override {enable | disable}
 <configuration commands>
end
config log fortianalyzer override-setting
 set override {enable | disable}
 <configuration commands>
end
config log syslogd override-filter
 set override {enable | disable}
 <configuration commands>
end
config log syslogd override-setting
 set override {enable | disable}
 <configuration commands>
end
```

## **Using Automatic Discovery to connect to a FortiAnalyzer unit**

Automatic Discovery can be used if the FortiAnalyzer unit is on the same network.



### To connect using automatic discovery

1. Log in to the CLI.
2. Enter the following command syntax:

```
config log fortianalyzer setting
 set status enable
 set server <ip_address>
 set gui-display enable
 set address-mode auto-discovery
end
```

If your FortiGate unit is in Transparent mode, the interface using the automatic discovery feature will not carry traffic. For more information about how to enable the interface to also carry traffic when using the automatic discovery feature, see the Fortinet Knowledge Base article, [Fortinet Discovery Protocol in Transparent mode](#).



The FortiGate unit searches within the same subnet for a response from any available FortiAnalyzer units.

---

## Activating a FortiCloud account for logging purposes

When you subscribe to FortiCloud, you can configure to send logs to the FortiCloud server. The account activation can be done within the web-based manager, from the *License Information* widget located in *System > Dashboard*.

From this widget, you can easily create a new account, or log in to the existing account. From within the License Information widget, after the account is activated, you can go directly to the FortiCloud web portal, or log out of the service if you are already logged in.

### To activate a FortiCloud account for logging purposes:

The following assumes that you are already at *System > Dashboard* and that you have located the License Information widget.

1. In the License Information widget, select *Activate* in the *FortiCloud* section.  
The Registration window appears. From this window, you create the login credentials that you will use to access the account.
2. Select *Create Account* and enter then information for the login credentials.  
After entering the login credentials, you are automatically logged in to your FortiCloud account.
3. Check that the account has been activated by viewing the account status from the License Information widget.

If you need more space, you can subscribe to the 200Gb FortiCloud service by selecting *Upgrade* in the *FortiCloud* section of the widget.

## Viewing log storage space

The `diag sys logdisk usage` command allows you to view detailed information about how much space is currently being used for logs. This is useful when you see a high percentage, such as 92 percent for the disk's capacity. The FortiGate unit uses only 75 percent of the available disk capacity to avoid a high storage amount so when there is a high percentage, it

refers to the percentage of the 75 percent that is available. For example, 92 percent of the 75 percent is available.

The following is an example of what you may see when you use `diag sys logdisk usage` command on a unit with no VDOMs configured:

```
diag sys logdisk usage
```

The following appears:

```
Total HD usage: 176MB/3011 MB
Total HD logging space: 22583MB
Total HD logging space for each vdom: 22583MB
HD logging space usage for vdom "root": 30MB/22583MB
```

If you want to view the size of the log database, as well as the log database table entries, use the `get log sql status` command. This command displays the amount of free space that is available as well as the first log database entry time and date and the last log database entry time and date.

The output of the `get log sql status` command contains information similar to the following:

```
Database size: 104856576
Free size in database: 670004416
Entry number:
 Event: 1263
 Traffic: 254039
 Attack: 4
 Antivirus: 8
 WebFilter: 5291
 DLP: 76544
 Application Control: 68103
 Netscan: 75
Total: 405331
First entry time: 2011-03-21 08:25:55
Last entry time: 2011-04-21 09:10:55
```

## Customizing and filtering log messages

When viewing log messages, you may want to customize and filter the information that you are seeing in the Log & Report menu (for example, Log & Report > Traffic Log > Forward Traffic). Filtering and customizing the display provides a way to view specific log information without scrolling through pages of log messages to find the information.

Customizing log messages is the process of removing or adding columns to the log display page, allowing you to view certain desired information. The most columns represent the fields from within a log message, for example, the user column represents the user field, as well as additional information. If you want to reset the customized columns on the page back to their defaults, you need to select *Reset All Columns* within the column title right-click menu.

Filtering information is similar to customizing, however, filtering allows you to enter specific information that indicates what should appear on the page. For example, including only log messages that appeared on February 24, between the hours of 8:00 and 8:30 am.

## To customize and filter log messages

The following is an example that displays all traffic log messages that originate from the source IP address 172.20.120.24, as well as displaying only the columns:

- OS Name
- OS Version
- Policy ID
- Src (Source IP)

The following assumes that you are already on the page of the log messages you want to customize and filter. In this example, the log messages that we are customizing and filtering are in *Log & Report > Traffic Log > Forward Traffic*.

1. On the Forward Traffic page, right click anywhere on a column title.
2. Right click on a column title, and mouse over *Column Settings* to open the list.
3. Select each checkmarked title to uncheck it and remove them all from the displayed columns.
4. Scroll down to the list of unchecked fields and select 'OS Name', 'OS Version', 'Policy ID', and 'Src' to add checkmarks next to them.
5. Click outside the menu, and wait for the page to refresh with the new settings in place.
6. Select the funnel icon next to the word Src in the title bar of the Src column.
7. Enter the IP you want displayed (in this example, 172.20.120.24) in the text box.
8. Press Apply, and wait for the page to reload.

## Viewing logs from the CLI

You can easily view log messages from within the CLI. In this example, we are viewing DLP log messages.

1. Log in to the CLI and then enter the following to configure the display of the DLP log messages.

```
execute log filter category 9
execute log filter start-line 1
execute log filter view-lines 20
```

The customized display of log messages in the CLI is similar to how you customize the display of log messages in the web-based manager. For example, `category 9` is the DLP log messages, and the `start-line` is the first line in the log database table for DLP log messages, and there will be 20 lines (`view-lines 20`) that will display.

2. Enter the following to view the log messages:

```
execute log display
```

The following appears below `execute log display`:

```
600 logs found
```

```
20 logs returned
```

along with the 20 DLP log messages.

## Configuring NAC quarantine logging

NAC quarantine log messages provide information about what was banned and quarantined by a DLP sensor. The following explains how to configure NAC quarantine logging and enable it on

a policy. This procedure assumes the DLP sensor is already in place. View the UTM Handbook for more details on DLP.

### To configure NAC quarantine logging

1. Go to *Policy > Policy > Policy*.
2. Select the security policy that you want to apply the DLP profile to, and then select *Edit*.
3. Within the Security Profiles section, enable *DLP Sensor* and then select the profile from the drop-down list.
4. Select *OK*.
5. Log in to the CLI.
6. Enter the following to enable NAC quarantine in the DLP sensor:

```
config dlp sensor
 edit <dlp_sensor_name>
 set nac-quar-log enable
 end
```

## Logging local-in policies

Local-in security policies are policies that control the flow of internal traffic, and can be used to broaden or restrict an administrator's access privileges. These local-in policies can also be configured to log traffic and activity that the policies control.

You can enable logging of local-in policies from either the web-based manager or the CLI. In the web-based manager, you must first enable local-in policies in *System > Settings > Features* by enabling *Local-In Policy* and selecting *Apply*. The Local-In Policy page will then be available in *Policy > Policy > Local-In Policy*.

Use the following table when deciding what to log when local-in policy activity is occurring.

**Table 83: Local-in Policy Options**

Log options for Local-in policies	Description
Enable Logging for Denied Traffic	This records all implicit local deny or a local-in policy that has the action deny. For example, someone trying to log in to a port 80 that is not allowed by the local-in policy.
Enable Logging for Allowed Traffic	This records all administrator, system, user, and FortiGuard traffic.
Enable Logging for Local Out Traffic	This records all traffic leaving the FortiGate.

When deciding what local-in policy traffic you want logged, consider the following:

**Table 84: Special Traffic**

Traffic activity	Traffic Direction	Description
<b>FortiGuard update announcements</b>	IN	All push announcements of updates that are coming from the FortiGuard system. For example, IPS or AV updates.
FortiGuard update requests	OUT	All updates that are checking for antivirus or IPS as well as other FortiGuard service updates.
Firewall authentication	IN	The authentication made using either the web-based manager or CLI.
Central management (a FortiGate unit being managed by a FortiManager unit)	IN	The access that a FortiManager has managing the FortiGate unit.
DNS	IN	All DNS traffic.
DHCP/DHCP Relay	IN	All DHCP and/or DHCP Relay traffic.
HA (heart beat sync policy)	IN/OUT	For high-end platforms with a backplane heart beat port.
HA (Session sync policy)	IN/OUT	This will get information from the CMDB and updated by session sync daemon.
CAPWAP	IN	This activity is logged only when a HAVE_CAPWAP is defined.
Radius	IN	This is recorded only within FortiCarrier.
NETBIOS forward	IN	Any interface that NETBIOS forward is enabled on.
RIP	IN	
OSPF	IN	
VRRP	IN	
BFD	IN	
IGMP	IN	This is recorded only when PIM is enabled.
PIM	IN	This is recorded only when PIM is enabled.
BGP	IN	This is recorded only when config bgp and bgp neighbor is enabled in the CLI.
WCCP policy	IN	Any interface that WCCP is enabled; however, if in Cache mode, this is not recorded because it is not available.
WAN Opt/ Web Cache	IN	Any interface where WAN Opt is enabled.

**Table 84: Special Traffic**

Traffic activity	Traffic Direction	Description
WANOpt Tunnel	IN	This is recorded when HAVE_WANOPT is defined.
SSL-VPN	IN	Any interface from a zone where the action in the policy is SSL VPN.
IPSEC	IN	
L2TP	IN	
PPTP	IN	
VPD	IN	This is recorded only when FortiClient is enabled.
Web cache db test facility	IN	This is recorded only when WA_CS_REMOTE_TEST is defined.
GDBserver	IN	This is recorded only when debug is enabled.

## Tracking specific search phrases in reports

It is possible to use the Web Filter to track specific search keywords and phrases and record the results for display in the report.

You should verify that the web filter profile you are using indicates what search phrases you want to track and monitor, so that the report includes this information.

**1. Log in to the CLI and enter show webfilter profile default.**

This provides details about the webfilter profile being used by the security policy. In this example, the details (shown in the following in bold) indicate that safe search is enabled, but not specified or being logged.

```
show webfilter profile default
 config webfilter profile
 edit "default"
 set comment "default web filtering"
 set inspection-mode flow-based
 set options https-scan
 set post-action comfort
 config web
 set safe-search url
 end
 config ftgd-wf
 config filters
 edit 1
 set action block
 set category 2
 next
 edit 2
 set action block
 set category 7
 next
 edit 3
 set action block
 set category 8
```

**2. Enter the following command syntax so that logging and the keyword for the safe search will be included in logging.**

```
config webfilter profile
 edit default
 config web
 set log-search enable
 set keyword-match "fortinet" "easter" "easter bunny"
 end
 end
```

**3. To test that the keyword search is working, go to a web browser and begin searching for the words that were included in the webfilter profile, such as easter.**

You can tell that the test works by going to *Log & Report > Traffic Log > Forward Traffic* and viewing the log messages.

## Creating a dataset containing attack name instead of attack ID

If you want to create a dataset that contains the attack name instead of the attack ID, use the following as a basis.

```
config report dataset
 edit top_attacks_1hr
 set log-type attack
 set time-period last-n-hours
 set period-last-n 1
 set query "SELECT attack_id, COUNT(*) AS totalnum FROM $log
 WHERE $filter and attack_id IS NOT NULL GROUP BY attack_id
 ORDER BY totalnum DESC LIMIT 10"
 end
```

## Reverting modified report settings to default settings

If you need to go back to the original default report settings, you can easily revert to those settings in the Report menu. Reverting to default settings means that your previously modified report settings will be lost.

To revert back to default report settings, in *Log & Report > Report > Local*, select *Restore Defaults* from the top navigation. This may take a minute or two. You can also use the CLI command `execute report-config reset` to reset the report to defaults.

If you are having problems with report content being outdated or incorrect, especially after a firmware update, you can recreate the report database using your current log information with the CLI command `execute report recreate-db`.

## Customizing FortiOS reports with CLI

You may want to configure a FortiOS report to create a specific report that contains only specific charts and images. A FortiOS report is created in the CLI. A FortiOS report is time-consuming, so allow plenty of time for creating this type of report.

Reports are made up of four components:

- Charts, which are the content of the report
- Datasets, which determine what information is contained in the charts
- Styles, which determine the appearance (text style, etc) of the content
- and Themes, which determine the structure of the report

You should also review the [FortiGate CLI Reference](#) since it contains detailed information about each command used in the following.

### Configuring a style

There are default styles and summary styles to choose from; however, you may want to create your own styles. You can also customize the default styles.



## To customize default styles

1. Log in to the CLI.

Enter the following command syntax:

```
config report style
 edit style <style_name>
```

For example default.graph.

To view a list of available styles, enter ? after entering edit.

## To create a new style

1. Log in to the CLI.

Enter the following command syntax:

```
config report style
 edit <new_style_name>
 set options {align | border | color | column | font | margin |
 padding | size | text}
 set align {center | justify | left | right }
 set bg-color {color_name1 | color_name2 | color_name3 | ...}
 set border-bottom <border_width_pixels> <border_style_{solid |
 dotted | dashed}> <border_color>
 set border-left <border_width_pixels> <border_style_{solid |
 dotted | dashed}> <border_color>
 set border-top <border_width_pixels> <border_style_{solid |
 dotted | dashed}> <border_color>
 set column-gap <pixels>
 set column-span {all | none}
 set fg-color <color>
 set font-family {Arial | Courier | Helvetica | Times | Verdana}
 set font-size {xx-small | x-small | small | medium | large |
 x-large | xx-large | <pixels>}
 set font-style {italic | normal}
 set font-weight {bold | normal}
 set height <pixels or percentage>
 set line-height <pixels or percentage>
 set margin-bottom <pixels>
 set margin-left <pixels>
 set margin-right <pixels>
 set margin-top <pixels>
 set padding-bottom <pixels>
 set padding-left <pixels>
 set padding-right <pixels>
 set padding-top <pixels>
 set width <pixels or percentage>
 end
```

## Example

Since you need to configure several different styles for different parts of the report, such as footers and headers, the following is grouped into parts.

### **To configure footers and headers**

```
config report style
 edit web-footerheader
 set options align font
 set align center
 set font-size xx-small
 set font-family Arial
 set font-weight normal
 set font-style normal
 next
```

### **To configure pages**

```
edit web-pages
 set options align margin column
 set align justify
 set margin-top 5
 set margin-bottom 5
 set margin-right 6
 set margin-left 7
 set column-gap 3
 set column-span all
next
```

## To configure the table of contents, headings and title

```
edit web-toc-title
 set options align font
 set font-family Arial
 set font-weight bold
 set font-style normal
 set font-size x-large
 set align center
next
edit web-toc-heading1
 set options align font
 set font-weight bold
 set font-size large
 set font-style normal
 set align left
next
edit web-toc-heading2
 set options align font
 set font-family Arial
 set font-size medium
 set font-style normal
 set font-weight bold
 set align left
next
edit web-toc-heading3
 set options align font
 set font-family Arial
 set font-size medium
 set font-style italic
 set font-weight bold
 set align left
next
```

### To configure the page headings style

```
edit web-page-heading1
set options align font
set font-family Arial
set font-size large
set font-style normal
set font-weight bold
set align left
next
edit web-page-heading2
set options align font
set font-family Arial
set font-size medium
set font-style normal
set font-weight bold
next
edit web-page-heading3
set options align font
set font-family Arial
set font-size medium
set font-style italic
set font-weight bold
next
```

### To configure the chart style

```
edit web-chart
set options align font
set font-family Arial
set font-size small
set font-style normal
set font-weight bold
set align font
next
```

### To configure the cover page

```
edit web-cover
set options align font
set font-family Arial
set font-size xx-large
set font-style normal
set font-weight bold
set align center
end
end
```

## Configuring a theme

A theme is a group of settings that creates the general style of a report. For example, the styles that are applied to the table of contents section of the report.

## To configure a theme for a report

### 1. Log in to the CLI.

Enter the following command syntax:

```
config report theme
 edit <theme_name>
 set column-count [1 | 2 | 3]
 set default-html-style <string>
 set default-pdf-style <string>
 set graph-chart-style <string>
 set heading1-style <string>
 set heading2-style <string>
 set heading3-style <string>
 set heading4-style <string>
 set hline-style <string>
 set image-style <string>
 set normal-text-style <string>
 set page-footer-style <string>
 set page-header-style <string>
 set page-orient [landscape | portrait]
 set page-style <string>
 set report-subtitle-style <string>
 set report-title-style <string>
 set table-chart-caption-style <string>
 set table-chart-even-row-style <string>
 set table-chart-head-style <string>
 set table-chart-odd-row-style <string>
 set table-chart-style <string>
 set toc-heading1-style <string>
 set toc-heading2-style <string>
 set toc-heading3-style <string>
 set toc-heading4-style <string>
 set toc-title-style <string>
 end
```

To choose a style for any one of the above commands, except for column-count and page-orient, enter ? to view the available choices.

## Example of a theme

```
config report theme
 edit web-theme
 set column-count 2
 set default-pdf-style web-cover
 set graph-chart-style web-chart
 set page-orient landscape
 set page-style web-pages
 set table-chart-style web-chart
 set toc-title-style web-toc-title
 set toc-heading1-style web-toc-heading1
 set toc-heading2-style web-toc-heading2
 set toc-heading3-style web-toc-heading3
 set heading1-style web-heading1
 set heading2-style web-heading2
 set heading3-style web-heading3
 set page-footer-style web-footerheader
 set page-header-style web-footerheader
 set report-title-style web-cover
 end
```

## Configuring charts

Charts display the log information in a clear and concise way using a graph. You can find out what information is in each default chart by using the following get command:

```
config report chart
 edit <chart_name>
 get
```

The information displays similar to the following:

```
name: web.allowed-request.sites.user
policy: 0
type: graph
period: last24h
comments: (null)
dataset: web.allowed-request.sites.user
category: webfilter
favorite: no
graph-type: bar
style: auto
dimension: 3D
 x-series
 caption: (null)
 databind: field(1)
 is-category: yes
 label-angle: 45-degree
 unit: (null)
 y-series
 caption: Requests
 databind: field(2)
 extra-y: disable
 group: (null)
 label-angle: horizontal
 unit: (null)
title: Top Allowed Web Sites for User by Request legend: enable
```

### Example of a new chart

```
config report chart
 edit chart-web
 set type table
 set period last24h
 set comments "To view web activity information for the past 24
 hours"
 set dataset dataset-web-last24h
 set category webfilter
 set favorite no
 set style auto
 set title "Internet searches in the past 24 hours"
 end
```

## Adding a chart

The following is an example of how to add a Client Reputation Summary chart to the report using the CLI:

```
config report layout
 edit default
 config body-item
 edit 701
 set type chart
 set chart cr.summary
 end
 end
end
```



# Troubleshooting and logging

This section explains how to troubleshoot logging configuration issues, as well as connection issues, that you may have with your FortiGate unit and a log device. This section also contains information about how to use log messages when troubleshooting issues that are about other FortiGate features, such as VPN tunnel errors.

The following topics are included in this section:

- [Using log messages to help in troubleshooting issues](#)
- [Connection issues between FortiGate unit and logging devices](#)
- [Log database issues](#)
- [Logging daemon \(Miglogd\)](#)

## Using log messages to help in troubleshooting issues

Log messages can help when troubleshooting issues that occur, since they can provide details about what is occurring. The uses and methods for involving logging in troubleshooting vary depending on the problem. The following are examples of how log messages can assist when troubleshooting networking issues.

### Using IPS packet logging in diagnostics

This type of logging should only be enabled when you need to know about specific diagnostic information, for example, when you suspect a signature is triggered by a false positive. These log messages can help troubleshoot individual problems with misidentified or missing packets and network intrusions involving malicious packets.

#### To configure IPS packet logging

1. Go to *Security Profiles > Intrusion Protection > IPS Sensors*.
2. Select the IPS sensor that you want to enable IPS packet logging on, and then select *Edit*.
3. Within the sensor, select *Filter Options*.
4. In the filter options, enable *Packet Logging*.
5. Select *OK*.

If you want to configure the packet quota, number of packets that are recorded before alerts and after attacks, use the following procedure.

#### To configure additional settings for IPS packet logging

1. Log in to the CLI.
2. Enter the following to start configuring additional settings:

```
config ips settings
 set ips-packet-quota <integer>
 set packet-log-history <integer>
 set packet-log-post-attack <integer>
end
```

## Using HA log messages to determine system status

When the FortiGate unit is in HA mode, you may see the following log message content within the event log:

```
type=event subtype=ha level=critical msg= "HA slave heartbeat interface
internal lost neighbor information"
```

OR

```
type=event subtype=ha level=critical msg= "Virtual cluster 1 of group 0
detected new joined HA member"
```

OR

```
type=event subtype=ha level=critical msg= "HA master heartbeat
interface internal get peer information"
```

The log messages occur within a given time, and indicate that the units within the cluster are not aware of each other anymore. These log messages provide the information you need to fix the problem.

## Connection issues between FortiGate unit and logging devices

If external logging devices are not recording the log information properly or at all, the problem will likely be due to one of two situations: no data is being received because the log device cannot be reached, or no data is being sent because the FortiGate unit is no longer logging properly.

### Unable to connect to a supported log device

After configuring logging to a supported log device, and testing the connection, you may find you cannot connect. To determine whether this is the problem:

1. Verify that the information you entered is correct; it could be a simple mistake within the IP address or you may have not selected *Apply* on the Log Settings page after changing them, which would prevent them from taking effect.
2. Use `execute ping` to see if you can ping to the log device.
3. If you are unable to ping to the log device, check to see if the log device itself working and that it is on the network and assigned an appropriate address.

### FortiGate unit has stopped logging

If the FortiGate unit stopped logging to a device, test the connection between both the FortiGate unit and device using the `execute ping` command. The log device may have been turned off, is upgrading to a new firmware version, or just not working properly.

The FortiGate unit may also have a corrupted log database. When you log into the web-based manager and you see an SQL database error message, it is because the SQL database has become corrupted. View "[SQL database errors](#)" in the next section before taking any further actions, to avoid losing your current logs.

## Log database issues

If attempting to troubleshoot issues with the SQL log database, use the following to help guide you to solving issues that occur.

## SQL statement syntax errors

There may be errors or inconsistencies in the SQL used to maintain the database. Here are some example error messages and possible causes:

You have an error in your SQL syntax (remote/MySQL) or

ERROR: syntax error at or near... (local/PostgreSQL)

- Verify that the SQL keywords are spelled correctly, and that the query is well-formed.
- Table and column names are demarked by grave accent (`) characters. Single (') and double (") quotation marks will cause an error.

No data is covered.

- The query is correctly formed, but no data has been logged for the log type. Verify that you have configured the FortiGate unit to save that log type. On the Log Settings page, make sure that the log type is checked.

## Connection problems

If well-formed SQL queries do not produce results, and logging is turned on for the log type, there may be a database configuration problem with the remote database.


Ensure that:

- MySQL is running and using the default port 3306.
- You have created an empty database and a user who has read/write permissions for the database.
- Here is an example of creating a new MySQL database named fazlogs, and adding a user for the database:
  1. #Mysql -u root -p
  2. mysql> Create database fazlogs;
  3. mysql> Grant all privileges on fazlogs.\* to 'fazlogger'@'\*' identified by 'fazpassword';
  4. mysql> Grant all privileges on fazlogs.\* to 'fazlogger'@'localhost' identified by 'fazpassword';

## SQL database errors

If the database seems inaccessible, you may encounter the following error message (Figure 295) after upgrading or downgrading the FortiGate unit's firmware image.

**Figure 295:**Example of an SQL database error message



The error message indicates that the SQL database is corrupted and cannot be updated with the SQL schemas any more. When you see this error message, you can do one of the following:

- select *Cancel* and back up all log files; then select *Rebuild* to blank and rebuild the database.
- select *Rebuild* immediately, which will blank the database and previous logs will be lost.

Until the database is rebuilt, no information will be logged by the FortiGate unit regardless of the log settings that are configured on the unit. When you select *Rebuild*, all logs are lost because the SQL database is erased and then rebuilt again. Logging resumes automatically according to your settings after the SQL database is rebuilt.

To view the status of the database, use the `diagnose debug sqlldb-error status` command in the CLI. This command will inform you whether the database has errors present.

If you want to view the database's errors, use the `diagnose debug sqlldb-error read` command in the CLI. This command indicates exactly what errors occurred, and what tables contain those errors.

Log files are backed up using the `execute backup {disk | memory } {alllogs | logs}` command in the CLI. You must use the text variable when backing up log files because the text variable allows you to view the log files outside the FortiGate unit. When you back up log files, you are really just copying the log files from the database to a specified location, such as a TFTP server.

## Logging daemon (Miglogd)

The number of logging daemon child processes has been made available for editing. A higher number can affect performance, and a lower number can affect log processing time, although no logs will be dropped or lost if the number is decreased.

If you are suffering from performance issues, you can alter the number of logging daemon child processes, from 0 to 15, using the following syntax. The default is 8.

```
config system global
 set miglogd-children <integer>
end
```

# Appendix: FortiGate report charts

The following tables explain the charts that are available for reports in FortiOS 5.0, which includes what log file or files the FortiGate unit gathers the information needed for the chart. This appendix can help you when determining which charts to include in reports, without having to view them from the CLI.

This appendix also includes the dataset and drill-down chart that is included within the chart itself. Drill-down charts are not included in every chart.

This section includes the following topics:

- [Traffic charts](#)
- [Web filter charts](#)
- [IPS \(or attack\) charts](#)
- [Antivirus charts](#)
- [Email filter charts](#)
- [VPN charts](#)

## Traffic charts

Traffic charts contain information about the traffic activity, which includes email, terminal and instant messaging traffic. Terminal activity is activity that is SSH or Telnet, which is used to log in to the FortiGate unit's from a remote location.

The information from these charts are taken from the traffic log table in the SQLite database; however, the traffic.bandwidth.wanopt chart contains information from the traffic log's subtype wanopt which is stored within the traffic log table.

The information displayed in these charts is in a bar chart style; however, the traffic.bandwidth.wanopt displays information in a pie chart style.

**Table 85:** Explanation of the traffic charts

Chart	Dataset included in chart	Explanation of the chart
traffic.bandwidth.apps.app_cat	traffic.bandwidth.apps.app_cat	The top applications for an application category by bandwidth.
traffic.bandwidth.app_cats.user	traffic.bandwidth.app_cats.user	The top application categories for a user by bandwidth.
traffic.bandwidth.users	traffic.bandwidth.users	The top user traffic by bandwidth.
traffic.sessions.apps.app_cat	traffic.sessions.apps.app_cat	The top applications for sessions by application category.
traffic.sessions.app_cats.user	traffic.sessions.app_cats.user	The top sessions by application categories for a user.

**Table 85:** Explanation of the traffic charts (Continued)

traffic.session.users	traffic.sessions.users	The top users by sessions.
traffic.bandwidth.apps.user	traffic.bandwidth.apps.user	The top applications for a user by bandwidth.
traffic.bandwidth.users.app	traffic.bandwidth.users.app	The top users using applications by bandwidth.
traffic.bandwidth.app_cats	traffic.bandwidth.app_cats	The top application categories by bandwidth.
traffic.sessions.apps.user	traffic.sessions.apps.user	The top applications that are used by a user by sessions.
traffic.sessions.users.app	traffic.sessions.users.app	The top users for application by sessions.
traffic.sessions.app_cats	traffic.sessions.app_cats	The top application categories by session.
traffic.bandwidth.wanopt	traffic.bandwidth.wanopt	The WAN Optimization and web cache performance information.

## Web filter charts

Web filter charts contain information about the web activity on the network. The information from these charts are taken from the web filter table in the SQLite database. All web filter charts display their information within a bar chart style.

**Table 86:** Web filter charts and matching datasets

Chart	Dataset included in chart	Explanation of the chart
web.allowed-requests.sites.user	web.allowed-request.sites.user	The top allowed web sites a user has accessed by requests.
web.allowed-requests.users.web_cat	web.allowed-request.users.web_cat	The top allowed users for web category by requests.
web.allowed-requests.web_cats	web.allowed-request.web_cats	The top allowed web categories by request.
web.blocked-requests.sites.user	web.blocked-request.sites.user	The top blocked web sites that a user tried to access by request.
web.blocked-requests.users.web_cat	web.blocked-request.users.web_cat	The top blocked users for web category by requests.
web.blocked-requests.web_cats	web.blocked-request.web_cats	The top blocked web categories by requests.
web.requests.phrases.user	web.requests.phrases.user	The top search phrases for a user.

**Table 86:** Web filter charts and matching datasets

web.requests.phrases	web.requests.phrases	The top search phrases.
web.requests.users.phrase	web.requests.users.phrase	The top users that used search phrase.
web.allowed-requests.users.site	web.allowed-request.users.site	The top allowed users for web sites by requests.
web.allowed-requests.sites	web.allowed-request.sites	The top allowed web sites by requests.
web.blocked-requests.users.site	web.blocked-request.users.site	The top blocked users for web site by request.
web.blocked-requests.sites	web.blocked-request.sites	The top blocked web sites by requests.
web.bandwidth.sites.user	web.bandwidth.sites.user	The top web sites that a user has accessed by bandwidth.
web.bandwidth.users.site	web.bandwidth.users.site	The top users for web site by bandwidth.
web.bandwidth.sites	web.bandwidth.sites	The top web sites by bandwidth.
web.bandwidth.stream-sites.user	web.bandwidth.stream-sites.user	The top video streaming web sites for user by bandwidth.
web.bandwidth.users.stream-site.	web.bandwidth.users.stream-site	The top users for video streaming web sites by bandwidth.
web.bandwidth.stream-sites	web.bandwidth.stream-sites	The top video streaming web sites by bandwidth.

## IPS (or attack) charts

The IPS, or attack, charts contain information about attacks that occurred on your network. The information from these charts is taken from the attack table in the SQLite database. All IPS charts display their information in a bar chart style.

**Table 87:** Attack charts and matching datasets

Chart	Dataset included in chart	Explanation of the chart
attack.count.critical-attacks.user	attack.count.critical-attacks.user	The top critical or high level attacks for a user.
attack.count.users.critical-attack	attack.count.users.critical-attack	The top users for an attack.
attack.count.critical-attacks	attack.count.critical-attacks	The top high or critical attacks.

**Table 87:** Attack charts and matching datasets

attack.count.attacks.user	attack.count.attacks.user	The top attacks for a user.
attack.count.users.attack	attack.count.users.attack	The top ten users for an attack.
attack.count.attacks	attack.count.attacks	The top attacks that occurred.

## Antivirus charts

The antivirus charts contain information about any potential or suspicious viruses that were detected by the FortiGate unit, as well as information about files that may have contained viruses.

Antivirus chart information is taken from the virus chart table in the SQLite database. All antivirus chart information is displayed using a bar chart style.

**Table 88:** Antivirus charts and matching datasets

Chart	Dataset included in chart	Explanation of the chart
virus.count.viruses.user	virus.count.viruses.user	The top viruses for user.
virus.count.users.virus	virus.count.users.virus	The top users targeted by a virus.
virus.count.viruses	virus.count.viruses	The top viruses that were detected.
virus.count.viruses.protocol	virus.count.viruses.protocol	The top viruses per protocol.
virus.count.protocols	virus.count.protocols	The virus protocol distribution.
virus.count.users	virus.count.users	The top users targeted by viruses.

## Email filter charts

Email filter charts contain information about email filtering that occurred, such as how many spam emails versus clean emails were detected by the FortiGate unit. These charts can help in determining the amount of emails that are coming into and going out of your network.

The email filter information is taken from the email filter table in the SQLite database. All email filter charts display the information in a bar chart style except for email.bandwidth.timeperiods.sender and email.request.timeperiods.receiver, which display their information using a line chart style.



**Table 89:** Email filter charts and matching datasets

Chart	Dataset included in chart	Explanation of the chart
email.request.timeperiods.sender	email.request.timeperiods.sender	The number of emails that were sent from the sender email address.
email.request.senders	email.request.senders	The top email senders.
email.bandwidth.timeperiods.sender	email.bandwidth.timeperiods.sender	The bandwidth used by email messages that were sent from the sender email address.
email.bandwidth.senders	email.bandwidth.senders	The top email senders by bandwidth. This chart displays information using bar chart.
email.request.timeperiods.receiver	email.request.timeperiods.receiver	The number of emails that were sent to the recipient.
email.request.receivers	email.request.receivers	The top email receivers.

## VPN charts

VPN charts contain information about VPN activity. This information is gathered from the event-vpn log messages and includes VPN user activity, IPsec and SSL tunnel activity, and tunnel activity volume.

**Table 90:** VPN charts and matching datasets

Chart	Dataset included in chart	Explanation of the chart
vpn.bandwidth.static-tunnels.user	vpn.bandwidth.static-tunnels.user	The bandwidth for static VPN tunnels per user.
vpn.bandwidth.users.static-tunnel	vpn.bandwidth.users.static-tunnel	The bandwidth used per user in static VPN tunnels.
vpn.bandwidth.static-tunnels	vpn.bandwidth.static-tunnels	The bandwidth used in static VPN tunnels.
vpn.bandwidth.ssl-sources.user	vpn.bandwidth.ssl-sources.user	The bandwidth used for source IP addresses of SSL VPN users
vpn.bandwidth.users-ssl-source	vpn.bandwidth.users-ssl-source	The bandwidth used for users that use SSL VPN tunnels and those source IP addresses of the users.
vpn.bandwidth.ssl-sources	vpn.bandwidth.ssl-sources	The bandwidth of SSL VPN connections and their source IP addresses.
vpn.bandwidth.dynamic-tunnels.user	vpn.bandwidth.dynamic-tunnels.user	The bandwidth of dynamic VPN tunnels per user.

**Table 90:** VPN charts and matching datasets

vpn.bandwidth.users.dynamic-tunnel	vpn.bandwidth.users.dynamic-tunnel	The bandwidth used per user for dynamic VPN tunnels.
vpn.bandwidth.dynamic-tunnels	vpn.bandwidth.dynamic-tunnels	The bandwidth used for dynamic VPN tunnels.

# Chapter 14 Managing Devices for FortiOS 5.0

This FortiOS Handbook chapter contains the following sections:

[Managing “bring your own device”](#) describes device monitoring, devices, device groups, and device policies. The administrator can monitor all types of devices and control their access to network resources.

[Endpoint Protection](#) describes how you can enforce the use of FortiClient Endpoint Control and apply an endpoint profile to users’ devices. Endpoint profiles include real-time antivirus protection, application control, web category filtering, and VPN provisioning.

[Vulnerability Scan](#) describes how perform network vulnerability scanning to look for security weaknesses in your servers and workstations.

# Managing “bring your own device”

FortiOS can control network access for different types of personal mobile devices that your employees bring onto your premises. You can:

- identify and monitor the types of devices connecting to your networks, wireless or wired
- use MAC address based access control to allow or deny individual devices
- create policies based on device type
- enforce endpoint control on devices that can run FortiClient Endpoint Control software

## Device monitoring

The FortiGate unit can monitor your networks and gather information about the devices operating on those networks. Collected information includes:

- MAC address
- IP address
- operating system
- hostname
- user name
- how long ago the device was detected and on which FortiGate interface

You can go to *User & Device > Device > Device Definitions* to view this information.

Device	OS	User	Hostname	IP Address	Custom Group	FortiClient State	Last Seen	Alias
18:03:73:b6:f9:e9				172.20.120.100		N/A	2 minutes ago (wan1)	
00:0b:82:17:a2:de						N/A	yesterday (wan1)	
18:03:73:89:1b:25			Marc-PC	172.20.120.235		N/A	2 minutes ago (wan1)	
78:2b:cb:d8:36:68				172.20.120.71		N/A	1 second ago (wan1)	
00:09:0f:fe:d0:67				172.20.120.136		N/A	1 minute ago (wan1)	
00:26:ab:9b:9e:63				172.20.120.111		N/A	6 minutes ago (wan1)	
a8:20:66:14:fa:da			wd-mb	172.20.120.226		N/A	5 seconds ago (wan1)	
c4:2c:03:21:a9:8e				172.20.120.83		N/A	7 seconds ago (wan1)	
00:0c:29:ba:54:2e				172.20.120.54		N/A	3 minutes ago (wan1)	
00:09:0f:4e:70:b1				172.20.120.122		N/A	yesterday (wan1)	
Bob	iPhone / iOS	A		10.10.82.4		N/A		Bob
jcoles-mac	Mac OS X / 10.x			172.20.120.51	Employees	N/A	1 second ago (wan1)	jcoles-mac
00:0c:29:92:7f:4a				172.20.120.52		N/A	24 seconds ago (wan1)	
00:0c:29:73:1e:df				172.20.120.13		N/A	1 minute ago (wan1)	
00:09:0f:15:04:86						N/A	1 minute ago (wan1)	
f0:4d:a2:f1:d3:4a				172.20.120.36		N/A	11 seconds ago (wan1)	
00:24:e8:e0:98:66			akaye-notebook	172.20.120.223		N/A	yesterday (wan1)	
f0:4d:a2:f1:d6:60				172.20.120.46		N/A	1 minute ago (wan1)	
c4:2c:03:21:af:04				172.20.120.14		N/A	1 second ago (wan1)	
00:09:0f:99:4b:e4			FG100D3G12804410			N/A	yesterday (wan1)	
00:0c:29:df:22:b0			bill-0b2i3pig5	172.20.120.222		N/A	yesterday (wan1)	
00:0c:29:93:6d:bd			FortiGate-VM			N/A	yesterday (wan1)	
00:09:0f:35:6d:41			FAP22B3U11022065	172.20.120.230		N/A	yesterday (wan1)	
f0:4d:a2:f1:bf:a3				172.20.120.26		N/A	yesterday (wan1)	
00:09:0f:67:2d:58	Android / 2.2, 2.3			172.20.120.2		N/A	1 minute ago (wan1)	

Device monitoring is enabled separately on each interface. Device detection is intended for devices directly connected to your LAN ports. If enabled on a WAN port, device detection may be unable to determine some devices' operating system.

### To configure device monitoring

1. Go to *System > Network > Interfaces*.
2. Edit the interface that you want to monitor devices on.
3. In *Device Management*, select *Detect and Identify Devices*.
4. Select *OK*.
5. Repeat steps 2 through 4 for each interface that will monitor devices.

### To assign an alias to a detected device or change device information

1. Go to *User & Device > Device > Device Definitions*.
2. Double-click the device entry or right-click it and select *Edit*.
3. Enter an *Alias* such as the user's name to identify the device.  
This step is compulsory. The alias replaces the MAC address in the device list.
4. Change other information as needed.
5. Select *OK*.

### To add a device manually

1. Go to *User & Device > Device > Device Definitions* and select *Create New*.
2. Enter the following information.
  - Alias (required)
  - MAC address
  - Device Type
3. Optionally, select *Custom Groups* or enter *Comments*.
4. Select *OK*.

## Device Groups

Device Groups are used in device policies to specify which devices match the policy. FortiOS automatically adds detected devices of well-known device types to predefined device groups. You can also create custom device groups so that you can create a different policy for devices that you know than for devices in general.

Go to *User & Device > Device > Device Groups* to view the list of device groups. To view all groups, select *Show Empty Groups* at the top right of the list.

**Table 91:** Predefined Device Groups

Group	Devices
All	All devices.
Android Phone	All Android-based phones in the Device Visibility database.
Android Tablet	Tablets running Android OS.
BlackBerry Phone	All BlackBerry-based phones in the Device Visibility database.
BlackBerry PlayBook	All BlackBerry PlayBook devices in the Device Visibility database.
Collected Emails	All devices from which FortiOS has collected a user email address.
Fortinet Device	FortiGate, FortiManager, FortiAnalyzer, FortiMail, etc.

**Table 91:** Predefined Device Groups

Group	Devices
Gaming Console	All Gaming consoles listed in the Device Visibility database. This includes Xbox, PS2, PS3, Wii, PSP.
IP Phone	All IP phones.
iPad	All IOS-based tablets in the Device Visibility database.
iPhone	All IOS-based phones in the Device Visibility database.
Linux PC	PCs running a Linux-based OS.
Mac	Apple Macintosh computers.
Media Streaming	Media streaming devices such as Apple TV.
Router/NAT Device	Router.
Windows PC	PCs running a Windows OS.
Windows Phone	All Windows OS based phones.
Windows Tablet	All Windows-based tablets.
Other Network Device	All other network devices not categorized under any other group.

## Creating a custom device group

In addition to the predefined device groups, you can create custom device groups, where you choose the member devices.

For ease of identifying devices, Fortinet recommends that you assign each device an Alias. For previously detected devices, you can edit the existing device definition to assign an alias. For devices that have not yet been detected, you can add a device definition if you know the device's MAC address. At that time you can also assign an alias for the device.

### To create a custom device group and add devices to it

1. Go to *User & Device > Device > Device Groups* and select *Create New*.
2. Enter a *Name* for the group, Employees for example.
3. Click in the *Members* field and click a device to add. Repeat to add other devices.
4. Select *OK*.

The devices are added to the custom device group.

## Controlling access with a MAC Address Access Control List

A MAC Address Access Control List is best used to handle exceptions. If you want to limit network access to a larger group such as your employees, it is better to create a custom device group and specify that group in your device-based security policies.

A MAC Address Access Control List functions as either a list of blocked devices or a list of allowed devices. This is determined by the *Unknown MAC Address* entry.

- By default, unknown MAC addresses are allowed: *Action* is *Assign IP*. You add an entry for each MAC address that you want to block and set its *Action* to *Block*.
- If you want to restrict access to a limited set of devices, you set the *Unknown MAC Address* entry to *Block* and add an entry for each allowed MAC address with *Action* set to *Assign IP*.

### To create a MAC Address Access Control List

1. In the SSID or other interface configuration, select *Enable DHCP Server*.
2. Enter the required *Address Range* and *Netmask*.
3. Expand *Advanced*.
4. In *MAC Address Access Control List*, select *Create New* and enter the device's *MAC Address*.
5. Select *Assign IP* to allow the device or *Block* to block the device and then select *OK*.
6. Repeat Steps 4 and 5 for each additional MAC address entry.
7. If needed, edit the *Unknown MAC Address* entry to set the correct *Action*.

## Device policies

Policies based on device identity enable you to implement policies according to device type. For example:

- Gaming consoles cannot connect to the company network or the Internet.
- Personal tablet and phone devices can connect to the Internet but not to company servers.
- Company-issued laptop computers can connect to the Internet and company servers. Web filtering and antivirus are applied.
- Employee laptop computers can connect to the Internet, but web filtering is applied. They can also connect to company networks, but only if FortiClient Endpoint Security is installed to protect against viruses.

Figure 296 and Figure 297 show these policies implemented for WiFi to the company network and to the Internet.

**Figure 296:**Device policies for WiFi access to the company network

**Edit Policy**

Policy Type:  Firewall  VPN

Policy Subtype:  Address  User Identity  Device Identity

Incoming Interface: wifi (SSID: fortinet)

Source Address: all

Outgoing Interface: internal

Enable NAT

Use Destination Interface Address  Fixed Port

Use Dynamic IP Pool

**Configure Authentication Rules**

Destination Address	Device	Endpoint Compliance	Service	Schedule	UTM Security	Traffic Shaping	Logging	Action
all	Gaming Console	-	ALL	always	-	⊗	⊗	⊗ DENY
all	Android Phone Android Tablet BlackBerry Phone BlackBerry PlayBook iPad iPhone	-	ALL	always	-	⊗	⊗	⊗ DENY
all	company laptop	⊗	ALL	always	🚫🚫🚫	⊗	⊗	✓ ACCEPT
all	employee laptop	✓	ALL	always	-	⊗	⊗	✓ ACCEPT
all	employee laptop	-	ALL	always	-	⊗	⊗	🚫 Captive Portal - Enforce FortiClic

Customize Authentication Messages

Comments:  0/255

**OK** **Cancel**

**Figure 297:**Device policies for WiFi access to the Internet

**Edit Policy**

Policy Type:  Firewall  VPN

Policy Subtype:  Address  User Identity  Device Identity

Incoming Interface: wifi (SSID: fortinet)

Source Address: all

Outgoing Interface: wan1

Enable NAT

Use Destination Interface Address  Fixed Port

Use Dynamic IP Pool

**Configure Authentication Rules**

Destination Address	Device	Endpoint Compliance	Service	Schedule	UTM Security	Traffic Shaping	Logging	Action
all	Gaming Console	-	ALL	always	-	⊗	⊗	⊗ DENY
all	Android Phone Android Tablet BlackBerry Phone BlackBerry PlayBook iPad iPhone	⊗	ALL	always	-	⊗	⊗	✓ ACCEPT
all	company laptop	⊗	ALL	always	🚫🚫🚫	⊗	⊗	✓ ACCEPT
all	employee laptop	⊗	ALL	always	🚫	⊗	⊗	✓ ACCEPT

Customize Authentication Messages

Comments:  0/255

**OK** **Cancel**

The next section explains device policy creation in detail.



## Creating device policies

Device-based security policies are similar to policies based on user identity:

- The policy enables traffic to flow from one network interface to another.
- NAT can be enabled.
- Authentication rules can allow or deny specific devices or device groups.
- UTM protection can be applied.

### To create a device identity policy

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. In *Policy Subtype*, select *Device Identity*.
3. Choose *Incoming Interface*, *Source Address*, and *Outgoing Interface* as you would for any security policy.
4. Select *Enable NAT* if appropriate.  
You are now ready to create authentication rules.

### To create an authentication rule

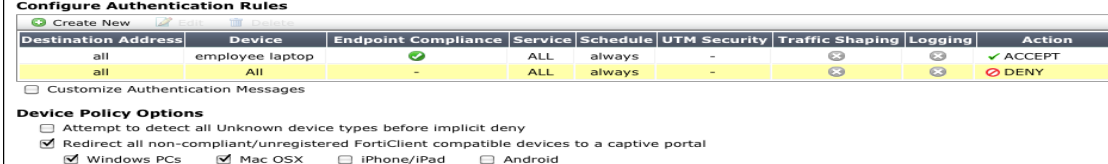
1. Select *Create New*.
2. Enter *Destination*, *Schedule*, and *Service* as you would for any security policy.
3. In *Device*, select the devices or device groups to which this policy applies.  
You can select multiple devices or groups.
4. Select *Compliant with Endpoint Profile* if you want to enforce use of FortiClient Endpoint Security by the client devices. This is available here only if Action is ACCEPT. See [“Adding endpoint protection”](#) next.
5. Select either ACCEPT or DENY as the policy Action.
6. Configure *UTM Security Profiles* as you would for any security policy.
7. Select *OK*.
8. Select *OK* again to complete creation of the security policy.

## Adding endpoint protection

Optionally, you can require that users’ devices have FortiClient Endpoint Security software installed. The software provides FortiOS more detailed information about the applications being used. FortiOS pushes a FortiClient profile out to the FortiClient software, configuring network protection such as antivirus, application control, and web category filtering. Devices without an up-to-date installation of FortiClient software are restricted to a captive portal from which the user can download a FortiClient installer.

If you have already created an ACCEPT rule for particular device groups, you simply edit this rule and enable *Compliant with Endpoint Profile*. Then select the device policy option that directs FortiClient-compatible devices to a captive portal.

**Figure 298:**Authentication rule with Endpoint compliance and captive portal enabled



Destination Address	Device	Endpoint Compliance	Service	Schedule	UTM Security	Traffic Shaping	Logging	Action
all	employee laptop	✓	ALL	always	-	⊗	⊗	✓ ACCEPT
all	All	-	ALL	always	-	⊗	⊗	⊗ DENY

Customize Authentication Messages

**Device Policy Options**

- Attempt to detect all Unknown device types before implicit deny
- Redirect all non-compliant/unregistered FortiClient compatible devices to a captive portal
  - Windows PCs
  - Mac OSX
  - iPhone/iPad
  - Android

For more information, see [“Endpoint Protection” on page 2003](#).

## Setting Device Policy Options

1. Optionally, enable *Attempt to detect all Unknown device types before implicit deny*.
2. *Redirect all non-compliant/unregistered FortiClient compatible devices to a captive portal* enables the captive portal. Select which device platforms to include.
3. Optionally, enable *Prompt Email Address Collection Portal for all devices*. This requests an email address from the device user. See [“Guest access in a retail environment” on page 512](#).

# Endpoint Protection

This section describes the Endpoint Protection feature and how to configure it.

The following topics are included in this section:

- [Endpoint Protection overview](#)
- [Configuration overview](#)
- [Creating a FortiClient profile](#)
- [Enabling Endpoint Protection in security policies](#)
- [Configuring endpoint registration over a VPN](#)
- [Monitoring endpoints](#)
- [Modifying the Endpoint Protection replacement messages](#)

## Endpoint Protection overview

Endpoint Protection enforces the use of up-to-date FortiClient Endpoint Security software on endpoints (workstation computers and mobile devices). It pushes a FortiClient profile to the FortiClient application, specifying security settings, including:

- Real-time antivirus protection - on or off
- FortiClient web category filtering based on web filters defined in a FortiGate web filter profile
- FortiClient application control (application firewall) using application sensors defined in the FortiGate application control feature
- Endpoint vulnerability scanning daily, weekly, or monthly

The FortiClient profile can also specify:

- VPN configurations
- Uploading of logs to the FortiGate unit hourly or daily
- Configuration profile (.mobileconfig file for iOS)
- Dashboard banner

You enable Endpoint Security in device identity security policies by enabling *Compliant with FortiClient Profile*. Optionally, the security policy can redirect non-compliant endpoints to a captive portal to download FortiClient software. Otherwise, non-compliant endpoints are blocked.

## User experience

Endpoint Protection applies to users attempting to make a connection that is controlled by a firewall policy that specifies a FortiClient profile. The user of a non-compliant endpoint using a web browser receives a replacement message HTML page from the FortiGate unit. The message explains the non-compliance. Depending on the FortiClient profile, the user may be allowed to continue or is blocked from further access. For information about modifying these replacement pages, see [“Modifying the Endpoint Protection replacement messages” on page 2010](#).

## FortiClient non-compliance

If the authentication rule in a device identity policy requires endpoint protection, a device without the latest version of FortiClient Endpoint Security installed is denied access. Optionally, the user is redirected to a captive and sees a message like this:

**Figure 299:**Default FortiClient non-compliance message for Windows



After installing FortiClient Endpoint Security, the user will receive an invitation to register with the FortiGate unit. If the user accepts the invitation, the FortiClient profile is sent to the device's FortiClient application. Now the user can pass the authentication rule and connect to the network. FortiClient Endpoint Security registered with a FortiGate unit does not need to be separately licensed with FortiGuard.

The FortiGate unit can also register endpoints who connect over the Internet through a VPN. The user can accept an invitation to register with the FortiGate unit. See [“Configuring endpoint registration over a VPN” on page 2009](#).

## FortiGate endpoint registration limits

To view the number of endpoints that are registered and the total that can be registered, go to *System > Dashboard > Status*. Under *License Information*, find *FortiClient Software*. You will see a line like “Registered/Allowed 4 of 10”. This means that there are four registered endpoints and a total of ten are allowed.

When the registration limit is reached, the next FortiClient-compatible device will not be able to register with the FortiGate unit. The user sees a message in FortiClient application about this. The FortiClient profile is not sent to client and the client cannot connect through the FortiGate unit.

For all FortiGate models, the maximum number of registered endpoints is ten. For all models except 20C and 40C, you can purchase an endpoint license to increase this capacity:

### To add an endpoint license - web-based manager

1. Go to *System > Dashboard > Status*.
2. In the *License Information* widget, under *FortiClient Software*, select *[Enter License]*, enter the license key, and select *OK*.

**Table 92:** Maximum registered endpoints with endpoint license

Model type	Max Registered Endpoints
Desktop	200
Rack - 1U	2000
Rack - 2U+	8000

## Configuration overview

Endpoint Protection requires that all hosts using the firewall policy have the FortiClient Endpoint Security application installed. Make sure that all hosts affected by this policy are able to install this application. Currently, FortiClient Endpoint Security is available for Microsoft Windows (2000 and later) and Apple Mac OSX only.

To set up Endpoint Protection, you need to

- Enable Central Management by the FortiGuard Analysis & Management Service if you will use FortiGuard Services to update the FortiClient application or antivirus signatures. You do not need to enter account information. See “Centralized Management” in the System Administration chapter of this Handbook.
- By default, the FortiGuard service provides the FortiClient installer. If you prefer to host it on your own server, see “[Changing the FortiClient installer download location](#)” on page 2005.
- In Security Profiles, configure application sensors and web filters profiles as needed to monitor or block applications. See the Security Profiles Guide chapter of this Handbook for details.
- Create a FortiClient profile or use a predefined profile. See “[Creating a FortiClient profile](#)” on page 2006. Enable the application sensor and web category filtering profiles that you want to use.
- Enable *Compliant with FortiClient Profile* in the authentication rules of Device Identity security policies that the endpoints will use.
- Optionally, configure the FortiGate unit to support endpoint registration by IPsec or SSL VPN.

### Changing the FortiClient installer download location

By default, FortiClient installers are downloaded from the FortiGuard network. You can also host these installers on a server for your users to download. In that case, you must configure FortiOS with this custom download location. For example, to set the download location to a customer web server with address `custom.example.com`, enter the following command:

```
config endpoint-control settings
 set download-location custom
 set download-custom-link "http://custom.example.com"
end
```

## Creating a FortiClient profile

Each FortiClient profile is assigned to particular device groups, user groups, or individual users. When *Compliant with FortiClient Profile* is selected in Device Identity policy authentication rule, all users of that rule must have FortiClient Endpoint Security installed. The FortiGate unit pushes the FortiClient profile settings to the FortiClient application on the client.

There is a default FortiClient profile for Windows and Mac OS that enables only AntiVirus, Web Filtering, and VPN. You can also create your own FortiClient profiles.

### To create a FortiClient profile - web-based manager

1. If you will use the Application Firewall feature, go to *Security Profiles > Application Control > Application Sensors* to create the Application Sensors that you will need.
2. If you will use Web Category Filtering, go to *Security Profiles > Web Filter > Profiles* to create the web filter profile that you will need.
3. Go to *User & Device > Endpoint Protection > FortiClient Profiles*.  
The list of FortiClient profiles is displayed.
4. Select *Create New* or select an existing profile and *Edit* it.
5. In *Assign Profile To*, select the device groups, user groups, and users to which this FortiClient profile applies. This is not available for the default profile.
6. Enter the *FortiClient Configuration Deployment* settings for *Windows and Mac*:

<b>Antivirus Protection</b>	ON — enable the FortiClient realtime AntiVirus feature.
<b>Web Category Filtering</b>	ON — enable web category filtering. Select the web filter profile to use.
<b>Disable Web Category Filtering when protected by this FortiGate</b>	Disables FortiClient web category filtering when client traffic is filtered by the FortiGate unit. Selected by default.
<b>Client VPN Provisioning</b>	Enable to configure the FortiClient VPN client. Enter the VPN configuration details.
<b>Application Firewall</b>	ON — enable application control. Select the application sensor to use.
<b>Endpoint Vulnerability Scan on Client</b>	ON — FortiGate unit will perform vulnerability scan on client. Select the desired schedule.
<b>Initiate Scan After Client Registration</b>	Enables scan following registration, regardless of schedule. Selected by default.
<b>Upload logs to FortiAnalyzer /FortiManager</b>	ON — FortiClient software will upload its logs to the specified FQDN or IP address. Select the desired schedule.
<b>Use FortiManager for client software/signature update</b>	ON — FortiClient software obtain AV signatures and software updates from the specified FQDN or IP address. <i>Failover to FDN when FortiManager is not available</i> is enabled by default.
<b>Dashboard Banner</b>	ON — Display dashboard banner.

7. Enter the *FortiClient Configuration Deployment* settings for *iOS*:

---

<b>Web Category Filtering</b>	ON — enable web category filtering. Select the web filter profile to use.
<b>Disable Web Category Filtering when protected by this FortiGate</b>	Disables FortiClient web category filtering when client traffic is filtered by the FortiGate unit. Selected by default.
<b>Client VPN Provisioning</b>	Enable to configure the FortiClient VPN client. You can enter multiple VPN configurations by selecting the “+” button.
<b>VPN Name</b>	Enter a name to identify this VPN configuration in the FortiClient application.
<b>Type</b>	Select <i>IPsec</i> or <i>SSL-VPN</i> .  If you select <i>IPsec</i> , select a <i>VPN Configuration File</i> that contains the required IPsec VPN configuration. The Apple iPhone Configuration Utility produces .mobileconfig files which contain configuration information for an iOS device.  If you select <i>SSL-VPN</i> , enter the VPN configuration details.
<b>Distribute Configuration Profile</b>	ON — Distribute configuration information to iOS devices running FortiClient Endpoint Security. Select <i>Browse</i> and locate the file to be distributed.  The Apple iPhone Configuration Utility produces .mobileconfig files which contain configuration information for an iOS device.

---

8. Enter the *FortiClient Configuration Deployment* settings for *Android*:

---

<b>Web Category Filtering</b>	ON — enable web category filtering. Select the web filter profile to use.
<b>Disable Web Category Filtering when protected by this FortiGate</b>	Disables FortiClient web category filtering when client traffic is filtered by the FortiGate unit. Selected by default.
<b>Client VPN Provisioning</b>	Enable to configure the FortiClient VPN client. You can enter multiple VPN configurations by selecting the “+” button.
<b>VPN Name</b>	Enter a name to identify this VPN configuration in the FortiClient application.
<b>Type</b>	Select <i>IPsec</i> or <i>SSL-VPN</i> . Enter the VPN configuration details.

---

9. Select *OK*.

### To create a FortiClient profile - CLI

This example creates a profile for Windows and Mac computers.

```
config endpoint-control profile
 edit ep-profile1
 set device-groups mac windows-pc
 config forticlient-winmac-settings
 set forticlient-av enable
 set forticlient-wf enable
 set forticlient-wf-profile default
 end
 end
end
```

## Enabling Endpoint Protection in security policies

Endpoint Protection is applied to any traffic where the controlling firewall policy has Endpoint Security enabled. The device group to which the device belongs determines which FortiClient profile is applied. The policy searches the list of FortiClient profiles starting from the top and applies the first profile assigned to the device group.

### To enable Endpoint Protection - web-based manager

1. Go to *Policy > Policy > Policy* and edit the device identity firewall policy where you want to enable Endpoint Protection.
2. Create or edit an authentication rule.
3. Select *Compliant with FortiClient profile*.
4. Select *OK*.

### To configure the firewall policy - CLI

In this example, the LAN connects to Port 2 and the Internet is connected to Port 1. a FortiClient profile is applied.

```
config firewall policy
 edit 0
 set srcintf port2
 set dstintf port1
 set srcaddr LANUsers
 set dstaddr all
 set action accept
 set identity-based enable
 set identity-from device
 set nat enable
 config identity-based-policy
 edit 1
 set schedule always
 set service ALL
 set devices all
 set endpoint-compliance enable
 end
 end
end
```



## Configuring endpoint registration over a VPN

FortiGate units can register FortiClient-equipped endpoints over either an interface-based IPsec VPN or a tunnel-mode SSL VPN. After the user authenticates, the FortiGate unit sends the FortiClient application the IP address and port to be used for registration. If the user accepts the FortiGate invitation to register, registration proceeds and the FortiClient profile is downloaded to the client.

Users without FortiClient Endpoint Security connecting to the SSL VPN through a browser can be redirected to a captive portal to download and install the FortiClient software. The security policy must enable *Redirect all non-compliant/unregistered FortiClient compatible devices to a captive portal*, but not select any specific device types.

### Endpoint registration on an IPsec VPN

You can enable endpoint registration when you configure the FortiClient VPN or you can enable it on an existing FortiClient VPN.

#### To enable endpoint registration while configuring the VPN

- Enable *Allow Endpoint Registration* on the Network page of the VPN Wizard when creating the FortiClient VPN.

#### To enable endpoint registration on an existing VPN

1. Go to *System > Network > Interfaces* and edit the VPN's tunnel interface.  
The tunnel is a subinterface of the physical network interface.
2. In *Administrative Access*, make sure that *FCT-Access* is enabled.
3. Select *OK*.

### Endpoint registration on the SSL VPN

#### To enable endpoint registration on the SSL VPN

1. Go to *VPN > SSL > Portal*.
2. Make sure *Enable Tunnel Mode* is enabled.
3. Optionally, enable *Include FortiClient Download*.  
Users who access the VPN with a browser will be able to download FortiClient Endpoint Security for their device.
4. Select *Apply*.
5. Go to *VPN > SSL > Config*, make sure *Allow Endpoint Registration (Tunnel Mode Only)* is enabled, then select *Apply*.

This procedure does not include all settings needed to configure a working SSL VPN.

### Synchronizing endpoint registrations

To support roaming users in a network with multiple FortiGate units, you need to configure synchronization of the endpoint registration databases between the units. The registered endpoints are then recognized on all of the FortiGate units. This is configured in the CLI. For

example, to synchronize this FortiGate unit's registered endpoint database with another unit named `other1` at IP address 172.20.120.4, enter:

```
config endpoint-control forticlient-registration-sync
 edit other1
 set peer-ip 172.20.120.4
 end
```

## Monitoring endpoints

Go to *User & Device > Monitor > FortiClient* to monitor endpoints.

## Modifying the Endpoint Protection replacement messages

If the security policy has *Redirect all non-compliant/unregistered FortiClient compatible devices to a captive portal* enabled, users of non-compliant devices are redirected to a captive portal that is defined by the *Endpoint NAC Download Portal* replacement message. There are different portals for Android, iOS, Mac, Windows, and “other” devices. Optionally, you can modify them.

### To modify the Endpoint NAC Download Portal

1. Go to *System > Config > Replacement Message Group* and select *Extended View*.
2. In the *Endpoint Control* section select the message that you want to edit.  
The replacement message and its HTML code appear in a split screen in the lower half of the page.
3. Modify the text as needed and select *Save*.

# Vulnerability Scan

The Network Vulnerability Scan helps you to protect your network assets (servers and workstations) by scanning them for security weaknesses. You can scan on-demand or on a scheduled basis. Results are viewable on the FortiGate unit, but results are also sent to an attached FortiAnalyzer unit. The FortiAnalyzer unit can collect the results of vulnerability scans from multiple FortiGate units at different locations on your network, compiling a comprehensive report about network security.

This section describes how to configure a single FortiGate unit for network scanning and how to view the results of the scan.

The following topics are included in this section:

- [Configuring vulnerability scans](#)
- [Running a vulnerability scan and viewing scan results](#)
- [Requirements for authenticated scanning and ports scanned](#)

## Configuring vulnerability scans

You can configure the scan schedule and the assets to be scanned.

### To configure scanning - web-based manager

1. Go to *User & Device > Vulnerability Scan > Scan Definition*.
2. Beside *Schedule* select *Change* to set the scan schedule and mode:

---

<b>Recurrence</b>	Select <i>Daily</i> , <i>Weekly</i> , or <i>Monthly</i> and configure the details for the option you have selected.
<b>Suspend Scan between</b>	Set a time during which the scan should be paused if its running.
<b>Vulnerability Scan Mode</b>	<b>Quick</b> — check only the most commonly used ports <b>Standard</b> — check the ports used by most known applications <b>Full</b> — check all TCP and UDP ports  For a detailed list of the TCP and UDP ports examined by each scan mode, see <a href="#">Table 93 on page 2016</a> .

---

3. Select *Apply* to save the schedule and scan type.
4. Select *Create New* under *Asset Definitions* to select the devices on the network to scan.  
An asset can be a single server or workstation computer on your network or a range of addresses on your network. You must add assets to the vulnerability scan before you can run a scan.

To scan an entire network or part of a network you can just add the appropriate IP address range to the asset configuration. You can also add the IP addresses of Windows and Linux computers to include the user names and passwords for these machines. The vulnerability scanner will use these credentials to log into the computers and do more detailed vulnerability scanning.

Even if the asset is an address range you can add Windows and Linux credentials. The vulnerability scanner will attempt to log into all network device it finds using these credentials.

5. Enter the following information and select *OK*:

<b>Name</b>	Enter a name for this asset.
<b>Type</b>	Select <i>IP Address</i> to add a single IP address. Select <i>Range</i> to add a range of IP addresses to scan.
<b>IP Address</b>	Enter the IP address of the asset. ( <i>Type is IP Address.</i> )
<b>Range</b>	Enter the start and end of the IP address range. ( <i>Type is Range.</i> )
<b>Enable Scheduled Vulnerability Scanning</b>	Select to allow this asset to be scanned according to the schedule. Otherwise the asset is not scanned during a scheduled vulnerability scan.
<b>Windows Authentication</b>	Select to use authentication on a Windows operating system. Enter the username and password in the fields provided.  For more information, see <a href="#">“Requirements for authenticated scanning and ports scanned” on page 2013.</a>
<b>Unix Authentication</b>	Select to use authentication on a Unix operating system. Enter the username and password in the fields provided.  For more information, see <a href="#">“Requirements for authenticated scanning and ports scanned” on page 2013.</a>

6. Select *Apply* to save the configuration.

#### To configure scanning - CLI

To configure, for example, a standard scan to be performed every Sunday at 2:00am, you would enter:

```
config netscan settings
 set scan-mode standard
 set schedule enable
 set time 02:00
 set recurrence weekly
 set day-of-week sunday
end
```

#### To add an asset - CLI

This example adds a single computer to the Asset list:

```
config netscan assets
 edit 0
 set name "server1"
 set addr-type ip
 set start-ip 10.11.101.20
 set auth-windows enable
 set win-username admin
 set win-password zxcvbnm
```

```
 set scheduled enable
end
```

This example adds an address range to the Asset list. Authentication is not used:

```
config netscan assets
 edit 0
 set name "fileservers"
 set addr-type range
 set start-ip 10.11.101.160
 set end-ip 10.11.101.170
 set scheduled enable
 end
```

## Running a vulnerability scan and viewing scan results

### To run a vulnerability scan - web-based manager

1. Go to *User & Device > Vulnerability Scan > Scan Definition* and select *Start Scan*.  
When the scan is running you can pause or stop it at any time. You can also watch the progress of the scan.
2. When the scan is complete go to *User & Device > Vulnerability Scan > Vulnerability Result* to view the results of the scan.

### To run a vulnerability scan - CLI

Use the following CLI commands:

```
execute netscan start scan
execute netscan status
execute netscan pause
execute netscan resume
execute netscan stop
```

### To view vulnerability scan results

1. To view vulnerability scan results go to *User & Device > Vulnerability Scan > Vulnerability Result*.
2. Select any log entry to view log details.

## Requirements for authenticated scanning and ports scanned

The effectiveness of an authenticated scan is determined by the level of access the FortiGate unit obtains to the host operating system. Rather than use the system administrator's account, it might be more convenient to set up a separate account for the exclusive use of the vulnerability scanner with a password that does not change.

The following sections detail the account requirements for various operating systems.

## Microsoft Windows hosts - domain scanning

The user account provided for authentication must

- have administrator rights
- be a Security type of account
- have global scope
- belong to the Domain Administrators group
- meet the Group Policy requirements listed below:

### Group Policy - Security Options

In the Group Policy Management Editor, go to Computer Configuration > Windows Settings > Security Settings > Local Policies > Security Options.

Setting	Value
Network access: Sharing and security model for local accounts	Classic
Accounts: Guest account status	Disabled
Network access: Let Everyone permissions apply to anonymous users	Disabled

### Group Policy - System Services

In the Group Policy Management Editor, go to Computer Configuration > Windows Settings > Security Settings > System Services.

Setting	Value
Remote registry	Automatic
Server	Automatic
Windows Firewall	Automatic

### Group Policy - Administrative Templates

In the Group Policy Management Editor, go to Computer Configuration > Administrative Templates > Network > Network Connections > Windows Firewall > Domain Profile.

Setting	Value
Windows Firewall: Protect all network connections	Disabled

or

Setting	Value
Windows Firewall: Protect all network connections	Enabled
Windows Firewall: Allow remote administration exception	Enabled
Allow unsolicited messages from <sup>1</sup>	*

Windows Firewall: Allow file and printer sharing exception	Enabled
Allow unsolicited messages from <sup>1</sup>	*
Windows Firewall: Allow ICMP exceptions	Enabled
Allow unsolicited messages from <sup>1</sup>	*

<sup>1</sup>Windows prompts you for a range of IP addresses. Enter either “\*” or the IP address of the Fortinet appliance that is performing the vulnerability scan.

## Microsoft Windows hosts - local (non-domain) scanning

The user account provided for authentication must

- be a local account
- belong to the Administrators group

The host must also meet the following requirements:

- Server service must be enabled. (Windows 2000, 2003, XP)
- Remote Registry Service must be enabled.
- File Sharing must be enabled.
- Public folder sharing must be disabled. (Windows 7)
- Simple File Sharing (SFS) must be disabled. (Windows XP)

## Windows firewall settings

- Enable the *Remote Administration Exception* in Windows Firewall. (Windows 2003, Windows XP)
- Allow *File and Print sharing* and *Remote Administration* traffic to pass through the firewall. Specify the IP address or subnet of the Fortinet appliance that is performing the vulnerability scan. (Windows Vista, 2008)
- For each of the active *Inbound Rules* in the *File and Printer Sharing* group, set the *Remote IP address* under *Scope* to either *Any IP address* or to the IP address or subnet of the Fortinet appliance that is performing the vulnerability scan. (Windows 7)

## Unix hosts

The user account provided for authentication must be able at a minimum to execute these commands:

- The account must be able to execute “uname” in order to detect the platform for packages.
- If the target is running Red Hat, the account must be able to read /etc/redhat-release and execute “rpm”.
- If the target is running Debian, the account must be able to read /etc/debian-version and execute “dpkg”.

**Table 93:** Ports scanned in each scan mode

Scan Type	Ports scanned
<b>Standard Scan</b>	<p><b>TCP:</b> 1-3, 5, 7, 9, 11, 13, 15, 17-25, 27, 29, 31, 33, 35, 37-39, 41-223, 242-246, 256-265, 280-282, 309, 311, 318, 322-325, 344-351, 363, 369-581, 587, 592-593, 598, 600, 606-620, 624, 627, 631, 633-637, 666-674, 700, 704-705, 707, 709-711, 729-731, 740-742, 744, 747-754, 758-765, 767, 769-777, 780-783, 786, 799-801, 860, 873, 886-888, 900-901, 911, 950, 954-955, 990-993, 995-1001, 1008, 1010-1011, 1015, 1023-1100, 1109-1112, 1114, 1123, 1155, 1167, 1170, 1207, 1212, 1214, 1220-1222, 1234-1236, 1241, 1243, 1245, 1248, 1269, 131t3-1314, 1337, 1344-1625, 1636-1774, 1776-1815, 1818-1824, 1901-1909, 1911-1920, 1944-1951, 1973, 1981, 1985-2028, 2030, 2032-2036, 2038, 2040-2049, 2053, 2065, 2067, 2080, 2097, 2100, 2102-2107, 2109, 2111, 2115, 2120, 2140, 2160-2161, 2201-2202, 2213, 2221-2223, 2232-2239, 2241, 2260, 2279-2288, 2297, 2301, 2307, 2334, 2339, 2345, 2381, 2389, 2391, 2393-2394, 2399, 2401, 2433, 2447, 2500-2501, 2532, 2544, 2564-2565, 2583, 2592, 2600-2605, 2626-2627, 2638-2639, 2690, 2700, 2716, 2766, 2784-2789, 2801, 2908-2912, 2953-2954, 2998, 3000-3002, 3006-3007, 3010-3011, 3020, 3047-3049, 3080, 3127-3128, 3141-3145, 3180-3181, 3205, 3232, 3260, 3264, 3267-3269, 3279, 3306, 3322-3325, 3333, 3340, 3351-3352, 3355, 3372, 3389, 3421, 3454-3457, 3689-3690, 3700, 3791, 3900, 3984-3986, 4000-4002, 4008-4009, 4080, 4092, 4100, 4103, 4105, 4107, 4132-4134, 4144, 4242, 4321, 4333, 4343, 4443-4454, 4500-4501, 4567, 4590, 4626, 4651, 4660-4663, 4672, 4899, 4903, 4950, 5000-5005, 5009-5011, 5020-5021, 5031, 5050, 5053, 5080, 5100-5101, 5145, 5150, 5190-5193, 5222, 5236, 5300-5305, 5321, 5400-5402, 5432, 5510, 5520-5521, 5530, 5540, 5550, 5554-5558, 5569, 5599-5601, 5631-5632, 5634, 5678-5679, 5713-5717, 5729, 5742, 5745, 5755, 5757, 5766-5767, 5800-5802, 5900-5902, 5977-5979, 5997-6053, 6080, 6103, 6110-6112, 6123, 6129, 6141-6149, 6253, 6346, 6387, 6389, 6400, 6455-6456, 6499-6500, 6515, 6558, 6588, 6660-6670, 6672-6673, 6699, 6767, 6771, 6776, 6831, 6883, 6912, 6939, 6969-6970, 7000-7021, 7070, 7080, 7099-7100, 7121, 7161, 7174, 7200-7201, 7300-7301, 7306-7308, 7395, 7426-7431, 7491, 7511, 7777-7778, 7781, 7789, 7895, 7938, 7999-8020, 8023, 8032, 8039, 8080-8082, 8090, 8100, 8181, 8192, 8200, 8383, 8403, 8443, 8450, 8484, 8732, 8765, 8886-8894, 8910, 9000-9001, 9005, 9043, 9080, 9090, 9098-9100, 9400, 9443, 9535, 9872-9876, 9878, 9889, 9989-10000, 10005, 10007, 10080-10082, 10101, 10520, 10607, 10666, 11000, 11004, 11223, 12076, 12223, 12345-12346, 12361-12362, 12456, 12468-12469, 12631, 12701, 12753, 13000, 13333, 14237-14238, 15858, 16384, 16660, 16959, 16969, 17007, 17300, 18000, 18181-18186, 18190-18192, 18194, 18209-18210, 18231-18232, 18264, 19541, 20000-20001, 20011, 20034, 20200, 20203, 20331, 21544, 21554, 21845-21849, 22222, 22273, 22289, 22305, 22321, 22555, 22800, 22951, 23456, 23476-23477, 25000-25009, 25252, 25793, 25867, 26000, 26208, 26274, 27000-27009, 27374, 27665, 29369, 29891, 30029, 30100-30102, 30129, 30303, 30999, 31336-31337, 31339, 31554, 31666, 31785, 31787-31788, 32000, 32768-32790, 33333, 33567-33568, 33911, 34324, 37651, 40412, 40421-40423, 42424, 44337, 47557, 47806, 47808, 49400, 50505, 50766, 51102, 51107, 51112, 53001, 54321, 57341, 60008, 61439, 61466, 65000, 65301, 65512</p> <p><b>UDP:</b> 7, 9, 13, 17, 19, 21, 37, 53, 67-69, 98, 111, 121, 123, 135, 137-138, 161, 177, 371, 389, 407, 445, 456, 464, 500, 512, 514, 517-518, 520, 555, 635, 666, 858, 1001, 1010-1011, 1015, 1024-1049, 1051-1055, 1170, 1243, 1245, 1434, 1492, 1600, 1604, 1645, 1701, 1807, 1812, 1900, 1978, 1981, 1999, 2001-2002, 2023, 2049, 2115, 2140, 2801, 3024, 3129, 3150, 3283, 3527, 3700, 3801, 4000, 4092, 4156, 4569, 4590, 4781, 5000-5001, 5036, 5060, 5321, 5400-5402, 5503, 5569, 5632, 5742, 6073, 6502, 6670, 6771, 6912, 6969, 7000, 7300-7301, 7306-7308, 7778, 7789, 7938, 9872-9875, 9989, 10067, 10167, 11000, 11223, 12223, 12345-12346, 12361-12362, 15253, 15345, 16969, 20001, 20034, 21544, 22222, 23456, 26274, 27444, 30029, 31335, 31337-31339, 31666, 31785, 31789, 31791-31792, 32771, 33333, 34324, 40412, 40421-40423, 40426, 47262, 50505, 50766, 51100-51101, 51109, 53001, 61466, 65000</p>



**Table 93:** Ports scanned in each scan mode

Scan Type	Ports scanned
<b>Full Scan</b>	All TCP and UDP ports (1-65535)
<b>Quick Scan</b>	<p><b>TCP:</b> 11, 13, 15, 17, 19-23, 25, 37, 42, 53, 66, 69-70, 79-81, 88, 98, 109-111, 113, 118-119, 123, 135, 139, 143, 220, 256-259, 264, 371, 389, 411, 443, 445, 464-465, 512-515, 523-524, 540, 548, 554, 563, 580, 593, 636, 749-751, 873, 900-901, 990, 992-993, 995, 1080, 1114, 1214, 1234, 1352, 1433, 1494, 1508, 1521, 1720, 1723, 1755, 1801, 2000-2001, 2003, 2049, 2301, 2401, 2447, 2690, 2766, 3128, 3268-3269, 3306, 3372, 3389, 4100, 4443-4444, 4661-4662, 5000, 5432, 5555-5556, 5631-5632, 5634, 5800-5802, 5900-5901, 6000, 6112, 6346, 6387, 6666-6667, 6699, 7007, 7100, 7161, 7777-7778, 8000-8001, 8010, 8080-8081, 8100, 8888, 8910, 9100, 10000, 12345-12346, 20034, 21554, 32000, 32768-32790</p> <p><b>UDP:</b> 7, 13, 17, 19, 37, 53, 67-69, 111, 123, 135, 137, 161, 177, 407, 464, 500, 517-518, 520, 1434, 1645, 1701, 1812, 2049, 3527, 4569, 4665, 5036, 5060, 5632, 6502, 7778, 15345</p>

# Chapter 15 Unified Threat Management

## for FortiOS 5.0

This FortiOS Handbook chapter contains the following sections:

[Security Profiles overview](#) describes Security Profiles components and their relation to firewall policies, as well as SSL content scanning and inspection. We recommend starting with this section to become familiar with the different features in your FortiGate unit.

[Client Reputation](#) explains how to track client behavior and report on activities that you determine are risky or otherwise noteworthy.

[AntiVirus](#) explains how the FortiGate unit scans files for viruses and describes how to configure the antivirus options.

[Email filter](#) explains how the FortiGate unit filters email, describes how to configure the filtering options and the action to take with email detected as spam.

[Intrusion protection](#) explains basic Intrusion Protection System (IPS) concepts and how to configure IPS options; includes guidance and a detailed table for creating custom signatures as well as several examples.

[Web filter](#) and [FortiGuard Web Filter](#) The first of these sections describes basic web filtering concepts, the order in which the FortiGate unit performs web filtering, and configuration. The second section describes enhanced features of the subscription-based FortiGuard Web Filtering service and explains how to configure them. We recommend reading both sections if you are using FortiGuard Web Filtering because settings you configure in one feature may affect the other.

[Data leak prevention](#) describes the DLP features that allow you to prevent sensitive data from leaving your network and explains how to configure the DLP rules, compound rules, and sensors.

[Application control](#) describes how your FortiGate unit can detect and take action against network traffic based on the application generating the traffic.

# Security Profiles overview

Ranging from the FortiGate®-30 series for small businesses to the FortiGate-5000 series for large enterprises, service providers and carriers, the FortiGate line combines a number of security features to protect your network from threats. As a whole, these features, when included in a single Fortinet security appliance, are referred to as Security Profiles. The Security Profiles features your FortiGate model includes are:

- AntiVirus
- Intrusion Prevention System (IPS)
- Web filtering
- E-mail filtering, including protection against spam and grayware
- Data Leak Prevention (DLP)
- Application Control
- ICAP

Firewall policies limit access, and while this and similar features are a vital part of securing your network, they are not covered in this document.

The following topics are included in this section:

- [Traffic inspection](#)
- [Content inspection and filtering](#)
- [Security Profiles components](#)
- [Security Profiles/lists/sensors](#)

## Traffic inspection

When the FortiGate unit examines network traffic one packet at a time for IPS signatures, it is performing traffic analysis. This is unlike content analysis where the traffic is buffered until files, email messages, web pages, and other files are assembled and examined as a whole.

DoS policies use traffic analysis by keeping track of the type and quantity of packets, as well as their source and destination addresses.

Application control uses traffic analysis to determine which application generated the packet.

Although traffic inspection doesn't involve taking packets and assembling files they are carrying, the packets themselves can be split into fragments as they pass from network to network. These fragments are reassembled by the FortiGate unit before examination.

No two networks are the same and few recommendations apply to all networks. This topic offers suggestions on how you can use the FortiGate unit to help secure your network against content threats.

## IPS signatures

IPS signatures can detect malicious network traffic. For example, the Code Red worm attacked a vulnerability in the Microsoft IIS web server. Your FortiGate's IPS system can detect traffic attempting to exploit this vulnerability. IPS may also detect when infected systems communicate with servers to receive instructions.

## IPS recommendations

- Enable IPS scanning at the network edge for all services.
- Use FortiClient endpoint IPS scanning for protection against threats that get into your network.
- Subscribe to FortiGuard IPS Updates and configure your FortiGate unit to receive push updates. This will ensure you receive new IPS signatures as soon as they are available.
- Your FortiGate unit includes IPS signatures written to protect specific software titles from DoS attacks. Enable the signatures for the software you have installed and set the signature action to *Block*.
- You can view these signatures by going to *Security Profiles > Intrusion Protection > Predefined* and sorting by, or applying a filter to, the *Group* column.
- Because it is critical to guard against attacks on services that you make available to the public, configure IPS signatures to block matching signatures. For example, if you have a web server, configure the action of web server signatures to *Block*.

## Suspicious traffic attributes

Network traffic itself can be used as an attack vector or a means to probe a network before an attack. For example, SYN and FIN flags should never appear together in the same TCP packet. The SYN flag is used to initiate a TCP session while the FIN flag indicates the end of data transmission at the end of a TCP session.

The FortiGate unit has IPS signatures that recognize abnormal and suspicious traffic attributes. The SYN/FIN combination is one of the suspicious flag combinations detected in TCP traffic by the `TCP.BAD.FLAGS` signature.

The signatures that are created specifically to examine traffic options and settings, begin with the name of the traffic type they are associated with. For example, signatures created to examine TCP traffic have signature names starting with TCP.

## Application control

While applications can often be blocked by the ports they use, application control allows convenient management of all supported applications, including those that do not use set ports.

### Application control recommendations

- Some applications behave in an unusual manner in regards to application control. For more information, see [“Application considerations” on page 2153](#).
- By default, application control allows the applications not specified in the application control list. For high security networks, you may want to change this behavior so that only the explicitly allowed applications are permitted.

## Content inspection and filtering

When the FortiGate unit buffers the packets containing files, email messages, web pages, and other similar files for reassembly before examining them, it is performing content inspection. Traffic inspection, on the other hand, is accomplished by the FortiGate unit examining individual packets of network traffic as they are received.

No two networks are the same and few recommendations apply to all networks. This topic offers suggestions on how you can use the FortiGate unit to help secure your network against content threats. Be sure to understand the effects of the changes before using the suggestions.

## AntiVirus

The FortiGate antivirus scanner can detect viruses and other malicious payloads used to infect machines. The FortiGate unit performs deep content inspection. To prevent attempts to disguise viruses, the antivirus scanner will reassemble fragmented files and uncompress content that has been compressed. Patented Compact Pattern Recognition Language (CPRL) allows further inspection for common patterns, increasing detection rates of virus variations in the future.

### AntiVirus recommendations

- Enable antivirus scanning at the network edge for all services.
- Use FortiClient endpoint antivirus scanning for protection against threats that get into your network.
- Subscribe to FortiGuard AntiVirus Updates and configure your FortiGate unit to receive push updates. This will ensure you receive new antivirus signatures as soon as they are available.
- Enable the Extended Virus Database if your FortiGate unit supports it.
- Examine antivirus logs periodically. Take particular notice of repeated detections. For example, repeated virus detection in SMTP traffic could indicate a system on your network is infected and is attempting to contact other systems to spread the infection using a mass mailer.
- The *builtin-patterns* file filter list contains nearly 20 file patterns. Many of the represented files can be executed or opened with a double-click. If any of these file patterns are not received as a part of your normal traffic, blocking them may help protect your network. This also saves resources since files blocked in this way do not need to be scanned for viruses.
- To conserve system resources, avoid scanning email messages twice. Scan messages as they enter and leave your network or when clients send and retrieve them, rather than both.

## FortiGuard Web Filtering

The web is the most popular part of the Internet and, as a consequence, virtually every computer connected to the Internet is able to communicate using port 80, HTTP. Botnet communications take advantage of this open port and use it to communicate with infected computers. FortiGuard Web Filtering can help stop infections from malware sites and help prevent communication if an infection occurs.

### FortiGuard Web Filtering recommendations

- Enable FortiGuard Web Filtering at the network edge.
- Install the FortiClient application and use FortiGuard Web Filtering on any systems that bypass your FortiGate unit.
- Block categories such as Pornography, Malware, Spyware, and Phishing. These categories are more likely to be dangerous.
- In the email filter profile, enable *IP Address Check* in *FortiGuard Email Filtering*. Many IP addresses used in spam messages lead to malicious sites; checking them will protect your users and your network.

## Email filter

Spam is a common means by which attacks are delivered. Users often open email attachments they should not, and infect their own machine. The FortiGate email filter can detect harmful spam and mark it, alerting the user to the potential danger.

## Email filter recommendations

- Enable email filtering at the network edge for all types of email traffic.
- Use FortiClient endpoint scanning for protection against threats that get into your network.
- Subscribe to the FortiGuard AntiSpam Service.

## DLP

Most security features on the FortiGate unit are designed to keep unwanted traffic out of your network while DLP can help you keep sensitive information from leaving your network. For example, credit card numbers and social security numbers can be detected by DLP sensors.

### DLP recommendations

- Rules related to HTTP posts can be created, but if the requirement is to block all HTTP posts, a better solution is to use application control or the *HTTP POST Action* option in the web filter profile.
- While DLP can detect sensitive data, it is more efficient to block unnecessary communication channels than to use DLP to examine it. If you don't use instant messaging or peer-to-peer communication in your organization, for example, use application control to block them entirely.

## Security Profiles components

### AntiVirus

Your FortiGate unit stores a virus signature database that can identify more than 15,000 individual viruses. FortiGate models that support additional virus databases are able to identify hundreds of thousands of viruses. With a FortiGuard AntiVirus subscription, the signature databases are updated whenever a new threat is discovered.

AntiVirus also includes file filtering. When you specify files by type or by file name, the FortiGate unit will stop the matching files from reaching your users.

FortiGate units with a hard drive or configured to use a FortiAnalyzer unit can store infected and blocked files for that you can examine later.

### Intrusion Protection System (IPS)

The FortiGate Intrusion Protection System (IPS) protects your network against hacking and other attempts to exploit vulnerabilities of your systems. More than 3,000 signatures are able to detect exploits against various operating systems, host types, protocols, and applications. These exploits can be stopped before they reach your internal network.

You can also write custom signatures, tailored to your network.

### Web filtering

Web filtering includes a number of features you can use to protect or limit your users' activity on the web.

FortiGuard Web Filtering is a subscription service that allows you to limit access to web sites. More than 60 million web sites and two billion web pages are rated by category. You can choose to allow or block each of the 77 categories.

URL filtering can block your network users from access to URLs that you specify.

Web content filtering can restrict access to web pages based on words and phrases appearing on the web page itself. You can build lists of words and phrases, each with a score. When a web content list is selected in a web filter profile, you can specify a threshold. If a user attempts to load a web page and the score of the words on the page exceeds the threshold, the web page is blocked.

## Email filtering

FortiGuard AntiSpam is a subscription service that includes an IP address black list, a URL black list, and an email checksum database. These resources are updated whenever new spam messages are received, so you do not need to maintain any lists or databases to ensure accurate spam detection.

You can use your own IP address lists and email address lists to allow or deny addresses, based on your own needs and circumstances.

## Data Leak Prevention (DLP)

Data leak prevention allows you to define the format of sensitive data. The FortiGate unit can then monitor network traffic and stop sensitive information from leaving your network. Rules for U.S. social security numbers, Canadian social insurance numbers, as well as Visa, Mastercard, and American Express card numbers are included.

## Application Control

Although you can block the use of some applications by blocking the ports they use for communications, many applications do not use standard ports to communicate. Application control can detect the network traffic of more than 1000 applications, improving your control over application communication.

## ICAP

This module allows for the offloading of certain processes to a separate server so that your FortiGate firewall can optimize its resources and maintain the best level of performance possible.

## Security Profiles/lists/sensors

A profile is a group of settings that you can apply to one or more firewall policies. Each Security Profile feature is enabled and configured in a profile, list, or sensor. These are then selected in a security policy and the settings apply to all traffic matching the policy. For example, if you create an antivirus profile that enables antivirus scanning of HTTP traffic, and select the antivirus profile in the security policy that allows your users to access the World Wide Web, all of their web browsing traffic will be scanned for viruses.

Because you can use profiles in more than one security policy, you can configure one profile for the traffic types handled by a set of firewall policies requiring identical protection levels and types, rather than repeatedly configuring those same profile settings for each individual security policy.

For example, while traffic between trusted and untrusted networks might need strict protection, traffic between trusted internal addresses might need moderate protection. To provide the different levels of protection, you might configure two separate sets of profiles: one for traffic between trusted networks, and one for traffic between trusted and untrusted networks.

The Security Profiles include:

- antivirus profile
- IPS sensor
- Web filter profile
- Email filter profile
- Data Leak Prevention profile
- Application Control list
- VoIP profile

Although they're called profiles, sensors, and lists, they're functionally equivalent. Each is used to configure how the feature works.



# Client Reputation

The Security scan types available on FortiGate units are varied and tailored to detect specific attacks. However, sometimes user/client behavior can increase the risk of attack or infection. For example, if one of your network clients receives email viruses on a daily basis while no other clients receive these attachments, extra measures may be required to protect the client, or a discussion with the user about this issue may be worthwhile.

Before you can decide on a course of action, you need to know the problem is occurring. Client reputation can provide this information by tracking client behavior and reporting on activities that you determine are risky or otherwise noteworthy.

To learn how to enable and customize Client Reputation on your FortiGate go to the following FortiGate Cookbook video:

[Client Reputation Video](#)

## Summary of the Client Reputation features

Activities you can track include:

- **Bad Connection Attempts:** A typical BOT behavior is to connect to some hosts that do not exist on the Internet. This is because the BOT home needs to constantly change itself to dodge legislative enforcement or to hide from AV vendors. Bad connection attempts are tracked by:
  - Look ups for a DNS name that does not exist.
  - Connection attempts to an IP address that has no route.
  - HTTP 404 errors
- Packets that are blocked by security policies.
- **Intrusion protection:** Attack detected. The effect on reputation increases with severity of attack. A subscription to FortiGuard IPS updates is required.
- **Malware protection:** Malware detected. This requires a subscription to FortiGuard Antivirus updates.
- **Web activity:** Visit to web site in risky categories, including Potentially Liable, Adult/Mature Content, Bandwidth Consuming and Security Risk. A subscription to FortiGuard Web Filtering is required.
- **Application protection:** Client uses software in risky categories, including Botnet, P2P, Proxy, and Games applications. A subscription to FortiGuard IPS updates is required.
- **Geographical locations** that clients are communicating with. Access to the FortiGuard geographic database and a valid Fortinet support contract is required.

You can configure how severely each type of tracked activity will impact the reputation of the client in a sliding scale of Low, Medium, High or Critical. You can also choose to ignore an activity by setting it to Off. When an activity is turned off, it will have no effect on reputation.

You can enable client reputation tracking for your FortiGate unit by going to *Security Profiles > Client Reputation > Threat Level Definition*. Turning on client reputation tracking turns on traffic logging for all security policies, for all DoS policies and for all sniffer policies. While client

reputation is enabled, logging cannot be turned off for these policies. Traffic logging must be enabled for data to be added to the client reputation database.



Client reputation only highlights risky activity and does not include tools to stop it. Instead, client reputation is a tool that exposes risky behavior. When you uncover risky behavior that you are concerned about, you can take additional action to stop it. That action could include adding more restrictive security policies to block the activity or increase Security Profiles protection. You can also taking other measures outside your FortiGate unit to stop the activity.

To support client reputation your FortiGate unit must be registered, have a valid support contract and be licensed for FortiGuard antivirus, IPS and Web Filtering.

After client reputation is turned on, the FortiGate unit tracks recent behavior using a sliding window and displays current data for this window. The client reputation monitor displays clients and their activities in charts ordered according to how risky the behavior exhibited by the client is.

Client Reputation data is stored in traffic log messages in the newly added client reputation fields (*crscore* and *craction*). When you enable client reputation *Log Security Events* or *Log all Sessions* is enabled in all security policies. *Log Security Events* records traffic log messages for Security Profile sessions and *Log all Sessions* records traffic logs for all sessions. When Client Reputation is enabled you cannot select *No Log* in a security policy. Using client reputation data in log messages, you can configure FortiAnalyzer to produce a client reputation report.

Enabling client reputation can affect system performance if you had not been using traffic logging.

This chapter describes:

- [Applying client reputation monitoring to your network](#)
- [Viewing client reputation results](#)
- [Setting the client reputation profile/definition](#)
- [Expanding client reputation to include more types of behavior](#)
- [Client reputation execute commands](#)
- [Client reputation diagnose commands](#)

## Applying client reputation monitoring to your network

Client reputation monitoring is applied to network traffic by going to *Security Profiles > Client Reputation > Threat Level Definition* turning on *Client Reputation Tracking* and selecting *Apply*.

You can then either change the client reputation profile used by your FortiGate unit or you can accept the default profile. The client reputation profile indicates how risky you consider different types of client behavior to be. See [“Setting the client reputation profile/definition” on page 2028](#) for details.

## Viewing client reputation results


To view Client Reputation results go to *Security Profiles > Client Reputation > Reputation Score* to view the client reputation monitor. The monitor displays information about risky behavior as it was found. You can drill down into individual items to get more information about the behavior found and the client that caused it.

The client reputation monitor updates every 2 minutes. You can also select *Refresh* to manually update the display.

Select *Reset* to clear all client reputation data and restart the reporting window.

Figure 300 shows example client reputation results that shows activity from for different IP addresses that matched the kinds of traffic to be monitored according to the client reputation profile. You can see the IP address or name of each client and the amount of risky activity detected. The list at the bottom of the display shows more information about each device. The device information is gathered from enabling device monitoring by going to *User & Device > Device > Device Definition*.

**Figure 300:**Example client reputation results



You can select any of the bars in the graph to view information for each time the risky behavior was detected during the past 7 days (or whatever the Client Reputation window is). Information for each event detected includes the date and time the event was detected, the destination address, the application, and the client reputation score.

## Changing the client reputation reporting window and database size

By default, client reputation reports on activity for the last seven days. You can change this reporting window using the following command:

```
config client-reputation profile
 set window-size <interval_int>
end
```

Where <interval-int> is the reporting window in days. Range 1 to 30 days, default 7 days.

Enter the following command to set the client reputation report size:

```
config client-reputation profile
 set max-rep-db-size <size>
end
```

Where <size> can be from 10 to 2000 MBytes (2 GBytes). The default size is 100 MBytes.

## Client reputation data update and maintenance intervals

Client reputation updates its database every 2 minutes by querying the log database for client reputation information. This means that data displayed in the client reputation monitor is very current, at the most 2 minutes old.

Client reputation includes a data maintenance routine that runs every 12 hours to perform maintenance functions on the client reputation database. This routine:

- Checks the number of tracked hosts. If the number is at the maximum of 5000, the maintenance routine removes the oldest ten percent (500) of hosts from the list. If the number is less than the maximum, nothing changes.
- Deletes any reputation data associated with a host that is not in the tracking list (usually this only occurs if hosts are removed).
- Deletes any reputation data that is older than the current time minus the window-size in days.

## Setting the client reputation profile/definition

Configure the client reputation profile by going to *Security Profiles > Client Reputation > Threat Level Definition*. You configure one client reputation profile for all of the activity monitored by the FortiGate unit. The profile sets the risk levels for the types of behavior that client reputation monitors. You can set the risk to off, low, medium, high and critical for the following types of behavior:

- Application Protection
  - Botnet applications
  - P2P applications
  - Proxy applications
  - Games applications
- Intrusion protection (IPS)
  - Critical severity attack detected
  - High severity attack detected
  - Medium severity attack detected
  - Low severity attack detected
  - Informational severity attack detected
- Malware Protection
  - Malware detected
  - Botnet connection detected
- Packet based inspection
  - Blocked by firewall policy
  - Failed connection attempts
- Web Activity
  - All blocked URLs
  - Visit to security risk sites
  - Visit to potentially liable sites
  - Visit to adult/mature content sites
  - Visit to bandwidth consuming sites

Figure 301: Default client reputation profile

**Threat Level Definition**

**ON Client Reputation Tracking**

**Application Protection**

- Botnet Applications
- P2P Applications
- Proxy Applications
- Games Applications

**Intrusion Protection**

- Critical Severity Attack Detected
- High Severity Attack Detected
- Medium Severity Attack Detected
- Low Severity Attack Detected
- Informational Severity Attack Detected

**Malware Protection**

- Malware Detected
- Botnet Connection Detected

**Packet Based Inspection**

- Blocked by Firewall Policy
- Failed Connection Attempts

**Web Activity**

- All Blocked URLs
- Visit to Security Risk Sites
- Visit to Potentially Liable Sites
- Visit to Adult/Mature Content Sites
- Visit to Bandwidth Consuming Sites

**Risk Level Values**

LOW 5 MED 10 HIGH 30 CRIT 50

**Apply**

To configure the profile, decide how risky or dangerous each of the types of behavior are to your network and rate them accordingly. The higher you rate a type of behavior, the more visible clients engaging in this behavior will become in the client reputation monitor and the more easily you can detect this behavior.

For example, if you consider malware a high risk for your network, you can set the client reputation profile for malware to high or critical (as it is in the default client reputation profile). Then, whenever any amount of malware is detected, clients that originated the malware will be very visible in the client reputation monitor.

Set the risk to off for types of activity that you do not want client reputation to report on. This does not reduce the performance requirements or the amount of data gathered by client reputation, just the report output.

You can change a profile setting at any time and data that has already been collected will be used.

It is normally not necessary to change the *Risk Level Values* but it can be done if you need to alter the relative importance of the risk settings.

## Expanding client reputation to include more types of behavior

You can use the following command to change the client reputation profile from the CLI to include client reputation reporting about more settings:

```
config client-reputation profile
```

In addition to the settings configurable from the web-based manager, you can also set the following options:

- `geolocation` to enable reporting on connections to and from different countries (geographical locations). For example, use the following command to indicate that you consider communication with Aruba to be medium risk:

```
config client-reputation profile
 config geolocation
 edit 0
 set country AW
 set level medium
 end
 end
```

- `url-block-detected` to report on connections blocked by web filtering. Use the following command to enable reporting about blocked URLs and set the risk level to medium:

```
config client-reputation profile
 set url-block-detected medium
end
```

From the CLI you can configure client reputation to report more FortiGuard web filtering categories and more types of applications. For example, to report on social network activity (application control category 23):

```
config client-reputation-profile
 config application
 edit 0
 set category 23
 set level medium
 end
 end
```

To report on the local web filtering category (category 22):

```
config client-reputation-profile
 config web
 edit 0
 set group 22
 set level medium
 end
 end
```

## Client reputation execute commands

The `execute client-reputation` command includes the following options:

- `erase`, deletes all client reputation data.
- `host-count`, lists the clients that started sessions recorded by client reputation
- `host-detail`, for a specified client's IP address, displays the client reputation traffic log messages saved for that client.
- `host-summary`, for a specified client's IP address, displays the client's IP address, total entries, and total score.
- `purge`, deletes all data from the client reputation database.
- `topN`, display the top N clients identified by client reputation.

## Client reputation diagnose commands

The `diagnose client-reputation` command includes the following options

- `convert-timestamp` convert a client reputation database timestamp to date and time
- `test-all` adds log messages from multiple sources to the client reputation database for testing
- `test-app` adds application control log messages to the client reputation database for testing
- `test-ips` adds Intrusion Protection log messages to the client reputation database for testing
- `test-webfilter` adds webfilter log messages to the client reputation database for testing

# AntiVirus

This section describes how to configure the antivirus options. From an antivirus profile you can configure the FortiGate unit to apply antivirus protection to HTTP, FTP, IMAP, POP3, SMTP, IM, and NNTP sessions. If your FortiGate unit supports SSL content scanning and inspection, you can also configure antivirus protection for HTTPS, IMAPS, POP3S, SMTPS, and FTPS sessions. You can also apply flow-based antivirus protection to SMB or CIFS (Microsoft file sharing) traffic.

In many cases you can just customize the default antivirus profile and apply it to the security policy that accepts the traffic to be virus scanned. You can also create custom antivirus profiles if want to apply different types of virus protection to different traffic.

The following topics are included in this section:

- [Antivirus concepts](#)
- [Enable antivirus scanning](#)
- [Grayware scanning](#)
- [Windows file sharing \(CIFS\) flow-based antivirus scanning](#)
- [Advanced Persistent Threat \(APT\) protection](#)
- [Testing your antivirus configuration](#)
- [Antivirus examples](#)

## Antivirus concepts

The word “antivirus” refers to a group of features that are designed to prevent unwanted and potentially malicious files from entering your network. These features all work in different ways, which include checking for a file size, name, or type, or for the presence of a virus or grayware signature.

The antivirus scanning routines your FortiGate unit uses are designed to share access to the network traffic. This way, each individual feature does not have to examine the network traffic as a separate operation, and the overhead is reduced significantly. For example, if you enable file filtering and virus scanning, the resources used to complete these tasks are only slightly greater than enabling virus scanning alone. Two features do not require twice the resources.

### How antivirus scanning works

Antivirus scanning examines files for viruses, worms, trojans, and malware. The antivirus scan engine has a database of virus signatures it uses to identify infections. If the scanner finds a signature in a file, it determines that the file is infected and takes the appropriate action.

The most thorough scan requires that the FortiGate unit have the whole file for the scanning procedure. To achieve this, the antivirus proxy buffers the file as it arrives. Once the transmission is complete, the virus scanner examines the file. If no infection is present, it is sent to the destination. If an infection is present, a replacement message is set to the destination.

During the buffering and scanning procedure, the client must wait. With a default configuration, the file is released to the client only after it is scanned. You can enable client comforting in the Proxy Options profile to feed the client a trickle of data to prevent them from thinking the transfer is stalled, and possibly cancelling the download.



Buffering the entire file allows the FortiGate unit to eliminate the danger of missing an infection due to fragmentation because the file is reassembled before examination. Archives can also be expanded and the contents scanned, even if archives are nested.

Since the FortiGate unit has a limited amount of memory, files larger than a certain size do not fit within the memory buffer. The default buffer size is 10 MB. You can use the `uncompsizelimit` CLI command to adjust the size of this memory buffer.

Files larger than the buffer are passed to the destination without scanning. You can use the *Oversize File/Email* setting to block files larger than the antivirus buffer if allowing files that are too large to be scanned is an unacceptable security risk.

## Flow-based antivirus scanning

If your FortiGate unit supports flow-based antivirus scanning, you can choose to select it instead of proxy-based antivirus scanning. Flow-based antivirus scanning uses the FortiGate IPS engine to examine network traffic for viruses, worms, trojans, and malware, without the need to buffer the file being checked.

The advantages of flow-based scanning include faster scanning and no maximum file size. Flow-based scanning doesn't require the file be buffered so it is scanned as it passes through the FortiGate unit, packet-by-packet. This eliminates the maximum file size limit and the client begins receiving the file data immediately. Also, flow-based scanning does not change packets as they pass through the FortiGate unit, while proxy-based scanning can change packet details such as sequence numbers. The changes made by proxy-based scanning do not affect most networks.

The trade-off for these advantages is that flow-based scans detect a smaller number of infections. Viruses in documents, packed files, and some archives are less likely to be detected because the scanner can only examine a small portion of the file at any moment. Also, the file archive formats flow-based scanning will examine are limited to ZIP and GZIP.

## Antivirus scanning order

The antivirus scanning function includes various modules and engines that perform separate tasks.

### Proxy-based antivirus scanning order

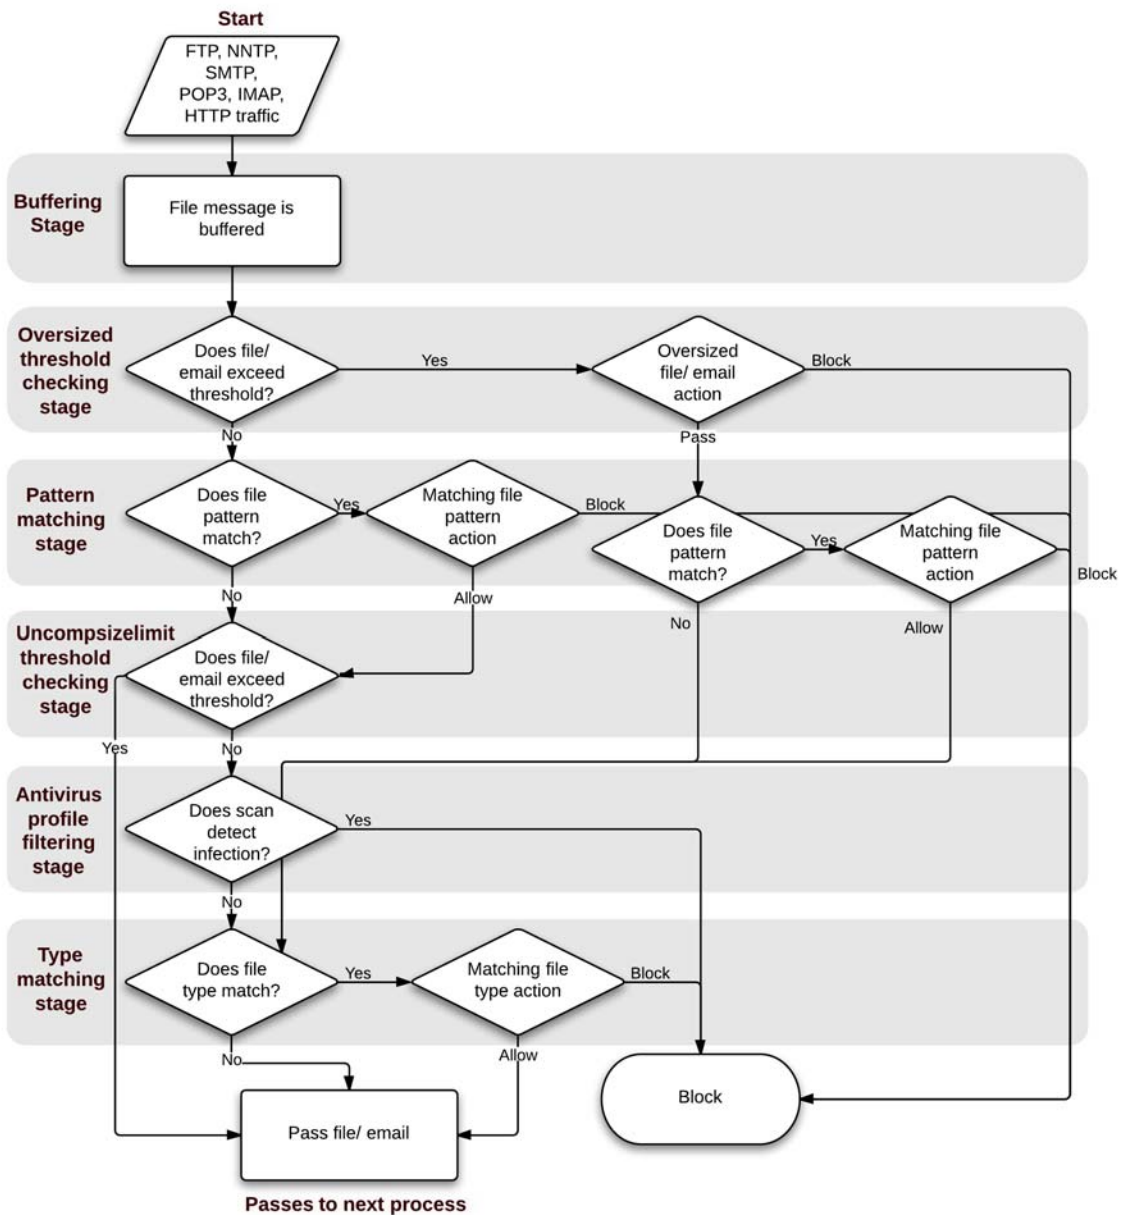
Figure 302 on page 2034 illustrates the antivirus scanning order when using proxy-based scanning. The first check for oversized files/email is to determine whether the file exceeds the configured size threshold. The `uncompsizelimit` check is to determine if the file can be buffered for file type and antivirus scanning. If the file is too large for the buffer, it is allowed to pass without being scanned. For more information, see the `config antivirus service` command. The antivirus scan includes scanning for viruses, as well as for grayware and heuristics if they are enabled.



File filtering includes file pattern and file type scans which are applied at different stages in the antivirus process.

---

**Figure 302:**Antivirus scanning order when using the normal, extended, or extreme database



If a file fails any of the tasks of the antivirus scan, no further scans are performed. For example, if the file *fakefile.EXE* is recognized as a blocked file pattern, the FortiGate unit will send the end user a replacement message, and delete or quarantine the file. The unit will not perform virus scan, grayware, heuristics, and file type scans because the previous checks have already determined that the file is a threat and have dealt with it.

### Flow-based antivirus scanning order

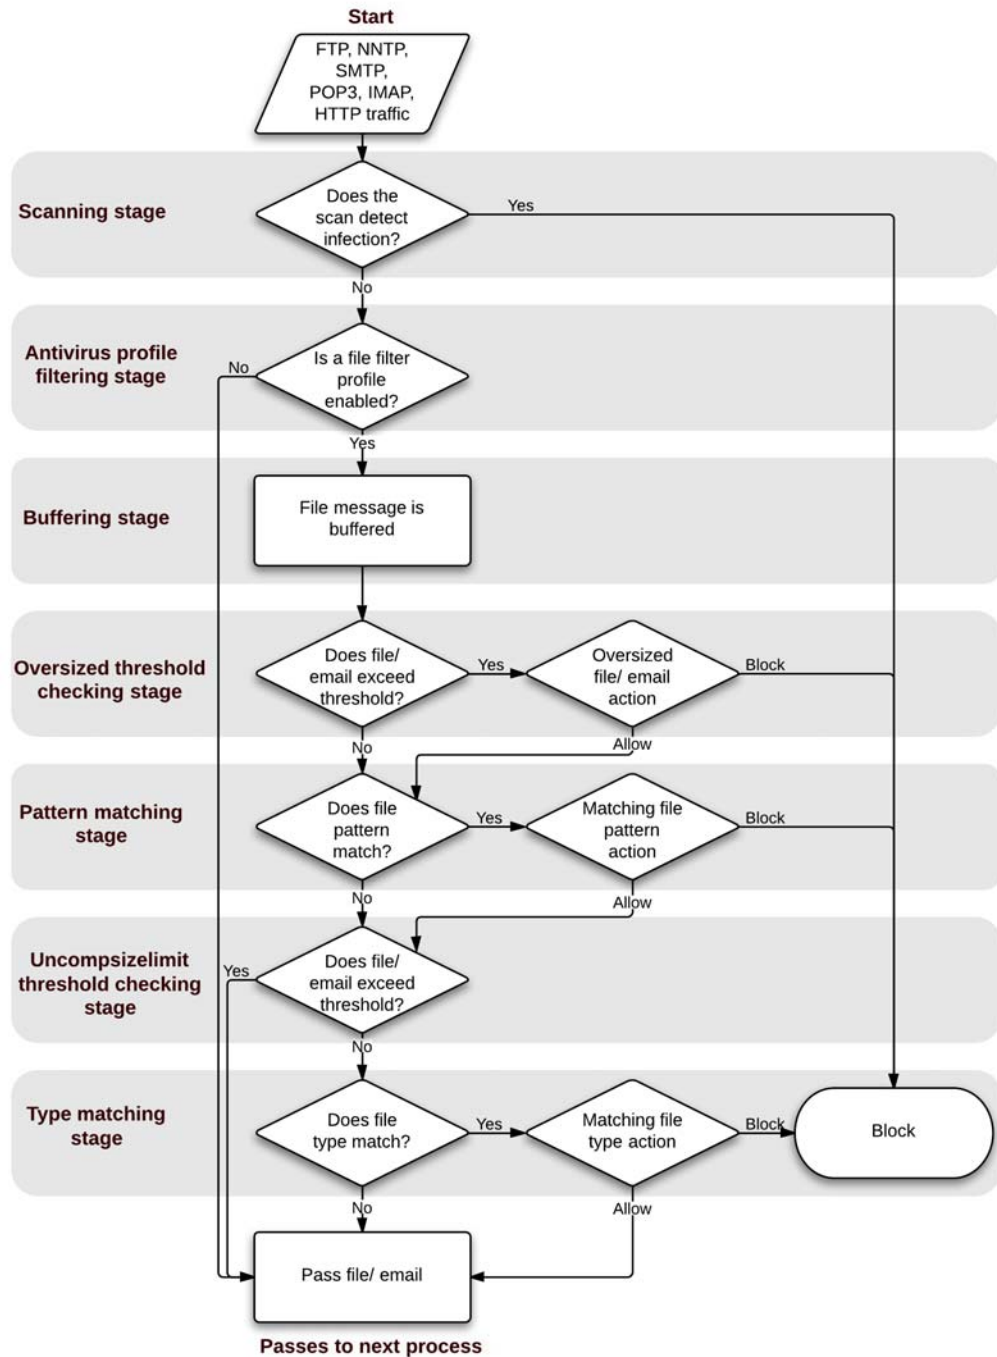
Figure 303 on page 2035 illustrates the antivirus scanning order when using flow-based scanning (i.e. the flow-based database). The antivirus scan takes place before any other

antivirus-related scan. If file filter is not enabled, the file is not buffered. The antivirus scan includes scanning for viruses, as well as for grayware and heuristics if they are enabled.



File filtering includes file pattern and file type scans which are applied at different stages in the antivirus process.

**Figure 303:**Antivirus scanning order when using flow-based scanning



## Antivirus databases

The antivirus scanning engine relies on a database of virus signatures to detail the unique attributes of each infection. The antivirus scan searches for these signatures, and when one is discovered, the FortiGate unit determines the file is infected and takes action.

All FortiGate units have the normal antivirus signature database but some models have additional databases you can select for use. Which you choose depends on your network and security needs.

<b>Normal</b>	Includes viruses currently spreading as determined by the FortiGuard Global Security Research Team. These viruses are the greatest threat. The Normal database is the default selection and it is available on every FortiGate unit.
<b>Extended</b>	Includes the normal database in addition to recent viruses that are no-longer active. These viruses may have been spreading within the last year but have since nearly or completely disappeared.
<b>Extreme</b>	Includes the extended database in addition to a large collection of 'zoo' viruses. These are viruses that have not spread in a long time and are largely dormant today. Some zoo viruses may rely on operating systems and hardware that are no longer widely used.

## Antivirus techniques

The antivirus features work in sequence to efficiently scan incoming files and offer your network optimum antivirus protection. The first four features have specific functions, the fifth, heuristics, protects against new, or previously unknown virus threats. To ensure that your system is providing the most protection available, all virus definitions and signatures are updated regularly through the FortiGuard antivirus services. The features are discussed in the order that they are applied, followed by FortiGuard antivirus.

### Virus scan

If the file passes the file pattern scan, the FortiGate unit applies a virus scan to it. The virus definitions are kept up-to-date through the FortiGuard Distribution Network (FDN). For more information, see [“FortiGuard Antivirus” on page 2036](#).

### Grayware

If the file passes the virus scan, it will be checked for grayware. Grayware configurations can be turned on and off as required and are kept up to date in the same manner as the antivirus definitions. For more information, see [“Grayware scanning” on page 2041](#).

### Heuristics

After an incoming file has passed the grayware scan, it is subjected to the heuristics scan. The FortiGate heuristic antivirus engine, if enabled, performs tests on the file to detect virus-like behavior or known virus indicators. In this way, heuristic scanning may detect new viruses, but may also produce some false positive results. You configure heuristics from the CLI.

## FortiGuard Antivirus

FortiGuard Antivirus services are an excellent resource which includes automatic updates of virus and IPS (attack) engines and definitions, as well as the local spam DNS black list (DNSBL),

through the FDN. The [FortiGuard Center](#) web site also provides the FortiGuard Antivirus virus and attack encyclopedia.

The connection between the FortiGate unit and FortiGuard Center is configured in *System > Config > FortiGuard*.

## Enable antivirus scanning

Antivirus scanning is configured in an antivirus profile, but it is enabled in a firewall policy. Once the use of an antivirus profile is enabled and selected in one or more firewall policies, all the traffic controlled by those firewall policies will be scanned according to your settings.

### Antivirus Profiles

From *Security Profiles > Antivirus > Profile* you can configure antivirus profiles that are then applied to firewall policies. A profile is specific configuration information that defines how the traffic within a policy is examined and what action may be taken based on the examination.

You can create multiple antivirus profiles for different antivirus scanning requirements. For example, you create an antivirus profile that specifies only virus scanning for POP3 which you then apply to the out-going firewall policy. You can also choose specific protocols, such as POP3, that will be blocked and then archived by the unit. This option is available only in the CLI.

Within antivirus profiles, you can also choose specific protocols to be blocked and then archive them. This is available only in the CLI.

#### To enable antivirus scanning – web-based manager

1. Go to *Security Profiles > AntiVirus > Profile*.
2. View and optionally change the *default* antivirus profile.
  - You can also select *Create New* to create a new antivirus profile, or select an existing antivirus profile and choose *Edit*.
3. Select the inspection and the traffic you want scanned for viruses.
4. Select *OK*.
5. Go to *Policy > Policy > Policy* and either add or select the security policy that accepts the traffic to be virus scanned.
6. Turn on antivirus and select the profile that you configured.
7. Select *OK* to save the security policy.

#### To enable antivirus scanning – CLI

You need to configure the scan option for each type of traffic you want scanned. In this example, antivirus scanning of HTTP traffic is enabled in the profile.

```
config antivirus profile
 edit default
 config http
 set options scan
 end
 end
```

Then enter a command similar to the following to add the default antivirus profile to a security policy.

```
config firewall policy
 edit 0
 set srcintf internal
 set dstintf wan1
 set srcaddr all
 set dstaddr all
 set schedule always
 set service ALL
 set action allow
 set utm-status enable
 set av-profile default
 end
```

## Changing the default antivirus database

If your FortiGate unit supports extended, extreme, or flow-based virus database definitions, you can select the virus database most suited to your needs.

In most circumstances, the regular virus database provides sufficient protection. Viruses known to be active are included in the regular virus database. The extended database includes signatures of the viruses that have become rare within the last year in addition to those in the normal database. The extreme database includes legacy viruses that have not been seen in the wild in a long time in addition to those in the extended database.

The flow-based database contains a subset of the virus signatures in the extreme database. Unlike the other databases, selecting the flow-based database also changes the way the FortiGate unit scans your network traffic for viruses. Instead of the standard proxy-based scan, network traffic is scanned as it streams through the FortiGate unit. For more information on the differences between flow-based and proxy-based antivirus scanning, see [“How antivirus scanning works” on page 2032](#).

If you require the most comprehensive antivirus protection, enable the extended virus database. The additional coverage comes at a cost, however, because the extra processing requires additional resources.

### To change the antivirus database

```
config antivirus settings
 set default-db extended
end
```

## Configuring the scan buffer size

When checking files for viruses using the proxy-based scanning method, there is a maximum file size that can be buffered. Files larger than this size are passed without scanning. The default size for all FortiGate models is 10 megabytes.

Archived files are extracted and email attachments are decoded before the FortiGate unit determines if they can fit in the scan buffer. For example, a 7 megabyte ZIP file containing a 12 megabyte EXE file will be passed without scanning with the default buffer size. Although the archive would fit within the buffer, the uncompressed file size will not.

In this example, the `uncompsizelimit` CLI command is used to change the scan buffer size to 20 megabytes for files found in HTTP traffic:

```
config antivirus service http
 set uncompsizelimit 20
end
```

The maximum buffer size varies by model. Enter `set uncompsizelimit?` to display the buffer size range for your FortiGate unit.



Flow-based scanning does not use a buffer and therefore has no file-size limit. File data is scanned as it passes through the FortiGate unit. The `uncompsizelimit` setting has no effect for flow-based scanning.

## Configuring archive scan depth

The antivirus scanner will open archives and scan the files inside. Archives within other archives, or nested archives, are also scanned to a default depth of twelve nestings. You can adjust the number of nested archives to which the FortiGate unit will scan with the `uncompnestlimit` CLI command. Further, the limit is configured separately for each traffic type.

For example, this CLI command sets the archive scan depth for SMTP traffic to 5. That is, archives within archives will be scanned five levels deep.

```
config antivirus service smtp
 set uncompnestlimit 5
end
```

You can set the nesting limit from 2 to 100.

## Configuring a maximum allowed file size

Proxy options allow you to enforce a maximum allowed file size for each of the network protocols in the profile. They are HTTP, FTP, IMAP, POP3, SMTP, IM, and NNTP. If your FortiGate unit supports SSL content scanning and inspection, you can also configure a maximum file size for HTTPS, IMAPS, POP3S, SMTPS, and FTPS.

The action you set determines what the FortiGate unit does with a file that exceeds the oversized file threshold. Two actions are available:

<b>Block</b>	Files that exceed the oversize threshold are dropped and a replacement message is sent to the user instead of the file.
<b>Pass</b>	Files exceed the oversized threshold are allowed through the FortiGate unit to their destination. Note that passed files are not scanned for viruses. File Filtering, both file pattern and file type, are applied, however.

You can also use the maximum file size to help secure your network. If you're using a proxy-based virus scan, the proxy scan buffer size limits the size of the files that can be scanned for infection. Files larger than this limit are passed without scanning. If you configure the maximum file size to block files larger than the scan buffer size, large infected files will not by-pass antivirus scanning.

In this example, the maximum file size will be configured to block files larger than 10 megabytes, the largest file that can be antivirus scanned with the default settings. You will need to configure a proxy options profile and add it to a security policy.

### **Set proxy options profile to block files larger than 10 MB**

1. Go to *Policy > Policy > Proxy Options*.
2. Edit the default or select *Create New* to add a new one.
3. Scroll down to the common Options Section and place a check in the box next to *BlockOversized File/Email*
4. The sub line *Threshold (MB)* will appear with a value field. Enter 10.
5. Select *OK* or *Apply*.

The proxy options profile is configured, but to block files, you must select it in the firewall policies handling the traffic that contains the files you want blocked.

### **To select the Proxy Options profile in a security policy**

1. Go to *Policy > Policy > Policy*.
2. Edit or create a security policy.
3. Select a proxy-based security profile. You will know that there is a proxy component to the Security Profile because when a Security Profile is Proxy based the *Proxy Options* field will be visible (for example, select an Antivirus profile that includes proxy scanning).
4. Beside *Proxy Options* select the name of the MTU proxy options protocol.
5. Select *OK* to save the security policy.

Once you complete these steps, any files in the traffic subject to Security Profile scanning handled by this policy that are larger than 10MB will be blocked. If you have multiple firewall policies, examine each to determine if you want to apply similar file blocking the them as well.

## **Configuring client comforting**

When proxy-based antivirus scanning is enabled, the FortiGate unit buffers files as they are downloaded. Once the entire file is captured, the FortiGate unit scans it. If no infection is found, the file is sent along to the client. The client initiates the file transfer and nothing happens until the FortiGate finds the file clean, and releases it. Users can be impatient, and if the file is large or the download slow, they may cancel the download, not realizing that the transfer is in progress.

The client comforting feature solves this problem by allowing a trickle of data to flow to the client so they can see the file is being transferred. The default client comforting transfer rate sends one byte of data to the client every ten seconds. This slow transfer continues while the FortiGate unit buffers the file and scans it. If the file is infection-free, it is released and the client will receive the remainder of the transfer at full speed. If the file is infected, the FortiGate unit caches the URL and drops the connection. The client does not receive any notification of what happened because the download to the client had already started. Instead, the download stops and the user is left with a partially downloaded file.

If the user tries to download the same file again within a short period of time, the cached URL is matched and the download is blocked. The client receives the Infection cache message



replacement message as a notification that the download has been blocked. The number of URLs in the cache is limited by the size of the cache.



Client comforting can send unscanned and therefore potentially infected content to the client. You should only enable client comforting if you are prepared to accept this risk. Keeping the client comforting interval high and the amount low will reduce the amount of potentially infected data that is downloaded.

Client comforting is available for HTTP and FTP traffic. If your FortiGate unit supports SSL content scanning and inspection, you can also configure client comforting for HTTPS and FTPS traffic.

#### **Enable and configure client comforting**

1. Go to *Policy > Policy > Proxy Options*.
2. Select a Proxy Options profile and choose *Edit*, or select *Create New* to make a new one.
3. Scroll down to the *Common Options* section and check the box next to *Comfort Clients*. This will set the option on all of the applicable protocols. The ability to set this feature on a protocol by protocol basis exists in the CLI
4. Select *OK* or *Apply* to save the changes.
5. Select this Proxy Options profile in any security policy for it to take effect on all traffic handled by the policy.

The default values for Interval and Amount are 10 and 1, respectively. This means that when client comforting takes effect, 1 byte of the file is sent to the client every 10 seconds. You can change these values to vary the amount and frequency of the data transferred by client comforting.

## **Grayware scanning**

Grayware programs are unsolicited software programs installed on computers, often without the user's consent or knowledge. Grayware programs are generally considered an annoyance, but they can also cause system performance problems or be used for malicious purposes.

To allow the FortiGate unit to scan for known grayware programs, you must enable both antivirus scanning and grayware detection.

Enter the following command to enable grayware detection:

```
config antivirus settings
 set grayware enable
end
```

With grayware detection enabled, the FortiGate unit will scan for grayware any time it checks for viruses.

## **Windows file sharing (CIFS) flow-based antivirus scanning**

FortiOS 5.0 now supports virus scanning of Windows file sharing traffic. This includes CIFS, SMB, and SAMBA traffic. This feature is applied by enabling SMB scanning in an antivirus profile and then adding this profile to a security policy that accepts CIFS traffic. CIFS virus scanning is available only through flow-based antivirus scanning.

FortiOS 5.0 flow-based virus scanning can detect the same number of viruses in CIFS/SMB/SAMBA traffic as it can for all supported content protocols.

**Figure 304:**Configuring CIFS/SMB/SAMBA virus scanning

**New AntiVirus Profile**

Name: SMB-CIFS-SAMBA-only

Comments: AV scanning of SMB, CIFS, and SAMBA traffic only 48/255

Inspection Mode:  Proxy  Flow-based

Protocol	Virus Scan and Removal
<b>Web</b>	
HTTP	<input type="checkbox"/>
<b>Email</b>	
SMTP	<input type="checkbox"/>
POP3	<input type="checkbox"/>
IMAP	<input type="checkbox"/>
MAPI	<input type="checkbox"/>
<b>File Transfer</b>	
FTP	<input type="checkbox"/>
SMB	<input checked="" type="checkbox"/>
<b>IM</b>	
ICQ, Yahoo, MSN Messenger	<input type="checkbox"/>

OK Cancel

Use the following command to enable CIFS/SMB/SAMBA virus scanning in an antivirus profile:

```
config antivirus profile
 edit smb-profile
 config smb
 set options scan
 set avdb flow-based
 end
```

Then add this antivirus profile to a security policy that accepts the traffic to be virus scanned. In the security policy the service can be set to ANY, SAMBA, or SMB.

```
config firewall policy
 edit 0
 set service ANY
 ...
 set utm-status enable
 set av-profile smb-profile
 end
```

Note the following about CIFS/SMB/SAMBA virus scanning:

- Some newer version of SAMBA clients and SMB2 can spread one file across multiple sessions, preventing some viruses from being detected if this occurs.
- Enabling CIFS/SMB/SAMBA virus scanning can affect FortiGate performance.
- SMB2 is a new version of SMB that was first partially implemented in Windows Vista.
- Currently SMB2 is supported by Windows Vista or later, and partly supported by Samba 3.5 and fully support by Samba 3.6.
- The latest version of SMB2.2 will be introduced with Windows 8.
- Most clients still use SMB as default setting.

## Advanced Persistent Threat (APT) protection

New advanced persistent threat (APT) protection features in FortiOS 5.0 include botnet protection, phishing protection, and zero-day threat protection using FortiGuard Analytics for sandboxing.

### Botnet and phishing protection

In an antivirus profile you can configure the FortiGate unit to detect and block botnet connection attempts. This feature also blocks attempted access to phishing URLs.

The antivirus database includes a constantly updated database of known command and control (C&C) sites that Botnet clients attempt to connect too as well as a database of phishing URLs.

To enable Botnet and phishing protection in an antivirus profile select Block Connections to Botnet Servers. Botnet protection is available for proxy and flow-based antivirus profiles.

**Figure 305:** Adding Botnet and phishing protection.

**Edit Antivirus Profile** default

Name: default

Comments: scan and delete virus 21/255

Inspection Mode:  Proxy  Flow-based

Block Connections to Botnet Servers

Protocol	Virus Scan and Removal
<b>Web</b>	
HTTP	<input checked="" type="checkbox"/>
<b>Email</b>	
SMTP	<input checked="" type="checkbox"/>
POP3	<input checked="" type="checkbox"/>
IMAP	<input checked="" type="checkbox"/>
MAPI	<input type="checkbox"/>
<b>File Transfer</b>	
FTP	<input checked="" type="checkbox"/>
SMB	<input type="checkbox"/>
<b>IM</b>	
ICQ, Yahoo, MSN Messenger	<input checked="" type="checkbox"/>

Apply

### FortiGuard Sandbox (in the cloud sandboxing, zero day threat analysis and submission)

In a Proxy Mode antivirus profile, enabling Send Files to *FortiGuard Sandbox for Inspection* causes your FortiGate unit to upload files to FortiGuard where the file will be executed and the resulting behavior analyzed for risk. You have the choice of uploading all files or only the suspicious ones. If the file exhibits risky behavior or is found to contain a virus, a new virus signature is created and added to the FortiGuard antivirus signature database. The next time your FortiGate unit updates its antivirus database it will have the new signature.

Currently, a file is considered suspicious if it does not contain a known virus and if it has some suspicious characteristics. The suspicious characteristics can change depending on the current threat climate and other factors. Fortinet optimizes how files are uploaded as required.



The FortiGuard Sandbox feature is available if you have a valid FortiCloud subscription. To verify whether or not a subscription is associated with your FortiGate go to *System > Dashboard > Status* and check the License Information widget in the FortiCloud subsection.

**Figure 306:**Enabling FortiGuard Sandbox in an Antivirus Profile

Protocol	Virus Scan and Removal
<b>Web</b>	
HTTP	<input checked="" type="checkbox"/>
<b>Email</b>	
SMTP	<input checked="" type="checkbox"/>
POP3	<input checked="" type="checkbox"/>
IMAP	<input checked="" type="checkbox"/>
MAPI	<input type="checkbox"/>
<b>File Transfer</b>	
FTP	<input checked="" type="checkbox"/>
<b>IM</b>	
ICQ, Yahoo, MSN Messenger	<input checked="" type="checkbox"/>

On the FortiGate, there are two ways to verify that files are being uploaded to the FortiCloud Sandbox. The first is to go to *System > Config > FortiSandbox*. The window is for configuring whether or not the FortiGate unit is to use the FortiCloud Sandbox or a FortiSandbox Appliance but it also shows the statistics of files submitted to the Sandbox over the last seven days.

The second method is to go to *System > Dashboard > Status* and view the *Advanced Threat Protection Statistics* dashboard widget. This widget will show essentially the same information. This widget is not one of the default ones so you will have to add it to the Dashboard.

**Figure 307:**Example Advanced Threat Protection Statistics widget showing Sandbox submissions

FortiGate Statistics	
Number of Files Scanned	6311
Malicious	40
Detected Zero-Day Malware Variants	0
Suspicious Files	0
Clean	6271

FortiGuard Sandbox Statistics (Last 7 Days)	
# of Files Submitted to FortiGuard Sandbox	1798
Malicious	0
Clean	1735

To view information relating to the Antivirus function from the FortiCloud side, go to *System > Dashboard > Status* and look at the *License Information* widget. In the *FortiCloud* subsection in the *Account* line, select the *Launch Portal* link. Once at the portal select the icon for the specific FortiGate that you view the information for.

Under the Logs & Archives tab of the menu bar you will find the UTM option. Once this option is selected, you will have the option of choosing AntiVirus. The site will display records within the designated time frame that refer to AntiVirus events recorded by the logs.

**Figure 308:**Example view of FortiCloud's AntiVirus logs

#	Time	Level	Source	Destination	Source Interface	Destination Interface	Message	Reference
1	04-01 21:23	2	192.168.10.100	172.20.120.111	port10	wan1	File submitted to Sandbox.	
2	04-01 21:23	2	192.168.10.100	172.20.120.111	port10	wan1	File submitted to Sandbox.	
3	04-01 21:23	2	192.168.10.100	172.20.120.111	port10	wan1	File submitted to Sandbox.	
4	04-01 21:23	2	192.168.10.100	172.20.120.111	port10	wan1	File submitted to Sandbox.	
5	04-01 21:23	2	192.168.10.100	172.20.120.111	port10	wan1	File submitted to Sandbox.	
6	04-01 21:23	2	192.168.10.100	172.20.120.111	port10	wan1	File submitted to Sandbox.	
7	04-01 21:23	2	192.168.10.100	172.20.120.111	port10	wan1	File submitted to Sandbox.	
8	04-01 21:23	2	192.168.10.100	172.20.120.111	port10	wan1	File submitted to Sandbox.	
9	04-01 21:23	2	192.168.10.100	172.20.120.111	port10	wan1	File submitted to Sandbox.	
10	04-01 21:23	2	192.168.10.100	172.20.120.111	port10	wan1	File submitted to Sandbox.	

In addition to the normal UTM logs, there is a new menu item in that top menu bar that appears when your FortiGate is configured to submit files to the FortiSandbox. This page on the site will

display more granular information on files with viruses that are submitted by your FortiGate unit. This information will include:

- Date and Time
- File Name
- User Name
- Service
- Source IP
- Destination IP
- Vdom
- Analysis
- URL

## Testing your antivirus configuration

You have configured your FortiGate unit to stop viruses, but you'd like to confirm your settings are correct. Even if you have a real virus, it would be dangerous to use for this purpose. An incorrect configuration will allow the virus to infect your network.

To solve this problem, the European Institute of Computer Anti-virus Research has developed a test file that allows you to test your antivirus configuration. The EICAR test file is not a virus. It can not infect computers, nor can it spread or cause any damage. It's a very small file that contains a sequence of characters. Your FortiGate unit recognizes the EICAR test file as a virus so you can safely test your FortiGate unit antivirus configuration.

Go to <http://www.fortiguard.com/antivirus/eicartest.html> to download the test file (eicar.com) or the test file in a ZIP archive (eicar.zip).

If the antivirus profile applied to the security policy that allows you access to the Web is configured to scan HTTP traffic for viruses, any attempt to download the test file will be blocked. This indicates that you are protected.

## Antivirus examples

The following examples provide a sample antivirus configuration scenario for a fictitious company.

### Configuring simple antivirus protection

Small offices, whether they are small companies, home offices, or satellite offices, often have very simple needs. This example details how to enable antivirus protection on a FortiGate unit located in a satellite office. The satellite office does not have an internal email server. To send and retrieve email, the employees connect to an external mail server.

#### Creating an antivirus profile

Most antivirus settings are configured in an antivirus profile. Antivirus profiles are selected in firewall policies. This way, you can create multiple antivirus profiles, and tailor them to the traffic controlled by the security policy in which they are selected. In this example, you will create one antivirus profile.

#### **To create an antivirus profile – web-based manager**

1. Go to *Security Profiles > AntiVirus > Profiles*.

2. Select *Create New*.
3. In the *Name* field, enter `basic_antivirus`.
4. In the *Comments* field, enter `Antivirus protection for web and email traffic`.
5. Select the *Virus Scan* check boxes for the *HTTP*, *IMAP*, *POP3*, and *SMTP* traffic types.
6. Select *OK* to save the antivirus profile.

#### To create an antivirus profile — CLI

```
config antivirus profile
 edit basic_antivirus
 set comment "Antivirus protection for web and email traffic"
 config http
 set options scan
 end
 config imap
 set options scan
 end
 config pop3
 set options scan
 end
 config smtp
 set options scan
 end
 end
end
```

#### Selecting the antivirus profile in a security policy

An antivirus profile directs the FortiGate unit to scan network traffic only when it is selected in a security policy. When an antivirus profile is selected in a security policy, its settings are applied to all the traffic the security policy handles.

#### To select the antivirus profile in a security policy — web-based manager

1. Go to *Policy > Policy > Policy*.
2. Create a new or edit a security policy.
3. Turn on *Antivirus*.
4. Select an antivirus profile.
5. Select *OK* to save the security policy.

#### To select the antivirus profile in a security policy — CLI

```
config firewall policy
 edit 1
 set utm-status enable
 set profile-protocol-options default
 set av-profile basic_antivirus
 end
end
```

HTTP, IMAP, POP3, and SMTP traffic handled by the security policy you modified will be scanned for viruses. A small office may have only one security policy configured. If you have multiple policies, consider enabling antivirus scanning for all of them.

## Protecting your network against malicious email attachments

Grayware is commonly delivered by email or the web. The Example.com corporation has been the victim of multiple greyware infections in the past. Now that the company has a FortiGate unit protecting its network, you (Example.com's system administrator) can configure the unit to scan email and web traffic to filter out greyware attachments.

### Enabling antivirus scanning in the antivirus profile

The primary means to avoid viruses is to configure the FortiGate unit to scan email and web traffic for virus signatures. You enable virus scanning in the antivirus profile and then select the antivirus profile in firewall policies that control email traffic.

#### To enable antivirus scanning in the antivirus profile

1. Go to *Security Profiles > AntiVirus > Profiles*.
2. Create a new or edit an antivirus profile.
3. Select *Virus Scan and Removal for HTTP* to scan web traffic for viruses.
4. Select the *Virus Scan* check box for *IMAP*, *POP3*, and *SMTP* to scan all email protocols for viruses.
5. Select *OK* or *Apply* to save the antivirus profile.

### Selecting the antivirus profile in a security policy

An antivirus profile directs the FortiGate unit to scan network traffic only when it is selected in a security policy. When an antivirus profile is selected in a security policy, its settings are applied to all the traffic the security policy handles.

#### To select the antivirus profile in a security policy

1. Go to *Policy > Policy > Policy*.
2. Create or edit a policy that controls the network traffic controlling email traffic.
3. Turn on Antivirus.
4. Select an antivirus profile.
5. Select *OK* to save the security policy.



# Email filter

This section describes how to configure FortiGate email filtering for IMAP, POP3, and SMTP email. Email filtering includes both spam filtering and filtering for any words or files you want to disallow in email messages. If your FortiGate unit supports SSL content scanning and inspection, you can also configure spam filtering for IMAPS, POP3S, and SMTPS email traffic.

The following topics are included in this section:

- [Email filter concepts](#)
- [Enable email filtering](#)
- [Configure email traffic types to inspect](#)
- [Configure the spam action](#)
- [Configure the tag location](#)
- [Configure the tag format](#)
- [Configure FortiGuard email filters](#)
- [Configure local email filters](#)
- [Email filter examples](#)

## Email filter concepts

You can configure the FortiGate unit to manage unsolicited commercial email by detecting and identifying spam messages from known or suspected spam servers.

The FortiGuard Antispam Service uses both a sender IP reputation database and a spam signature database, along with sophisticated spam filtering tools, to detect and block a wide range of spam messages. Using FortiGuard Antispam email filter profile settings, you can enable IP address checking, URL checking, email checksum checking, and spam submission. Updates to the IP reputation and spam signature databases are provided continuously via the global FortiGuard Distribution Network.

From the [FortiGuard Antispam Service](#) page in the FortiGuard Center, you can find out whether an IP address is blacklisted in the FortiGuard antispam IP reputation database, or whether a URL or email address is in the signature database.

## Email filter techniques

The FortiGate unit has a number of techniques available to help detect spam. Some use the FortiGuard Antispam Service and require a subscription. The remainder use your DNS servers or use lists that you must maintain.

### FortiGuard IP address check

The FortiGate unit queries the FortiGuard Antispam Service to determine if the IP address of the client delivering the email is blacklisted. A match will cause the FortiGate unit to treat delivered messages as spam.

The default setting of the `smtp-spamhdrip` CLI command is `disable`. If enabled, the FortiGate unit will check all the IP addresses in the header of SMTP email against the FortiGuard Antispam Service. For more information, see the [FortiGate CLI Reference](#).

## FortiGuard URL check

The FortiGate unit queries the FortiGuard Antispam service to determine if any URL in the message body is associated with spam. If any URL is blacklisted, the FortiGate unit determines that the email message is spam.

## Detect phishing URLs in email

The FortiGate unit sends the URL links in email messages to FortiGuard to determine if the links are associated with a known phishing site. If such a link is detected, the link is removed from the message. The URL remains, but it is no longer a selectable hyperlink.

## FortiGuard email checksum check

The FortiGate unit sends a hash of an email to the FortiGuard Antispam server, which compares the hash to hashes of known spam messages stored in the FortiGuard Antispam database. If the hash results match, the email is flagged as spam.

## FortiGuard spam submission

Spam submission is a way you can inform the FortiGuard AntiSpam service of non-spam messages incorrectly marked as spam. When you enable this setting, the FortiGate unit adds a link to the end of every message marked as spam. You then select this link to inform the FortiGuard AntiSpam service when a message is incorrectly marked.

## IP address black/white list check

The FortiGate unit compares the IP address of the client delivering the email to the addresses in the IP address black/white list specified in the email filter profile. If a match is found, the FortiGate unit will take the action configured for the matching black/white list entry against all delivered email.

The default setting of the `smtp-spamhdrip` CLI command is `disable`. If enabled, the FortiGate unit will check all the IP addresses in the header of SMTP email against the specified IP address black/white list. For more information, see the [FortiGate CLI Reference](#).

## HELO DNS lookup

The FortiGate unit takes the domain name specified by the client in the HELO greeting sent when starting the SMTP session and does a DNS lookup to determine if the domain exists. If the lookup fails, the FortiGate unit determines that any messages delivered during the SMTP session are spam.

## Email address black/white list check

The FortiGate unit compares the sender email address, as shown in the message envelope MAIL FROM, to the addresses in the email address black/white list specified in the email filter profile. If a match is found, the FortiGate unit will take the action configured for the matching black/white list entry.

## Return email DNS check

The FortiGate unit performs a DNS lookup on the reply-to domain to see if there is an A or MX record. If no such record exists, the message is treated as spam.

## Banned word check

The FortiGate unit blocks email messages based on matching the content of the message with the words or patterns in the selected spam filter banned word list. This feature is only available in the CLI.

## Order of spam filtering

The FortiGate unit checks for spam using various filtering techniques. The order in which the FortiGate unit uses these filters depends on the mail protocol used.

Filters requiring a query to a server and a reply (FortiGuard Antispam Service and DNSBL/ORDBL) are run simultaneously. To avoid delays, queries are sent while other filters are running. The first reply to trigger a spam action takes effect as soon as the reply is received.

Each spam filter passes the email to the next if no matches or problems are found. If the action in the filter is *Mark as Spam*, the FortiGate unit tags the email as spam according to the settings in the email filter profile.

For SMTP and SMTPS, if the action is discard, the email message is discarded or dropped.

If the action in the filter is *Mark as Clear*, the email is exempt from any remaining filters. If the action in the filter is *Mark as Reject*, the email session is dropped. Rejected SMTP or SMTPS email messages are substituted with a configurable replacement message.

## Order of SMTP and SMTPS spam filtering

The FortiGate unit scans SMTP and SMTPS email for spam in the order given below. SMTPS spam filtering is available on FortiGate units that support SSL content scanning and inspection.

1. IP address black/white list (BWL) check on last hop IP
2. DNSBL & ORDBL check on last hop IP, FortiGuard Antispam IP check on last hop IP, HELO DNS lookup
3. MIME headers check, E-mail address BWL check
4. Banned word check on email subject
5. IP address BWL check (for IPs extracted from "Received" headers)
6. Banned word check on email body
7. Return email DNS check, FortiGuard Antispam email checksum check, FortiGuard Antispam URL check, DNSBL & ORDBL check on public IP extracted from header.

## Order of IMAP, POP3, IMAPS and POP3S spam filtering

The FortiGate unit scans IMAP, POP3, IMAPS and POP3S email for spam in the order given below. IMAPS and POP3S spam filtering is available on FortiGate units that support SSL content scanning and inspection.

1. MIME headers check, E-mail address BWL check
2. Banned word check on email subject
3. IP BWL check
4. Banned word check on email body
5. Return email DNS check, FortiGuard Antispam email checksum check, FortiGuard Antispam URL check, DNSBL & ORDBL check.

## Enable email filtering

Unlike antivirus protection, no single control enables all email filtering. Your FortiGate unit uses many techniques to detect spam; some may not be appropriate for every situation. For this reason, when you enable email filtering, you must then choose when techniques are applied to email traffic.

### To enable email filtering

1. Go to *Security Profiles > Email Filter > Profile*.

The default email filter profile is presented. You can edit this profile or create a new one.

2. Select the *Inspection Mode*.

Proxy detection involves buffering the file and examining it as a whole. Advantages of proxy-based detection include a more thorough examination of attachments, especially archive formats and nesting.

Flow-based detection examines the file as it passes through the FortiGate unit without any buffering. Advantages of flow-based detection include speed and no interruption of detection during conserve mode.

3. Select *Enable Spam Detection and Filtering*.

4. If you wish to leave everything in its default setting you can select *OK* or *Apply*.

Once you have enabled the email filter you can further specify what protocols to inspect.

## Configure email traffic types to inspect

The FortiGate unit examines IMAP, POP3, and SMTP email traffic. If your FortiGate unit supports content inspection, it can also examine IMAPS, POP3S, and SMTPS traffic. The options that you will see in the profile window are IMAP, POP3 and SMTP

### To select the email traffic types to inspect

1. Go to *Security Profiles > Email Filter > Profile*.

2. The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.

3. Select *Enable Spam Detection and Filtering*.

4. Select the types of email that you want the FortiGate unit to examine when using this email filter profile.

5. Select *Apply*.

The traffic types you enable will be examined according to the settings in the email filter profile.

## Configure the spam action

When spam is detected, the FortiGate unit will deal with it according to the *Spam Action* setting in the email filter profile. Note that POP3S, IMAPS and SMTPS spam filtering is available only on FortiGate units that support SSL content scanning and inspection. POP3, IMAP, POP3S and IMAPS mail can only be tagged. SMTP and SMTPS mail can be set to *Discard* or *Tagged*:

- **Discard:** When the spam action is set to *Discard*, messages detected as spam are deleted. No notification is sent to the sender or recipient.
- **Tagged:** When the spam action is set to *Tagged*, messages detected as spam are labelled and delivered normally. The text used for the label is set in the *Tag Format* field and the label is placed in the subject or the message header, as set with the *Tag Location* option.

### To configure the spam action

1. Go to *Security Profiles > Email Filter > Profile*.
2. The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
3. Select *Enable Spam Detection and Filtering*.
4. Select the types of email that you want the FortiGate unit to examine when using this email filter profile.
5. Select *Apply*.
6. The *Spam Action* row has a drop-down selection under the SMTP traffic type. Select *Discard* or *Tagged*.

No selection is available for POP3 or IMAP traffic. *Tagged* is the only applicable action for those traffic types.

By default, the tag location for any traffic set to *Tagged* is *Subject* and the tag format is *Spam*. If you want to change these settings, continue with [“Configure the tag location” on page 2053](#) and [“Configure the tag format” on page 2053](#).

7. Select *Apply*.

Select the edited email filter profile in a security policy, and the traffic controlled by the security policy will be scanned according to the settings you configured. You may select the email filter profile in more than one security policy if required.

## Configure the tag location

When the spam action is set to *Tagged*, the *Tag Location* setting determines where the tag is applied in the message.

### To configure the tag location

1. Go to *Security Profiles > Email Filter > Profile*.
2. The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
3. Select *Enable Spam Detection and Filtering*.
4. Select the types of email that you want the FortiGate unit to examine when using this email filter profile.
5. Select *Apply*.
6. The *Tag Location* row has two options for each traffic type. Note that if the spam action for SMTP traffic is set to discard, the tag location will not be available. Select the tag location:
  - *Subject*: The FortiGate unit inserts the tag at the beginning of the message subject. For example, if the message subject is “Buy stuff!” and the tag is “[spam]”, the new message subject is “[spam] Buy stuff!” if the message is detected as spam.
  - *MIME*: The FortiGate unit inserts the tag into the message header. With most mail readers and web-based mail services, the tag will not be visible. Despite this, you can still set up a rule based on the presence or absence of the tag.
7. Select *Apply*.

## Configure the tag format

When the spam action is set to *Tagged*, the *Tag Format* setting determines what text is used as the tag applied to the message.

### To configure the tag format

1. Go to *Security Profiles > Email Filter > Profile*.
2. The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
3. Select *Enable Spam Detection and Filtering*.
4. Select the types of email that you want the FortiGate unit to examine when using this email filter profile.
5. Select *Apply*.
6. The *Tag Format* row has a field for each traffic type. Note that if the spam action for SMTP traffic is set to discard, the tag format will not be available. Enter the text the FortiGate unit will use as the tag for each traffic type.
7. Select *Apply*.

## Configure FortiGuard email filters

FortiGuard email filtering techniques use FortiGuard services to detect the presence of spam among your email. A FortiGuard subscription is required to use the FortiGuard email filters. You can enable the following types of FortiGuard email filtering:

<b>FortiGuard IP address checking</b>	When you enable FortiGuard IP address checking, your FortiGate unit will submit the IP address of the client to the FortiGuard service for checking. If the IP address exists in the FortiGuard IP address black list, your FortiGate unit will treat the message as spam.
<b>FortiGuard URL checking</b>	When you enable FortiGuard URL checking, your FortiGate unit will submit all URLs appearing in the email message body to the FortiGuard service for checking. If a URL exists in the FortiGuard URL black list, your FortiGate unit will treat the message as spam.
<b>FortiGuard phishing URL detection</b>	When you enable FortiGuard phishing URL detection, your FortiGate unit will submit all URL hyperlinks appearing in the email message body to the FortiGuard service for checking. If a URL exists in the FortiGuard URL phishing list, your FortiGate unit will remove the hyperlink from the message. The URL will remain in place, but it will no longer be a selectable hyperlink.
<b>FortiGuard email checksum checking</b>	When you enable FortiGuard email checksum checking, your FortiGate unit will submit a checksum of each email message to the FortiGuard service for checking. If a checksum exists in the FortiGuard checksum black list, your FortiGate unit will treat the message as spam.
<b>FortiGuard spam submission</b>	When you enable FortiGuard email checksum checking, your FortiGate unit will append a link to the end of every message detected as spam. This link allows email users to “correct” the FortiGuard service by informing it that the message is not spam.



Carefully consider the use of the *Spam submission* option on email leaving your network. Users not familiar with the feature may click the link on spam messages because they are curious. This will reduce the accuracy of the feature.

### To enable FortiGuard email filtering

1. Go to *Security Profiles > Email Filter > Profile*.
2. The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
3. Select *Enable Spam Detection and Filtering*.
4. Select the types of email that you want the FortiGate unit to examine when using this email filter profile.
5. Select *Apply*.
6. Under the heading *FortiGuard Spam Filtering*, select one or more of the following options:
  - *IP Address Check*.
  - *URL Check*.
  - *Detect Phishing URLs in Email*.
  - *E-mail Checksum Check*.
  - *Spam Submission*.
7. Select *Apply*.

Select the edited email filter profile in a security policy, and the traffic controlled by the security policy will be scanned according to the settings you configured. You may select the email filter profile in more than one security policy if required.

## Configure local email filters

Local email filtering techniques use your own resources, whether DNS checks or IP address and email address lists that you maintain. You can enable three types of local filtering:

- Black and white list (BWL) checking (includes email addresses and IP addresses)
- HELO DNS lookup
- Return email DNS checking

### Enabling IP address and email address black/white list checking

When you enable black/white list (BWL) checking, your FortiGate unit will perform IP address BWL checking and email address BWL checking.

IP address BWL checking matches client IP addresses with IP addresses in the selected email BWL list and acts according to the action configured for the IP address in the list: allow the message, reject it, or mark it as spam.

Email address BWL checking matches sender email addresses with email addresses in the selected email BWL list acts according to the action configured for the email address in the list: allow the message or mark it as spam.

Before you can enable IP address and email address black/white list spam filtering you must create an email black/white list.

#### To create an email black/white list

1. Go to *Security Profiles > Email Filter > Email List*.
2. Select *Create New*.
3. Enter a name for the BWL list.
4. Optionally, enter a description or comments about the list.
5. Select *OK* to save the list.

When a new black/white list is created, it is empty. To perform any actions, you must add IP and email addresses to the list.

#### **To add an IP address to an email black/white list**

1. Go to *Security Profiles > Email Filter > Email List*.
2. Edit a list.
3. Select *Create New*.
4. Select *IP/Netmask*.
5. Enter the IP address or netmask in the IP/netmask field.
6. Select the action:
  - *Mark as Clear*: Messages from clients with matching IP addresses will be allowed, bypassing further email filtering.
  - *Mark as Reject*: Messages from clients with matching IP addresses will be rejected. The FortiGate unit will return a reject message to the client. *Mark as Reject* only applies to mail delivered by SMTP. If an IP address black/white list is used with POP3 or IMAP mail, addresses configured with the *Mark as Reject* action will be marked as spam.
  - *Mark as Spam*: Messages from clients with matching IP addresses will be treated as spam, subject to the action configured in the applicable email filter profile. For more information, see [“Configure the spam action” on page 2052](#).
7. By default, the address is enabled and the FortiGate unit will perform the action if the address is detected. To disable checking for the address, clear the *Enable* check box.
8. Select *OK*.

#### **To add an email address to an email black/white list**

1. Go to *Security Profiles > Email Filter > Email List*.
2. Edit a list.
3. Select *Create New*.
4. Select *Email Address*.
5. Enter the email address in the *Email Address* field.
6. If you need to enter a pattern in the *Email Address* field, select whether to use wildcards or regular expressions to specify the pattern.

Wildcard uses an asterisk (“\*”) to match any number of any character. For example, \*@example.com will match all addresses ending in @example.com.

Regular expressions use Perl regular expression syntax. See <http://perldoc.perl.org/perlretut.html> for detailed information about using Perl regular expressions.
7. Select the action:
  - *Mark as Spam*: Messages with matching reply-to email addresses will be treated as spam, subject to the action configured in the applicable email filter profile. For more information, see [“Configure the spam action” on page 2052](#).
  - *Mark as Clear*: Messages with matching reply-to addresses will be allowed, bypassing further email filtering.
8. By default, the address is enabled and the FortiGate unit will perform the action if the address is detected. To disable checking for the address, clear the *Enable* check box.
9. Select *OK* to save the address.

#### **To enable IP address black/white list checking**

1. Go to *Security Profiles > Email Filter > Profile*.



2. The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
3. Select *Enable Spam Detection and Filtering* and select *Apply*.
4. Under the heading *Local Spam Filtering*, select *BWL Check*.
5. Select the IP address black/white list to use from the drop-down list.
6. Select *Apply*.

Select the email filter profile in a security policy, and the traffic accepted by the security policy will be scanned according to the settings you configured.

## Enabling HELO DNS lookup

Whenever a client opens an SMTP session with a server, the client sends a HELO command with the client domain name. When you enable HELO DNS lookup, your FortiGate unit will take the domain the client submits as part of the HELO greeting and send it to the configured DNS. If the domain does not exist, your FortiGate unit will treat all messages the client delivers as spam.

The HELO DNS lookup is available only for SMTP traffic.

### To enable HELO DNS lookup

1. Go to *Security Profiles > Email Filter > Profile*.
2. The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
3. Select *Enable Spam Detection and Filtering* and select *Apply*.
4. Under the heading *Local Spam Filtering*, select *HELO DNS Lookup*.
5. Select *Apply*.

Select the edited email filter profile in a security policy, and the traffic controlled by the security policy will be scanned according to the settings you configured. You may select the email filter profile in more than one security policy if required.

## Enabling return email DNS checking

When you enable return email DNS checking, your FortiGate unit will take the domain in the reply-to email address and send it to the configured DNS. If the domain does not exist, your FortiGate unit will treat the message as spam.

### To enable return email DNS check

1. Go to *Security Profiles > Email Filter > Profile*.
2. The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
3. Select *Enable Spam Detection and Filtering* and select *Apply*.
4. Under the heading *Local Spam Filtering*, select *Return E-mail DNS Check*.
5. Select *Apply*.

Select the edited email filter profile in a security policy, and the traffic controlled by the security policy will be scanned according to the settings you configured. You may select the email filter profile in more than one security policy if required.

## Enabling banned word checking

When you enable banned word checking, your FortiGate unit will examine the email message for words appearing in the banned word list specified in the email filter profile. If the total score of the banned word discovered in the email message exceeds the threshold value set in the email filter profile, your FortiGate unit will treat the message as spam.

When determining the banned word score total for an email message, each banned word score is added once no matter how many times the word appears in the message. Use the command `config spamfilter bword` to add an email banned word list. Use the command `config spamfilter profile` to add a banned word list to an email filtering profile.

## How content is evaluated

Every time the banned word filter detects a pattern in an email message, it adds the pattern score to the sum of scores for the message. You set this score when you create a new pattern to block content. The score can be any number from zero to 99999. Higher scores indicate more offensive content. When the total score equals or exceeds the threshold, the email message is considered as spam and treated according to the spam action configured in the email filter profile. The score for each pattern is counted only once, even if that pattern appears many times in the email message. The default score for banned word patterns is 10 and the default threshold is 10. This means that by default, an email message is blocked by a single match.

A pattern can be part of a word, a whole word, or a phrase. Multiple words entered as a pattern are treated as a phrase. The phrase must appear as entered to match. You can also use wildcards or regular expressions to have a pattern match multiple words or phrases.

For example, the FortiGate unit scans an email message that contains only this sentence: "The score for each word or phrase is counted only once, even if that word or phrase appears many times in the email message."

Banned word pattern	Pattern type	Assigned score	Score added to the sum for the entire page	Comment
word	Wildcard	20	20	The pattern appears twice but multiple occurrences are only counted once.
word phrase	Wildcard	20	0	Although each word in the phrase appears in the message, the words do not appear together as they do in the pattern. There are no matches.
word*phrase	Wildcard	20	20	The wildcard represents any number of any character. A match occurs as long as "word" appears before "phrase" regardless of what is in between them.
mail*age	Wildcard	20	20	Since the wildcard character can represent any characters, this pattern is a match because "email message" appears in the message.

In this example, the message is treated as spam if the banned word threshold is set to 60 or less.

### Adding words to a banned word list

Each banned word list contains a number of words, each having a `score`, and specifying where the FortiGate unit will search for the word (in the message subject, message body, or `all` which means both)

When the FortiGate unit accepts an email message containing one or more words in the banned word list specified in the active email filter profile, it totals the scores of the banned words in the email message. If the total is higher than the threshold set in the email filter profile, the email message will be detected as spam. If the total score is lower than the threshold, the message will be allowed to pass as normal.

The score of a banned word present in the message will be counted toward the score total only once, regardless of how many times the word appears in the message.

When you enter a word, set the `Pattern-type` to wildcards or regular expressions.

*Wildcard* uses an asterisk (“\*”) to match any number of any character. For example, `re*` will match all words starting with “re”.

*Regular expression* uses Perl regular expression syntax. See <http://perldoc.perl.org/perlretut.html> for detailed information about using Perl regular expressions.

## Email filter examples

### Configuring simple antispam protection

Small offices, whether they are small companies, home offices, or satellite offices, often have very simple needs. This example details how to enable antispam protection on a FortiGate unit located in a satellite office.

#### Creating an email filter profile

Most email filter settings are configured in an email filter profile. Email filter profiles are selected in firewall policies. This way, you can create multiple email filter profiles, and tailor them to the traffic controlled by the security policy in which they are selected. In this example, you will create one email filter profile.

##### To create an email filter profile – web-based manager

1. Go to *Security Profiles > Email Filter > Profile*.
2. Select the *Create New* icon in the Edit Email Filter Profile window title.
3. In the *Name* field, enter `basic_emailfilter`.
4. Select *Enable Spam Detection and Filtering*.
5. Ensure that *IMAP*, *POP3*, and *SMTP* are selected in the header row.

These header row selections enable or disable examination of each email traffic type. When disabled, the email traffic of that type is ignored by the FortiGate unit and no email filtering options are available.

6. Under *FortiGuard Spam Filtering*, enable *IP Address Check*.
7. Under *FortiGuard Spam Filtering*, enable *URL Check*.
8. Under *FortiGuard Spam Filtering*, enable *E-mail Checksum Check*.
9. Select *OK* to save the email filter profile.

### To create an email filter profile – CLI

```
config spamfilter profile
 edit basic_emailfilter
 set options spamfsip spamfsurl spamfschksum
 end
```

### Selecting the email filter profile in a security policy

An email filter profile directs the FortiGate unit to scan network traffic only when it is selected in a security policy. When an email filter profile is selected in a security policy, its settings are applied to all the traffic the security policy handles.

### To select the email filter profile in a security policy – web-based manager

1. Go to *Policy > Policy > Policy*.
2. Create a new or edit a policy.
3. Turn on email filtering.
4. Select the `basic_emailfilter` profile from the list.
5. Select *OK* to save the security policy.

### To select the email filter profile in a security policy – CLI

```
config firewall policy
 edit 1
 set utm-status enable
 set profile-protocol-options default
 set spamfilter-profile basic_emailfilter
 end
```

IMAP, POP3, and SMTP email traffic handled by the security policy you modified will be scanned for spam. Spam messages have the text “Spam” added to their subject lines. A small office may have only one security policy configured. If you have multiple policies, consider enabling spam scanning for all of them.

## Blocking email from a user

Employees of the Example.com corporation have been receiving unwanted email messages from a former client at a company called example.net. The client’s email address is client@example.net. All ties between the company and the client have been severed, but the messages continue. The FortiGate unit can be configured to prevent these messages from being delivered.

### To create the email address list

1. Go to *Security Profiles > Email Filter > Email List*.
2. Select *Create New*.
3. Enter a name for the new email address list.
4. Optionally, enter a descriptive comment for the email address list.
5. Select *OK* to create the list.
6. Select *Create New* to add a new entry to the email address list.
7. Select *Email Address*.
8. Enter `client@example.net` in the *E-mail Address* field.
  - If you wanted to prevent everyone’s email from the client’s company from getting through you could have used `*@example.net` instead.

9. Leave *Pattern Type* set to the default, *Wildcard*.
10. Leave *Action* as *Mark as Spam* to have the FortiGate unit mark all messages from example.net as spam.

Now that the email address list is created, you must enable the email filter in the email filter profile.

#### **To enable Email Filter**

1. Go to *Security Profiles > Email Filter > Profile*.
2. Select the email filter profile that is used by the firewall policies handling email traffic from the email filter profile drop down list.
3. In the row *Tag Location*, select *Subject* for all three mail protocols.
4. In the row *Tag Format*, enter *SPAM:* in all three fields.
5. Select *Enable Spam Detection and Filtering*.
6. Ensure that the check boxes labeled *IMAP*, *POP3*, and *SMTP* in the header row are selected.
7. Under *Local Spam Filtering*, enable *BWL Check* and select the email address list you created in the previous procedure from the drop down list.
8. Select *OK*.

When this email filter profile is selected in a security policy, the FortiGate unit will add "SPAM:" to the subject of any email message from an address ending with @example.net for all email traffic handled by the security policy. Recipients can ignore the message or they can configure their email clients to automatically delete messages with "SPAM:" in the subject.

# Intrusion protection

The FortiGate Intrusion Protection system combines signature detection and prevention with low latency and excellent reliability. With intrusion protection, you can create multiple IPS sensors, each containing a complete configuration based on signatures. Then, you can apply any IPS sensor to any security policy.

This section describes how to configure the FortiGate Intrusion Protection settings.

If you enable virtual domains (VDOMs) on the FortiGate unit, intrusion protection is configured separately for each virtual domain.

The following topics are included:

- [IPS concepts](#)
- [Enable IPS scanning](#)
- [Configure IPS options](#)
- [Enable IPS packet logging](#)
- [IPS examples](#)

## IPS concepts

The FortiGate intrusion protection system protects your network from outside attacks. Your FortiGate unit has two techniques to deal with these attacks: anomaly- and signature-based defense.

### Anomaly-based defense

Anomaly-based defense is used when network traffic itself is used as a weapon. A host can be flooded with far more traffic than it can handle, making the host inaccessible. The most common example is the denial of service (DoS) attack, in which an attacker directs a large number of computers to attempt normal access of the target system. If enough access attempts are made, the target is overwhelmed and unable to service genuine users. The attacker does not gain access to the target system, but it is not accessible to anyone else.

The FortiGate DoS feature will block traffic above a certain threshold from the attacker and allow connections from other legitimate users. The DoS policy configuration information can be found in the Firewall Handbook.

### Signature-based defense

Signature-based defense is used against known attacks or vulnerability exploits. These often involve an attacker attempting to gain access to your network. The attacker must communicate with the host in an attempt to gain access and this communication will include particular commands or sequences of commands and variables. The IPS signatures include these command sequences, allowing the FortiGate unit to detect and stop the attack.

### Signatures

IPS signatures are the basis of signature-based intrusion protection. Every attack can be reduced to a particular string of commands or a sequence of commands and variables.

Signatures include this information so your FortiGate unit knows what to look for in network traffic.

Signatures also include characteristics about the attack they describe. These characteristics include the network protocol in which the attack will appear, the vulnerable operating system, and the vulnerable application.

To view the complete list of signatures, go to *Security Profiles > Intrusion Protection > IPS Signatures*. This will include the predefined signatures and any custom signatures that you may have created.

## Protocol decoders

Before examining network traffic for attacks, the IPS engine uses protocol decoders to identify each protocol appearing in the traffic. Attacks are protocol-specific, so your FortiGate unit conserves resources by looking for attacks only in the protocols used to transmit them. For example, the FortiGate unit will only examine HTTP traffic for the presence of a signature describing an HTTP attack.

## IPS engine

Once the protocol decoders separate the network traffic by protocol, the IPS engine examines the network traffic for the attack signatures.

## IPS sensors

The IPS engine does not examine network traffic for all signatures, however. You must first create an IPS sensor and specify which signatures are included. Add signatures to sensors individually using signature entries, or in groups using IPS filters.

To view the IPS sensors, go to *Security Profiles > Intrusion Protection > IPS Sensor*.

## IPS filters

IPS sensors contain one or more IPS filters. A filter is a collection of signature attributes that you specify. The signatures that have all of the attributes specified in a filter are included in the IPS filter.

For example, if your FortiGate unit protects a Linux server running the Apache web server software, you could create a new filter to protect it. By setting *OS* to *Linux*, and *Application* to *Apache*, the filter will include only the signatures that apply to both Linux and Apache. If you wanted to scan for all the Linux signatures and all the Apache signatures, you would create two filters, one for each.

To view the filters in an IPS sensor, go to *Security Profiles > Intrusion Protection > IPS Sensor*, select the IPS sensor containing the filters you want to view, and choose *Edit*.

## Custom/predefined signature entries

Signature entries allow you to add an individual custom or predefined IPS signature. If you need only one signature, adding a signature entry to an IPS sensor is the easiest way. Signature entries are also the only way to include custom signatures in an IPS sensor.

Another use for signature entries are to change the settings of individual signatures that are already included in a filter within the same IPS sensor. Add a signature entry with the required settings above the filter, and the signature entry will take priority.

## Policies

To use an IPS sensor, you must select it in a security policy or an interface policy. An IPS sensor that it not selected in a policy will have no effect on network traffic.

IPS is most often configured as part of a security policy. Unless stated otherwise, discussion of IPS sensor use will be in regards to firewall policies in this document.

## Enable IPS scanning

Enabling IPS scanning involves two separate parts of the FortiGate unit:

- The security policy allows certain network traffic based on the sender, receiver, interface, traffic type, and time of day. Firewall policies can also be used to deny traffic, but those policies do not apply to IPS scanning.
- The IPS sensor contains filters, signature entries, or both. These specify which signatures are included in the IPS sensor.

When IPS is enabled, an IPS sensor is selected in a security policy, and all network traffic matching the policy will be checked for the signatures in the IPS sensor.

### General configuration steps

For best results in configuring IPS scanning, follow the procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

1. Create an IPS sensor.
2. Add filters and/or predefined signatures and custom signatures to the sensor. The filters and signatures specify which signatures the IPS engine will look for in the network traffic.
3. Select a security policy or create a new one.
4. In the security policy, turn on *IPS*, and choose the IPS sensor from the list.

All the network traffic controlled by this security policy will be processed according to the settings in the policy. These settings include the IPS sensor you specify in the policy.

### Creating an IPS sensor

You need to create an IPS sensor and save it before configuring it with filters and entries.

#### To create a new IPS sensor

1. Go to *Security Profiles > Intrusion Protection > IPS Sensors*.
2. Select the *Create New* icon in the top of the Edit IPS Sensor window.
3. Enter the name of the new IPS sensor.
4. Optionally, you may also enter a comment. The comment will appear in the IPS sensor list and serves to remind you of the details of the sensor.
5. Select *OK*.

The IPS sensor is created and the sensor configuration window appears. A newly created sensor is empty and contains no filters or signatures. You need to add one or more filters or signatures before the sensor can take effect.

### Creating an IPS filter

While individual signatures can be added to a sensor, a filter allows you to add multiple signatures to a sensor by specifying the characteristics of the signatures to be added.

#### To create a new IPS filter

1. Go to *Security Profiles > Intrusion Protection > IPS Sensors*.



2. Select the IPS sensor to which you want to add the filter using the drop-down list in the top row of the Edit IPS Sensor window.
3. Select the *Create New* icon
4. For *Sensor Type* chose *Filter Based*.
5. Configure the filter that you require. Signatures matching all of the characteristics you specify in the filter will be included in the filter. Select *Specify* and choose the filter option that have the appropriate parameters.

### **Basic**

#### *Severity*

Refers to the level of threat possessed by the attack.

The options include:

- critical
- high
- medium
- low
- info

#### *Target*

Refers to the type of device targeted by the attack.

The options include:

- client
- server

#### *OS*

Refers to the Operating System affected by the attack.

The options include:

BSD	Linux	MacOS
Other	Solaris	Windows

### **Advanced**

#### *Application*

Refers to the vendor or or type of application affected by the attack.

The options include:.

Adobe	Apache	Apple
CGI_app	Cisco	HP
IBM	IE	IIS
Mozilla	MS_Office	Novel
Oracle	PHP_app	Sun

This list can be expanded to include more options by selecting the [show more...] link. The additional options include:

ASP_app	CA	DB2
IM	Ipswitch	MailEnable
MediaPlayer	MS_Exchange	MSSQL
MySQL	Netscape	P2P
PostgreSQL	Real	Samba
SAP	SCADA	Sendmail
Veritas	Winamp	Other

*Protocol*

Refers to the protocol that is the vector for the attack.

The options include:

DNS	FTP	HTTP
ICMP	IMAP	LDAP
POP3	SCCP	SIP
SMTP	SNMP	SSH
SSL	TCP	UDP

This list can be expanded to include more options by selecting the [show more...] link. The additional options include:

BO	DCERPC	DHCP
DNP3	H323	IM
MSSQL	NBSS	NNTP
P2P	RADIUS	RDT
RPC	TRCP	RTP
RTSP	TELNET	TFN
Other		

6. Choose an action for when a signature is triggered.

Action	Description
<b>Signature Default</b>	All predefined signatures have an <i>Action</i> attribute that is set to Pass or Drop. This means that if a signature included in the filter has an <i>Action</i> setting of Pass, traffic matching the signature will be detected and then allowed to continue to its destination. Select <i>Accept signature defaults</i> use the default action for each included signature.  Note: to see what the default for a signature is, go to the <i>IPS Signatures</i> page and enable the column <i>Action</i> , then find the row with the signature name in it.
<b>Monitor All</b>	Select <i>Monitor all</i> to pass all traffic matching the signatures included in the filter, regardless of their default <i>Action</i> setting.
<b>Block All</b>	Select <i>Block all</i> to drop traffic matching any the signatures included in the filter.
<b>Reset</b>	Select <i>Reset</i> to reset the session whenever the signature is triggered. In the CLI this action is referred to as <i>Reject</i> .
<b>Quarantine</b>	Has 2 fields the need to be configured:  1. Method: <ul style="list-style-type: none"><li>• Attacker's IP Address - Traffic from the Attacker's IP address is refused until the expiration time from the trigger is reached.</li><li>• Attacker and Victim Address - All traffic from the Attacker's address to the Victim's address will be blocked.</li><li>• Attack's incoming interface - the interface that experienced the attack will refuse further traffic.</li></ul> 2. Expires (time frame that the quarantine will be in effect): <ul style="list-style-type: none"><li>• 5 Minute(s)</li><li>• 30 Minutes(s)</li><li>• 1 Hour(s)</li><li>• 1 Day(s)</li><li>• Week(s)</li><li>• Month(s)</li><li>• Year(s)</li></ul>
<b>Packet Logging</b>	Select to enable packet logging for the filter.  When you enable packet logging on a filter, the unit saves a copy of the packets that match any signatures included in the filter. The packets can be analyzed later.  For more information about packet filtering, see <a href="#">"Monitoring Security Profiles activity" on page 2178</a>

7 Select *OK*.

The filter is created and added to the filter list.

## Updating predefined IPS signatures

The FortiGuard Service periodically updates the pre-defined signatures and adds new signatures to counter emerging threats as they appear.

Because the signatures included in filters are defined by specifying signature attributes, new signatures matching existing filter specifications will automatically be included in those filters. For example, if you have a filter that includes all signatures for the Windows operating system, your filter will automatically incorporate new Windows signatures as they are added.

## Viewing and searching predefined IPS signatures

Go to *Security Profiles > Intrusion Protection > IPS Signatures* to view the list of existing IPS signatures. You may find signatures by paging manually through the list, apply filters, or by using the search field.

### Searching manually

Signatures are displayed in a paged list, with 50 signatures per page. The bottom of the screen shows the current page and the total number of pages. You can enter a page number and press enter, to skip directly to that page. Previous Page and Next Page buttons move you through the list, one page at a time. The First Page and Last Page button take you to the beginning or end of the list.

### Applying filters

You can enter criteria for one or more columns, and only the signatures matching all the conditions you specify will be listed.

#### To apply filters

1. Go to *Security Profiles > Intrusion Protection > IPS Signatures*.
2. Select column by which to filter.
3. Select the funnel/filter icon and enter the value or values to filter by.
4. Use additional columns as needed to refine search.

The available options vary by column. For example, Enable allows you to choose between two options, while OS has multiple options, and you may select multiple items together. Filtering by name allows you to enter a text string and all signature names containing the string will be displayed.

## IPS processing in an HA cluster

IPS processing in an HA cluster is no different than with a single FortiGate unit, from the point of view of the network user. The difference appears when a secondary unit takes over from the primary, and what happens depends on the HA mode.

### Active-passive

In an active-passive HA cluster, the primary unit processes all traffic just as it would in a stand-alone configuration. Should the primary unit fail, a secondary unit will assume the role of the primary unit and begin to process network traffic. By default, the state of active communication sessions are not shared with secondary units and will not survive the fail-over condition. Once the sessions are reestablished however, traffic processing will continue as normal.

If your network requires that active sessions are taken over by the new primary unit, select *Enable Session Pick-up* in your HA configuration. Because session information must be sent to all subordinate units on a regular basis, session pick-up is a resource-intensive feature and is not enabled by default.

## Active-active

The fail-over process in an active-active cluster is similar to an active-passive cluster. When the primary unit fails, a secondary unit takes over and traffic processing continues. The load-balancing schedule used to distribute sessions to the cluster members is used by the new primary unit to redistribute sessions among the remaining subordinate units. If session pick-up is not enabled, the sessions active on the failed primary are lost, and the sessions redistributed among the secondary units may also be lost. If session pick-up is enabled, all sessions are handled according to their last-known state.

For more information about HA options and settings, see [“High Availability for FortiOS 5.0” on page 1116](#).

## Configure IPS options

There are a number of CLI commands that influence how IPS functions.

### Hardware Acceleration

In order to provide control over the hardware’s processing of IPS there are commands to configure and control the hardware acceleration of IPS. There are two settings that can be chosen, one for the network processor and one for the content processor.

Network processor acceleration can be disabled or set to enable basic acceleration.

Content processor acceleration can be disabled or set to either basic or advanced acceleration.

These Settings are only found in the CLI:

```
config ips global
 set np-accel-mode {none | basic}
 set cp-accel-mode {none | basic | advanced}
end
```

### Extended IPS Database.

Some models have access to an extended IPS Database. The extended database may affect the performance of the FortiGate unit so depending on the model of the FortiGate unit the extended database package may not be enabled by default. For example, the D-series Desktop model have this option disabled by default.

This feature can only be enabled through the CLI.

```
config ips global
 set database extended
end
```

## Configuring the IPS engine algorithm

The IPS engine is able to search for signature matches in two ways. One method is faster but uses more memory, the other uses less memory but is slower. Use the `algorithm` CLI command to select one method:

```
config ips global
 set algorithm {super | high | low | engine-pick}
end
```

Specify `high` to use the faster more memory intensive method or `low` for the slower memory efficient method. The setting `super` improves the performance for FortiGate units with more than 4GB of memory. The default setting is `engine-pick`, which allows the IPS engine to choose the best method on the fly.

## Configuring the IPS engine-count

FortiGate units with multiple processors can run more than one IPS engine concurrently. The `engine-count` CLI command allows you to specify how many IPS engines are used at the same time:

```
config ips global
 set engine-count <int>
end
```

The recommended and default setting is 0, which allows the FortiGate unit to determine the optimum number of IPS engines.

## Configuring fail-open

If the IPS engine fails for any reason, it will fail open by default. This applies for inspection of all the protocols inspected by FortiOS IPS protocol decoders, including but not limited to HTTP, HTTPS, FTP, SMTP, POP3, IMAP, etc. This means that traffic continues to flow without IPS scanning. If IPS protection is more important to your network than the uninterrupted flow of network traffic, you can disable this behavior using the `fail-open` CLI command:

```
config ips global
 set fail-open {enable | disable}
end
```

The default setting is `enable`.

## Configuring the session count accuracy

The IPS engine can keep track of the number of open session in two ways. An accurate count uses more resources than a less accurate heuristic count.

```
config ips global
 set session-limit-mode {accurate | heuristic}
end
```

The default is `heuristic`.

## Configuring the IPS buffer size

Set the size of the IPS buffer.

```
config ips global
 set socket-size <int>
end
```

The acceptable range is from 1 to 64 megabytes. The default size varies by model.

## Configuring protocol decoders

The FortiGate Intrusion Protection system uses protocol decoders to identify the abnormal traffic patterns that do not meet the protocol requirements and standards. For example, the HTTP decoder monitors traffic to identify any HTTP packets that do not meet the HTTP protocol standards.

To change the ports a decoder examines, you must use the CLI. In this example, the ports examined by the DNS decoder are changed from the default 53 to 100, 200, and 300.

```
config ips decoder dns_decoder
 set port_list "100,200,300"
end
```

You cannot assign specific ports to decoders that are set to *auto* by default. These decoders can detect their traffic on any port. Specifying individual ports is not necessary.

## Configuring security processing modules

FortiGate Security Processing Modules, such as the CE4, XE2, and FE8, can increase overall system performance by accelerating some security and networking processing on the interfaces they provide. They also allow the FortiGate unit to offload the processing to the security module, thereby freeing up its own processor for other tasks. The security module performs its own IPS and firewall processing, but you can configure it to favor IPS in hostile high-traffic environments.

If you have a security processing module, use the following CLI commands to configure it to devote more resources to IPS than firewall. This example shows the CLI commands required to configure a security module in slot 1 for increased IPS performance.

```
config system amc-slot
 edit sw1
 set optimization-mode fw-ips
 set ips-weight balanced
 set ips-p2p disable
 set ips-fail-open enable
 set fp-disable none
 set ipsec-inb-optimization enable
 set syn-proxy-client-timer 3
 set syn-proxy-server-timer 3
 end
```

In addition to offloading IPS processing, security processing modules provide a hardware accelerated SYN proxy to defend against SYN flood denial of service attacks. When using a security module, configure your DoS anomaly check for `tcp_syn_flood` with the *Proxy* action. The *Proxy* action activates the hardware accelerated SYN proxy.

## IPS signature rate count threshold

The IPS signature threshold can allow configuring a signature so that it will not be triggered until a rate count threshold is met. This provides a more controlled recording of attack activity. For example, if multiple login attempts produce a failed result over a short period of time then an alert would be sent and perhaps traffic blocked. This would be a more rational response than sending an alert every time a login failed.

The syntax for this configuration is as follows:

```
config ips sensor
 edit default
 config entries
 edit <Filter ID number>
 set rule <*id>
 set rate-count <integer between 1 - 65535>
 set rate-duration <integer between 1 - 65535>
```

The value of the rate-duration is an integer for the time in seconds.

```
set rate-mode <continuous | periodical>
```

The rate-mode refers to how the count threshold is met.

If the setting is “continuous”, and the action is set to block, as soon as the rate-count is reached the action is engaged. For example, if the count is 10, as soon as the signature is triggered 10 times the traffic would be blocked.

If the setting is “periodical”, the FortiGate allows up to the value of the rate-count incidents where the signature is triggered during the rate-duration. For example, if the rate count is 100 and the duration is 60, the signature would need to be triggered 100 times in 60 seconds for the action to be engaged.

```
set rate-track <dest-ip | dhcp-client-mac | dns-domain |
 none | src-ip>
```

This setting allow the tracking of one of the protocol fields within the packet.

```
end
end
```

## Enable IPS packet logging

Packet logging saves the network packets containing the traffic matching an IPS signature to the attack log. The FortiGate unit will save the logged packets to wherever the logs are configured to be stored, whether memory, internal hard drive, a FortiAnalyzer unit, or the FortiGuard Analysis and Management Service.

You can enable packet logging in the filters. Use caution in enabling packet logging in a filter. Filters configured with few restrictions can contain thousands of signatures, potentially resulting in a flood of saved packets. This would take up a great deal of space, require time to sort through, and consume considerable system resources to process. Packet logging is designed as a focused diagnostic tool and is best used with a narrow scope.



Although logging to multiple FortiAnalyzer units is supported, packet logs are not sent to the secondary and tertiary FortiAnalyzer units. Only the primary unit receives packet logs.

---



### To enable packet logging for a filter

1. Create a filter in an IPS sensor. For more information, see [“Creating an IPS filter” on page 2064](#).
2. Before saving the filter, select *Enable All for Packet Logging*.
3. Select the IPS sensor in the security policy that allows the network traffic the FortiGate unit will examine for the signature.

For information on viewing and saving logged packets, see [“Monitoring Security Profiles activity” on page 2178](#).

## IPS examples

### Configuring basic IPS protection

Small offices, whether they are small companies, home offices, or satellite offices, often have very simple needs. This example details how to enable IPS protection on a FortiGate unit located in a satellite office. The satellite office contains only Windows clients.

#### Creating an IPS sensor

Most IPS settings are configured in an IPS sensor. IPS sensors are selected in firewall policies. This way, you can create multiple IPS sensors, and tailor them to the traffic controlled by the security policy in which they are selected. In this example, you will create one IPS sensor.

#### To create an IPS sensor— web-based manager

1. Go to *Security Profiles > Intrusion Protection > IPS Sensors*.
2. Select the *Create New* icon in the top of the Edit IPS Sensor window.
3. In the *Name* field, enter `basic_ips`.
4. In the *Comments* field, enter `IPS protection for Windows clients`.
5. Select *OK*.
6. Select the *Create New* drop-down to add a new component to the sensor and for the *Sensor Type* choose *Filter Based*.
7. In the Filter Options choose the following:
  - a. For *Severity*: select all of the options
  - b. For *Target*: select *Client* only.
  - c. For *OS*: select *Windows* only.
8. For the *Action* leave as the default.
9. Select *OK* to save the filter.
10. Select *OK* to save the IPS sensor.

### To create an IPS sensor — CLI

```
config ips sensor
 edit basic_ips
 set comment "IPS protection for Windows clients"
 config entries
 edit 1
 set location client
 set os windows
 end
 end
 end
```

### Selecting the IPS sensor in a security policy

An IPS sensor directs the FortiGate unit to scan network traffic only when it is selected in a security policy. When an IPS sensor is selected in a security policy, its settings are applied to all the traffic the security policy handles.

### To select the IPS sensor in a security policy — web-based manager

1. Go to *Policy > Policy > Policy*.
2. Select a policy.
3. Select the *Edit* icon.
4. Enable the *IPS* option.
5. Select the `basic_ips` profile from the list.
6. Select *OK* to save the security policy.

### To select the IPS sensor in a security policy — CLI

```
config firewall policy
 edit 1
 set utm-status enable
 set ips-sensor basic_ips
 end
```


All traffic handled by the security policy you modified will be scanned for attacks against Windows clients. A small office may have only one security policy configured. If you have multiple policies, consider enabling IPS scanning for all of them.

## Using IPS to protect your web server

Many companies have web servers and they must be protected from attack. Since web servers must be accessible, protection is not as simple as blocking access. IPS is one tool your FortiGate unit has to allow you to protect your network.

In this example, we will configure IPS to protect a web server. As shown in [Figure 309 on page 2075](#), a FortiGate unit protects a web server and an internal network. The internal network will have its own policies and configuration but we will concentrate on the web server in this example.

**Figure 309:**A simple network configuration



The FortiGate unit is configured with:

- a virtual IP to give the web server a unique address accessible from the Internet.
- a security policy to allow access to the web server from the Internet using the virtual IP.

To protect the web server using intrusion protection, you need to create an IPS sensor, populate it with filters, then enable IPS scanning in the security policy.

#### **To create an IPS sensor**

1. Go to *Security Profiles > Intrusion Protection > IPS Sensors* and select *Create New*.
2. Enter `web_server` as the name of the new IPS sensor.
3. Select *OK*.

The new IPS sensor is created but it has no filters, and therefore no signatures are included.

The web server operating system is Linux, so you need to create a filter for all Linux server signatures.

#### **To create the Linux server filter**

1. Go to *Security Profiles > Intrusion Protection > IPS Sensors* and select the `web_server` IPS sensor and select the *Edit* icon.
2. Select *Add Filter*.
3. Enter `Linux Server` as the name of the new filter.
4. For *Target*, select *Specify* and choose *server*.
5. In the Filter Options choose the following:
  - a. For *Severity*: select all of the options
  - b. For *Target*: select *server* only.
  - c. For *OS*: select *Linux* only.
6. Select *OK*.

The filter is saved and the IPS sensor page reappears. In the filter list, find the *Linux Server* filter and look at the value in the *Count* column. This shows how many signatures match the current filter settings. You can select the *View Rules* icon to see a listing of the included signatures.

#### **To edit the security policy**

1. Go to *Policy > Policy > Policy*, select security policy that allows access to the web server, and select the *Edit* icon.
2. Enable IPS option and choose the `web_server` IPS sensor from the list.
3. Select *OK*.

Since IPS is enabled and the `web_server` IPS sensor is specified in the security policy controlling the web server traffic, the IPS sensor examines the web server traffic for matches to the signatures it contains.

## **Create and test a packet logging IPS sensor**

In this example, you create a new IPS sensor and include a filter that detects the EICAR test file and saves a packet log when it is found. This is an ideal first experience with packet logging because the EICAR test file can cause no harm, and it is freely available for testing purposes.

#### **Create an IPS sensor**

1. Go to *Security Profiles > Intrusion Protection > IPS Sensors*.
2. Select *Create New*.
3. Name the new IPS sensor `EICAR_test`.
4. Select *OK*.

#### **Create an entry**

1. Select the *Create New* drop down menu and for *Sensor Type* choose *Specify Signatures*.
2. Rather than search through the signature list, use the name filter by selecting the search icon over the header of the *Signature* column.
3. Enter `EICAR` in the Search field.
4. Highlight the `Eicar.Virus.Test.File` signature by clicking on it.
5. Select *Block All* as the *Action*.
6. Select *Enable, Packet Logging*.
7. Select *OK* to save the IPS sensor.

You are returned to the IPS sensor list. The `EICAR_test` sensor appears in the list.

#### **Add the IPS sensor to the security policy allowing Internet access**

1. Go to *Policy > Policy > Policy*.
2. Select the security policy that allows you to access the Internet.
3. Select the *Edit* icon.
4. Enable *Log Allowed Traffic*.
5. Enable the *IPS* option.
6. Choose `EICAR_test` from the available IPS sensors.
7. Select *OK*.

With the IPS sensor configured and selected in the security policy, the FortiGate unit blocks any attempt to download the EICAR test file.

### Test the IPS sensor

1. Using your web browser, go to [http://www.eicar.org/anti\\_virus\\_test\\_file.htm](http://www.eicar.org/anti_virus_test_file.htm).
2. Scroll to the bottom of the page and select *ecar.com* from the row labeled as using the standard HTTP protocol.
3. The browser attempts to download the requested file and,
  - If the file is successfully downloaded, the custom signature configuration failed at some point. Check the custom signature, the IPS sensor, and the firewall profile.
  - If the download is blocked with a high security alert message explaining that you're not permitted to download the file, the EICAR test file was blocked by the FortiGate unit antivirus scanner before the IPS sensor could examine it. Disable antivirus scanning and try to download the EICAR test file again.
  - If no file is downloaded and the browser eventually times out, the custom signature successfully detected the EICAR test file and blocked the download.

### Viewing the packet log

1. Go to *Log&Report > Log & Archive Access > Security Log*.
2. Locate the log entry that recorded the blocking of the EICAR test file block. The Message field data will be `tools: EICAR.AV.Test.File.Download`.
3. Select the *View Packet Log* icon in the *Packet Log* column.
4. The packet log viewer is displayed.

## Configuring a Fortinet Security Processing module

The Example Corporation has a web site that is the target of SYN floods. While they investigate the source of the attacks, it's very important that the web site remain accessible. To enhance the ability of the company's FortiGate-620B to deal with SYN floods, the administrator will install an ASM-CE4 Fortinet Security Processing module and have all external access to the web server come through it.

The security processing modules not only accelerate and offload network traffic from the FortiGate unit's processor, but they also accelerate and offload security and content scanning. The ability of the security module to accelerate IPS scanning and DoS protection greatly enhances the defense capabilities of the FortiGate-620B.

### Assumptions

As shown in other examples and network diagrams throughout this document, the Example Corporation has a pair of FortiGate-620B units in an HA cluster. To simplify this example, the cluster is replaced with a single FortiGate-620B.

An ASM-CE4 is installed in the FortiGate-620B.

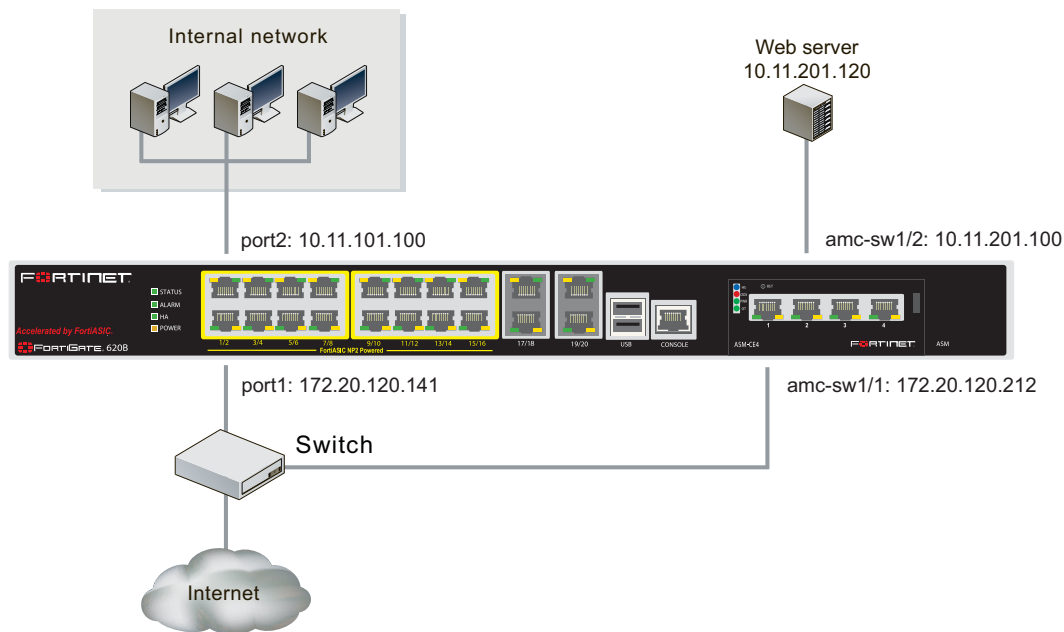
The network is configured as shown in [Figure 310](#).

### Network configuration

The Example Corporation network needs minimal changes to incorporate the ASM-CE4. Interface `amc-sw1/1` of the ASM-CE4 is connected to the Internet and interface `amc-sw1/1` is connected to the web server.

Since the main office network is connected to `port2` and the Internet is connected to `port1`, a switch is installed to allow both `port1` and `amc-sw1/1` to be connected to the Internet.

**Figure 310:**The FortiGate-620B network configuration



The switch used to connect port1 and amc-sw1/1 to the Internet must be able to handle any SYN flood, all of the legitimate traffic to the web site, and all of the traffic to and from the Example Corporation internal network. If the switch can not handle the bandwidth, or if the connection to the service provider can not provide the required bandwidth, traffic will be lost.

### Security module configuration

The Fortinet security modules come configured to give equal priority to content inspection and firewall processing. The Example Corporation is using a ASM-CE4 module to defend its web server against SYN flood attacks so firewall processing is a secondary consideration.

Use these CLI commands to configure the security module in ASM slot 1 to devote more resources to content processing, including DoS and IPS, than to firewall processing.

```
config system amc-slot
 edit sw1
 set optimization-mode fw-ips
 set ips-weight balanced
 set ips-p2p disable
 set ips-fail-open enable
 set fp-disable none
 set ipsec-inb-optimization enable
 set syn-proxy-client-timer 3
 set syn-proxy-server-timer 3
 end
```

These settings do not disable firewall processing. Rather, when the security module nears its processing capacity, it will chose to service content inspection over firewall processing.

### IPS Sensor

You can group signatures into IPS sensors for easy selection when applying to firewall policies. You can define signatures for specific types of traffic in separate IPS sensors, and then select those sensors in profiles designed to handle that type of traffic. For example, you can specify all

of the web-server related signatures in an IPS sensor, and that sensor can then be applied to a firewall policy that controls all of the traffic to and from a web server protected by the unit.

The FortiGuard Service periodically updates the pre-defined signatures, with signatures added to counter new threats. Since the signatures included in filters are defined by specifying signature attributes, new signatures matching existing filter specifications will automatically be included in those filters. For example, if you have a filter that includes all signatures for the Windows operating system, your filter will automatically incorporate new Windows signatures as they are added.

Each IPS sensor consists of two parts: filters and overrides. Overrides are always checked before filters.

Each filter consists of a number of signatures attributes. All of the signatures with those attributes, and only those attributes, are checked against traffic when the filter is run. If multiple filters are defined in an IPS Sensor, they are checked against the traffic one at a time, from top to bottom. If a match is found, the unit takes the appropriate action and stops further checking.

A signature override can modify the behavior of a signature specified in a filter. A signature override can also add a signature not specified in the sensor's filters. Custom signatures are included in an IPS sensor using overrides.

The signatures in the overrides are first compared to network traffic. If the IPS sensor does not find any matches, it then compares the signatures in each filter to network traffic, one filter at a time, from top to bottom. If no signature matches are found, the IPS sensor allows the network traffic.

The signatures included in the filter are only those matching every attribute specified. When created, a new filter has every attribute set to *all* which causes every signature to be included in the filter. If the severity is changed to high, and the target is changed to server, the filter includes only signatures checking for high priority attacks targeted at servers.

# Custom Application & IPS Signatures

## Creating a custom IPS signature

The FortiGate predefined signatures cover common attacks. If you use an unusual or specialized application or an uncommon platform, add custom signatures based on the security alerts released by the application and platform vendors.

You can add or edit custom signatures using the web-based manager or the CLI.

### To create a custom signature

1. Go to *Security Profiles > Intrusion Protection > IPS Signatures*.
2. Select *Create New* to add a new custom signature.
3. Enter a *Name* for the custom signature.
4. Enter the *Signature*. For information about completing this field, see [“All custom signatures follow a particular syntax. Each begins with a header and is followed by one or more keywords. The syntax and keywords are detailed in the next two topics.”](#).
5. Select *OK*.

All custom signatures follow a particular syntax. Each begins with a header and is followed by one or more keywords. The syntax and keywords are detailed in the next two topics.

## Custom signature syntax

A custom signature definition is limited to a maximum length of 512 characters. A definition can be a single line or span multiple lines connected by a backslash (\) at the end of each line.

A custom signature definition begins with a header, followed by a set of keyword/value pairs enclosed by parenthesis [( )]. The keyword and value pairs are separated by a semi colon (;) and consist of a keyword and a value separated by a space. The basic format of a definition is `HEADER (KEYWORD VALUE;)`

You can use as many keyword/value pairs as required within the 512 character limit. To configure a custom signature, go to *Security Profiles > Intrusion Protection > IPS Signatures*, select *Create New* and enter the data directly into the *Signature* field, following the guidance in the next topics.



**Table 94:** Valid syntax for custom signature fields

Field	Valid Characters	Usage
<b>HEADER</b>	F-SBID	The header for an attack definition signature. Each custom signature must begin with this header.
<b>KEYWORD</b>	Each keyword must start with a pair of dashes (--), and consist of a string of 1 to 19 characters.  Normally, keywords are an English word or English words connected by an underscore (_). Keywords are case insensitive.	The keyword is used to identify a parameter.
<b>VALUE</b>	Double quotes (") must be used around the value if it contains a space and/or a semicolon (;).  If the value is NULL, the space between the KEYWORD and VALUE can be omitted.  Values are case sensitive.  Note: If double quotes are used for quoting the value, the double quotes are not considered as part of the value string.	The value is set specifically for a parameter identified by a keyword.

## Custom signature keywords

### Information keywords

#### **attack\_id**

Syntax: `--attack_id <id_int>;`

Description:

Use this optional value to identify the signature. It cannot be the same value as any other custom rules. If an attack ID is not specified, the FortiGate automatically assigns an attack ID to the signature. If you are using VDOMs, custom signatures appear only in the VDOM in which you create them. You can use the same attack ID for signatures in different VDOMs.

An attack ID you assign must be between 1000 and 9999.

Example: `--attack_id 1234;`

#### **name**

Syntax: `--name <name_str>;`

Description:

Enter the name of the rule. A rule name must be unique. If you are using VDOMs, custom signatures appear only in the VDOM in which you create them. You can use the same rule name

for signatures in different VDOMs. The name you assign must be a string greater than 0 and less than 64 characters in length.

Example: `--name "Buffer_Overflow";`

## Session keywords

### flow

**Syntax:** `--flow {from_client[,reversed] | from_server[,reversed] | bi_direction };`

**Description:**

Specify the traffic direction and state to be inspected. They can be used for all IP traffic.

Example: `??src_port 41523; ??flow bi_direction;`

The signature checks traffic to and from port 41523.

If you enable “quarantine attacker”, the optional reversed keyword allows you to change the side of the connection to be quarantined when the signature is detected.

For example, a custom signature written to detect a brute-force log in attack is triggered when “Login Failed” is detected from\_server more than 10 times in 5 seconds. If the attacker is quarantined, it is the server that is quarantined in this instance. Adding reversed corrects this problem and quarantines the actual attacker.

Previous FortiOS versions used to\_client and to\_server values. These are now deprecated, but still function for backwards compatibility.

### service

**Syntax:** `--service {HTTP | TELNET | FTP | DNS | SMTP | POP3 | IMAP | SNMP | RADIUS | LDAP | MSSQL | RPC | SIP | H323 | NBSS | DCERPC | SSH | SSL};`

**Description:**

Specify the protocol type to be inspected. This keyword allows you to specify the traffic type by protocol rather than by port. If the decoder has the capability to identify the protocol on any port, the signature can be used to detect the attack no matter what port the service is running on. Currently, HTTP, SIP, SSL, and SSH protocols can be identified on any port based on the content.

## Content keywords

### byte\_jump

**Syntax:** `--byte_jump <bytes_to_convert>, <offset>[, multiplier][, relative] [, big] [, little] [, string] [, hex] [, dec] [, oct] [, align];`

**Description:**

Use the byte\_jump option to extract a number of bytes from a packet, convert them to their numeric representation, and jump the match reference up that many bytes (for further pattern matching or byte testing). This keyword allows relative pattern matches to take into account numerical values found in network data.

The available keyword options include:

- `<bytes_to_convert>`: The number of bytes to examine from the packet.
- `<offset>`: The number of bytes into the payload to start processing.
- `[multiplier]`: multiplier is optional. It must be a numerical value when present. The converted value multiplied by the number is the result to be skipped.
- `relative`: Use an offset relative to last pattern match.
- `big`: Process the data as big endian (default).
- `little`: Process the data as little endian.string: The data is a string in the packet.
- `hex`: The converted string data is represented in hexadecimal notation.
- `dec`: The converted string data is represented in decimal notation.
- `oct`: The converted string data is represented in octal notation.
- `align`: Round up the number of converted bytes to the next 32-bit boundary.

## byte\_test

**Syntax:** `--byte_test <bytes_to_convert>, <operator>, <value>, <offset>[, relative] [, big] [, little] [, string] [, hex] [, dec] [, oct];`

### Description:

Use the `byte_test` keyword to compare a byte field against a specific value (with operator). This keyword is capable of testing binary values or converting representative byte strings to their binary equivalent and testing them. The available keyword options include:

- `<bytes_to_convert>`: The number of bytes to compare.
- `<operator>`: The operation to perform when comparing the value (`<`, `>`, `=`, `!`, `&`).
- `<value>`: The value to compare the converted value against.
- `<offset>`: The number of bytes into the payload to start processing.
- `relative`: Use an offset relative to last pattern match.
- `big`: Process the data as big endian (default).
- `little`: Process the data as little endian.string: The data is a string in the packet.
- `hex`: The converted string data is represented in hexadecimal notation.
- `dec`: The converted string data is represented in decimal notation.
- `oct`: The converted string data is represented in octal notation.

## depth

**Syntax:** `--depth <depth_int>;`

### Description:

Use the `depth` keyword to search for the contents within the specified number of bytes after the starting point defined by the `offset` keyword. If no offset is specified, the offset is assumed to be equal to 0.

If the value of the `depth` keyword is smaller than the length of the value of the `content` keyword, this signature will never be matched.

The `depth` must be between 0 and 65535.

## distance

**Syntax:** `--distance <dist_int>;`

### Description:

Use the distance keyword to search for the contents within the specified number of bytes relative to the end of the previously matched contents. If the within keyword is not specified, continue looking for a match until the end of the payload.

The distance must be between 0 and 65535.

## content

**Syntax:** `--content [!]"<content_str>"`;

### Description:

Deprecated, see pattern and context keywords. Use the content keyword to search for the content string in the packet payload. The content string must be enclosed in double quotes.

To have the FortiGate search for a packet that does not contain the specified context string, add an exclamation mark (!) before the content string.

Multiple content items can be specified in one rule. The value can contain mixed text and binary data. The binary data is generally enclosed within the pipe (|) character.

The double quote ("), pipe sign(|) and colon(:) characters must be escaped using a back slash if specified in a content string.

If the value of the content keyword is greater than the length of the value of the depth keyword, this signature will never be matched.

## context

**Syntax:** `--context {uri | header | body | host}`;

### Description:

Specify the protocol field to look for the pattern. If context is not specified for a pattern, the FortiGate unit searches for the pattern anywhere in the packet buffer.

The available context variables are:

`uri`: Search for the pattern in the HTTP URI line.

`header`: Search for the pattern in HTTP header lines or SMTP/POP3/SMTP control messages.

`body`: Search for the pattern in HTTP body or SMTP/POP3/SMTP email body.

`host`: Search for the pattern in HTTP HOST line.

## no\_case

**Syntax:** `--no_case`;

### Description:

Use the no-case keyword to force the FortiGate unit to perform a case-insensitive pattern match.

## offset

**Syntax:** `--offset <offset_int>`;

### Description:

Use the offset keyword to look for the contents after the specified number of bytes into the payload. The specified number of bytes is an absolute value in the payload. Follow the offset keyword with the depth keyword to stop looking for a match after a specified number of bytes. If no depth is specified, the FortiGate unit continues looking for a match until the end of the payload. The offset must be between 0 and 65535.

## pattern

**Syntax:** `--pattern [!] "<pattern_str>"`;

### Description:

The FortiGate unit will search for the specified pattern. A pattern keyword normally is followed by a context keyword to define where to look for the pattern in the packet. If a context keyword is not present, the FortiGate unit looks for the pattern anywhere in the packet buffer. To have the FortiGate search for a packet that does not contain the specified URI, add an exclamation mark (!) before the URI.

Example: `--pattern "/level/" --pattern "|E8 D9FF FFFF|/bin/sh" --pattern "!|20|RTSP/"`

## pcre

**Syntax:** `--pcre [!] "<regex>/[ismxAEGRUB]"`;

### Description:

Similarly to the pattern keyword, use the pcre keyword to specify a pattern using Perl-compatible regular expressions (PCRE). A pcre keyword can be followed by a context keyword to define where to look for the pattern in the packet. If no context keyword is present, the FortiGate unit looks for the pattern anywhere in the packet buffer. For more information about PCRE syntax, go to <http://www.pcre.org>.

The switches include:

- **i:** Case insensitive.
- **s:** Include newlines in the dot metacharacter.
- **m:** By default, the string is treated as one big line of characters. **^** and **\$** match at the beginning and ending of the string. When **m** is set, **^** and **\$** match immediately following or immediately before any newline in the buffer, as well as the very start and very end of the buffer.
- **x:** White space data characters in the pattern are ignored except when escaped or inside a character class.
- **A:** The pattern must match only at the start of the buffer (same as **^**).
- **E:** Set **\$** to match only at the end of the subject string. Without **E**, **\$** also matches immediately before the final character if it is a newline (but not before any other newlines).
- **G:** Invert the “greediness” of the quantifiers so that they are not greedy by default, but become greedy if followed by **?**.
- **R:** Match relative to the end of the last pattern match. (Similar to `distance:0`).
- **U:** Deprecated, see the context keyword. Match the decoded URI buffers.

## uri

**Syntax:** `--uri [!] "<uri_str>"`;

### Description:

Deprecated, see pattern and context keywords. Use the uri keyword to search for the URI in the packet payload. The URI must be enclosed in double quotes ("). To have the FortiGate unit search for a packet that does not contain the specified URI, add an exclamation mark (!) before the URI. Multiple content items can be specified in one rule. The value can contain mixed text and binary data. The binary data is generally enclosed within the pipe (|) character. The double quote ("), pipe sign (|) and colon (:) characters must be escaped using a back slash (\) if specified in a URI string.

## within

**Syntax:** `--within <within_int>;`

**Description:**

Use this together with the distance keyword to search for the contents within the specified number of bytes of the payload. The within value must be between 0 and 65535.

## IP header keywords

### dst\_addr

**Syntax:** `--dst_addr [!]<ipv4>;`

**Description:**

Use the `dst_addr` keyword to search for the destination IP address. To have the FortiGate search for a packet that does not contain the specified address, add an exclamation mark (!) before the IP address. You can define up to 28 IP addresses or CIDR blocks. Enclose the comma separated list in square brackets.

Example: `dst_addr [172.20.0.0/16, 10.1.0.0/16, 192.168.0.0/16]`

### ip\_id

**Syntax:** `--ip_id <field_int>;`

**Description:**

Check the IP ID field for the specified value.

### ip\_option

**Syntax:** `--ip_option {rr | eol | nop | ts | sec | lsrr | ssrr | satid | any};`

**Description:**

Use the `ip_option` keyword to check various IP option settings.

The available options include:

- `rr`: Check if IP RR (record route) option is present.
- `eol`: Check if IP EOL (end of list) option is present.
- `nop`: Check if IP NOP (no op) option is present.
- `ts`: Check if IP TS (time stamp) option is present.
- `sec`: Check if IP SEC (IP security) option is present.
- `lsrr`: Check if IP LSRR (loose source routing) option is present.
- `ssrr`: Check if IP SSRR (strict source routing) option is present.
- `satid`: Check if IP SATID (stream identifier) option is present.
- `any`: Check if IP any option is present.

### ip\_tos

**Syntax:** `--ip_tos <field_int>;`

**Description:**

Check the IP TOS field for the specified value.

## ip\_ttl

**Syntax:** `--ip_ttl [< | >] <t1l_int>;`

### Description:

Check the IP time-to-live value against the specified value. Optionally, you can check for an IP time-to-live greater-than (>) or less-than (<) the specified value with the appropriate symbol.

## protocol

**Syntax:** `--protocol {<protocol_int> | tcp | udp | icmp};`

### Description:

Check the IP protocol header.

Example: `--protocol tcp;`

## src\_addr

**Syntax:** `--src_addr [!]<ipv4>;`

### Description:

Use the `src_addr` keyword to search for the source IP address. To have the FortiGate unit search for a packet that does not contain the specified address, add an exclamation mark (!) before the IP address. You can define up to 28 IP addresses or CIDR blocks. Enclose the comma separated list in square brackets.

Example: `src_addr 192.168.13.0/24`

## TCP header keywords

### ack

**Syntax:** `--ack <ack_int>;`

### Description:

Check for the specified TCP acknowledge number.

### dst\_port

**Syntax:** `--dst_port [!]{<port_int> | :<port_int> | <port_int>: | <port_int>:<port_int>;}`

### Description:

Use the `dst_port` keyword to specify the destination port number.

You can specify a single port or port range:

- `<port_int>` is a single port.
- `:<port_int>` includes the specified port and all lower numbered ports.
- `<port_int>:` includes the specified port and all higher numbered ports.
- `<port_int>:<port_int>` includes the two specified ports and all ports in between.

### seq

**Syntax:** `--seq <seq_int>;`

### Description:

Check for the specified TCP sequence number.

### src\_port

**Syntax:** `--src_port [!]{<port_int> | :<port_int> | <port_int>: | <port_int>:<port_int>};`

#### Description:

Use the `src_port` keyword to specify the source port number.

You can specify a single port or port range:

- `<port_int>` is a single port.
- `:<port_int>` includes the specified port and all lower numbered ports.
- `<port_int>:` includes the specified port and all higher numbered ports.
- `<port_int>:<port_int>` includes the two specified ports and all ports in between.

### tcp\_flags

**Syntax:** `--tcp_flags <SAFRUP120>[!|*|+] [,<SAFRUP120>];`

#### Description:

Specify the TCP flags to match in a packet.

- `S`: Match the SYN flag.
- `A`: Match the ACK flag.
- `F`: Match the FIN flag.
- `R`: Match the RST flag.
- `U`: Match the URG flag.
- `P`: Match the PSH flag.
- `1`: Match Reserved bit 1.
- `2`: Match Reserved bit 2.
- `0`: Match No TCP flags set.
- `!`: Match if the specified bits are not set.
- `*`: Match if any of the specified bits are set.
- `+`: Match on the specified bits, plus any others.

The first part if the value (`<SAFRUP120>`) defines the bits that must be present for a successful match.

Example:

```
--tcp_flags AP only matches the case where both A and P bits are set.
```

The second part (`[, <SAFRUP120>]`) is optional, and defines the additional bits that can be present for a match.

For example `tcp_flags S, 12` matches the following combinations of flags: `S`, `S` and `1`, `S` and `2`, `S` and `1` and `2`. The modifiers `!`, `*` and `+` cannot be used in the second part.

### window\_size

**Syntax:** `--window_size [!]<window_int>;`

#### Description:



Check for the specified TCP window size. You can specify the window size as a hexadecimal or decimal integer. A hexadecimal value must be preceded by 0x. To have the FortiGate search for the absence of the specified window size, add an exclamation mark (!) before the window size.

## UDP header keywords

### dst\_port

**Syntax:** `--dst_port [!]{<port_int> | :<port_int> | <port_int>: | <port_int>:<port_int>};`

**Description:**

Specify the destination port number.

You can specify a single port or port range:

- `<port_int>` is a single port.
- `:<port_int>` includes the specified port and all lower numbered ports.
- `<port_int>:` includes the specified port and all higher numbered ports.
- `<port_int>:<port_int>` includes the two specified ports and all ports in between.

### src\_port

**Syntax:** `--src_port [!]{<port_int> | :<port_int> | <port_int>: | <port_int>:<port_int>};`

**Description:**

Specify the destination port number.

You can specify a single port or port range:

- `<port_int>` is a single port.
- `:<port_int>` includes the specified port and all lower numbered ports.
- `<port_int>:` includes the specified port and all higher numbered ports.
- `<port_int>:<port_int>` includes the two specified ports and all ports in between.

## ICMP keywords

### icmp\_code

**Syntax:** `--icmp_code <code_int>;`

**Description:**

Specify the ICMP code to match.

### icmp\_id

**Syntax:** `--icmp_id <id_int>;`

**Description:**

Check for the specified ICMP ID value.

### icmp\_seq

**Syntax:** `--icmp_seq <seq_int>;`

Description:

Check for the specified ICMP sequence value.

### icmp\_type

**Syntax:** `--icmp_type <type_int>;`

**Description:**

Specify the ICMP type to match.

## Other keywords

### data\_size

**Syntax:** `--data_size {<size_int> | <<size_int> | ><size_int>;`

**Description:**

Test the packet payload size. With `data_size` specified, packet reassembly is turned off automatically. So a signature with `data_size` and `only_stream` values set is wrong.

- `<size_int>` is a particular packet size.
- `<<size_int>` is a packet smaller than the specified size.
- `><size_int>` is a packet larger than the specified size.

Example: `--data_size <300>;`

### data\_at

**Syntax:** `--data_at <offset_int>[, relative];`

**Description:**

Verify that the payload has data at a specified offset, optionally looking for data relative to the end of the previous content match.

### rate

**Syntax:** `--rate <matches_int>,<time_int>;`

**Description:**

Instead of generating log entries every time the signature is detected, use this keyword to generate a log entry only if the signature is detected a specified number of times within a specified time period.

- `<matches_int>` is the number of times a signature must be detected.
- `<time_int>` is the length of time in which the signature must be detected, in seconds.

For example, if a custom signature detects a pattern, a log entry will be created every time the signature is detected. If `--rate 100,10;` is added to the signature, a log entry will be created if the signature is detected 100 times in the previous 10 seconds. Use this command with `--track` to further limit log entries to when the specified number of detections occur within a certain time period involving the same source or destination address rather than all addresses.

### rpc\_num

**Syntax:** `--rpc_num <app_int>[, <ver_int> | *][, <proc_int> | *];`

**Description:**

Check for RPC application, version, and procedure numbers in SUNRPC CALL requests. The `*` wild card can be used for version and procedure numbers.

## same\_ip

**Syntax:** `--same_ip;`

**Description:**

Check that the source and the destination have the same IP addresses.

## track

**Syntax:** `--track {SRC_IP |DST_IP |DHCP_CLIENT |DNS_DOMAIN} [,block_int];`

**Description:**

When used with `--rate`, this keyword narrows the custom signature rate totals to individual addresses.

- `SRC_IP`: tracks the packet's source IP.
- `DST_IP`: tracks the packet's destination IP.
- `DHCP_CLIENT`: tracks the DHCP client's MAC address.
- `DNS_DOMAIN`: counts the number of any specific domain name.
- `block_int` has the FortiGate unit block connections for the specified number of seconds, from the client or to the server, depending on which is specified.

For example, if `--rate 100,10` is added to the signature, a log entry will be created if the signature is detected 100 times in the previous 10 seconds. The FortiGate unit maintains a single total, regardless of source and destination address.

If the same custom signature also includes `--track client`; matches are totaled separately for each source address. A log entry is added when the signature is detected 100 times in 10 seconds within traffic from the same source address.

The `--track` keyword can also be used without `--rate`. If an integer is specified, the client or server will be blocked for the specified number of seconds every time the signature is detected.

## Creating a custom signature to block access to example.com

In this first example, you will create a custom signature to block access to the example.com URL.

This example describes the use of the custom signature syntax to block access to a URL. To create the custom signature entry in the FortiGate unit web-based manager, see [“Creating a custom IPS signature” on page 2080](#).

**1. Enter the custom signature basic format**

All custom signatures have a header and at least one keyword/value pair. The header is always the same:

```
F-SBID()
```

The keyword/value pairs appear within the parentheses and each pair is followed by a semicolon.

## 2. Choose a name for the custom signature

Every custom signature requires a name, so it is a good practice to assign a name before adding any other keywords.

Use the `--name` keyword to assign the custom signature a name. The name value follows the keyword after a space. Enclose the name value in double-quotes:

```
F-SBID(--name "Block.example.com";)
```

The signature, as it appears here, will not do anything if you try to use it. It has a name, but does not look for any patterns in network traffic. You must specify a pattern that the FortiGate unit will search for.

## 3. Add a signature pattern

Use the `--pattern` keyword to specify what the FortiGate unit will search for:

```
F-SBID(--name "Block.example.com"; --pattern "example.com";)
```

The signature will now detect the `example.com` URL appearing in network traffic. The custom signature should only detect the URL in HTTP traffic, however. Any other traffic with the URL should be allowed to pass. For example, an email message to or from `example.com` should not be stopped.

## 4. Specify the service

Use the `--service` keyword to limit the effect of the custom signature to only the HTTP protocol.

```
F-SBID(--name "Block.example.com"; --pattern "example.com";
--service HTTP;)
```

The FortiGate unit will limit its search for the pattern to the HTTP protocol. Even though the HTTP protocol uses only TCP traffic, the FortiGate will search for HTTP protocol communication in TCP, UDP, and ICMP traffic. This is a waste of system resources that you can avoid by limiting the search further, as shown below.

## 5. Specify the traffic type.

Use the `--protocol tcp` keyword to limit the effect of the custom signature to only TCP traffic. This will save system resources by not unnecessarily scanning UDP and ICMP traffic.

```
F-SBID(--name "Block.example.com"; --pattern "example.com";
--service HTTP; --protocol tcp;)
```

The FortiGate unit will limit its search for the pattern to TCP traffic and ignore UDP and ICMP network traffic.

## 6. Ignore case sensitivity

By default, patterns are case sensitive. If a user directed his or her browser to `Example.com`, the custom signature would not recognize the URL as a match.

Use the `--no_case` keyword to make the pattern matching case insensitive.

```
F-SBID(--name "Block.example.com"; --pattern "example.com";
--service HTTP; --no_case;)
```

Unlike all of the other keywords in this example, the `--no_case` keyword has no value. Only the keyword is required.

## 7. Limit pattern scans to only traffic sent from the client

The `--flow` command can be used to further limit the network traffic being scanned to only that sent by the client or by the server.

```
F-SBID(--name "Block.example.com"; --pattern "example.com";
--service HTTP; --no_case; --flow from_client;)
```

Web servers do not contact clients until clients first open a communication session. Therefore, using the `--flow from_client` command will force the FortiGate to ignore all

traffic from the server. Since the majority of HTTP traffic flows from the server to the client, this will save considerable system resources and still maintain protection.

#### 8. Specify the context

When the client browser tries to contact example.com, a DNS is first consulted to get the example.com server IP address. The IP address is then specified in the URL field of the HTTP communication. The domain name will still appear in the host field, so this custom signature will not function without the `--context host` keyword/value pair.

```
F-SBID(--name "Block.example.com"; --pattern "example.com";
 --service HTTP; --no_case; --flow from_client;
 --context host;)
```

## Creating a custom signature to block the SMTP “vrfy” command

The SMTP “vrfy” command can be used to verify the existence of a single email address or to list all of the valid email accounts on an email server. A spammer could potentially use this command to obtain a list of all valid email users and direct spam to their inboxes.

In this example, you will create a custom signature to block the use of the vrfy command. Since the custom signature blocks the vrfy command from coming through the FortiGate unit, the administrator can still use the command on the internal network.

This example describes the use of the custom signature syntax to block the vrfy command. To create the custom signature entry in the FortiGate unit web-based manager, see [“Creating a custom IPS signature” on page 2080](#).

#### 1. Enter the custom signature basic format

All custom signatures have a header and at least one keyword/value pair. The header is always the same:

```
F-SBID()
```

The keyword/value pairs appear within the parentheses and each pair is followed by a semicolon.

#### 2. Choose a name for the custom signature

Every custom signature requires a name, so it is a good practice to assign a name before you add any other keywords.

Use the `--name` keyword to assign the custom signature a name. The name value follows the keyword after a space. Enclose the name value in double-quotes:

```
F-SBID(--name "Block.SMTP.VRFY.CMD";)
```

The signature, as it appears here, will not do anything if you try to use it. It has a name, but does not look for any patterns in network traffic. You must specify a pattern that the FortiGate unit will search for.

#### 3. Add a signature pattern

Use the `--pattern` keyword to specify what the FortiGate unit will search for:

```
F-SBID(--name "Block.SMTP.VRFY.CMD"; --pattern "vrfy";)
```

The signature will now detect the vrfy command appearing in network traffic. The custom signature should only detect the command in SMTP traffic, however. Any other traffic with the pattern should be allowed to pass. For example, an email message discussing the vrfy command should not be stopped.

#### 4. Specify the service

Use the `--service` keyword to limit the effect of the custom signature to only the HTTP protocol.

```
F-SBID(--name "Block.SMTP.VRFY.CMD"; --pattern "vrfy";
 --service SMTP;)
```

The FortiGate unit will limit its search for the pattern to the SMTP protocol.

Even though the SMTP protocol uses only TCP traffic, the FortiGate will search for SMTP protocol communication in TCP, UDP, and ICMP traffic. This is a waste of system resources that you can avoid by limiting the search further, as shown below.

#### 5. Specify the traffic type.

Use the `--protocol tcp` keyword to limit the effect of the custom signature to only TCP traffic. This will save system resources by not unnecessarily scanning UDP and ICMP traffic.

```
F-SBID(--name "Block.SMTP.VRFY.CMD"; --pattern "vrfy";
 --service SMTP; --protocol tcp;)
```

The FortiGate unit will limit its search for the pattern to TCP traffic and ignore the pattern in UDP and ICMP network traffic.

#### 6. Ignore case sensitivity

By default, patterns are case sensitive. If a user directed his or her browser to Example.com, the custom signature would not recognize the URL as a match.

Use the `--no_case` keyword to make the pattern matching case insensitive.

```
F-SBID(--name "Block.SMTP.VRFY.CMD"; --pattern "vrfy";
 --service SMTP; --no_case;)
```

Unlike all of the other keywords in this example, the `--no_case` keyword has no value. Only the keyword is required.

#### 7. Specify the context

The SMTP vrfy command will appear in the SMTP header. The `--context host` keyword/value pair allows you to limit the pattern search to only the header.

```
F-SBID(--name "Block.SMTP.VRFY.CMD"; --pattern "vrfy";
 --service SMTP; --no_case; --context header;)
```

# Web filter

This section describes FortiGate web filtering for HTTP traffic. The three main parts of the web filtering function, the Web Content Filter, the URL Filter, and the FortiGuard Web Filtering Service interact with each other to provide maximum control over what the Internet user can view as well as protection to your network from many Internet content threats. Web Content Filter blocks web pages containing words or patterns that you specify. URL filtering uses URLs and URL patterns to block or exempt web pages from specific sources. FortiGuard Web Filtering provides many additional categories you can use to filter web traffic.

This section describes the Web Content Filter and URL Filter functions. For information on FortiGuard Web Filtering, see [“FortiGuard Web Filter” on page 133](#)

The following topics are included in this section:

- [Web filter concepts](#)
- [Inspections Modes](#)
- [FortiGuard Web Filtering Service](#)
- [Overriding FortiGuard website categorization](#)
- [SafeSearch](#)
- [YouTube Education Filter](#)
- [Web Site Filter](#)
- [Web content filter](#)
- [Advanced web filter configurations](#)
- [Working with the Interface](#)
- [Web filtering example](#)

## Web filter concepts

Web filtering is a means of controlling the content that an Internet user is able to view. With the popularity of web applications, the need to monitor and control web access is becoming a key component of secure content management systems that employ antivirus, web filtering, and messaging security. Important reasons for controlling web content include:

- lost productivity because employees are accessing the web for non-business reasons
- network congestion — when valuable bandwidth is used for non-business purposes, legitimate business applications suffer
- loss or exposure of confidential information through chat sites, non-approved email systems, instant messaging, and peer-to-peer file sharing
- increased exposure to web-based threats as employees surf non-business-related web sites
- legal liability when employees access/download inappropriate and offensive material
- copyright infringement caused by employees downloading and/or distributing copyrighted material.

As the number and severity of threats increase on the World Wide Web, the risk potential increases within a company's network as well. Casual non-business related web surfing has caused many businesses countless hours of legal litigation as hostile environments have been created by employees who download and view offensive content. Web-based attacks and

threats are also becoming increasingly sophisticated. Threats and web-based applications that cause additional problems for corporations include:

- spyware/grayware
- phishing
- pharming
- instant messaging
- peer-to-peer file sharing
- streaming media
- blended network attacks.

**Spyware**, also known as grayware, is a type of computer program that attaches itself to a user's operating system. It does this without the user's consent or knowledge. It usually ends up on a computer because of something the user does such as clicking on a button in a pop-up window. Spyware can track the user's Internet usage, cause unwanted pop-up windows, and even direct the user to a host web site. For further information, visit the [FortiGuard Center](#).

Some of the most common ways of grayware infection include:

- downloading shareware, freeware, or other forms of file-sharing services
- clicking on pop-up advertising
- visiting legitimate web sites infected with grayware.

**Phishing** is the term used to describe attacks that use web technology to trick users into revealing personal or financial information. Phishing attacks use web sites and email that claim to be from legitimate financial institutions to trick the viewer into believing that they are legitimate. Although phishing is initiated by spam email, getting the user to access the attacker's web site is always the next step.

**Pharming** is a next generation threat that is designed to identify and extract financial, and other key pieces of information for identity theft. Pharming is much more dangerous than phishing because it is designed to be completely hidden from the end user. Unlike phishing attacks that send out spam email requiring the user to click to a fraudulent URL, pharming attacks require no action from the user outside of their regular web surfing activities. Pharming attacks succeed by redirecting users from legitimate web sites to similar fraudulent web sites that have been created to look and feel like the authentic web site.

**Instant messaging** presents a number of problems. Instant messaging can be used to infect computers with spyware and viruses. Phishing attacks can be made using instant messaging. There is also a danger that employees may use instant messaging to release sensitive information to an outsider.

**Peer-to-peer (P2P)** networks are used for file sharing. Such files may contain viruses. Peer-to-peer applications take up valuable network resources and may lower employee productivity but also have legal implications with the downloading of copyrighted or sensitive company material.

**Streaming media** is a method of delivering multimedia, usually in the form of audio or video to Internet users. Viewing streaming media impacts legitimate business by using valuable bandwidth.

**Blended network threats** are rising and the sophistication of network threats is increasing with each new attack. Attackers learn from each previous successful attack and enhance and update attack code to become more dangerous and fast spreading. Blended attacks use a combination of methods to spread and cause damage. Using virus or network worm techniques combined with known system vulnerabilities, blended threats can quickly spread through email, web sites, and Trojan applications. Examples of blended threats include Nimda, Code Red, Slammer, and Blaster. Blended attacks can be designed to perform different types of attacks,



which include disrupting network services, destroying or stealing information, and installing stealthy backdoor applications to grant remote access.

## Different ways of controlling access

The methods available for monitoring and controlling Internet access range from manual and educational methods to fully automated systems designed to scan, inspect, rate and control web activity.

Common web access control mechanisms include:

- establishing and implementing a well-written usage policy in the organization on proper Internet, email, and computer conduct
- installing monitoring tools that record and report on Internet usage
- implementing policy-based tools that capture, rate, and block URLs.

The final method is the focus of this topic. The following information shows how the filters interact and how to use them to your advantage.

## Order of web filtering

The FortiGate unit applies web filters in a specific order:

1. URL filter
2. FortiGuard Web Filter
3. web content filter
4. web script filter
5. antivirus scanning.

If you have blocked a FortiGuard Web Filter category but want certain users to have access to URLs within that pattern, you can use the *Override* within the FortiGuard Web Filter. This will allow you to specify which users have access to which blocked URLs and how long they have that access. For example, if you want a user to be able to access [www.example.com](http://www.example.com) for one hour, you can use the override to set up the exemption. Any user listed in an override must fill out an online authentication form that is presented when they try to access a blocked URL before the FortiGate unit will grant access to it. For more information, see [“FortiGuard Web Filter” on page 133](#).

## Inspections Modes

### Proxy

Proxy-based inspection involves buffering the traffic and examining it as a whole before determining an action. The process of having the whole of the data to analyze allow this process to include more points of data to analyze than the flow-based or DNS methods.

The advantage of a proxy-based method is that the inspection can be more thorough than the other methods, resulting in fewer false positive or negative results in the analysis of the data.

### Flow-based

The Flow-based inspection method examines the file as it passes through the FortiGate unit without any buffering. As each packet of the traffic arrives it is process and forwarded without waiting for the complete file or web page, etc.

The advantage of the flow-based method is that the user sees a faster response time for HTTP requests and there is less chance of a time-out error due to the server at the other end responding slowly.

The disadvantages of this method are that there is a higher probability of a false positive or negative in the analysis of the data and that a number of points of analysis that can be used in the proxy-based method are not available in the flow-based inspection method. There is also fewer actions available to choose from based on the categorization of the website by FortiGuard services.

## DNS

The DNS inspection method uses the same categories as the FortiGuard Service. It is lightweight in terms of resource usage because it doesn't involve any proxy-based or flow-based inspection.

A DNS request is typically the first part of any new session to a new website. This inspection method takes advantage of that and places the results of the categorization of websites right on the FortiGuard DNS servers. When the FortiGate resolves a URL, in addition to the IP address of the website it also receives a domain rating.

In the same way that the flow-based inspection method had fewer filters and points of analysis than the proxy-based inspection method, DNS has fewer settings still. All of its inspection is based on the IP address, the domain name and the rating provided by the FortiGuard DNS server.

## FortiGuard Web Filtering Service

FortiGuard Web Filter is a managed web filtering solution available by subscription from Fortinet. FortiGuard Web Filter enhances the web filtering features supplied with your FortiGate unit by sorting billions of web pages into a wide range of categories users can allow or block. The FortiGate unit accesses the nearest FortiGuard Web Filter Service Point to determine the category of a requested web page, and then applies the security policy configured for that user or interface.

FortiGuard Web Filter includes over 45 million individual ratings of web sites that apply to more than two billion pages. Pages are sorted and rated into several dozen categories administrators can allow or block. Categories may be added or updated as the Internet evolves. To make configuration simpler, you can also choose to allow or block entire groups of categories. Blocked pages are replaced with a message indicating that the page is not accessible according to the Internet usage policy.

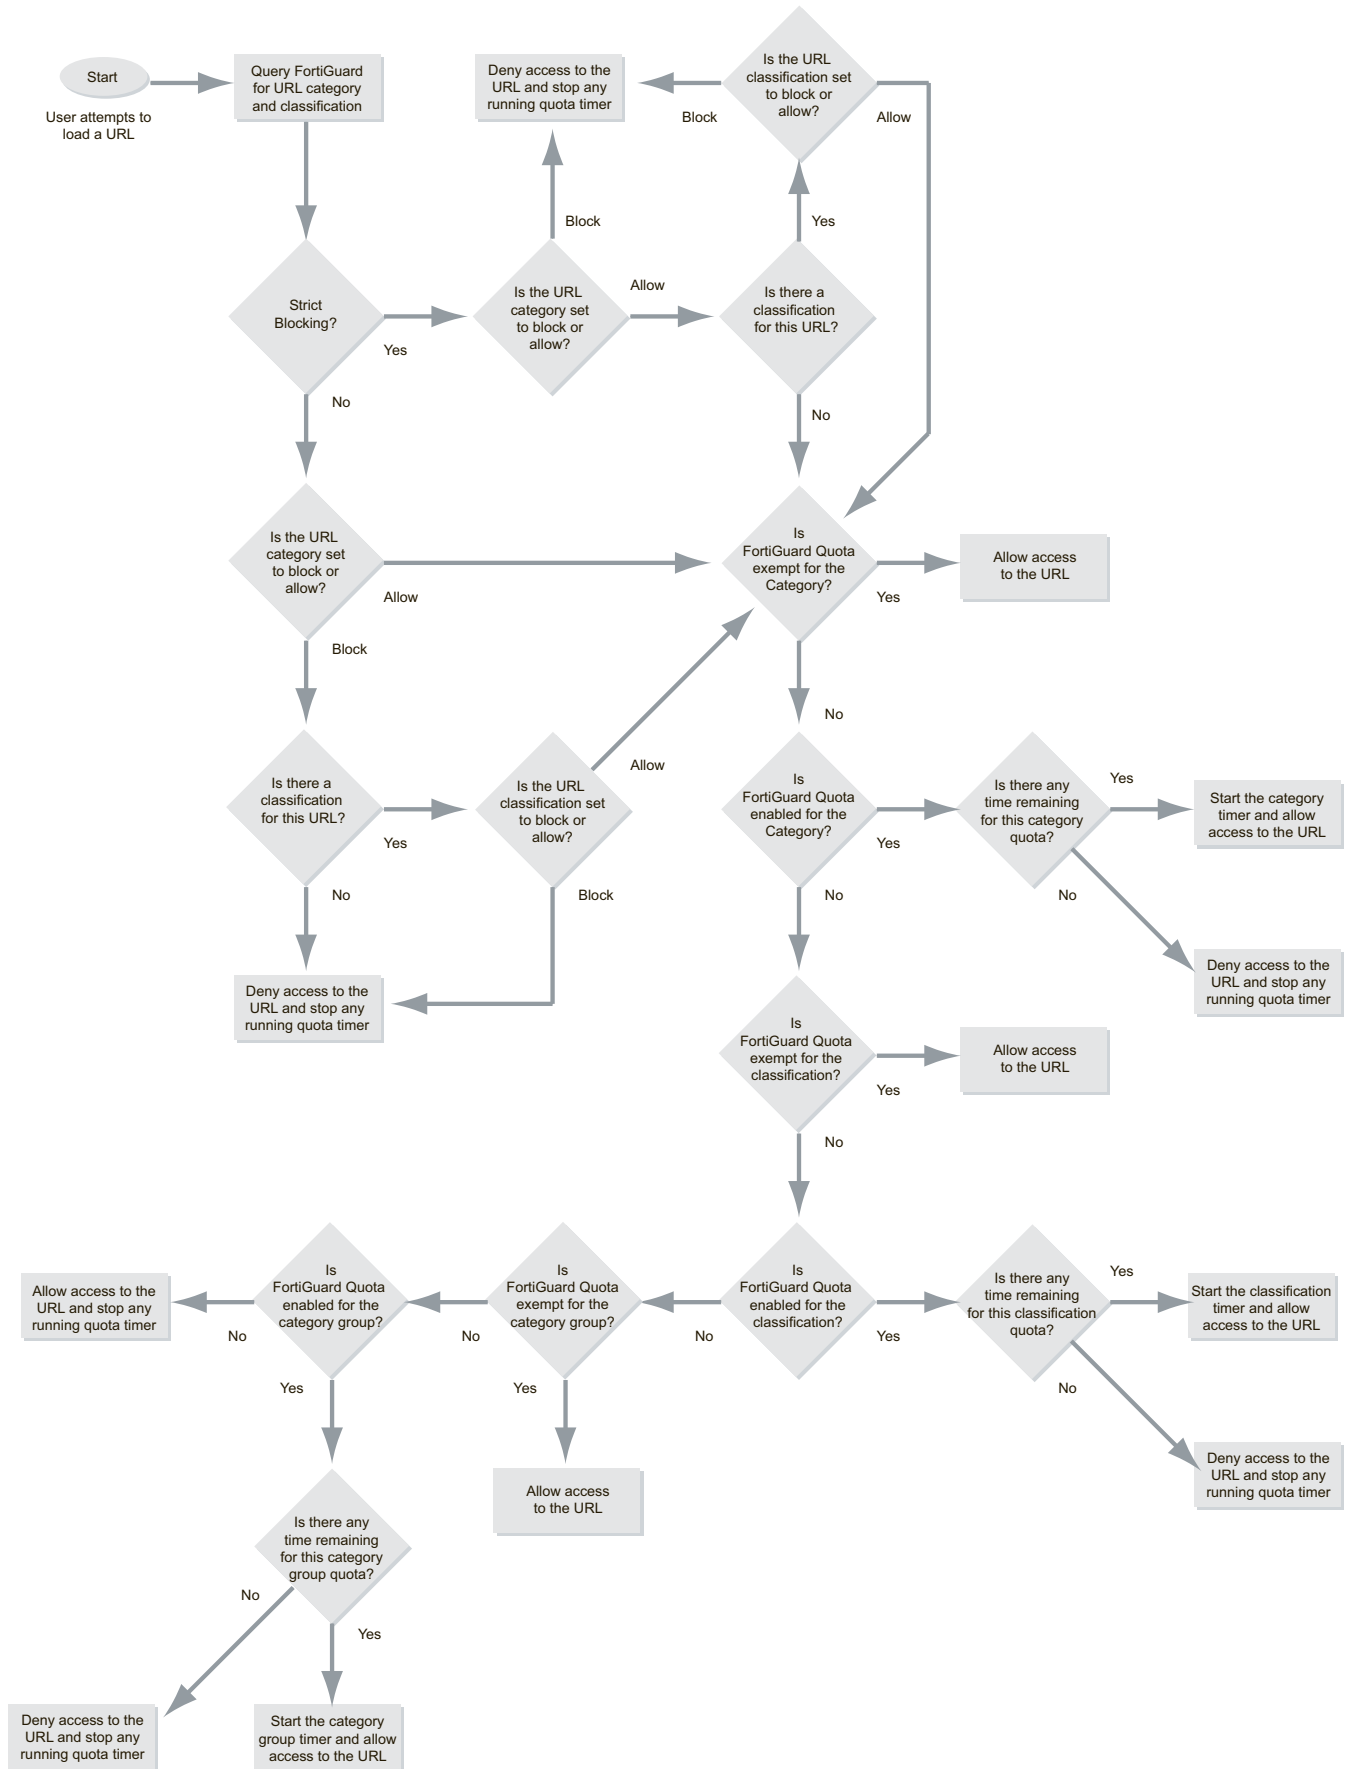
FortiGuard Web Filter ratings are performed by a combination of proprietary methods including text analysis, exploitation of the web structure, and human raters. Users can notify the FortiGuard Web Filter Service Points if they feel a web page is not categorized correctly, so that the service can update the categories in a timely fashion.

Before you begin to use the FortiGuard Web Filter options you should verify that you have a valid subscription to the service for your FortiGate firewall.

## FortiGuard Web Filter and your FortiGate unit

When FortiGuard Web Filter is enabled in a web filter profile, the setting is applied to all firewall policies that use this profile. When a request for a web page appears in traffic controlled by one of these firewall policies, the URL is sent to the nearest FortiGuard server. The URL category is returned. If the category is blocked, the FortiGate unit provides a replacement message in place of the requested page. If the category is not blocked, the page request is sent to the requested URL as normal.

Figure 311: Webfiltering flowchart



## Enabling FortiGuard Web Filter

FortiGuard Web Filter is enabled and configured within web filter profiles by enabling FortiGuard Categories. The service is engaged by turning on the Web Filter profile and selecting a profile that has FortiGuard Categories enabled on one or more active policies being run by the firewall.

There is also a system wide setting for the enabling or disabling of FortiGuard Web Filter that is only in the CLI.

```
config system fortiguard
 set webfilter-force-off
```

The two options on this setting are enable or disable. The syntax of the settings name is “force-off” so in order to enable FortiGuard Webfilter you have to choose disable for the setting and enable if you want to turn it off.

## General configuration steps

1. Go to Security Profiles > Web Filter > Profile.
2. Select the Edit icon of the web filter profile in which you want to enable FortiGuard Web Filter, or select Create New to add a new web filter profile.
3. Select an Inspection Mode.
4. If you are using FortiGuard Categories, enable the feature, select the categories and select the action to be performed.
5. The categories allow you to block or allow access to general or more specific web site categories. Configure access as required.
6. Save the filter and web filter profile.
7. To complete the configuration, you need to select the security policy controlling the network traffic you want to restrict. Then, in the security policy, enable Web Filter and select the appropriate web filter profile from the list.

## Configuring FortiGuard Web Filter settings

FortiGuard Web Filter includes a number of settings that allow you to determine various aspects of the filtering behavior.

### To configure FortiGuard Web Filter settings

1. Go to *Security Profiles > Web Filter > Profile*.
2. Select the web filter profile in which you want to enable FortiGuard Web Filter from the drop down list in the Edit Web Filter Profile window title bar, or select *Create New* to add a new web filter profile.
3. The category groups are listed in a table. You can expand each category group to view and configure every category within the groups. If you change the setting of a category group, all categories within the group inherit the change.
4. Select the category groups and categories to which you want to apply an action.

5. Select an action from the *Change Action for Selected Categories* drop-down list immediately below the category table. Five actions are available:
  - *Allow* permits access to the sites within the category.
  - *Monitor* permits and logs access to sites in the category. You may also enable user quotas when enabling the monitor action.
  - *Warning* presents the user with a message, allowing them to continue if they choose.
  - *Authenticate* requires a user authenticate with the FortiGate unit before being allowed access to the category or category group.
  - *Block* prevents access to sites within the category. Users attempting to access a blocked site will receive a replacement message explaining that access to the site is blocked.
6. Select *OK*.



In older versions of FortiOS there was a character limitation for the URL of 2048 bytes or approximately 321 characters. If the URL you were trying to reach was longer the URL sent to FortiGuard would be truncated and the service would be unable to categorize the site. Starting in version 5 of the firmware the parsed URL has been increase to 4Kilobytes, effectively doubling the length of a URL capable of being categorized.

---

## To configure the FortiGuard Web Filter categories

1. Go to *Security Profiles > Web Filter > Profiles*.
2. Select the web filter profile in which you want to enable FortiGuard Web Filter from the drop down list in the Edit Web Filter Profile window title bar, or select *Create New* to add a new web filter profile.
3. Select *Create New*.
4. Select a *Filter Type* of *Category*.
5. Select the required category groups. You may also expand the category groups to select individual categories.
6. Select the *Monitor* action.
7. Enable *Enforce Quota* to activate the quota for the selected categories and category groups.
8. Select *Hours*, *Minutes*, or *Seconds* and enter the number of hours, minutes, or seconds. This is the daily quota allowance for each user.
9. Select *OK*.
10. Select *Apply*.

Apply the web filter profile to an identity-based security policy. All the users subject to that policy are restricted by the quotas.



If you look at your logs carefully, you may notice that not every URL connection in the log shows a category. They are left blank. If you take one of those URL and enter it in the FortiGuard website designed to show the category for a URL it will successfully categorize it.

The reason for this is that to optimize speed throughput and reduce the load on the FortiGuard servers the FortiGate does not determine a category rating on scripts and css files.

---

## Configuring FortiGuard Web Filter usage quotas

In addition to using category and classification blocks and overrides to limit user access to URLs, you can set a daily timed access quota by category, category group, or classification.

Quotas allow access for a specified length of time, calculated separately for each user. Quotas are reset every day at midnight.

Users must authenticate with the FortiGate unit. The quota is applied to each user individually so the FortiGate must be able to identify each user. One way to do this is to configure a security policy using the identity based policy feature. Apply the web filter profile in which you have configured FortiGuard Web Filter and FortiGuard Web Filter quotas to such a security policy.



The use of FortiGuard Web Filter quotas requires that users authenticate to gain web access. The quotas are ignored if applied to a security policy in which user authentication is not required.

---

When a user first attempts to access a URL, they're prompted to authenticate with the FortiGate unit. When they provide their user name and password, the FortiGate unit recognizes them, determines their quota allowances, and monitors their web use. The category and classification of each page they visit is checked and FortiGate unit adjusts the user's remaining available quota for the category or classification.



Editing the web filter profile resets the quota timers for all users.

1. Select the *Monitor* action.
2. Enable *Enforce Quota* to activate the quota for the selected categories and category groups.
3. Select *Hours*, *Minutes*, or *Seconds* and enter the number of hours, minutes, or seconds. This is the daily quota allowance for each user.
4. Select *OK*.
5. Select *Apply*.

Apply the web filter profile to an identity-based security policy. All the users subject to that policy are restricted by the quotas.

### Quota hierarchy

You can apply quotas to categories and category groups. Only one quota per user can be active at any one time. The one used depends on how you configure the FortiGuard Web Filter.

When a user visits a URL, the FortiGate unit queries the FortiGuard servers for the category of the URL. From highest to lowest, the relative priority of the quotas are:

1. Category
2. Category group

## Overriding FortiGuard website categorization

In most things there is an exception to the rule. When it comes to the rules about who is allowed to go to which websites in spite of the rules or in this case, policies, it seems that there are more exceptions than to most rules. There are numerous valid reasons and scenarios for exceptions so it follows that there needs to be a way to accommodate this exceptions.

## The different methods of override

There are actually two different ways to override web filtering behavior based on FortiGuard categorization of a websites. The second method has 2 variations in implementation and each of the three has a different level of granularity.

### 1. Using Alternate Categories

#### *Rating Override*

This method manually assigns a specific website to a different Fortinet category or a locally created category.

### 2. Using Alternate Profiles

#### *Administrative Override or Allow Blocked Override*

In this method all of the traffic going through the FortiGate unit, using identity based policies and a Web Filtering profile has the option where configured users or IP addresses can use an alternative Web Filter profile when attempting to access blocked websites.

## Using Alternate Categories

### Rating Overrides

There are two approached to overriding the FortiGuard Web Filtering. The first is an identity based method that can be configured using a combination of identity based policies and specifically designed webfilter profiles. This has been addressed in the Firewall Handbook.

The second method is the system wide approach that locally (on the FortiGate Firewall) reassigns a URL to a different FortiGuard Category and even subcategory. This is where you can set assign a specific URL to the FortiGuard Category that you want to you can also set the URL to one of the Local Categories that you have created

The Rating Overrides option is available because different people will have different criteria for how they categorize websites. Even is the criteria is the same an organization may have reason to block the bulk of a category but need to be able to access specific URLs that are assigned to that category.

A hypothetical example could be that a website, example.com is categorized as being in the Sub-Category Pornography. The law offices of Barrister, Solicitor and Lawyer do not want their employees looking at pornography at work so they have used the FortiGuard Webfilter to block access to sites that have been assigned to the Category "Pornography". However, the owners of example.com are clients of the law office and they are aware that example.com is for artists that specialize in nudes and erotic images. In this case to approaches can be taken. The first is that the Rating Override function can be used to assign example.com to Nudity and Risque instead of Pornography for the purposes of matching the criteria that the law office goes by or the site can be assigned to a Custom Category that is not blocked because the site belongs to one of their clients and they always want to be able to access the site.

Another hypothetical example from the other side of the coin. A private school has decided that a company that specializes in the online selling of books that could be considered inappropriate for children because of their violent subject matter, should not be accessible to anyone in the school. The categorization by Fortinet of the site example2.com is General Interest - Business with the subcategory of Shopping and Auction, which is a category that is allowed at the school. In this case they school could reassign the site to the Category Adult Material which is a blocked category.

## Local Categories

User-defined categories can be created to allow users to block groups of URLs on a per-profile basis. The categories defined here appear in the global URL category list when configuring a web filter profile. Users can rate URLs based on the local categories.

Users can create user-defined categories then specify the URLs that belong to the category. This allows users to block groups of web sites on a per profile basis. The ratings are included in the global URL list with associated categories and compared in the same way the URL block list is processed.

The local assignment of a category overrides the FortiGuard server ratings and appear in reports as “Local” Categories or “Custom” Categories depending on the context.

Local categories are configured in the CLI.

To create a Local Category:

```
config webfilter ftgd-local-cat
 edit local_category_1
 set id 140
 end
```

There is also a way to create a new category in the Web Based Manager.

Go to *Security Profiles > Web Filter > Rating Overrides*. In the process creating a new override, you can choose to override that URL to the *Category* “Custom Categories”. When you are choosing the *Sub-Category* from the drop down menu, the last item in the list will be *Create New*. If you select this item you will be given the option to fill out a field called *Category Name*. Just type in the name of the new category from this point on it will be listed in the *Custom Categories* of the *Rating Overrides* and as one of the Local Categories in the FortiGuard Webfilter.

## Configuring Rating Overrides

1. Go to *Security Profiles > Web Filter > Rating Overrides*.
2. Select Create New
3. Type in the URL field the URL of the Website that you wish to recategorize.
4. Select the Lookup Rating button to verify the current categorization that is being assigned to the URL.
5. Change the Category field to one of the more applicable options from the drop down menu.
6. Change the Sub-Category field to a more narrowly defined option within the main category.
7. Select OK.



It is usually recommended that you choose a category that you know will be addressed in existing Webfilter profiles so that you will not need to engage in further configuration.

---

## Using Alternate Profiles

### Allow Blocked Overrides or Web Overrides

Depending on which patch level you are on of Firmware Version 5 you will have one of 2 options for the implementation of the alternate profile approach to overriding a Web Filter profile that blocks access to a URL.



## Before FortiOS 5 Patch 4

The only feature available was referred to as "Allow Blocked Override". The configuration settings for this feature were found in the individual Web Filter profiles by going to the Advanced section near the bottom of the configuration window.

## Starting with FortiOS Patch 4

The Administrative Override feature for Web Filtering was added and is found by going to *Security Profiles > Web Filter > Web Overrides*. This opening window will display a listing of all of the overrides of this type. The editing window referred to the configuration as an Administrative Override.

## The Concept

When a Web filter profile is overridden it does not necessarily remove all control and restrictions that were previously imposed by the Web Filter. The idea is to replace a restrictive filter with a different one. In practice, it makes sense that this will likely be a profile that is less restrictive than the original one but there is nothing that forces this. The degree to which that the alternate profile is less restrictive is open. It can be as much as letting the user access everything on the Internet or as little as allowing only one additional website. The usual practice though is to have as few alternate profiles as are needed to allow approved people to access what they need during periods when an exception to the normal rules is needed but still having enough control that the organizations web usage policies are not compromised.

You are not restricted to having only one alternative profile as an option to the existing profile. The new profile depends on the credentials or IP address making the connection. For example, John connecting through the "Standard" profile could get the "Allow\_Streaming\_Video" profile while George would get the "Allow\_Social\_Networking\_Sites" profile.

The other thing to take into account is the time factor on these overrides. They are not indefinite. The longest that an override can be enabled is for 1 year less a minute. Often these overrides are set up for short periods of time for specific reasons such as a project. Having the time limitation means that the System Administrator does not have to remember to go back and turn the feature off after the project is finished.

## Identity or Address

In either case what these override features do is, for specified users, user groups or IP addresses, allow sites blocked by Web Filtering profiles to be overridden for a specified length of time. The drawback of this method of override is that it takes more planning and preparation than the rating override method. The advantage is that once this has been set up, this method requires very little in the way of administrative overhead to maintain.

When planning to use the alternative profile approach keep in mind the following: In Boolean terms, one of the following "AND" conditions has to be met before overriding the Web Filter is possible

### Based on the IP address:

- The Web Filter profile must be specified as allowing overrides
- AND the user's computer is one of the IP addresses specified
- AND the time is within the expiration time frame.

While the conditions are fewer for this situation there is less control over who has the ability to bypass the filtering configured for the site. All someone has to do is get on a computer that is allowed to override the Web Filter and they have access.

### **Based on user or user group:**

- The Web Filter profile must be specified as allowing overrides
- AND the policy the traffic is going through must be identity based
- AND the user's credentials matches the identity credentials specified
- AND the time is within the expiration time frame.

This method is the one most likely to be used as it gives more control in that the user has to have the correct credential and more versatile because the user can use the feature from any computer that uses the correct policy to get out on the Internet.

## **Settings**

When using an alternate profile approach to Web Filter overrides the following settings are used to determine authentication and outcome. Not every setting is used in both methods but enough of them are common to describe them collectively.

### **Apply to Group(s)**

This is found in the Allow Blocked Overrides configuration. Individual users can not be selected. You can select one or more of the User Groups that are recognized by the FortiGate unit, whether they are local to the system or from a third part authentication device such as a AD server through FSSO.

### **Original Profile**

This is found in the Administrative Override configuration. In the Allow Blocked Overrides setting the configuration is right inside the profile so there was no need to specify which profile was the original one, but the Administrative Override setup is done separately from the profiles themselves.

### **Assign to Profile or New Profile**

Despite the difference in the name of the field, this is the same thing in both variations of the feature. You select from the drop down menu the alternate Web Filter Profile that you wish to set up for this override.

### **Scope or Scope Range**

When setting up the override in the "Allow Blocked Overrides" variation you are given a drop down menu next to the field name Scope while in the Administrative Override configuration you are asked to select a radio button next to the same options. In both cases this is just a way of selecting which form of credentials will be required to approve the overriding of the existing Web Filter profile.

When the Web Filter Block Override message page appears it will display a field named "Scope:" and depending on the selection, it will show the type of credentials used to determine whether or not the override is allowed. The available options are:

#### **User**

This means that the authentication for permission to override will be based on whether or not the user is using a specific user account.

#### **User Group**

This means that the authentication for permission to override will be based on whether or not the user account supplied as a credential is a member of the specified User Group.

#### **IP**

This means that the authentication for permission to override will be based on the IP address of the computer that was used to authenticate. This would be used with computers that have multiple users. Example: If Paul logs on to the computer, engages the override using

his credentials and then logs off, if the scope was based on the IP address of the computer, anybody logging in with any account on that computer would now be using the alternate override Web Filter profile.

When entering an IP address in the Administrative Override version, only individual IP addresses are allowed.

#### **Differences between IP and Identity based scope**

- Using the IP scope does not require the use of an Identity based policy.
- When using the Administrative Override variation and IP scope, you may not see a warning message when you change from using the original Web Filter profile to using the alternate profile. There is no requirement for credentials from the user so, if allowed, the page will just come up in the browser.

#### **Ask**

This option is available only in the "Allowed Blocked Overrides" variation and when used configures the message page to ask which scope the user wished to use. Normally, when the page appears the scope options are greyed out and not editable, but by using the ask option the option is dark and the user can choose from the choice of:

- User
- User Group
- IP Address

#### **Duration Mode**

This option is available only in the "Allowed Blocked Overrides" variation. The Administrative Override sets a specified time frame that is always used for that override. The available options from the drop down menu are:

##### **Constant**

Using this setting will mean that what ever is set as the duration will be the length of time that the override will be in effect. If the Duration variable is set to 15 minutes the length of the override will always be 15 minutes. The option will be visible in the Override message page but the setting will be greyed out.

##### **Ask**

Using this setting will give the person the option of setting the duration to the override when it is engaged. The duration time which is greyed out if the Constant setting is used will be dark and editable. The user can set the duration in terms of Day, Hours and or Minutes.

#### **Duration**

Duration is one of the areas where the two variations take a different approach, on two aspects of the setting. As already indicated the "Administrative Override" only uses a static time frame there is no option for the user to select on the fly how long it will last. The other way in which the two variations differ is that the "Allow Blocked Overrides" starts the clock when the user logs in with his credentials. For example, if the duration is 1 hour and John initiates an override at 2:00 p.m. on January 1, at the end of that hour he will revert back to using the original profile but he can go back and re-authenticate and start the process over again. The Administrative override variation starts the clock from when the override was configured, which is why it shows an expiration date and time when you are configuring it.

This option, which is available when the Duration Mode is set to Constant is the time in minutes that the override will last when engaged by the user.

When setting up a constant duration in the Web Based Interface, minutes is the only option for units of time. To set a longer time frame or to use the units of hours or days you can use the CLI.

```
config webfilter profile
 edit <name of webfilter profile>
 config override
 set ovr-dur <###d##h##m>
 end
```

When configuring the duration you don't have to set a value for a unit you are not using. If you are not using days or hours you can use

```
set ovr-dur 30m
```

instead of

```
set ovr-dur 0d0h30m
```

However, each of the units of time variable has their own maximum level

```
###d cannot be more than 364
##h cannot be more than 23
##m cannot be more than 59
```

So the maximum length that the override duration can be set to is 364 days, 23 hours, and 59 minutes(a minute shy of 1 year)

## SafeSearch

SafeSearch is a feature of popular search sites that prevents explicit web sites and images from appearing in search results. Although SafeSearch is a useful tool, especially in educational environments, the resourceful user may be able to simply turn it off. Enabling SafeSearch for the supported search sites enforces its use by rewriting the search URL to include the code to indicate the use of the SafeSearch feature. For example, on a Google search it would mean adding the string “&safe=active” to the URL in the search.

The search sites supported are:

- Google
- Yahoo
- Bing
- Yandex

### Enabling SafeSearch – CLI

```
config webfilter profile
 edit default
 config web
 set safe-search url
 end
 end
```

This enforces the use of SafeSearch in traffic controlled by the firewall policies using the web filter you configure.

## YouTube Education Filter

YouTube for Schools is a way to access educational videos from inside a school network. This YouTube feature gives schools the ability to access a broad set of educational videos on the YouTube EDU channel and to select the specific videos that are accessible from within the school network.

Before using this feature you must create a google account for your school and sign up for YouTube for Schools. The Youtube Schools account provides a unique ID registration key. This ID becomes part of the filter that is used to access YouTube EDU's educational videos in schools even if other YouTube content is blocked by the policy.

More details can be found by going to:

<http://www.youtube.com/schools>

### Enabling YouTube Education Filter in CLI

```
config webfilter profile
 edit default
 config web
 set safe-search url header
 set youtube-edu-filter-id ABCD1234567890abcdef
 end
 end
```

## Deep Scanning Restrictions

This section doesn't have a label such as "Deep Scanning Restrictions" but there are two settings in the profile that relate to the topic. In the profile, they appear as:

- *Enable HTTPS URL Scan Only*
- *Categories Exempt from Deep Scanning...*

### Enable HTTPS URL Scan Only

When Deep Scanning is turned on traffic that is encrypted using SSL is scanned for issues just as unencrypted traffic is. However, scanning encrypted traffic puts a larger load on the resources of the FortiGate unit.

Even if the scanning of the contents of the traffic is not a requirement many administrator prefer to scan the URLs being sent over HTTPS so that users cannot bypass the blocking of access to a site by putting "https://" as a prefix to a URL. The setting restricts the deep scanning of the traffic to the URL destination which is in the header. This way the resources tied up in decrypting the traffic are minimized, yet the administrator can still enforce policy regarding access to prohibited websites

### Categories Exempt from Deep Scanning

For the purposes of personal privacy, there are 3 categories that can be exempted from deep scanning by the FortiGate unit. They are Banking, Health Care and Personal Privacy.

When HTTPS URL Scan Only is enabled you will notice that the option to exclude these categories from deep scanning is removed. This is because if only the URL is being scanned then the contents of the traffic is not being scanned anyway so there is no need to exclude it.

## Web Site Filter

You can allow or block access to specific URLs by adding them to the Web Site Filter list. You add the URLs by using patterns containing text and regular expressions. The FortiGate unit allows or blocks web pages matching any specified URLs or patterns and displays a replacement message instead.



URL blocking does not block access to other services that users can access with a web browser. For example, URL blocking does not block access to ftp://ftp.example.com. Instead, use firewall policies to deny ftp connections.

When adding a URL to the URL filter list, follow these rules:

- Type a top-level URL or IP address to control access to all pages on a web site. For example, `www.example.com` or `192.168.144.155` controls access to all pages at this web site.
- Enter a top-level URL followed by the path and file name to control access to a single page on a web site. For example, `www.example.com/news.html` or `192.168.144.155/news.html` controls access to the news page on this web site.
- To control access to all pages with a URL that ends with `example.com`, add `example.com` to the filter list. For example, adding `example.com` controls access to `www.example.com`, `mail.example.com`, `www.finance.example.com`, and so on.
- Control access to all URLs that match patterns using text and regular expressions (or wildcard characters). For example, `example.*` matches `example.com`, `example.org`, `example.net` and so on.



URLs with an action set to exempt or monitor are not scanned for viruses. If users on the network download files through the FortiGate unit from a trusted web site, add the URL of this web site to the URL filter list with an action to pass it so the FortiGate unit does not virus scan files downloaded from this URL.

### URL formats

When adding a URL to the URL filter list, follow these rules:

#### How URL formats are detected when using HTTPS

If your unit does not support SSL content scanning and inspection or if you have selected the *URL filtering* option in web content profile for *HTTPS content filtering mode* under *Protocol Recognition*, filter HTTPS traffic by entering a top level domain name, for example, `www.example.com`. HTTPS URL filtering of encrypted sessions works by extracting the CN from the server certificate during the SSL negotiation. Since the CN only contains the domain name of the site being accessed, web filtering of encrypted HTTPS sessions can only filter by domain names.

If your unit supports SSL content scanning and inspection and if you have selected Deep Scan, you can filter HTTPS traffic in the same way as HTTP traffic.

#### How URL formats are detected when using HTTP

URLs with an action set to exempt are not scanned for viruses. If users on the network download files through the unit from trusted web site, add the URL of this web site to the URL

filter list with an action set to exempt so the unit does not virus scan files downloaded from this URL.

- Type a top-level URL or IP address to control access to all pages on a web site. For example, `www.example.com` or `192.168.144.155` controls access to all pages at this web site.
- Enter a top-level URL followed by the path and filename to control access to a single page on a web site. For example, `www.example.com/news.html` or `192.168.144.155/news.html` controls the news page on this web site.
- To control access to all pages with a URL that ends with `example.com`, add `example.com` to the filter list. For example, adding `example.com` controls access to `www.example.com`, `mail.example.com`, `www.finance.example.com`, and so on.
- Control access to all URLs that match patterns created using text and regular expressions (or wildcard characters). For example, `example.*` matches `example.com`, `example.org`, `example.net` and so on.
- Fortinet URL filtering supports standard regular expressions.



If virtual domains are enabled on the unit, web filtering features are configured globally. To access these features, select *Global Configuration* on the main menu.

---

## Web Site Filter actions

You can select one of four actions for how traffic will be treated as it attempts to reach a site in the list.

### Block

Attempts to access any URLs matching the URL pattern are denied. The user will be presented with a replacement message.

### Allow

Any attempt to access a URL that matches a URL pattern with an allow action is permitted. The traffic is passed to the remaining antivirus proxy operations, including FortiGuard Web Filter, web content filter, web script filters, and antivirus scanning.

Allow is the default action. If a URL does not appear in the URL list, it is permitted.

### Monitor

Traffic to, and reply traffic from, sites matching a URL pattern with a monitor be allowed through in the same way as the “Allow” action. The difference with the Monitor action being that a log message will be generated each time a matching traffic session is established. The requests will also be subject to all other Security Profiles inspections that would normally be applied to the traffic.

### Exempt

*Exempt* allows trusted traffic to bypass the antivirus proxy operations, but it functions slightly differently. In general, if you’re not certain that you need to use the *Exempt* action, use *Monitor*.

HTTP 1.1 connections are persistent unless declared otherwise. This means the connections will remain in place until closed or the connection times out. When a client loads a web page, the client opens a connection to the web server. If the client follows a link to another page on

the same site before the connection times out, the same connection is used to request and receive the page data.

When you add a URL pattern to a URL filter list and apply the *Exempt* action, traffic sent to and replies traffic from sites matching the URL pattern will bypass all antivirus proxy operations. The connection itself inherits the exemption. This means that all subsequent reuse of the existing connection will also bypass all antivirus proxy operations. When the connection times out, the exemption is cancelled.

For example, consider a URL filter list that includes `example.com/files` configured with the *Exempt* action. A user opens a web browser and downloads a file from the URL `example.com/sample.zip`. This URL does not match the URL pattern so it is scanned for viruses. The user then downloads `example.com/files/beautiful.exe` and since this URL does match the pattern, the connection itself inherits the exempt action. The user then downloads `example.com/virus.zip`. Although this URL does not match the exempt URL pattern, a previously visited URL did, and since the connection inherited the exempt action and was re-used to download a file, the file is not scanned.

If the user next goes to an entirely different server, like `example.org/photos`, the connection to the current server cannot be reused. A new connection to `example.org` is established. This connection is not exempt. Unless the user goes back to `example.com` before the connection to that server times out, the server will close the connection. If the user returns after the connection is closed, a new connection to `example.com` is created and it is not exempt until the user visits a URL that matches the URL pattern.

Web servers typically have short time-out periods. A browser will download multiple components of a web page as quickly as possible by opening multiple connections. A web page that includes three photos will load more quickly if the browser opens four connections to the server and downloads the page and the three photos at the same time. A short time-out period on the connections will close the connections faster, allowing the server to avoid unnecessarily allocating resources for a long period. The HTTP session time-out is set by the server and will vary with the server software, version, and configuration.

Using the *Exempt* action can have unintended consequences in certain circumstances. You have a web site at `example.com` and since you control the site, you trust the contents and configure `example.com` as exempt. But `example.com` is hosted on a shared server with a dozen other different sites, each with a unique domain name. Because of the shared hosting, they also share the same IP address. If you visit `example.com`, your connection your site becomes exempt from any antivirus proxy operations. Visits to any of the 12 other sites on the same server will reuse the same connection and the data you receive is exempt from scanned.

Use of the *Exempt* action is not suitable for configuration in which connections through the FortiGate unit use an external proxy. For example, you use `proxy.example.net` for all outgoing web access. Also, as in the first example, URL filter list that includes a URL pattern of `example.com/files` configured with the *Exempt* action. Users are protected by the antivirus protection of the FortiGate unit until a user visits a URL that matches the of `example.com/files` URL pattern. The pattern is configured with the *Exempt* action so the connection to the server inherits the exemption. With a proxy however, the connection is from the user to the proxy. Therefore, the user is entirely unprotected until the connection times out, no matter what site he visits.

Ensure you are aware of the network topology involving any URLs to which you apply the *Exempt* action.

## Status

The Web Site Filter has the option to either enable or disable individual web sites in the list. This allows for the temporary removal of the actions against a site so that it can be later reengaged without having to rewrite the configuration.



## Configuring a Web Site Filter

### To create a URL Filter list

1. Go to *Security Profiles > Web Filter > Profiles*.
2. Select the Web Filter Profile that you wish to add the Web Site Filter to.
3. About half way down the Edit Web Filter Profile page check the box next to Enable Web Site Filter.
4. Select *Create New*.
5. Enter a URL for the website.
6. Enter optional comments to describe it.
7. Select *OK*.

## Configuring a URL filter list

Each URL filter list can have up to 5000 entries. For this example, the URL `www.example*.com` will be used. You configure the list by adding one or more URLs to it.

### To add a URL to a URL filter list

1. Go to *Security Profiles > Web Filter > URL Filter*.
2. Select an existing list and choose *Edit*.
3. Select *Create New*.
4. Enter the URL, without the “http”, for example: `www.example*.com`.
5. Select a *Type: Simple, Wildcard or Regular Expression*.
6. In this example, select *Wildcard*.
7. Select the *Action* to take against matching URLs: *Exempt, Block, Allow, or Monitor*.
8. Select *Enable*.
9. Select *OK*.

## Web content filter

You can control web content by blocking access to web pages containing specific words or patterns. This helps to prevent access to pages with questionable material. You can also add words, phrases, patterns, wild cards and Perl regular expressions to match content on web pages. You can add multiple web content filter lists and then select the best web content filter list for each web filter profile.

Enabling web content filtering involves three separate parts of the FortiGate configuration.

- The security policy allows certain network traffic based on the sender, receiver, interface, traffic type, and time of day.
- The web filter profile specifies what sort of web filtering is applied.
- The web content filter list contains blocked and exempt patterns.

The web content filter feature scans the content of every web page that is accepted by a security policy. The system administrator can specify banned words and phrases and attach a numerical value, or score, to the importance of those words and phrases. When the web content filter scan detects banned content, it adds the scores of banned words and phrases in the page. If the sum is higher than a threshold set in the web filter profile, the FortiGate unit blocks the page.

## General configuration steps

Follow the configuration procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

1. Create a web content filter list.
2. Add patterns of words, phrases, wildcards, and regular expressions that match the content to be blocked or exempted.
3. You can add the patterns in any order to the list. You need to add at least one pattern that blocks content.
4. In a web filter profile, enable the web content filter and select a web content filter list from the options list.

To complete the configuration, you need to select a security policy or create a new one. Then, in the security policy, enable *Webfilter* and select the appropriate web filter profile from the list.

## Creating a web filter content list

You can create multiple content lists and then select the best one for each web filter profile. Creating your own web content lists can be accomplished only using the CLI.

This example shows how to create a web content list called inappropriate language, with two entries, offensive and rude.

### To create a web filter content list

```
config webfilter content
 edit 3
 set name "inappropriate language"
 config entries
 edit offensive
 set action block
 set lang western
 set pattern-type wildcard
 set score 15
 set status enable
 next
 edit rude
 set action block
 set lang western
 set pattern-type wildcard
 set score 5
 set status enable
 end
 end
end
```

## How content is evaluated

Every time the web content filter detects banned content on a web page, it adds the score for that content to the sum of scores for that web page. You set this score when you create a new pattern to block the content. The score can be any number from zero to 99999. Higher scores indicate more offensive content. When the sum of scores equals or exceeds the threshold score, the web page is blocked. The default score for web content filter is 10 and the default

threshold is 10. This means that by default a web page is blocked by a single match. Blocked pages are replaced with a message indicating that the page is not accessible according to the Internet usage policy.

Banned words or phrases are evaluated according to the following rules:

- The score for each word or phrase is counted only once, even if that word or phrase appears many times in the web page.
- The score for any word in a phrase without quotation marks is counted.
- The score for a phrase in quotation marks is counted only if it appears exactly as written.

The following table describes how these rules are applied to the contents of a web page. Consider the following, a web page that contains only this sentence: “The score for each word or phrase is counted only once, even if that word or phrase appears many times in the web page.”

**Table 95:** Banned Pattern Rules

Banned pattern	Assigned score	Score added to the sum for the entire page	Threshold score	Comment
word	20	20	20	Appears twice but only counted once. Web page is blocked.
word phrase	20	40	20	Each word appears twice but only counted once giving a total score of 40. Web page is blocked
word sentence	20	20	20	“word” appears twice, “sentence” does not appear, but since any word in a phrase without quotation marks is counted, the score for this pattern is 20. Web page is blocked.
“word sentence”	20	0	20	“This phrase does not appear exactly as written. Web page is allowed.
“word or phrase”	20	20	20	This phrase appears twice but is counted only once. Web page is blocked.

## Enabling the web content filter and setting the content threshold

When you enable the web content filter, the web filter will block any web pages when the sum of scores for banned content on that page exceeds the content block threshold. The threshold will be disregarded for any exemptions within the web filter list.

### To enable the web content filter and set the content block threshold

1. Go to *Security Profiles > Web Filter > Profiles*.
2. Select the *Create New* icon on the Edit Web Filter Profile window title bar.

3. In the *Name* field, enter the name of the new web filter profile.
4. Optionally, you may also enter a comment. The comment can remind you of the details of the sensor.
5. Select the *Inspection Method*.  
Proxy-based detection involves buffering the file and examining it as a whole. Advantages of proxy-based detection include a more thorough examination of attachments, especially archive formats and nesting.  
Flow-based detection examines the file as it passes through the FortiGate unit without any buffering. Advantages of flow-based detection include speed and no interruption of detection during conserve mode.
6. Expand the *Advanced Filter* heading.
7. Enable *Web Content Filter*.
8. Select the required web filter content list from the *Web Content Filter* drop-down list.
9. Select *Apply*.

The web filter profile configured with web content filtering is ready to be added to a firewall profile.

## Advanced web filter configurations

### Allow websites when a rating error occurs

Enable to allow access to web pages that return a rating error from the FortiGuard Web Filter service.

If your FortiGate unit cannot contact the FortiGuard service temporarily, this setting determines what access the FortiGate unit allows until contact is re-established. If enabled, users will have full unfiltered access to all web sites. If disabled, users will not be allowed access to any web sites.

### ActiveX filter

Enable to filter ActiveX scripts from web traffic. Web sites using ActiveX may not function properly with this filter enabled.

### Block HTTP redirects by rating

Enable to block HTTP redirects.

Many web sites use HTTP redirects legitimately but in some cases, redirects may be designed specifically to circumvent web filtering, as the initial web page could have a different rating than the destination web page of the redirect.

This option is not supported for HTTPS.

### Block Invalid URLs

Select to block web sites when their SSL certificate CN field does not contain a valid domain name.

FortiGate units always validate the CN field, regardless of whether this option is enabled. However, if this option is not selected, the following behavior occurs:

- If the request is made directly to the web server, rather than a web server proxy, the FortiGate unit queries for FortiGuard Web Filtering category or class ratings using the IP address only, not the domain name.
- If the request is to a web server proxy, the real IP address of the web server is not known. Therefore, rating queries by either or both the IP address and the domain name is not reliable. In this case, the FortiGate unit does not perform FortiGuard Web Filtering.

## Cookie filter

Enable to filter cookies from web traffic. Web sites using cookies may not function properly with this enabled.

## Provide Details for Blocked HTTP 4xx and 5xx Errors

Enable to have the FortiGate unit display its own replacement message for 400 and 500-series HTTP errors. If the server error is allowed through, malicious or objectionable sites can use these common error pages to circumvent web filtering.

## HTTP POST action

Select the action to take with HTTP POST traffic. HTTP POST is the command used by your browser when you send information, such as a form you have filled-out or a file you are uploading, to a web server.

The available actions include:

<b>Normal</b>	Allow use of the HTTP POST command as normal.
<b>Comfort</b>	<p>Use client comforting to slowly send data to the web server as the FortiGate unit scans the file. Use this option to prevent a server time-out when scanning or other filtering is enabled for outgoing traffic.</p> <p>The client comforting settings used are those defined in the Proxy Options profile selected in the security policy. For more information, see <a href="#">“Configuring client comforting” on page 2040</a>.</p>
<b>Block</b>	<p>Block the HTTP POST command. This will limit users from sending information and files to web sites.</p> <p>When the post request is blocked, the FortiGate unit sends the http-post-block replacement message to the web browser attempting to use the command.</p>

## Java applet filter

Enable to filter java applets from web traffic. Web sites using java applets may not function properly with this filter enabled.

## Rate Images by URL

Enable to have the FortiGate retrieve ratings for individual images in addition to web sites. Images in a blocked category are not displayed even if they are part of a site in an allowed category.

Blocked images are replaced on the originating web pages with blank place-holders. Rated image file types include GIF, JPEG, PNG, BMP, and TIFF.

## Rate URLs by Domain and IP Address

Enable to have the FortiGate unit request the rating of the site by URL and IP address separately, providing additional security against attempts to bypass the FortiGuard Web Filter.

If the rating determined by the domain name and the rating determined by the IP address defer the Action that is enforce will be determined by a weighting assigned to the different categories. The higher weighted category will take precedence in determining the action. This will have the side effect that sometimes the Action will be determined by the classification based on the domain name and other times it will be determined by the classification that is based on the IP address.



FortiGuard Web Filter ratings for IP addresses are not updated as quickly as ratings for URLs. This can sometimes cause the FortiGate unit to allow access to sites that should be blocked, or to block sites that should be allowed.

---

An example of how this would work would be if a URL's rating based on the domain name indicated that it belonged in the category Lingerie and Swimsuit, which is allowed but the category assigned to the IP address was Pornography which has an action of Block, because the Pornography category has a higher weight the effective action is Block.

## Web resume download block

Enable to prevent the resumption of a file download where it was previously interrupted. With this filter enabled, any attempt to restart an aborted download will download the file from the beginning rather than resuming from where it left off.

This prevents the unintentional download of viruses hidden in fragmented files.

Note that some types of files, such as PDF, fragment files to increase download speed and enabling this option can cause download interruptions. Enabling this option may also break certain applications that use the Range Header in the HTTP protocol, such as YUM, a Linux update manager.

## Working with the Interface

In order to find out the status of your configuration it helps to understand the interface. In the Web filtering section there are a number of pages that you will need to be able to read.

### Profile page

Lists each web filter profile that you created. On this page, you can edit, delete or create a new web filter profile. You are redirected to this page when you select *View List* on the Edit Web Filter Profile page.

**Note:** Web filtering overrides are profile-based, allowing a rule to be created that changes the web filter profile that applies to a user. An override link appears in all related blocked pages. This is available only in the CLI.

## New Web Filter Profile page

Provides settings for configuring a web filter profile. Advanced features, such as web content filtering and FortiGuard web filtering, is configured in the CLI.

This page appears when you select *Create New* on the Edit Web Filter Profile page. If you are on the Profile page, and you select *Create New*, you will be redirected to the New Web Filter Profile page.

**Note:** Logging is enabled in the CLI.

The following explains the web filtering options in the Web Filtering menu. If your unit supports SSL content scanning and inspection you can also configure web filtering for HTTPS traffic.

If you want to configure advanced settings, such as web content filter, you must configure them within the CLI. Advanced settings also includes overrides.

This topic includes the following:

- [Profile](#)
- [URL Filter](#)
- [Rating Overrides](#)

## Profile

The Profile menu allows you to configure a web filter profile to apply to a firewall policy. A profile is specific information that defines how the traffic within a policy is examined and what action may be taken based on the examination.

### Web profile configuration settings

The following are web filter profile configuration settings in *Security Profiles > Web Filter > Profiles*. If you want to configure advanced settings, such as FortiGuard web filtering overrides, you must configure these settings within the CLI.

---

#### Profile page

Lists each web filter profile that you created. On this page, you can edit, delete or create a new web filter profile. You are redirected to this page when you select *View List* on the Edit Web Filter Profile page.

**Note:** Web filtering overrides are profile-based, allowing a rule to be created that changes the web filter profile that applies to a user. An override link appears in all related blocked pages. This is available only in the CLI.

---

<b>Create New</b>	Creates a new web filter profile. When you select <i>Create New</i> , you are automatically redirected to the New Web Filter Profile page.
<b>Edit</b>	Modifies settings within a web filter profile. When you select <i>Edit</i> , you are automatically redirected to the Edit Web Filter Profile page.

---

<b>Delete</b>	<p>Removes a web filter profile from within the list on the Profile page.</p> <p>To remove multiple web filter profiles from within the list, on the Profile page, in each of the rows of the file filter lists you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all web filter profiles from the list, on the Profile page, select the check box in the check box column and then select <i>Delete</i>.</p>
<b>Name</b>	The name of the web filter profile.
<b>Comments</b>	A description given to the web filter profile. This is an optional setting.
<b>Ref.</b>	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page ( <i>Security Profiles &gt; Antivirus &gt; Profile</i>), 1 appears in <i>Ref.</i>.</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> <li>• <b>View the list page for these objects</b> – automatically redirects you to the list page where the object is referenced at.</li> <li>• <b>Edit this object</b> – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when this icon is selected, the user is redirected to the Edit Policy page.</li> <li>• <b>View the details for this object</b> – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.</li> </ul>

### New Web Filter Profile page

Provides settings for configuring a web filter profile. Advanced features, such as web content filtering and FortiGuard web filtering, is configured in the CLI.

This page appears when you select *Create New* on the Edit Web Filter Profile page. If you are on the Profile page, and you select *Create New*, you will be redirected to the New Web Filter Profile page.

**Note:** Logging is enabled in the CLI.

<b>Name</b>	<p>Enter a name for the web filter profile.</p> <p>If you want to edit the name at any time, select the profile and enter a new name in the <i>Name</i> field. Select <i>Apply</i> to save the change.</p>
-------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------



<b>Comments</b>	<p>Enter a description for the web filter profile. This is optional.</p> <p>If you want to edit the description at any time, select the profile and enter the new description in the <i>Comments</i> field. Select <i>Apply</i> to save the change.</p>
<b>Inspection mode</b>	<p>Select to enable either flow-based web filtering or proxy-based.</p> <p>Flow-based web filtering is a non-proxy solution, which provides high concurrent session, high session rate, and low-latency web filtering service.</p>
<b>FortiGuard Categories</b>	<p>A list of FortiGuard category groups and categories that are used to rate web sites. Selecting a category group will automatically select all of the categories within the group. For example, if you select Security Risk, you can see that all of the categories within are selected if you expand the group. You can however, select or deselect categories within groups as required.</p>
<b>Show</b>	<p>Select an action to view all of the categories that are currently configured with the selected action.</p>
<b>Change Action for Selected Categories to</b>	<p>Select an action, and all of the selected categories will have the selected action applied. Selected category groups will have the action applied to all categories within the group.</p>
<b>Quota on Categories</b>	<p>Users can have their web browsing time limited by category through the use of quotas. Quotas can be applied only to categories that are configured with the Monitor action.</p> <p>If you create a quota for a single category, every authenticated user subject to the security policy in which the web filter profile is applied is limited in browsing web sites in the category to the duration you specify. If you create a single quota that includes multiple categories, the quota will apply to the categories as a whole.</p> <p>Quotas are ignored for unauthenticated users. To enforce quotas, configure the security policy to require authentication.</p>
<b>Enable Safe Search (Support Search Engines: Google, Yahoo and Bing)</b>	<p>When enabled, the supported search engines exclude offensive material from search results.</p>
<b>HTTPS Scanning</b>	<p>Available only on models that support HTTPS.</p> <p>Select to have all of the web filtering specified in the web filter profile to HTTPS traffic as well as HTTP traffic.</p>
<b>Advanced Filter</b>	<p>Expand this heading for advanced web filtering options.</p>
<b>Web URL Filter</b>	<p>Enable to block access to URLs listed in the selected URL list.</p>

<b>Web Resume Download Block</b>	<p>Enable to prevent the resumption of a file download where it was previously interrupted. With this filter enabled, any attempt to restart an aborted download will download the file from the beginning rather than resuming from where it left off.</p> <p>This prevents the unintentional download of viruses hidden in fragmented files.</p> <p>Note that some types of files, such as PDF, fragment files to increase download speed and enabling this option can cause download interruptions. Enabling this option may also break certain applications that use the Range Header in the HTTP protocol, such as YUM, a Linux update manager.</p>
<b>Block Invalid URLs</b>	<p>Select to block web sites when their SSL certificate CN field does not contain a valid domain name.</p> <p>FortiGate units always validate the CN field, regardless of whether this option is enabled. However, if this option is not selected, the following behavior occurs:</p> <ul style="list-style-type: none"> <li>• If the request is made directly to the web server, rather than a web server proxy, the FortiGate unit queries for FortiGuard Web Filtering category or class ratings using the IP address only, not the domain name.</li> <li>• If the request is to a web server proxy, the real IP address of the web server is not known. Therefore, rating queries by either or both the IP address and the domain name is not reliable. In this case, the FortiGate unit does not perform FortiGuard Web Filtering.</li> </ul>
<b>HTTP POST Action</b>	<p>Select the action to take with HTTP POST traffic. HTTP POST is the command used by your browser when you send information, such as a form you have filled-out or a file you are uploading, to a web server.</p> <p>The available actions include:</p> <ul style="list-style-type: none"> <li>• Normal: Allow use of the HTTP POST command as normal.</li> <li>• Comfort: Use client comforting to slowly send data to the web server as the FortiGate unit scans the file. Use this option to prevent a server time-out when scanning or other filtering is enabled for outgoing traffic. The client comforting settings used are those defined in the Proxy Options profile selected in the security policy.</li> <li>• Block: Block the HTTP POST command. This will limit users from sending information and files to web sites. When the post request is blocked, the FortiGate unit sends the http-post-block replacement message to the web browser attempting to use the command.</li> </ul>
<b>Remove Java Applet Filter</b>	<p>Enable to filter java applets from web traffic. Web sites using java applets may not function properly with this filter enabled.</p>
<b>Remove ActiveX Filter</b>	<p>Enable to filter ActiveX scripts from web traffic. Web sites using ActiveX may not function properly with this filter enabled.</p>

<b>Remove Cookie Filter</b>	Enable to filter cookies from web traffic. Web sites using cookies may not function properly with this enabled.
<b>Search Engine Keyword Filter</b>	Enter the keywords that you want to monitor when users enter those same or similar keywords during a search within the supported search engines.
<b>Web Content Filter</b>	Enable to block access to web pages that include the words included in the selected web content filter list.
<b>Provide Details for Blocked HTTP 4xx and 5xx Errors</b>	Enable to have the FortiGate unit display its own replacement message for 400 and 500-series HTTP errors. If the server error is allowed through, malicious or objectionable sites can use these common error pages to circumvent web filtering.
<b>Rate Images by URL (Blocked images will be replaced with blanks)</b>	<p>Enable to have the FortiGate retrieve ratings for individual images in addition to web sites. Images in a blocked category are not displayed even if they are part of a site in an allowed category.</p> <p>Blocked images are replaced on the originating web pages with blank place-holders. Rated image file types include GIF, JPEG, PNG, BMP, and TIFF.</p>
<b>Allow Websites When a Rating Error Occurs</b>	<p>Enable to allow access to web pages that return a rating error from the FortiGuard Web Filter service.</p> <p>If your FortiGate unit cannot contact the FortiGuard service temporarily, this setting determines what access the FortiGate unit allows until contact is re-established. If enabled, users will have full unfiltered access to all web sites. If disabled, users will not be allowed access to any web sites.</p>
<b>Rate URLs by Domain and IP Address</b>	<p>Enable to have the FortiGate unit request the rating of the site by URL and IP address separately, providing additional security against attempts to bypass the FortiGuard Web Filter.</p> <p>FortiGuard Web Filter ratings for IP addresses are not updated as quickly as ratings for URLs. This can sometimes cause the FortiGate unit to allow access to sites that should be blocked, or to block sites that should be allowed.</p>
<b>Block HTTP Redirects by Rating</b>	<p>Enable to block HTTP redirects.</p> <p>Many web sites use HTTP redirects legitimately but in some cases, redirects may be designed specifically to circumvent web filtering, as the initial web page could have a different rating than the destination web page of the redirect.</p> <p>This option is not supported for HTTPS.</p>

## URL Filter

Allow or block access to specific URLs by adding them to the URL filter list. Add patterns using text and regular expressions (or wildcard characters) to allow or block URLs. The unit allows or blocks web pages matching any specified URLs or patterns and displays a replacement message.

You can add multiple URL filter lists and then select the best URL filter list for each profile.

You can add the following to block or exempt URLs:

- complete URLs
- IP addresses
- partial URLs to allow or block all sub-domains

Each URL filter list can have up to 5000 entries.

### URL filter configuration settings

The following are URL filter configuration settings in *Security Profiles > Web Filter > URL Filter*.



URL blocking does not block access to other services that users can access with a web browser. For example, URL blocking does not block access to `ftp://ftp.example.com`. Instead, use firewall policies to deny FTP connections.

---

#### URL Filter page

Lists each URL filter that you created. On this page, you can edit, delete or create a new URL filter.

---

**Create New** Creates a new URL filter list. When you select *Create New*, you are automatically redirected to the New List page. This page provides a name field and comment field. You must enter a name to go to the URL Filter Settings page.

---

**Edit** Modifies settings within a URL filter list. When you select *Edit*, you are automatically redirected to the URL Filter Settings page.

---

**Delete** Removes the URL filter list from the list on the URL Filter page. The *Delete* icon is only available if the URL filter list is not selected in any profiles.

To remove multiple URL filter list from within the list, on the URL Filter page, in each of the rows of the file filter lists you want removed, select the check box and then select *Delete*.

To remove all URL filter list from the list, on the URL Filter page, select the check box in the check box column and then select *Delete*.

---

**Name** The available URL filter lists.

---

**# Entries** The number of URL patterns in each URL filter list.

---

**MMS Profiles (FortiOS Carrier only)** The name of the MMS profile

---

**Comments** Optional description of each URL filter list.

---

<b>Ref.</b>	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page ( <i>Security Profiles &gt; AntiVirus &gt; Profiles</i>), 1 appears in <i>Ref.</i>.</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> <li>• <b>View the list page for these objects</b> – automatically redirects you to the list page where the object is referenced at.</li> <li>• <b>Edit this object</b> – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when this icon is selected, the user is redirected to the Edit Policy page.</li> <li>• <b>View the details for this object</b> – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.</li> </ul>
-------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

---

### URL Filter Settings page

Provides settings for configuring URLs that make up the URL filter, and also lists the URLs that you created. You are automatically redirected to this page from the New List Page. If you are editing a URL filter, you are automatically redirected to this page.

<b>Name</b>	If you are editing an existing URL filter setting and want to change the name, enter a new name in this field. You must select <i>OK</i> to save the change.
<b>Comments</b>	If you are editing an existing URL filter setting and want to change or add a description, enter the new text in this field. You must select <i>OK</i> to save these changes.
<b>Create New</b>	Adds a URL address and filter settings to the list. When you select <i>Create New</i> , you are automatically redirected to the New URL Filter list.
<b>Edit</b>	Modifies the settings within a URL filter.
<b>Delete</b>	<p>Removes an entry from the list.</p> <p>To remove multiple URL filters from within the list, on the URL Filter Settings page, in each of the rows of the filters you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all URL filters from the list, on the URL Filter Settings page, select the check box in the check box column and then select <i>Delete</i>.</p>
<b>Enable</b>	Enables a filter in the list.
<b>Disable</b>	Disables a filter in the list.

<b>Move To</b>	Moves the URL to any position in the list. When you select <i>Move To</i> , the Move URL Filter window appears.  To move a URL, select the new position <i>Before</i> or <i>After</i> , which will place the current URL entry before or after the entry you enter in the ( <i>URL</i> ) field. For example, 1example.com is being moved after 3example.com, so 3example.com is entered in the ( <i>URL</i> ) field.
<b>Remove All Entries</b>	Removes all filter entries within the list on the URL Filter Settings page.
<b>Enable</b>	Indicates whether the URL is enable or disabled. A green check mark indicates that the URL is enabled; a gray check mark indicates that the URL is disabled.
<b>URL</b>	The URL address.
<b>Action</b>	The type of action the unit will take when there is a match.
<b>Type</b>	The type of URL. For example, the type of URL is <i>Regex</i> .

### New URL Filter page

Provides settings for configuring a URL to add to the filter list.

<b>URL</b>	Enter the URL.
<b>Type</b>	Select a type from the drop-down list: <i>Simple</i> , <i>Regex</i> (regular expression), or <i>Wildcard</i> .
<b>Action</b>	Select an action the unit will take. <ul style="list-style-type: none"> <li>• <i>Allow</i> – any attempt to access a URL that matches a URL pattern with an allow action is permitted.</li> <li>• <i>Exempt</i> – similar to <i>Pass</i> in that it allows trusted traffic to bypass the antivirus proxy operations, but it functions slightly differently; ensure you are aware of the network topology involving URLs that you applied the Exemption action. Additional information about the Exempt action is found in the Security Profiles chapter of the FortiOS Handbook.</li> <li>• <i>Block</i> – attempts to access any URLs matching the URL pattern are denied; user is presented with a replacement message.</li> <li>• <i>Pass</i> – traffic to, and replay traffic from sites that match a URL pattern with a pass action will bypass all antivirus proxy operations, including FortiGuard Web Filter, web content filter, web script filters, and antivirus scanning. Make sure you trust the content of any site you pass, otherwise there may be a security risk.</li> </ul>
<b>Enable</b>	Select to enable the URL. By default, the URL is enabled.



Type a top-level domain suffix (for example, “com” without the leading period) to block access to all URLs with this suffix.

## Web filtering example

Web filtering is particularly important for protecting school-aged children. There are legal issues associated with improper web filtering as well as a moral responsibility not to allow children to view inappropriate material. The key is to design a web filtering system in such a way that students and staff do not fall under the same web filter profile in the FortiGate configuration. This is important because the staff may need to access websites that are off-limits to the students.

### School district

The background for this scenario is a school district with more than 2300 students and 500 faculty and staff in a preschool, three elementary schools, a middle school, a high school, and a continuing education center. Each elementary school has a computer lab and the high school has three computer labs with connections to the Internet. Such easy access to the Internet ensures that every student touches a computer every day.

With such a diverse group of Internet users, it was not possible for the school district to set different Internet access levels. This meant that faculty and staff were unable to view websites that the school district had blocked. Another issue was the students' use of proxy sites to circumvent the previous web filtering system. A proxy server acts as a go-between for users seeking to view web pages from another server. If the proxy server has not been blocked by the school district, the students can access the blocked website.

When determining what websites are appropriate for each school, the district examined a number of factors, such as community standards and different needs of each school based on the age of the students.

The district decided to configure the FortiGate web filtering options to block content of an inappropriate nature and to allow each individual school to modify the options to suit the age of the students. This way, each individual school was able to add or remove blocked sites almost immediately and have greater control over their students' Internet usage.

In this simplified example of the scenario, the district wants to block any websites with the word **example** on them, as well as the website `www.example.com`. The first task is to create web content filter lists for the students and the teachers.

#### To create a web content filter list for the students

```
config webfilter content
 edit 5
 set name "Student Web Content List"
 config entries
 edit example
 set action block
 set status enable
 end
 end
 end
```

It might be more efficient if the Teacher Web Content List included the same blocked content as the student list. From time to time a teacher might have to view a blocked page. It would then be a matter of changing the *Action* from *Block* to *Allow* as the situation required.

### To create a web content filter list for the teachers

```
config webfilter content
 edit 5
 set name "Teacher Web Content List"
 config entries
 edit example
 set action exempt
 set status enable
 end
 end
end
```

URL filter lists with filters to block unwanted web sites must be created for the students and teachers. For this example the URL *www.example.com* will be used.

### To create a URL filter for the students

1. Go to *Security Profiles > Web Filter > URL Filter*.
2. Select *Create New*.
3. Enter *Student URL List* as the URL filter *Name*.
4. Enter optional comments to describe the contents of the list.
5. Select *OK*.

The URL filter for the students has been created. Now it must be configured.

6. Select *Create New*.
7. Enter *example.com* in the URL field.
8. Select *Simple* from the *Type* list.
9. Select *Block* from the *Action* list.
10. Select *Enable*.
11. Select *OK*.
12. Select *OK*.

The teachers should be able to view the students' blocked content, however, so an additional URL filter is needed.

### To create a URL filter for the teachers

1. Go to *Security Profiles > Web Filter > URL Filter*.
2. Select *Create New*.
3. Enter *Teacher URL List* as the URL filter *Name*.
4. Enter optional comments to describe the list.
5. Select *OK*.

The URL filter for the students has been created. Now it must be configured.

6. Select *Create New*.
7. Enter *www.example.com* in the *URL* field.
8. Select *Simple* from the *Type* list.
9. Select *Exempt* from the *Action* list.
10. Select *Enable*.
11. Select *OK*.
12. Select *OK*.

A web filter profile must be created for the students and the teachers.



### **To create a web filter profile for the students**

1. Go to *Security Profiles > Web Filter > Profiles*.
2. Select the *Create New* icon in the Edit Web Filter window title bar.
3. Enter *Students* as the *Profile Name*.
4. Enter optional comments to identify the profile.
5. Expand the *Advanced Filter* heading.
6. Enable *Web Content Filter*.
7. Select *Student Web Content List* from the *Web Content Filter* drop-down list.
8. Enable *Web URL Filter*.
9. Select *Student URL List* from the *Web URL Filter* drop-down list.
10. Enable *Web Resume Download Block*.

Selecting this setting will block downloading parts of a file that have already been downloaded and prevent the unintentional download of virus files hidden in fragmented files. Note that some types of files, such as PDFs, are fragmented to increase download speed, and that selecting this option can cause download interruptions with these types.

11. Select *OK*.

### **To create a security policy for the students**

1. Go to *Policy > Policy > Policy*.
2. Select *Create New*.
3. Enable *Web Filter*.
4. Select *Students* from the web filter drop-down list.
5. Enter optional comments.
6. Select *OK*.

### **To create a web filter profile for the teachers**

1. Go to *Security Profiles > Web Filter > Profiles*.
2. Select the *Create New* icon in the Edit Web Filter window title bar.
3. Enter *Teachers* as the *Profile Name*.
4. Enter optional comments to identify the profile.
5. Expand the *Advanced Filter* heading.
6. Enable *Web Content Filter*.
7. Select *Teacher Web Content List* from the *Web Content Filter* drop-down list.
8. Enable *Web URL Filter*.
9. Select *Teacher URL List* from the *Web URL Filter* drop-down list.
10. Enable *Web Resume Download Block*.
11. Select *OK*.

### **To create a security policy for Teachers**

1. Go to *Policy > Policy > Policy*.
2. Select *Create New*.
3. Enable *Web Filter*.
4. Select *Teachers* from the web filter drop-down list.
5. Enter optional comments.

Select *OK*.

# Data leak prevention

The FortiGate data leak prevention (DLP) system allows you to prevent sensitive data from leaving your network. When you define sensitive data patterns, data matching these patterns will be blocked, or logged and allowed, when passing through the FortiGate unit. You configure the DLP system by creating individual filters based on file type, file size, a regular expression, an advanced rule, or a compound rule, in a DLP sensor and assign the sensor to a security policy.

Although the primary use of the DLP feature is to stop sensitive data from leaving your network, it can also be used to prevent unwanted data from entering your network and to archive some or all of the content passing through the FortiGate unit.

This section describes how to configure the DLP settings.

The following topics are included:

- [Data leak prevention concepts](#)
- [Enable data leak prevention](#)
- [Fingerprint](#)
- [File filter](#)
- [DLP archiving](#)
- [DLP examples](#)

## Data leak prevention concepts

Data leak prevention examines network traffic for data patterns you specify. You define whatever patterns you want the FortiGate unit to look for in network traffic. The DLP feature is broken down into a number of parts.

### DLP sensor

A DLP sensor is a package of filters. To use DLP, you must enable it in a security policy and select the DLP sensor to use. The traffic controlled by the security policy will be searched for the patterns defined in the filters contained in the DLP sensor. Matching traffic will be passed or blocked according to how you configured the filters.

### DLP filter

Each DLP sensor has one or more filters configured within it. Filters can examine traffic for known files using DLP fingerprints, for files of a particular type or name, for files larger than a specified size, for data matching a specified regular expression, or for traffic matching an advanced rule or compound rule.

You can configure the action taken when a match is detected. The actions include:

- None
- Log Only,
- Block
- Quarantine User,
- Quarantine IP address
- Quarantine Interface

Log Only is enabled by default.

## Fingerprint

Fingerprint scanning allows you to create a library of files for the FortiGate unit to examine. It will create checksum fingerprints so each file can be easily identified. Then, when files appear in network traffic, the FortiGate will generate a checksum fingerprint and compare it to those in the fingerprint database. A match triggers the configured action.

## File filter

File filters use file filter lists to examine network traffic for files that match either file names or file types. For example, you can create a file filter list that will find files called secret.\* and also all JPEG graphic files. You can create multiple file filter lists and use them in filters in multiple DLP sensors as required.

## File size

This filter-type checks for files exceeding a configured size. All files larger than the specified size are subject to the configured action.

## Regular expression

The FortiGate unit checks network traffic for the regular expression specified in a regular expression filter. The regular expression library used by Fortinet is a variation of a library called PCRE (Perl Compatible Regular Expressions). A number of these filters can be added to a sensor making a sort of 'dictionary' subset within the sensor.

Some other, more limited DLP implementations, use a list of words in a text file to define what words are searched for. While the format used here is slightly different than what some people are used to, the resulting effect is similar. Each Regular Expression filter can be thought of as a more versatile word to be searched against. In this dictionary (or sensor), the list of words is not limited to just predefined words. It can include expressions that can accommodate complex variations on those words and even target phrases. Another advantage of the individual filter model of this dictionary over the list is that each word can be assigned its own action, making this implementation much more granular.

## Watermark

Watermarking is essentially marking files with a digital pattern to mark the file as being proprietary to a specific company. Fortinet has a utility that will apply a digital watermark to files. The utility adds a small (approx. 100 byte) pattern to the file that is recognised by the DLP Watermark filter. the pattern is invisible to the end user.

When watermarking a file it should be verified that the pattern matches up to a category found on the FortiGate firewall. For example, if you are going to watermark a file with the sensitivity

level of “Secret” you should verify that “Secret” is a sensitivity level that has been assigned in the FortiGate unit.

## Software Versions

Before planning on using watermarking software it is always best to verify that the software will work with your OS. Currently the utility was only available for the Linux and Windows operating systems.

The Linux version can be found in one of 3 command line executable programs.

- watermark\_linux\_amd64
- watermark\_linux\_arm
- watermark\_linux\_x86

The Windows version is part of the FortiExplorer software.

## File types

The Watermark tool does not work with every file type. The following file types are supported by the watermark tool:

- .txt
- .pdf
- .doc
- .xls
- .ppt
- .docx
- .pptx
- .xlsx

Currently the DLP only works with Fortinet’s watermarking software.

## Using the FortiExplorer Watermark tool

The FortiExplorer software can be downloaded from the Fortinet Support Site.

**1. Choose whether to "Apply Watermark To:"**

- Select File
- Entire Directory

**2. Fill in the fields:**

**a. Select File**

This Field has a browse icon next to it which will allow the user to browse to and select a single file or directory to apply the water mark to.

**b. Sensitivity Level**

This field is a drop down menu that lists the available sensitivity levels that the FortiGate can scan for

**c. Identifier**

This is a unique identifier string of characters to identify the company that the document belongs to.

**d. Output Directory**

This Field has a browse icon next to it which will allow the user to browse to a directory where the altered file will be placed. If the output directory is the same as the source

directory the original file will be overwritten. If the output directory is different than the source directory then the watermarked version of the file will be placed there and the unaltered original will be left in the source directory.

**3. Select *Apply Watermark* to start the process.**

You should get output in the window similar to this:

```
> fortinet-watermark-win.exe -v -f "C:\Users\TestUser\Documents\test
document.txt" -i "123456ABC" -l "Private" -o
"C:\Users\TestUser\Watermarked Documents"
Creating watermark. Pattern:
=====identifier=123456ABC
sensitivity=Private=====
--> 'C:\Users\TestUser\Documents\test document.txt'
Inserted watermark size 231

1 file(s) processed. (success = 1, failure = 0)
```

## Installation of the watermark utility on Linux

**Add the watermark file to a location on the system that is in the \$PATH.**

To see what the path is use the command

```
echo $PATH
```

Example results:

```
/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin:/usr/ga
mes
```

for example you could move or copy the file to the `:/bin` directory.

**Permissions on the watermark file:**

Check the existing permissions:

The command in Linux for listing file along with the permissions is:

```
ls -l
```

Run the check to see if the permission status. The results may be something along these lines:

```
-rw-r--r-- 1 root root 2053868 Jan 10 11:44 watermark
```

You will see that in this case it has no executable permissions

To change the permissions on the watermark file:

It will be assumed for this command that the utility is in the `bin` directory and that you have ownership level access.

```
chmod o+x /bin/watermark
```

To verify the change:

```
ls -l wa*
```

```
-rw-r--r-x 1 root root 2053868 Jan 10 11:44 watermark
```

You can see how the `x` for executable has been added to the permissions for the *others* group.

## Syntax of the Watermark utility

The tool is executed in a Linux environment by passing in files or directories of files to insert a watermark.

## USAGE:

```
watermark <options> -f <file name> -i <identifier> -l <sensitivity
level>
watermark <options> -d <directory> -i <identifier> -l <sensitivity
level>
```

## Options:

```
-h print help
-v verbose information
-I inplace watermarking (don't copy file)
-o output directory
-e encode <to non-readable>
-a add additional watermark (by default replaces watermarks existing
watermarks)
-D delete all watermarks
```

## Using the watermark utility

Now if you are in your home directory and you want to watermark a file in the Documents directory you could plan out the command like this:

```
watermark [because that is the executable to be used]
-v [so that you can get as much feedback as possible]
-I [because you don't want a new file you just want to watermark the existing one]
-f [because you only want to change the one file not the entire directory]
filename.pdf [the name of the file]
-i 123456 [to set the identifier to 123456 - this is a required setting]
-l Private [to set the sensitivity level to "Private"]
```

Now at the command prompt enter all of these components in order:

```
watermark -v -I -f filename.pdf -i 12345 -l Private
Creating watermark. Pattern:
=====identifier=12345
sensitivity=Private=====
Watermarking file: 'filename.pdf'
Inserted watermark size 148
```

## Enable data leak prevention

DLP examines your network traffic for data patterns you specify. The FortiGate unit then performs an action based on the which pattern is found and a configuration set for each filter trigger.

### General configuration steps

Follow the configuration procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

#### 1. Create a DLP sensor.

New DLP sensors are empty. You must create one or more filters in a sensor before it can examine network traffic.

2. Add one or more filters to the DLP sensor.  
Each filter searches for a specific data pattern. When a pattern in the active DLP sensor appears in the traffic, the FortiGate unit takes the action configured in the matching filter. Because the order of filters within a sensor cannot be changed, you must configure DLP in sequence.
3. Add the DLP sensor to one or more firewall policies that control the traffic to be examined.

## Creating a DLP sensor

DLP sensors are collections of filters. You must also specify an action for the filter when you create it in a sensor. Once a DLP sensor is configured, you can select it a security policy profile. Any traffic handled by the security policy will be examined according to the DLP sensor configuration.

### To create a DLP sensor

1. Go to *Security Profiles > Data Leak Prevention > Sensor*.
2. Select the *Create New* icon on the Edit DLP Sensor window title bar.
3. In the *Name* field, enter the name of the new DLP sensor.
4. Optionally, you may also enter a comment. The comment appears in the DLP sensor list and can remind you of the details of the sensor.
5. Select *OK*.  
The DLP sensor is created and the sensor configuration window appears.
6. Select *OK*.

A newly created sensor is empty, containing no filters. Without filters, the DLP sensor will do nothing.

## Adding filters to a DLP sensor

Once you have created a DLP sensor, you need to add filters.

1. To add filters to a DLP sensor
2. Go to *Security Profiles > Data Leak Prevention > Sensor*.
3. Select the Sensor in the Edit DLP Sensor window title bar drop-down list.
4. Select *Create New*.
5. Enter a filter name.
6. Select the type of filter. You can choose either Messages or Files. Depending on which of these two are chosen different options will be available

Message filter will have these configuration options:

- Containing: [drop down menu including: Credit Card # or SSN]
- Regular Expression [input field]
- Encrypted

Examine the following Services:

Web Access

- HTTP-POST

Email

- SMTP
- POP3

- IMAP
- MAPI

Others

- NNTP

Action [from drop down menu]

- None
- Log Only,
- Block
- Quarantine User,
- Quarantine IP address
- Quarantine Interface

Files filter will have these options:

- Containing: drop down menu including: Credit Card # or SSN
- File Size >= [ ]kb
- File Type included in [drop down menu of File Filters]
- File Finger Print : [drop down menu]
- Watermark Sensitivity: [drop down menu] and Corporate Identifier [id field]
- Regular Expression [input field]
- Encrypted

Examine the following Services:

Web Access

- HTTP-POST
- HTTP-GET

Email

- SMTP
- POP3
- IMAP
- MAPI

Others

- FTP
- NNTP

Action [from drop down menu]

- None
- Log Only,
- Block
- Quarantine User,
- Quarantine IP address
- Quarantine Interface



**Table 96:** Option explanations

Option	Description
Containing..	the predefined settings for this filter are: <ul style="list-style-type: none"> <li>• Credit Card numbers - The number formats used by American Express, Visa, and Mastercard credit cards are detected.</li> <li>• Social Security Numbers.</li> </ul>
Regular Expression	Network traffic is examined for the pattern described by the regular expression.
Encrypted	This filter is triggered by encrypted files.
File Size	Enter a file size in kilobytes. Files larger than the specified size are treated according to the selected action.
File Type	Select a file filter list that includes the file patterns and file types the network traffic will be examined for. Files matching the types or patterns in the selected list are treated according to the selected action.  To create a file filter list, see <a href="#">“Creating a file filter list” on page 2141</a> .
File Finger Print	A fingerprint filter checks files in traffic against those in the FortiGate unit document fingerprint database. A match triggers the configured action.  You must configure a document source or uploaded documents to the FortiGate unit for fingerprint scanning to work. For more information about document fingerprinting, see <a href="#">“Fingerprint” on page 2131</a> .
Watermark Sensitivity	If you are using watermarking on your files you can use this filter to check for watermarks that correspond to sensitivity categories that you have set up. The Corporate Identifier is to make sure that you are only blocking watermarks that your company has place on the files, not watermarks with the same name by other companies.
Services	Configure the filter to examine the traffic over the selected services. This setting gives you a tool to optimized the resources of the FortiGate unit by only using processing cycles on the relevant traffic. Just check the boxes associated with the service / protocol that you want to have checked for filter triggers.

**Table 97:** Action Options

Action	Description
None	No action is taken if filter even if filter is triggered
Log Only	The FortiGate unit will take no action on network traffic matching a rule with this action. The filter match is logged, however. Other matching filters in the same sensor may still operate on matching traffic.
Block	Traffic matching a filter with the block action will not be delivered. The matching message or download is replaced with the data leak prevention replacement message.

**Table 97:** Action Options

Action	Description
Quarantine User	<p>If the user is authenticated, this action blocks all traffic to or from the user using the protocol that triggered the rule and adds the user to the Banned User list. If the user is not authenticated, this action blocks all traffic of the protocol that triggered the rule from the user's IP address.</p> <p>If the banned user is using HTTP, FTP, or NNTP (or HTTPS if the FortiGate unit supports SSL content scanning and inspection) the FortiGate unit displays the "Banned by data leak prevention" replacement message. If the user is using IM, the IM and P2P "Banned by data leak prevention" message replaces the banned IM message and this message is forwarded to the recipient. If the user is using IMAP, POP3, or SMTP (or IMAPS, POP3S, SMTPS if your FortiGate unit supports SSL content scanning and inspection) the Mail "Banned by data leak prevention" message replaces the banned email message and this message is forwarded to the recipient. These replacement messages also replace all subsequent communication attempts until the user is removed from the banned user list.</p> <p>If this action is chosen the additional field for [ ] minutes will appear so that a time limit can be set for the duration of the quarantine. This field cannot be left blank.</p>
Quarantine IP Address	<p>This action blocks access for any IP address that sends traffic matching a filter with this action. The IP address is added to the Banned User list. The FortiGate unit displays the "NAC Quarantine DLP Message" replacement message for all connection attempts from this IP address until the IP address is removed from the banned user list.</p> <p>If this action is chosen the additional field for [ ] minutes will appear so that a time limit can be set for the duration of the quarantine. This field cannot be left blank.</p>
Quarantine Interface	<p>This action blocks access to the network for all users connecting to the interface that received traffic matching a filter with this action. The FortiGate unit displays the "NAC Quarantine DLP Message" replacement message for all connection attempts to the interface until the interface is removed from the banned user list.</p> <p>If this action is chosen the additional field for [ ] minutes will appear so that a time limit can be set for the duration of the quarantine. This field cannot be left blank.</p>

*Quarantine User*, *Quarantine IP*, and *Quarantine Interface* provide functionality similar to NAC quarantine. However, these DLP actions block users and IP addresses at the application layer while NAC quarantine blocks IP addresses and interfaces at the network layer.

7. Select OK.
8. Repeat Steps 6 and 7 for each filter.
9. Select Apply to confirm the settings of the sensor.



If you have configured DLP to block IP addresses and if the FortiGate unit receives sessions that have passed through a NAT device, all traffic from that NAT device — not just traffic from individual users — could be blocked. You can avoid this problem by implementing authentication.



To view or modify the replacement message text, go to *System > Config > Replacement Message*.

## DLP document fingerprinting

One of the DLP techniques to detect sensitive data is fingerprinting (also called document fingerprinting). Most DLP techniques rely on you providing a characteristic of the file you want to detect, whether it's the file type, the file name, or part of the file contents. Fingerprinting is different in that you provide the file itself. The FortiGate unit then generates a checksum fingerprint and stores it. The FortiGate unit generates a fingerprint for all files detected in network traffic, and it is compared to all of the fingerprints stored in its fingerprint database. If a match is found, the configured action is taken.

The document fingerprint feature requires a FortiGate unit with internal storage. The document fingerprinting menu item does not appear on models without internal storage.

Any type of file can be detected by DLP fingerprinting and fingerprints can be saved for each revision of your files as they are updated.

To use fingerprinting you select the documents to be fingerprinted and then add fingerprinting filters to DLP sensors and add the sensors to firewall policies that accept the traffic to which to apply fingerprinting.

### Fingerprinted Documents

The FortiGate unit must have access to the documents for which it generates fingerprints. One method is to manually upload documents to be fingerprinted directly to the FortiGate unit. The other is to allow the FortiGate unit to access a network share that contains the documents to be fingerprinted.

If only a few documents are to be fingerprinted, a manual upload may be the easiest solution. If many documents require fingerprinting, or if the fingerprinted documents are frequently revised, using a network share makes user access easier to manage.

#### To configure manual document fingerprints

1. Go to *Security Profiles > Data Leak Prevention > Document Fingerprinting*.
2. In the Manual Document Fingerprints section, select *Create New*.
3. Select the file to be fingerprinted.
4. Choose a Sensitivity level. The default choices are *Critical*, *Private* and *Warning*, but more can be added in the CLI.
5. If the file is an archive containing other files, select *Process files inside archive* if you also want the individual files inside the archive to have fingerprints generated in addition to the archive itself.
6. Select *OK*.

The file is uploaded and a fingerprint generated.

#### To configure a fingerprint document source

1. Go to *Security Profiles > Data Leak Prevention > Document Fingerprinting*.
2. In the Document Sources section, select *Create New*.

### 3. Configure the settings:

<b>Name</b>	Enter a descriptive name for the document source.
<b>Server Type</b>	This refers to the type of server share that is being accessed. The default is Windows Share but this will also work on Samba shares.
<b>Server Address</b>	Enter the IP address of the server.
<b>User Name Password</b>	Enter the user name and password of the account the FortiGate unit uses to access the server network share.
<b>Path</b>	Enter the path to the document folder.
<b>Filename Pattern</b>	You may enter a filename pattern to restrict fingerprinting to only those files that match the pattern. To fingerprint all files, enter an asterisk ("**").
<b>Sensitivity Level</b>	Select a sensitivity level. The sensitivity is a tag for your reference that is included in the log files. It does not change how fingerprinting works.
<b>Scan Periodically</b>	To have the files on the document source scanned on a regular basis, select this option. This is useful if files are added or changed regularly. Once selected, you can choose Daily, Weekly, or Monthly update options, and enter the time of day the files are fingerprinted.
<b>Advanced</b>	Expand the Advanced heading for additional options.
<b>Fingerprint files in subdirectories</b>	By default, only the files in the specified path are fingerprinted. Files in subdirectories are ignored. Select this option to fingerprint files in subdirectories of the specified path.
<b>Remove fingerprints for deleted files</b>	Select this option to retain the fingerprints of files deleted from the document source. If this option is disabled, fingerprints for deleted files will be removed when the document source is rescanned.
<b>Keep previous fingerprints for modified files</b>	Select this option to retain the fingerprints of previous revisions of updated files. If this option is disabled, fingerprints for previous version of files will be deleted when a new fingerprint is generated.

### 4. Select OK.

## File filter

File filter is a DLP option that allows you to block files based on their file name or their type.

- **File patterns** are a means of filtering based purely on the names of files. They may include wildcards (\*). For example, blocking \*.scr will stop all files with an scr file extension, which is

commonly used for Windows screen saver files. Files trying to pass themselves off as Windows screen saver files by adopting the file-naming convention will also be stopped.

- Files can specify the full or partial file name, the full or partial file extension, or any combination. File pattern entries are not case sensitive. For example, adding \*.exe to the file pattern list also blocks any files ending with .EXE.
- Files are compared to the enabled file patterns from top to bottom, in list order.
- In addition to the built-in patterns, you can specify more file patterns to block. For details, see [“Creating a file filter list” on page 2141](#).
- **File types** are a means of filtering based on an examination of the file contents, regardless of the file name. If you block the file type *Archive (zip)*, all zip archives are blocked even if they are renamed with a different file extension. The FortiGate examines the file contents to determine what type of file it is and then acts accordingly.

The FortiGate unit can take either of the following actions toward the files that match a configured file pattern or type:

- **Block:** the file is blocked and a replacement message is sent to the user. If both file pattern filtering and virus scan are enabled, the FortiGate unit blocks files that match the enabled file filter and does not scan these files for viruses.
- **Allow:** the file is allowed to pass.

The FortiGate unit also writes a message to the Security log and sends an alert email message if configured to do so.



File filter does not detect files within archives. You can use file filter to block or allow the archives themselves, but not the contents of the archives.

---

## General configuration steps

The following steps provide an overview of file filter configuration. For best results, follow the procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

1. Create a file filter list.
2. Create one or more file patterns or file types to populate the file filter list.
3. Enable the file filter list by adding it to a filter in a DLP sensor.
4. Select the DLP sensor in a security policy.

## Creating a file filter list

Before your FortiGate unit can filter files by pattern or type, you must create a file filter list. The action triggered by the contents of a file filter list will be decided in the sensor so it is important to make sure that when you are building a list that you intend for the same action to be applied to all of the entries in the same list.

### To create a file filter list

1. Go to *Security Profiles > Data Leak Prevention > File Filter*.
2. Select *Create New*.
3. Enter a *Name* for the new file filter list.
4. Select *OK*.

The new list is created and the edit file filter list window appears. The new list is empty. You need to populate it with one or more file patterns or file types.

## Creating a file pattern

A file pattern allows you to block or allow files based on the file name. File patterns are created within file filter lists.

### To create a file pattern

1. Go to *Security Profiles > Data Leak Prevention > File Filter*.
2. Select a file filter list.
3. Select the *Edit* icon.
4. Select *Create New*.
5. Select *File Name Pattern* as the *Filter Type*.
6. Enter the pattern in the *Pattern* field. The file pattern can be an exact file name or can include wildcards (\*). The file pattern is limited to a maximum of 80 characters.
7. Select *OK*.

## Creating a file type

A file type allows you to block or allow files based on the kind of file. File types are created within file filter lists.

### To create a file type

1. Go to *Security Profiles > Data Leak Prevention > File Filter*.
2. Select the *Edit* icon of the file filter list to which you will add the file type.
3. Select *Create New*.
4. Select *File Type* as the *Filter Type*.
5. Select the kind of file from the *File Type* list.
6. Select *OK*.

DLP can detect the following file types:

- Archive (arj)
- Archive (bzip)
- Archive (bzip2)
- Archive (cab)
- Archive (gzip)
- Archive (Archive (lzh))
- Archive (rar)
- Archive (tar)
- Archive (zip)
- Audio (wav)
- Audio (wma)
- BMP (bmp)
- Batch File (bat)
- Common Console Document (msc)
- Encoded Data (base64)

- Encoded Data (binhex)
- Encoded Data (mime)
- Encoded Data (uue)
- Executable (elf)
- Executable (exe)
- GIF Image (gif)
- HTML Application (hta)
- HTML File (html)
- Ignored File Type (ignored)
- JPEG Image (jpeg)
- Java Application Descriptor (jad)
- Java Class File (class)
- Java Compiled Bytecode (cod)
- JavaScript File (javascript)
- Microsoft Office (msoffice)
- PDF (pdf)
- PNG Image (png)
- Packer (aspack)
- Packer (fsg)
- Packer (petite)
- Packer (upx)
- PalmOS Application (prc)
- Real Media Streaming (rm)
- Symbian Installer System File (sis)
- TIFF Image (tiff)
- Torrent (torrent)
- Unknown File Type (unknown)
- Video (mov)
- Video (mpeg)
- Windows Help File (hlp)
- activemime (activemime)



The “unknown” type is any file type that is not listed in the table. The “ignored” type is the traffic the unit typically does not scan. This includes primarily streaming audio and video.

---

## Preconfigured sensors

A number of preconfigured sensors are provided with your FortiGate unit. These can be edited or added to more closely match your needs.

Some of the preconfigured sensors with filters ready to go are:

- Credit-Card - This sensor logs the traffic, both files and messages, that contain credit card numbers in the formats used by American Express, MasterCard and Visa.
- Large-File - This sensor logs the traffic consisting of files larger than 5120 kB or approximately 5 MB.
- SSN-Sensor - This sensor logs the traffic, both files and messages, that contain Social Security Numbers with the exception of those that are WebEx invitation emails.



These rules affect only unencrypted traffic types. If you are using a FortiGate unit that can decrypt and examine encrypted traffic, you can enable those traffic types in these rules to extend their functionality if required.



Before using the rules, examine them closely to ensure you understand how they will affect the traffic on your network.

---

## DLP archiving

DLP is typically used to prevent sensitive information from getting out of your company network, but it can also be used to record network use. This is called DLP archiving. The DLP engine examines email, FTP, IM, NNTP, and web traffic. Enabling archiving for rules when you add them to sensors directs the FortiGate unit to record all occurrences of these traffic types when they are detected by the sensor.

Since the archive setting is configured for each rule in a sensor, you can have a single sensor that archives only the things you want.

You can archive Email, FTP, HTTP, IM, and session control content:

- Email content includes IMAP, POP3, and SMTP sessions. Email content can also include email messages tagged as spam by Email filtering. If your unit supports SSL content scanning and inspection, Email content can also include IMAPS, POP3S, and SMTPS sessions.
- HTTP content includes HTTP sessions. If your unit supports SSL content scanning and inspection HTTP content can also include HTTPS sessions.
- IM content includes AIM, ICQ, MSN, and Yahoo! sessions.

DLP archiving comes in two forms: *Summary Only*, and *Full*.

Summary archiving records information about the supported traffic types. For example, when an email message is detected, the sender, recipient, message subject, and total size are recorded. When a user accesses the Web, every URL the user visits recorded. The result is a summary of all activity the sensor detected.

For more detailed records, full archiving is necessary. When an email message is detected, the message itself, including any attachments, is archived. When a user accesses the Web, every page the user visits is archived. Far more detailed than a summary, full DLP archives require more storage space and processing.

Because both types of DLP archiving require additional resources, DLP archives are saved to a FortiAnalyzer unit or the FortiGuard Analysis and Management Service (subscription required).



You can use DLP archiving to collect and view historical logs that have been archived to a FortiAnalyzer unit or the FortiGuard Analysis and Management Service. DLP archiving is available for FortiAnalyzer when you add a FortiAnalyzer unit to the Fortinet configuration. The FortiGuard Analysis server becomes available when you subscribe to the FortiGuard Analysis and Management Service.

Two sample DLP sensors are provided with DLP archiving capabilities enabled. If you select the `Content_Summary` sensor in a security policy, it will save a summary DLP archive of all traffic the security policy handles. Similarly, the `Content_Archive` sensor will save a full DLP archive of all traffic handled the security policy you apply it to. These two sensors are configured to detect all traffic of the supported types and archive them.

DLP archiving is set in the CLI only.

To set the archive to Full

```
config dlp sensor
 edit <name of sensor>
 set full-archive-proto smtp pop3 imap http ftp nntp aim icq msn
 yahoo mapi
 end
```

To set the archive to Summary Only

```
config dlp sensor
 edit <name of sensor>
 set summary-proto smtp pop3 imap http ftp nntp aim icq msn yahoo
 mapi
 end
```

## DLP examples

### Blocking content with credit card numbers

When the objective is to block credit card numbers one of the important things to remember is that 2 filters will need to be used in the sensor.

In the default Credit-Card sensor, you will notice a few things.

- The Action is set to Log Only
- In the Files filter not all of the services are being examined.

If you wish to block as much content as possible with credit card numbers in it instead of just logging most the traffic that has it, the existing sensor will have to be edited.

*Security Profile > Data Leak Prevention > Sensor.*

Use the drop down menu to select Credit-Card.

1. Edit the first filter.
  - a. Change the Action to Block
  - b. Make sure all of the services are being examined.
  - c. Select OK
2. Repeat for the second filter
3. Select Apply
4. Edit the appropriate policies so that under *Security Profiles*, *DLP* is turned on and the *Credit-Card* sensor is selected.

## Blocking emails larger than 15 MB and logging emails from 5 MB to 15 MB

Because the designated size is over 10 MB the proxy options cannot be used to block the file size. Multiple filters will have to be used in this case and the order that they are used is important. Because there is no mechanism to move the filters within the sensor the order that they are added to the sensor is important.

*Security Profile > Data Leak Prevention > Sensor.*

Create a new sensor

Use the following values

<b>Name</b>	large_emails
<b>Comment</b>	<optional>

Once the Sensor has been created a new filter will need to be added.

Create New

Use the following values

Filter:

- Choice between *Meassages* and *Files*: choose *Files*.
- Choose radio button to the left of *File Size*
- In the field for the file size type 15360  
1MB = 1024kB, 15 MB = 15 x 1024kB = 15360kB

Examine the following Services

<b>SMTP</b>	enabled
<b>POP3</b>	enabled
<b>IMAP</b>	enabled
<b>HTTP</b>	not enabled
<b>FTP</b>	not enabled
<b>AIM</b>	not enabled
<b>ICQ</b>	not enabled
<b>MSN</b>	not enabled
<b>Yahoo!</b>	not enabled
<b>NNTP</b>	not enabled
<b>MAPI</b>	not enabled

Action

- From the drop down menu choose Block

Select OK

A second filter will need to be added.

*Create New*

Use the following values

Filter:

- Choice between Meassages and Files: choose Files
- Choose radio button to the left of *File Size*
- In the field for the file size type 1024

Examine the following Services

<b>SMTP</b>	enabled
<b>POP3</b>	enabled
<b>IMAP</b>	enabled
<b>HTTP</b>	not enabled
<b>FTP</b>	not enabled
<b>AIM</b>	not enabled
<b>ICQ</b>	not enabled
<b>MSN</b>	not enabled
<b>Yahoo!</b>	not enabled
<b>NNTP</b>	not enabled
<b>MAPI</b>	not enabled

Action

- From the drop down menu choose *Log Only*.

Select OK

Select Apply

Add the sensor to the appropriate policy.

The reason that the block filter is placed first is because the filters are applied in sequence and once the traffic triggers a filter the action is applied and then the traffic is passed on to the next test. If the Log Only filter which checks for anything over 1MB is triggered this would include traffic over 15MB, so a 16 MB file would only be logged. In the described order, the 16 MB file will be blocked and the 3 MB file will be logged.

## Selective blocking based on a finger print

The following is a fairly complex example but shows what can be done by combining various components in the correct configuration.

The company has a number of copyrighted documents that it does not want “escaping” to the Internet but it does want to be able to send those documents to the printers for turning into hardcopy.

The policies and procedures regarding this issue state that:

- Only members of the group *Senior\_Editors* can send copyrighted material to the printers.
- Every member of the company by default is included in the group *employees*.
- Even permitted transmission of copyrighted material should be recorded.
- All of the printers IP addresses are in a group called *approved\_printers*.
- There is a file share called *copyrighted* where any file that is copyrighted is required to have a copy stored.
- It doesn't happen often but for legal reasons sometimes these files can be changed, but all versions of a file in this directory need to be secured.
- All network connections to the Internet must have Antivirus enabled using at least the default profile.
- The SSL/SSH Inspection profile used will be *default*.

It is assumed for the purposes of this example that:

- Any addresses or address groups have been created.
- User accounts and groups have been created.
- The account used by the FortiGate is *fgtaccess*.
- The Copyrighted sensitivity level needs to be created.
- The copyrighted material is stored at `\\192.168.27.50\books\copyrighted\`

## Sensitivity Level Addition

```
config dlp fp-sensitivity
 edit copyrighted
end
```

## Finger print configuration

*Security Profile > Data Leak Prevention > Document Fingerprinting.*

In the *Document Sources* section select *Create New*

Use the following field values

<b>Name</b>	copyrighted_material
<b>Server Type</b>	Windows Share
<b>Server Address</b>	192.168.27.50
<b>User Name</b>	fgtaccess
<b>Password</b>	*****
<b>Path</b>	books/copyrighted/
<b>Filename Pattern</b>	*.pdf
<b>Sensitivity</b>	copyrighted
<b>Scan Periodically</b>	enabled
<b>&lt;Frequency&gt;</b>	Daily, Hour: 2, Min: 0
<b>Advanced</b>	

<b>Fingerprint files in subdirectories</b>	enabled
<b>Remove fingerprints for deleted files</b>	not enabled
<b>Keep previous fingerprints for modified files</b>	enabled

## Create DLP Sensors

*Security Profile > Data Leak Prevention > Sensor*

Create a new sensor. This can be done one of two ways.

- In the menu bar at the top on the right hand, use the Create New icon (circle with + symbol inside).
- In the menu bar at the top on the right hand, use the View List icon to go to the list window and use the Create New icon on the top left of that page.

Two Sensors need to be created. One for blocking the transmission of copyrighted material and a second for allowing the passing of copyrighted material under specific circumstances.

### Configuration for the first sensor that blocks transmission.

Use the following field values:

<b>Name</b>	block_copyrighted
<b>Comment</b>	<optional>

Once the Sensor has been created a new filter will need to be added.

Create New

Use the following values

Filter:

- Choice between Meessages and Files: choose Files
- Choose radio button to the left of File Finger Print
- From the drop down for File Finger Print choose “copyrighted”

Examine the following Services

<b>SMTP</b>	enabled
<b>POP3</b>	enabled
<b>IMAP</b>	enabled
<b>HTTP</b>	enabled
<b>FTP</b>	enabled
<b>AIM</b>	enabled
<b>ICQ</b>	enabled
<b>MSN</b>	enabled

<b>Yahoo!</b>	enabled
<b>NNTP</b>	enabled
<b>MAPI</b>	enabled

Action

- From the drop down menu choose Block

### Configuration for the second sensor that allows transmission.

Use the following field values:

<b>Name</b>	allow_copyrighted
<b>Comment</b>	<optional>

Once the Sensor has been created a new filter will need to be added.

This will be identical to the filter in the block\_copyrighted sensor except that the action will be *Log Only*.

## Create policies and attach DLP sensors

### Policy to allow transmission of copyrighted material

Policy > Policy > Policy

Create New

Use the following values in the Policy:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	User Identity
<b>Incoming Interface</b>	LAN
<b>Source Address</b>	all
<b>Outgoing Interface</b>	wan1
<b>Enable NAT</b>	enabled -- Use Destination Interface Address
<b>Enable Web cache</b>	<optional>
<b>Enable WAN Optimization</b>	<optional>
<b>Skip this policy for unauthenticated user</b>	do not enable
<b>Disclaimer</b>	<optional>
<b>Customize authentication Messages</b>	<optional>

Configure Authentication Rules:

<b>Destination Address</b>	approved_printers
<b>Group(s)</b>	Senior_Editors
<b>User(s)</b>	<optional>
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT
<b>Log Allowed Traffic</b>	<optional>
<b>Security Profiles</b>	
<b>Antivirus</b>	<ON> default
<b>Webfilter</b>	<optional>
<b>Application Control</b>	<optional>
<b>IPS</b>	<optional>
<b>Email Filter</b>	<optional>
<b>DLP Sensor</b>	<ON> Copyrighted
<b>VoIP</b>	<optional>
<b>ICAP</b>	<optional>
<b>Proxy Options</b>	
<b>SSL/SSH Inspection</b>	<ON>
<b>Traffic Shaping</b>	<optional>

This policy should be placed as close to the beginning of the list of policies so that it is among the first tested against.

### Policy to block transmission of copyrighted material

This will in effect be the default template for all following policies in that they will have to use the DLP profile that blocks the transmission of the copyrighted material.

*Policy > Policy > Policy*

*Create New* or *Edit* the existing policies.

The fields should include whatever values you need to accomplish your requirements, but each policy should include the DLP sensor `block_copyrighted` or if a different DLP configuration is required it should include a filter that blocks *copyrighted* fingerprinted files.

If you need to create a policy that is identity based, make sure that there is an Authentication rule for the group *employees* that uses the DLP sensor that blocks copyrighted material.

# Application control

Using the application control Security Profile feature, your FortiGate unit can detect and take action against network traffic depending on the application generating the traffic. Based on FortiGate Intrusion Protection protocol decoders, application control is a user-friendly and powerful way to use Intrusion Protection features to log and manage the behavior of application traffic passing through the FortiGate unit. Application control uses IPS protocol decoders that can analyze network traffic to detect application traffic even if the traffic uses non-standard ports or protocols.

The FortiGate unit can recognize the network traffic generated by a large number of applications. You can create application control sensors that specify the action to take with the traffic of the applications you need to manage and the network on which they are active, and then add application control sensors to the firewall policies that control the network traffic you need to monitor.

Fortinet is constantly increasing the list of applications that application control can detect by adding applications to the [FortiGuard Application Control Database](#). Because intrusion protection protocol decoders are used for application control, the application control database is part of the [FortiGuard Intrusion Protection System Database](#) and both of these databases have the same version number.

You can find the version of the application control database that is installed on your unit, by going to the *License Information* dashboard widget and find IPS Definitions version.

You can go to the [FortiGuard Application Control List](#) to see the complete list of applications supported by FortiGuard. This web page lists all of the supported applications. You can select any application name to see details about the application.

If you enable virtual domains (VDOMs) on the Fortinet unit, you need to configure application control separately for each virtual domain.

The following topics are included in this section:

- [Application control concepts](#)
- [Application considerations](#)
- [Application traffic shaping](#)
- [Application control monitor](#)
- [Enable application control](#)
- [Application control examples](#)

## Application control concepts

You can control network traffic generally by the source or destination address, or by the port, the quantity or similar attributes of the traffic itself in the security policy. If you want to control the flow of traffic from a specific application, these methods may not be sufficient to precisely define the traffic. To address this problem, the application control feature examines the traffic itself for signatures unique to the application generating it. Application control does not require knowledge of any server addresses or ports. The FortiGate unit includes signatures for over 1000 applications, services, and protocols.

Updated and new application signatures are delivered to your FortiGate unit as part of your FortiGuard Application Control Service subscription. Fortinet is constantly increasing the



number of applications that application control can detect by adding applications to the [FortiGuard Application Control Database](#). Because intrusion protection protocol decoders are used for application control, the application control database is part of the [FortiGuard Intrusion Protection System Database](#) and both of these databases have the same version number.

To view the version of the application control database installed on your FortiGate unit, go to the *License Information* dashboard widget and find the *IPS Definitions* version.

To see the complete list of applications supported by FortiGuard Application Control go to the [FortiGuard Application Control List](#). This web page lists all of the supported applications. You can select any application name to see details about the application.

## Application considerations

Some applications behave differently from most others. You should be aware of these differences before using application control to regulate their use.

### Automatically allowing basic applications

A common practice is to block applications by category, because the alternative is to list each specific traffic on an individual basis. While listing the applications individually gives a great deal of granularity it does tend to allow for missing some of them. On the other hand, blocking by category has the drawback of blocking some traffic that was not intended to be blocked.

There are a number of basic applications that you may want to be allowed on a default basis. For example, DNS. If you were to block the category Network Services you would end up blocking your web browsing, unless your users are members of a very limited group that do their web browsing by using IP addresses instead of URLs. Without DNS the systems will not be able to resolve URLs into IP addresses.

Using a set of options in the CLI the FortiGate unit can be configured to automatically allow the following types of traffic, regardless of whether or not their category is blocked:

- DNS
- ICMP
- Generic HTTP Web browsing
- Generic SSL communications

### Syntax

```
config application list
 edit appcontrol
 set options allow-dns allow-icmp allow-http allow-ssl
 end
```

As the example indicates, DNS is vitally important to multiple other types of traffic so by default it is set to be allowed, however the other settings must be specifically enabled.

## IM applications

The Application Control function for a number of IM application is not in the Web Based Manager, in the CLI of the FortiGate unit. These applications are:

- AIM
- ICQ
- MSN
- Yahoo

These applications are controlled by either permitting or denying the users from logging in to the service. Individual IM accounts are configured as to whether or not they are permitted and then there is a global policy for how to action unknown users, by the application, and whether to add the user to the black list or the white list.

The configuration details for these settings can be found in the CLI Reference guide under the heading of imp2p.

## Skype

Based on the NAT firewall type, Skype takes advantage of several NAT firewall traversal methods, such as STUN (Simple Traversal of UDP through NAT), ICE (Interactive Connectivity Establishment) and TURN (Traversal Using Relay NAT), to make the connection.

The Skype client may try to log in with either UDP or TCP, on different ports, especially well-known service ports, such as HTTP (80) and HTTPS (443), because these ports are normally allowed in firewall settings. A client who has previously logged in successfully could start with the known good approach, then fall back on another approach if the known one fails.

The Skype client could also employ Connection Relay. This means if a reachable host is already connected to the Skype network, other clients can connect through this host. This makes any connected host not only a client but also a relay server.

## Application traffic shaping

You can apply traffic shaping for application list entries you configure to pass. Traffic shaping enables you to limit or guarantee the bandwidth available to the application or applications specified in an application list entry. You can also prioritize traffic by using traffic shaping.

You can create or edit traffic shapers by going to *Firewall Objects > Traffic Shaper > Shared*. Per-IP traffic shapers are not available for use in application traffic shaping.

For more information about traffic shaping, see [“Traffic shaping methods” on page 2243](#).

## Direction of traffic shaping

When Traffic Shaping is enabled the direction that traffic shaping will be applied must also be chosen.

Forward direction traffic shaping refers to the direction of the initial connection. This would be the direction described by the policy that the Application Control Sensor is assigned to. If the policy has an Incoming Interface of LAN and an Outgoing Interface of wan1 then any Forward Direction Traffic Shaping profile will apply to network traffic heading in that direction only. If the connection used by that policy involved a response that included a download of Gigabytes of traffic the shaper would not be applied to that traffic.

Reverse Direction Traffic Shaping is applied to traffic that is flowing in the opposite direction indicated by the direction of the policy. If the policy has an Incoming Interface of LAN and an Outgoing Interface of wan1 then the shaper would only be applied to the traffic that was coming from the wan1 interface to the LAN interface.

For example, if you find that your network bandwidth is being overwhelmed by streaming HTTP video, one solution is to limit the bandwidth by applying a traffic shaper to an application control entry that allows the HTTP.Video application. Your users access the Web using a security policy that allows HTTP traffic from the internal interface to the external interface. Firewall policies are required to initiate communication so even though web sites respond to requests, a policy to allow traffic from the external interface to the internal interface is not required for your users to access the Web. The internal to external policy allows them to open communication sessions to web servers, and the external servers can reply using the existing session.

If you enable *Traffic Shaping* and select the Forward Direction shaper in an application sensor specified in the security policy, the problem will continue. The reason is the shaper you select for *Traffic Shaping* is applied only to the application traffic moving in the direction stated in the security policy. In this case, that is from the internal interface to the external interface. The security policy allows the user to visit the web site and start the video, but the video itself is streamed from the server to the user, or from the external interface to the internal interface. This is the reverse of the direction specified in the security policy. To solve the problem, you must enable *Reverse Direction Traffic Shaping* and select the appropriate shaper.

## Shaper re-use

Shapers are created independently of firewall policies and application sensors so you are free to reuse the same shapers in multiple list entries and policies. Shared shapers can be configured to apply separately to each security policy or across all policies. This means that if a shaper is configured to guaranteed 1000 KB/s bandwidth, each security policy using the shaper will have its own 1000 KB/s reserved, or all of the policies using the shaper will share a pool of 1000 KB/s, depending on how it is configured.

The same thing happens when a shaper is used in application sensors. If an application sensor using a shaper is applied to two separate policies, how the bandwidth is limited or guaranteed depends on whether the shaper is set to apply separately to each policy or across all policies. In fact, if a shaper is applied directly to one security policy, and it is also included in an application sensor that is applied to another security policy, the same issue occurs. How the bandwidth is limited or guaranteed depends on the shaper configuration.

If a shaper is used more than once within a single application sensor, all of the applications using the shaper are restricted to the maximum bandwidth or share the same guaranteed bandwidth.

For example, you want to limit the bandwidth used by Skype and Facebook chat to no more than 100 KB/s. Create a shaper, enable *Maximum Bandwidth*, and enter 100. Then create an application sensor with an entry for Skype and another entry for Facebook chat. Apply the shaper to each entry and select the application sensor in the security policy that allows your users to access both services.

This configuration uses the same shaper for each entry, so Skype **and** Facebook chat traffic are limited to no more than 100 KB/s in total. That is, traffic from both applications is added and the total is limited to 100 KB/s. If you want to limit Skype traffic to 100 KB/s and Facebook chat traffic to 100 KB/s, you must use separate shapers for each application control entry.

## Application control monitor

The application monitor enables you to gain an insight into the applications generating traffic on your network. When monitor is enabled in an application sensor entry and the list is selected in

a security policy, all the detected traffic required to populate the selected charts is logged to the SQL database on the FortiGate unit hard drive. The charts are available for display in the executive summary section of the log and report menu.



Because the application monitor relies on a SQL database, the feature is available only on FortiGate units with an internal hard drive.

---

While the monitor charts are similar to the top application usage dashboard widget, it offers several advantages. The widget data is stored in memory so when you restart the FortiGate unit, the data is cleared. Application monitor data is stored on the hard drive and restarting the system does not affect old monitor data.

Application monitor allows you to choose to compile data for any or all of three charts: top ten applications by bandwidth use, top ten media users by bandwidth, and top ten P2P users by bandwidth. Further, there is a chart of each type for the traffic handled by each security policy with application monitor enabled. The top application usage dashboard widget shows only the bandwidth used by the top applications since the last system restart.

## Application Control monitor

Once you have configured application control and associated the sensors with firewall policies, you can monitor the results. The applications that will be reported on the ones that are included in sensors that are assigned to firewall policies.

*Security Profile > Monitor > Application Monitor.*

Here you will find some widgets that include charts:

- Top Applications by Bandwidth
- Top Applications by Session Count
- Top IP/User for <application>

The number of “Top” can be set to the value of 5, 10 or 15 on any of these widgets.

## Enable application control

Application control examines your network traffic for traffic generated by the applications you want it to control.

### General configuration steps

Follow the configuration procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

1. Create an application sensor.
2. Configure the sensor to include the signatures for the application traffic you want the FortiGate unit to detect. Configure each entry to allow or pass the traffic.
3. Enable application control in a security policy and select the application sensor.

### Creating an application sensor

You need to create an application sensor before you can enable application control.

### To create an application sensor

1. Go to *Security Profiles > Application Control > Application Sensor*.
2. Select the *Create New* icon in the title bar of the *Edit Application Sensor* window.
3. In the *Name* field, enter the name of the new application sensor.
4. Optionally, you may also enter a comment.
5. Select *OK*.

The application sensor is created and the sensor configuration window appears. A newly created application sensor is empty. Without applications, the application sensor will have no effect.

## Adding applications to an application sensor

Once you have created an application sensor, you need to need to define the applications that you want to control.

You can add applications using application entries and application filters. Entries allow you to choose individual applications. Filters allow you to choose application attributes and all the applications with matching attributes are included in the filter.



The sequence of the entries in the table is significant. The entries are checked against the traffic in sequence, from top to bottom. If a match is found and the action is *Block* or *Reset*, the action is performed and further checking is stopped. If the action is *Monitor* the traffic is checked against all of the signatures in the sensor and the best match to the signature is the one that is logged.

### To add an application entry to an application sensor

1. Go to *Security Profiles > Application Control > Application Sensors*.
2. Select an application sensor from the drop-down list in the *Edit Application Sensor* window title bar.
3. Select the *Create New* icon in the sensor area and this will bring up a new window entitled *New Application Filter*.
4. Choose the format of the filter. There are two types of entries that can be added to a sensor. The type of entry is determined by the selection of the sensor type. The choices are either *Filter Based* or *Specify Applications*.

- **Filter Based**

This option is for choosing groups of similar applications based on the filters of Category, Popularity, Technology and Risk. Once the parameters of the 4 filter types have been chosen every application that falls in to that filtered list will be included in the list that the Application Control engine will use to filter the network traffic.

- **Specify Applications**

This option is good for a more granular approach to picking application to be filtered. It will allow for the use of the same filters that were used in the *Filter Based* option to develop a list of applications to be filtered but the *Specify Applications* option can be selective of which applications in that list are actually filtered. They are selected individually.

The difference in the Web-based Manager, when alternating between the Sensor Types, is that when the *Filter Based* option is chosen the *Filter Options* section will appear by default. If the *Specify Applications* sensor type is chosen you can click on the [Filter Options] link to make it appear and use it to narrow down the list of possible applications but it does not show up by default. The other difference is that with the *Specify Applications* option you are

given an additional field at the top of the Application List that allow you to type out the name of an application to search for it in that manner.

To use the search field, located above the application list, start typing any portion of the application name. The mail list of application will adjust accordingly.

5. Narrow down the list of applications to be filtered. This will depend a little on which *Sensor Type* was chosen. If the Filter Based option was chosen, by default, the top section of the window will show the properties by which the list of application filters can be filtered into a more manageable list. These properties are broken into 4 sections representing the properties of Category, Popularity, Technology and Risk. Between the property filter section and the Action section of the window there is a listing of the individual application filters that have been configured into the appliance.

Each of these individual application filters is assigned values in each of the 4 properties. The values that can be assigned to these properties are listed in the 4 sections. By enabling the check boxes next to the properties in the sections the list can be narrowed down until it only includes the subset of the individual application filters that you wish to make up the sensor entry or Application Filter.

When choosing a property, if the specific value is unknown do not disable the property section as this will cause the list of individual application filters to be empty.

The properties have been broken down into the following sections:

**a. Category**

These are the types of application that are available to filter by:

**Table 98:** Property Values listed in Category section along with ID#

Category Name	Category ID#
Botnet	19
eMail	21
File.Sharing	24
Game	8
General.Interest	12
IM	1
Media	5
Network.Service	15
P2P	2
Proxy	6
Remote.Access	7
Social.Networking	23
Storage.Backup	22
Update	17
VoIP	3
Web.Surfing	25

there is also a category designation reserved for future use.

These categories should cover the bulk of application based network traffic. If you wanted to disallow the use of Peer to Peer (P2P) applications because you didn't want your users tying up your bandwidth with torrent downloads you would select the P2P category and set the Action to Block

**b. Popularity**

Popularity is broken down into 5 levels of popularity represented by stars. 5 stars representing the most popular applications and 1 star representing applications that are the least popular. The Popularity property works well when trying to narrow down the list of one of the categories. Using the previous category example of P2P traffic but you wanted to monitor the activity of the most popular applications, which numbers about 30 as opposed to over 100, you would choose P2P from Category and the 5 star popularity.

**c. Technology**

Technology is broken down into 3 technology models as well as the more basic Network-Protocol which would can be used as a catch all for anything not covered by the more narrowly defined technologies of:

- Browser-Based
- Client-Server
- Peer -to-Peer

**d. Risk**

The Risk property does not indicate the level of risk but the type of impact that is likely to occur by allowing the traffic from that application to occur. The Risk list is broken down into the following

- Botnet
- Excessive-Bandwidth
- None

**6. Pick the individual applications if using the Specify Applications Sensor type.**

From the list of possible applications highlight the application by selecting the application. If you choose an application in error you can unhighlight or deselect the application by clicking on it again.

If the Filter Based sensor type is being used this will not be an option.

**7. Select the Action the FortiGate unit will take when it detects network traffic from the application:**

- *Monitor* allows the application traffic to flow normally and log all occurrences.

If you set the action to *Monitor*, you have the option of enabling traffic shaping for the application or applications specified in this application list entry. For more information about application control traffic shaping, see [“Enabling application traffic shaping” on page 2161](#)

- *Block* will stop all traffic from the application and log all occurrences.
- *Reset* will reset the network connection on the session that the specified application traffic was detected on.
- *Traffic Shaping* will allow a Traffic Shaping profile to be applied to the applicatin traffic that triggered the sensor.

Choosing the Traffic shaping action will cause to appear the secondary options of:

- Forward Direction Traffic Shaping with a checkbox
- Reverse Direction Traffic Shaping with a checkbox

If the checkbox is enable for these options a dropdown menu will appear next to that option that will allow you to choose one of the existing Traffic Shaping profiles. If you are

going to want to use Traffic Shaping as an action in Application Control it is best to set up any of the Traffic Shaping profiles that you will want in advance.

## Viewing and searching the application list

Go to *Security Profiles > Application Control > Application List* to view the list of applications the FortiGate unit recognizes. You may find applications by paging manually through the list, apply filters, or by using the search field.

### Searching manually

Applications are displayed in a paged list, with 50 applications per page. The bottom of the screen shows the current page and the total number of pages. You can enter a page number and press enter, to skip directly to that page. Previous Page and Next Page buttons move you through the list, one page at a time. The First Page and Last Page button take you to the beginning or end of the list.

### Applying application list filters

You can enter criteria for one or more columns, and only the applications matching all the conditions you specify will be listed.

#### To apply filters

1. Go to *Security Profiles > Application Control > Application List*.
2. Goto the column you intend to filter by and select the filter icon in the heading to the left of the column name.
3. A small window will appear which will have a field for the value to intend to filter by and a checkbox for *NOT* so that you can choose to view all of the values except the one you enter into the field. You can also input multiple values if you separate them using commas.
4. Select *Apply*.
5. Continue to add more filters to narrow your search, if required.
6. Select *OK*.

## Creating a New Custom Application Signature

If you have to deal with an application that is not already in the *Application List* you have the option to create a new one.

### Creating a new Application Filter

1. Go to *Security Profiles > Application Control > Application List*.
2. Select *Create New*.
3. The *New Custom Application Signature* window will appear.
4. In the *Name* field give the new signature a unique name.
5. In the *Comments* field give a brief discription of the application or what you intend to filter by.
6. In the *Signature* field include the signature that you intend to base your application filtering on. For more details on how to design a signature see [“Creating a custom IPS signature” on page 2080](#)
7. As an optional step you can select the Submit Signature link to submit your newly created signature to Fortiguard for possible inclusion in future predefined application lists.
- 8.



## Enabling application traffic shaping

Enabling traffic shaping in an application sensor involves selecting the required shaper. You can create or edit shapers in *Firewall Objects > Traffic Shaper > Shared*.

### To enable traffic shaping

1. Go to *Security Profiles > Application Control > Application Sensors*.
2. Select an application sensor from the drop-down list in the Edit Application Sensor window title bar.
3. Select the application control list entry and choose *Edit*.
4. Select *Traffic Shaping* and choose the required traffic shaper from the list.  
If the action is set to *Block*, the traffic shaping option is not available. Only allowed traffic can be shaped.
5. Select *Reverse Direction Traffic Shaping* and choose the required traffic shaper from the list if traffic flowing in the opposite direction also requires shaping.
6. Select *OK*.

Any security policy with this application sensor selected will shape application traffic according to the applications specified in the list entry and the shaper configuration.

## Application control examples

### Blocking all instant messaging

Instant messaging use is not permitted at the Example Corporation. Application control helps enforce this policy.

First you will create an application sensor with a single entry that includes all instant messaging applications. You will set the list action to block.

### To create the application sensor

1. Go to *Security Profiles > Application Control > Application Sensors*.
2. Select the *Create New* icon in the title bar of the *Edit Application Sensor* window.
3. In the *Name* field, enter `no_IM` for the application sensor name.
4. Select *OK*.
5. Select the *Create New* icon in the sensor.
6. For the Sensor Type select *Filter Based*.
7. For *Category*, select only *IM*.
8. For *Popularity*, *Technology* and *Risk*, make sure that all of the options are selected.
9. For *Action*, select *Block*.
10. Select *OK* to save the new filter.
11. Select *Apply* to save the sensor.

Next you will assign the sensor to a policy.

### To enable application control and select the application sensor

1. Go to *Policy > Policy > Policy*.
2. Select the security policy that allows the network users to access the Internet and choose *Edit*.

3. Under the heading *Security Profiles* toggle the button next to *Application Control* to turn it on.
4. In the drop down menu field next to the *Application Control* select the *no\_IM* application sensor.
5. Select *OK*.

No IM use will be allowed by the security policy. If other firewall policies handle traffic that users could use for IM, enable application control with the *no IM* application sensor for those as well.

## Allowing only software updates

Some departments at Example Corporation do not require access to the Internet to perform their duties. Management therefore decided to block their Internet access. Software updates quickly became an issue because automatic updates will not function without Internet access and manual application of updates is time-consuming.

The solution is configuring application control to allow only automatic software updates to access the Internet.

### To create an application sensor — web-based manager

1. Go to *Security Profiles > Application Control > Application Sensors*.
2. Select the *Create New* icon in the title bar of the *Edit Application Sensor* window.
3. In the *Name* field, enter `Updates_Only` as the application sensor name.
4. Select *OK*.
5. Select the *Create New* icon in the sensor.
6. For the *Sensor Type* select *Filter Based*.
7. Enable only *Update* in the *Category* list.
8. Select *Monitor* from the *Action* list.
9. Select *OK* to save the filter to the sensor.

The filter just finished filter will allow all software update application traffic.

10. Select the application filter *All Other Known Applications*.
11. Select *Edit*.
12. Select *Block* from the *Action* list.
13. Select *OK*.

The filter just finished filter will block all traffic from recognized applications that are not specified in this application sensor.

14. Select the *All Other Unknown Applications* entry.
15. Select *Edit*.
16. Select *Block* from the *Action* list.
17. Select *OK*.

The filter just finished filter will block all traffic from applications that are not recognized by the application control feature.

18. Select *Apply* to save the application sensor.

### To create an application sensor – CLI

```
config application list
 edit Updates_Only
 config entries
 edit 1
 set category 17
 set action pass
 end
 set other-application-action block
 set unknown-application-action block
 end
```



You will notice that there are some differences in the naming convention between the Web Based Interface and the CLI. For instance the *Action* in the CLI is “pass” and the *Action* in the Web Based Manager is “Monitor”.

---

### Selecting the application sensor in a security policy

An application sensor directs the FortiGate unit to scan network traffic only when it is selected in a security policy. When an application sensor is selected in a security policy, its settings are applied to all the traffic the security policy handles.

#### To select the application sensor in a security policy – web-based manager

1. Go to *Policy > Policy > Policy*.
2. Select a policy.
3. Select the *Edit* icon.
4. Under the heading *Security Profiles* toggle the button next to *Application Control* to turn it on.
5. In the drop down menu field next to the *Application Control* select the *Updates\_only* list.
6. Select *OK*.

#### To select the application sensor in a security policy – CLI

```
config firewall policy
 edit 1
 set utm-status enable
 set profile-protocol-options default
 set application-list Updates_Only
 end
```

Traffic handled by the security policy you modified will be scanned for application traffic. Software updates are permitted and all other application traffic is blocked.

# ICAP

ICAP is the acronym for Internet Content Adaptation Protocol. The purpose of the feature is to off-load work that would normally take place on the firewall to a separate server specifically set up for the specialized processing of the incoming traffic. This takes some of the resource strain off of the FortiGate firewall leaving it to concentrate its resources on things that only it can do.

Off-loading value-added services from Web servers to ICAP servers allows those same web servers to be scaled according to raw HTTP throughput versus having to handle these extra tasks.

ICAP servers are focused on a specific function, for example:

- Ad insertion
- Virus scanning
- Content translation
- HTTP header or URL manipulation
- Language translation
- Content filtering



ICAP does not appear by default in the web-based manager. You must enable it in *System > Config > Features* to display ICAP in the web-based manager.

---

The following topics are included in this section:

- [The Protocol](#)
- [Offloading using ICAP](#)
- [Configuration Settings](#)
- [Example ICAP sequence](#)
- [Example Scenerio](#)

## The Protocol

The protocol is a lightweight member of the TCP/IP suite of protocols. It is an Application layer protocol and its specifications are set out in RFC 3507. The default TCP that is assigned to it is 1344. Its purpose is to support HTTP content adaptation by providing simple object-based content vectoring for HTTP services. ICAP is usually used to implement virus scanning and content filters in transparent HTTP proxy caches. Content Adaptation refers to performing the particular value added service, or content manipulation, for an associated client request/response.

Essentially it allows an ICAP client, in this case the FortiGate firewall, to pass HTTP messages to an ICAP server like a remote procedure call for the purposes of some sort of transformation or other processing adaptation. Once the ICAP server has finished processing the the content, the modified content is sent back to the client.

The messages going back and forth between the client and server are typically HTTP requests or HTTP responses. While ICAP is a request/response protocol similar in semantics and usage

to HTTP/1.1 it is not HTTP nor does it run over HTTP, as such it cannot be treated as if it were HTTP. For instance ICAP messages can not be forwarded by HTTP surrogates.

## Offloading using ICAP

If you enable ICAP in a security policy, HTTP traffic intercepted by the policy is transferred to an ICAP server in the ICAP profile added to the policy. Responses from the ICAP server are returned to the FortiGate unit which forwards them to an HTTP client or server.

You can offload HTTP responses or HTTP requests (or both) to the same or different ICAP servers.

If the FortiGate unit supports HTTPS inspection, HTTPS traffic intercepted by a policy that includes an ICAP profile is also offloaded to the ICAP server in the same way as HTTP traffic.

When configuring ICAP on the FortiGate unit, you must configure an ICAP profile that contains the ICAP server information; this profile is then applied to a security policy.

## Configuration Settings

There are 2 sections where ICAP is configured:

### Servers

The available settings to be configured regarding the profile are

IP Type (in the GUI) or IP address version (in the CLI)

The options for this field in the GUI are 2 radio buttons labelled “IPv4” and “IPv6”. In the CLI the approach is slightly different. There is a field “ip-version” that can be set to “4” or “6”.

### IP address

depending on whether you’ve set the IP version to 4 or 6 will determine the format that the content of this field will be set into. In the GUI it looks like the same field with a different format but in the CLI it is actually 2 different fields named “ip-address” and ip6-address.

### Maximum Connections

This value refers to the maximum number of concurrent connections that can be made to the ICAP server. The default setting is 100. This setting can only be configured in the CLI.

The syntax is:

```
config icap server
 edit <icap_server_name>
 set max-connections <integer>
 end
```

### Port

this is the TCP port used for the ICAP traffic. The range can be from 1 to 65535. The default value is 1344.

## Profiles

### Enable Request Processing

Enabling this setting allows the ICAP server to process request messages.

If enabled this setting will also require:

- *Server* - This is the name of the ICAP server. It is chosen from the drop down menu in the field. The servers are configured in the Security Profiles > ICAP > Server section.
- *Path* - This is the path on the server to the processing component. For instance if the Windows share name was "Processes" and the directory within the share was "Content-Filter" the path would be "/Processes/Content-Filter"
- *On Failure* - There are 2 options. You can choose by the use of radio buttons either *Error* or *Bypass*.

### Enable Response Processing

Enabling this setting allows the ICAP server to process response messages.

If enabled this setting will also require:

- *Server* - This is the name of the ICAP server. It is chosen from the drop down menu in the field. The servers are configured in the Security Profiles > ICAP > Server section.
- *Path* - This is the path on the server to the processing component. For instance if the Windows share name was "Processes" and the directory within the share was "Content-Filter" the path would be "/Processes/Content-Filter"

*On Failure* - There are 2 options. You can choose by the use of radio buttons either *Error* or *Bypass*.

### Enable Streaming Media Bypass

Enabling this setting allows streaming media to ignore offloading to the ICAP server.

## Example ICAP sequence

This example is for an ICAP server performing web URL filtering on HTTP requests

1. A user opens a web browser and sends an HTTP request to connect to a web server.
2. The FortiGate unit intercepts the HTTP request and forwards it to an ICAP server.
3. The ICAP server receives the request and determines if the request is for a URL that should be blocked or allowed.
  - If the URL should be blocked the ICAP server sends a response to the FortiGate unit. The FortiGate unit returns this response to the user's web browser. This response could be a message informing the user that their request was blocked.
  - If the URL should be allowed the ICAP server sends a request to the FortiGate unit. The FortiGate unit forwards the request to the web server that the user originally attempted to connect to.
  - When configuring ICAP on the FortiGate unit, you must configure an ICAP profile that contains the ICAP server information; this profile is then applied to a security policy.

## Example Scenario

Information relevant to the following example:

- The ICAP server is designed to do proprietary content filtering specific to the organization so it will have to receive the messages and sent back appropriate responses.
- The content filter is a required security precaution so if the message cannot be processed it is not allowed through.
- Resources on both the Fortigate and the ICAP server are considerable so the maximum connections setting will set at a double the default value to analyse the impact on performance.
- The ICAP server's IP address is 172.16.100.55.
- The path to the processing component is "/proprietary\_code/content-filter/".
- Streaming media is not something that the filter considers, but is allowed through the policy so processing it would be a waste of resources.
- The ICAP profile is to be added to an existing firewall policy.
- It is assumed that the display of the policies has already been configured to show the column "ID".

1. Enter the following to configure the ICAP server:

Go to *Security Profiles > ICAP > Server*.

Use the following values:

<b>Name</b>	content-filtration-server4
<b>IP Type</b>	4
<b>IP Address</b>	172.16.100.55
<b>Port</b>	1344

Use the CLI to set the max-connections value.

```
config icap server
 edit content-filtration-server4
 set max-connections 200
 end
```

2. Enter the following to configure the ICAP profile to then apply to a security policy:

Use the following values:

<b>Name</b>	Prop-Content-Filtration
<b>Enable Request Processing</b>	enable
<b>Server</b>	content-filtration-server4
<b>Path</b>	/proprietary_code/content-filter/
<b>On Failure</b>	Error
<b>Enable Response Processing</b>	enable
<b>Server</b>	content-filtration-server4

<b>Path</b>	/proprietary_code/content-filter/
<b>On Failure</b>	Error
<b>Enable Streaming Media Bypass</b>	enable

**3.** Apply the ICAP profile to policy:

The purposes of this particular ICAP profile is to filter the content of the traffic coming through the firewall via policy ID#17

- a. Go to *Policy > Policy > Policy*.
- b. Open the existing policy ID# 17 for editing.
- c. Go to the section *Security Profiles*.
- d. Select the button next to *ICAP* so that it indicates that it's status is *ON*.
- e. Select the field with the profile name and use the drop down menu to select *Prop-Content-Filtration*.
- f. Select *OK*.



# Other Security Profiles considerations

The following topics are included in this section:

- [Profile Groups](#)
- [Security Profiles and Virtual domains \(VDOMs\)](#)
- [Conserve mode](#)
- [SSL content scanning and inspection](#)
- [Monitoring Security Profiles activity](#)
- [Using wildcards and Perl regular expressions](#)
- [Monitor interface reference](#)

## Profile Groups

One of the options when adding Security profiles to policies is the use of the Profile Groups feature. This works much the same way as an address group or a service group. You assign a selection of Security profiles to the Group and assign the group to a policy. This can be very convenient in an environment that has a large number of policies because instead of deciding each time you make a policy which Security profiles are going to be used you can have a small selection of Profile groups and every policy is assigned one of those groups. If changes need to be made, rather than going into each policy to make individual changes you only have to make changes to the group and the changes automatically propagate through to all of the policies that are using the Profile Group. It makes Security Profiles administration much simpler to implement, simpler to administrate and simpler to remember what Security Profiles features are being assigned to policies.

To refine the application of Security Profiles even further you can use the Profile Group in combination with Identity based policies and User Groups so that depending upon which User group a person belongs to that can be assigned a common set of Security profiles. A good example of this would a school environment. Staff and students are going to have significantly different permissions and restrictions associated with them. Staff will be allowed access to websites that children are not (Web Filter). Staff will be allowed to transmit certain data under certain circumstances while students cannot transmit that type of data at all (DLP). Staff might have access to applications to communicate with colleagues in real time while students might be denied social networking access to get them from being distracted from their studies (Application Control). There are a number of permutations and possibilities made simpler and easier to administrate using these features together.

## Creating a new group

### Security profiles that can be grouped

When setting up a Profile Group you can assign to a group, or not as you want, the following Profile types:

- AntiVirus
- Web Filter
- Application Control
- IPS
- Email Filter
- DLP Sensor
- VoIP
- ICAP

Because the Security profiles need to use one, if you are assigning a Security profile to a policy you must assign a Proxy Option profile.

### Using the Web-based Manager

To keep the interface simpler and less cluttered, by default, some versions of the firmware only display a default profile for each of the profile types and a default Profile Group. By going into the Admin Settings section and enabling the display of Multiple Security Profiles the option to have multiple Profile Groups in the Web Based Manager is also enabled.

1. Go to Security Profiles --> Profile Group --> Profile Group
2. Select Create New
3. Give the New Profile group a name.
4. Select the Security Profiles.
  - a. Use the check-boxes to determine whether or not a particular Security profile will be assigned.
  - b. Use the drop-down menu to determine which Security profile will be used.
  - c. Select a Proxy Option profile.

The Default Proxy Option Profile will be added by default if another profile is not selected.
5. Select OK.

## Using the CLI

In the CLI enter the commands:

```
config firewall profile-group
 edit <profile_group_name>
 set profile-protocol-options <protocol_options_name>
 set av-profile <name_of_av-profile>
 set webfilter-profile <name_of_webfilter-profile>
 set spamfilter-profile <name_of_spamfilter-profile>
 set dlp-sensor <name_of_dlp-sensor>
 set ips-sensor <name_of_ips-sensor>
 set application-list <name_of_application-list>
 set voip-profile <name_of_voip-profile>
 set icap-profile <name_of_icap-profile>
 set deep-inspection-options <name_of_deep-inspection-options>
 next
end
```

## Adding a Profile Group to a policy

### Using the CLI

1. Go to the Firewall policy that you wish to associate the Profile Group

- a. For an Address Firewall policy:

```
config firewall policy
 edit <policyID>
```

- b. For an Identity based policy

```
config firewall policy
 edit <policyID>
 config identity-based-policy
 edit <policy_id>
```

2. To assign a Profile Group to a security policy the following additional settings need to be added to the policy configuration.

```
set utm-status enable
set profile-type group
set profile-group <name of the profile group>
end
```

## When adding a Profile Group to a policy there are 2 potential points of confusion:

1. Depending on your interpretation, there may be some confusion on the profile-type setting.
  - `group` indicates the use of a profile group.
  - `single` indicates the use of individual Security profiles.
2. In the CLI, the context, or placement in the "syntax tree" of configuration settings, can make some options available or unavailable depending on other settings.

In an Address Policy you only have to go down 2 "levels" to have the options for configuring the Profile Groups available.

When an Identity policy is being used the Profile Group options are not available at the same level. You have to go down a further 2 levels, to inside the Authentication rule that is nested within the overall umbrella of the Firewall Policy. This is where the Profile Group settings will be available to you.

## Security Profiles and Virtual domains (VDOMs)

If you enable virtual domains (VDOMs) on your FortiGate unit, all Security Profiles configuration is limited to the VDOM in which you configure it.

While configuration is not shared, the various databases used by Security Profiles features are shared. The FortiGuard antivirus and IPS databases and database updates are shared. The FortiGuard web filter and spam filter features contact the FortiGuard distribution network and access the same information when checking email for spam and web site categories and classification.

## Conserve mode

FortiGate units perform all Security Profiles processing in physical RAM. Since each model has a limited amount of memory, conserve mode is activated when the remaining free memory is nearly exhausted or the AV proxy has reached the maximum number of sessions it can service. While conserve mode is active, the AV proxy does not accept new sessions.

### The AV proxy

Most content inspection the FortiGate unit performs requires that the files, email messages, URLs, and web pages be buffered and examined as a whole. The AV proxy performs this function, and because it may be buffering many files at the same time, it uses a significant amount of memory. Conserve mode is designed to prevent all the component features of the FortiGate unit from trying to use more memory than it has. Because the AV proxy uses so much memory, conserve mode effectively disables it in most circumstances. As a result, the content inspection features that use the AV proxy are also disabled in conserve mode.

All of the Security Profiles features use the AV proxy with the exception of IPS, application control, DoS as well as flow-based antivirus, DLP, and web filter scanning. These features continue to operate normally when the FortiGate unit enters conserve mode.

### Entering and exiting conserve mode

A FortiGate unit will enter conserve mode because it is nearly out of physical memory, or because the AV proxy has reached the maximum number of sessions it can service. The memory threshold that triggers conserve mode varies by model, but it is about 20% free memory. When memory use rises to the point where less than 20% of the physical memory is free, the FortiGate unit enters conserve mode.

The FortiGate unit will leave conserve mode only when the available physical memory exceeds about 30%. When exiting conserve mode, all new sessions configured to be scanned with features requiring the AV proxy will be scanned as normal, with the exception of a unit configured with the one-shot option.

### Conserve mode effects

What happens when the FortiGate unit enters conserve mode depends on how you have `av-failopen` configured. There are four options:

#### **off**

The off setting forces the FortiGate unit to stop all traffic that is configured for content inspection by Security Profiles features that use the AV proxy. New sessions are not allowed but

current sessions continue to be processed normally unless they request more memory. Sessions requesting more memory are terminated.

For example, if a security policy is configured to use antivirus scanning, the traffic it permits is blocked while in conserve mode. A policy with IPS scanning enabled continues as normal. A policy with both IPS and antivirus scanning is blocked because antivirus scanning requires the AV proxy.

Use the off setting when security is more important than a loss of access while the problem is rectified.

### pass

The pass setting allows traffic to bypass the AV proxy and continue to its destination. Since the traffic is bypassing the proxy, no Security Profiles scanning that requires the AV proxy is performed. Security Profiles scanning that does not require the AV proxy continues normally.

Use the pass setting when access is more important than security while the problem is rectified.

Pass is the default setting.

### one-shot

The one-shot setting is similar to pass in that traffic is allowed when conserve mode is active. The difference is that a system configured for one-shot will force new sessions to bypass the AV proxy even after it leaves conserve mode. The FortiGate unit resumes use of the AV proxy only when the `av-failopen` setting is changed or the unit is restarted.

### idledrop

The idledrop setting will recover memory and session space by terminating all the sessions associated with the host that has the most sessions open. The FortiGate may force this session termination a number of times, until enough memory is available to allow it to leave conserve mode.

The idledrop setting is primarily designed for situations in which malware may continue to open sessions until the AV proxy cannot accept more new sessions, triggering conserve mode. If your FortiGate unit is operating near capacity, this setting could cause the termination of valid sessions. Use this option with caution.

## Configuring the av-failopen command

You can configure the `av-failopen` command using the CLI.

```
config system global
 set av-failopen {off | pass | one-shot | idledrop}
end
```

The default setting is pass.

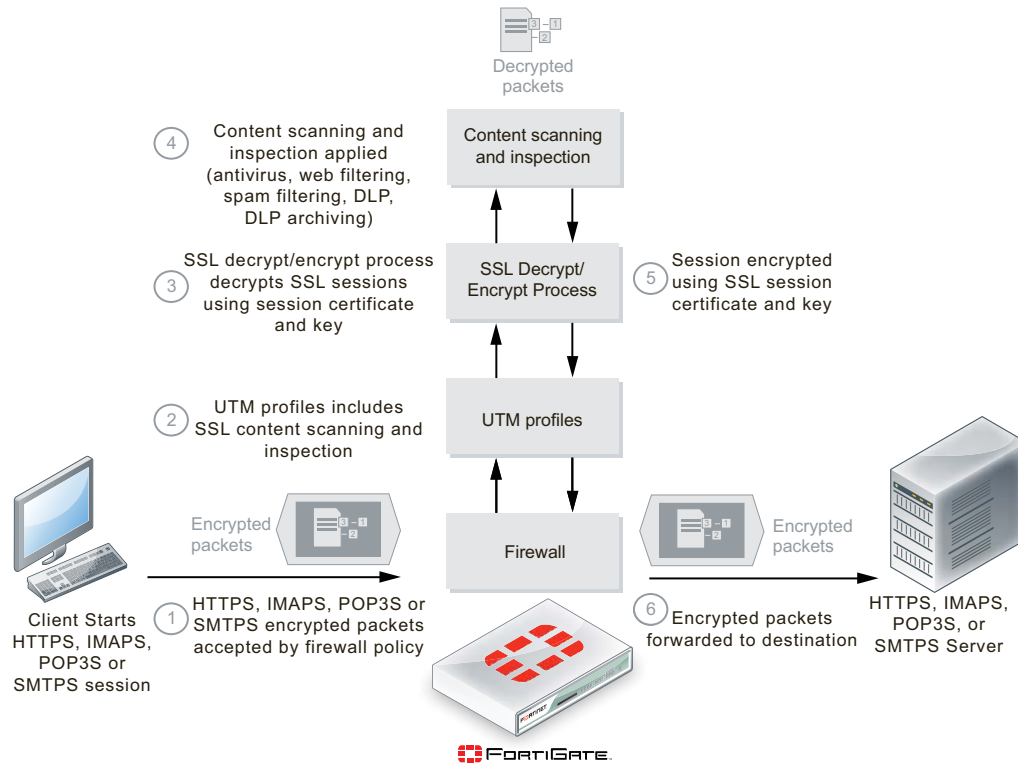
## SSL content scanning and inspection

If your FortiGate model supports SSL content scanning and inspection, you can apply antivirus scanning, web filtering, FortiGuard Web Filtering, and email filtering to encrypted traffic. You can

also apply DLP and DLP archiving to HTTPS, IMAPS, POP3S, and SMTPS traffic. To perform SSL content scanning and inspection, the FortiGate unit does the following:

- intercepts and decrypts HTTPS, IMAPS, POP3S, SMTPS, and FTPS sessions between clients and servers (FortiGate SSL acceleration speeds up decryption)
- applies content inspection to decrypted content, including:
- HTTPS, IMAPS, POP3S, and SMTPS Antivirus, DLP, and DLP archiving
- HTTPS web filtering and FortiGuard web filtering
- IMAPS, POP3S, and SMTPS email filtering
- encrypts the sessions and forwards them to their destinations.

**Figure 312:**FortiGate SSL content scanning and inspection packet flow



## Setting up certificates to avoid client warnings

To use SSL content scanning and inspection, you need to set up and use a certificate that supports it. FortiGate SSL content scanning and inspection intercepts the SSL keys that are passed between clients and servers during SSL session handshakes and then substitutes spoofed keys. Two encrypted SSL sessions are set up, one between the client and the FortiGate unit, and a second one between the FortiGate unit and the server. Inside the FortiGate unit the packets are decrypted.

While the SSL sessions are being set up, the client and server communicate in clear text to exchange SSL session keys. The session keys are based on the client and server certificates. The FortiGate SSL decrypt/encrypt process intercepts these keys and uses a built-in signing CA certificate named `Fortinet_CA_SSLProxy` to create keys to send to the client and the server. This signing CA certificate is used only by the SSL decrypt/encrypt process. The SSL decrypt/encrypt process then sets up encrypted SSL sessions with the client and server and uses these keys to decrypt the SSL traffic to apply content scanning and inspection.

Some client programs (for example, web browsers) can detect this key replacement and will display a security warning message. The traffic is still encrypted and secure, but the security warning indicates that a key substitution has occurred.

You can stop these security warnings by importing the signing CA certificate used by the server into the FortiGate unit SSL content scanning and inspection configuration. Then the FortiGate unit creates keys that appear to come from the server and not the FortiGate unit.



You can add one signing CA certificate for SSL content scanning and inspection. The CA certificate key size must be 1024 or 2048 bits. 4096-bit keys are not supported for SSL content scanning and encryption.

---

You can replace the default signing CA certificate, `Fortinet_CA_SSLProxy`, with another signing CA certificate. To do this, you need the signing CA certificate file, the CA certificate key file, and the CA certificate password.

### To add a signing CA certificate for SSL content scanning and inspection

1. Obtain a copy of the signing CA certificate file, the CA certificate key file, and the password for the CA certificate.
2. Go to *System > Certificates > Local Certificates* and select *Import*.
3. Set *Type* to *Certificate*.
4. For *Certificate file*, use the *Browse* button to select the signing CA certificate file.
5. For *Key file*, use the *Browse* button to select the CA certificate key file.
6. Enter the CA certificate *Password*.
7. Select *OK*.

The CA certificate is added to the *Local Certificates* list. In this example the signing CA certificate name is `Example_CA`. This name comes from the certificate file and key file name. If you want the certificate to have a different name, change these file names.

8. Add the imported signing CA certificate to the SSL content scanning and inspection configuration. Use the following CLI command if the certificate name is `Example_CA`.

```
config firewall ssl setting
 set caname Example_CA
end
```

The `Example_CA` signing CA certificate will now be used by SSL content scanning and inspection for establishing encrypted SSL sessions.

## SSL content scanning and inspection settings

If SSL content scanning and inspection is available on your FortiGate unit, you can configure SSL settings. The following table provides an overview of the options available and where to find further instruction:

**Table 99:** SSL content scanning and inspection settings

Setting	Description
<b>Predefined firewall services</b>	The IMAPS, POP3S and SMTPS predefined services. You can select these services in a security policy and a DoS policy.
<b>Protocol recognition</b>	<p>The TCP port numbers that the FortiGate unit inspects for HTTPS, IMAPS, POP3S, and SMTPS. Go to <i>Policy &gt; Policy &gt; Proxy Options</i>. Add or edit a Proxy Options profile, configure HTTPS, IMAPS, POP3S, SMTPS, and FTPS.</p> <p>Using <i>Proxy Options</i>, you can also configure the FortiGate unit to perform URL filtering of HTTPS or to use SSL content scanning and inspection to decrypt HTTPS so that the FortiGate unit can also apply antivirus and DLP content inspection and DLP archiving to HTTPS. Using SSL content scanning and inspection to decrypt HTTPS also allows you to apply more web filtering and FortiGuard Web Filtering options to HTTPS.</p> <p>To enable full SSL content scanning of web filtering, select <i>Enable Deep Scanning</i> under HTTPS in the Proxy Options profile.</p>
<b>Antivirus</b>	<p>Antivirus options including virus scanning and file filtering for HTTPS, IMAPS, POP3S, and SMTPS.</p> <p>Go to <i>AntiVirus &gt; Profile</i>. Add or edit a profile and configure <i>Virus Scan</i> for HTTPS, IMAPS, POP3S, and SMTPS.</p>
<b>Antivirus quarantine</b>	<p>Antivirus quarantine options to quarantine files in HTTPS, IMAPS, POP3S, SMTPS, and FTPS sessions.</p> <p>Go to <i>Security Profiles &gt; AntiVirus &gt; Quarantine</i>. You can quarantine infected files, suspicious files, and blocked files found in HTTPS, IMAPS, POP3S, SMTPS, and FTPS sessions.</p>
<b>Web filtering</b>	<p>Web filtering options for HTTPS:</p> <ul style="list-style-type: none"> <li>• Web Content Filter</li> <li>• Web URL Filter</li> <li>• ActiveX Filter</li> <li>• Cookie Filter</li> <li>• Java Applet Filter</li> <li>• Web Resume Download Block</li> <li>• Block invalid URLs</li> </ul> <p>Go to <i>Security Profiles &gt; Web Filter &gt; Profile</i>. Add or edit a web filter profile and configure web filtering for HTTPS.</p>



**Table 99:** SSL content scanning and inspection settings (continued)

Setting	Description
<b>FortiGuard Web Filtering</b>	<p>FortiGuard Web Filtering options for HTTPS:</p> <ul style="list-style-type: none"> <li>• Enable FortiGuard Web Filtering</li> <li>• Enable FortiGuard Web Filtering Overrides</li> <li>• Provide Details for Blocked HTTP 4xx and 5xx Errors</li> <li>• Rate Images by URL (Blocked images will be replaced with blanks)</li> <li>• Allow Websites When a Rating Error Occurs</li> <li>• Strict Blocking</li> <li>• Rate URLs by Domain and IP Address</li> <li>• Block HTTP Redirects by Rating</li> </ul> <p>Go to <i>Security Profiles &gt; Web Filter &gt; Profile</i>. Add or edit a profile and configure FortiGuard Web Filtering for HTTPS.</p>
<b>Email filtering</b>	<p>Email filtering options for IMAPS, POP3S, and SMTPS:</p> <ul style="list-style-type: none"> <li>• FortiGuard Email Filtering IP Address Check, URL check, E-mail Checksum Check, and Spam Submission</li> <li>• IP Address BWL Check</li> <li>• E-mail Address BWL Check</li> <li>• Return S-mail DNS Check</li> <li>• Banned Word Check</li> <li>• Spam Action</li> <li>• Tag Location</li> <li>• Tag Format</li> </ul> <p>Go to <i>Security Profiles &gt; Email Filter &gt; Profile</i>. Add or edit a profile and configure email filtering for IMAPS, POP3S, and SMTPS.</p>
<b>Data Leak Prevention</b>	<p>DLP for HTTPS, IMAPS, POP3S, and SMTPS. To apply DLP, follow the steps below:</p> <ul style="list-style-type: none"> <li>• Go to <i>Security Profiles &gt; Data Leak Prevention &gt; Sensor</i>, create a new DLP sensor or edit an existing one and then add any combination of the DLP advanced rules, DLP compound rules, file filters, a Regular Expressions, and file size limits to a DLP sensor.</li> <li>• Go to <i>Policy &gt; Policy &gt; Proxy Options</i>. Add or edit a profile and select <i>Enable Deep Scan</i> under HTTPS.</li> <li>• Go to <i>Policy &gt; Policy &gt; Policy</i>, edit the required policy, enable <i>DLP Sensor</i> and select the DLP sensor.</li> <li>• Go to <i>Policy &gt; Policy &gt; Policy</i>, edit the required policy, enable <i>Proxy Options</i> and select a profile that has <i>Enable Deep Scan</i> selected under HTTPS. <b>Note:</b> If no Proxy Options profile is selected, or if <i>Enable Deep Scan</i> is not selected within the Proxy Options profile, DLP rules <b>cannot</b> inspect HTTPS.</li> </ul>

**Table 99:** SSL content scanning and inspection settings (continued)

Setting	Description
<b>DLP archiving</b>	DLP archiving for HTTPS, IMAPS, POP3S, and SMTPS. Add DLP Rules for the protocol to be archived.
<b>Monitor DLP content information on the system dashboard</b>	<p>DLP archive information on the Log and Archive Statistics widget on the system dashboard for HTTPS, IMAPS, POP3S, and SMTPS.</p> <p>Go to <i>Policy &gt; Policy &gt; Proxy Options</i>. Add or edit a profile. For each protocol you want monitored on the dashboard, enable <i>Monitor Content Information for Dashboard</i>.</p> <p>These options display meta-information on the Statistics dashboard widget.</p>

## Exeptions

Periodically, you will come across situations where SSL and certificates will interfere with the smooth operation of an application or website. For instance, there is a popular application called Dropbox that does not work when deep SSL inspection is enabled. The reason for this is that the trusted certificate authority that is recognised by Dropbox is imbedded in the software and Dropbox cannot be reconfigured to recognise the FortiGate certificates that are used when deep SSL inspection is implemented.

One way to by-pass the deep inspection for Dropbox is to add dropbox.com to a local category in webfiltering and add that local category to the `ftgd-wf-ssl-exempt` list in the webfilter profile. This way any connections with dropbox.com will be exempt from deep SSL inspection.

Whenever an exception is found, the reason that it causes an issue will have to be determined in order to figure out a way to accommodate that application or website.

## Monitoring Security Profiles activity

The first two steps in monitoring activity covered by Security profiles is make sure that logging is enabled on the FortiGate and that the policies are configured to collect those logs as traffic goes through them.

Check the Logging and Reporting handbook for configuration of such details as to whether the logs are stored locally on a disk or in memory, or use a remote service of some kind such as a FortiAnalyzer or SNMP server. The important thing is that the storing of logs is taking place somewhere. This is configured by going to *Log & Report > Log Config > Log Setting*. If you are going to log locally you will also have to enable logging locally in the CLI.

The next step is to get the firewall policies to collect traffic logs. In the configuration of policies there are 3 logging options:

- No log
- Log Security events
- Log all Sessions

Make sure that either Log Security events or Log all sessions is selected.

There are two ways to view the Security Profiles activity based on the collected logs. The first gives you an overview based on a sampling of logs over time. This is good for spotting trends and giving you an idea of the overall impact of a type of Security Profiles threat. For instance you can see if you are a lot of your users are trying to get to sites that you have blocked or which email protocol is receiving the most blocked email.

Go to *Security Profiles > Monitor*. From here you can choose information from the different types of Security profile that you have running.

From the AV Monitor, you can see information relating to the Antivirus Profile.

- What are the Top Viruses coming through the FortiGate unit, listing:
  - Virus name,
  - Last time it was detected
  - A count of how many times it was detected.

From the Web Monitor you can find information relating to Web filtering. You can choose:

- Report by FortiGuard Webfilter Category
  - Top blocked Categories (pie chart and graph)
  - Total blocked requests
- Report by Webfilter Technique
  - Pie chart of requests (allowed, etc.)
  - Blocked Requests (Bar chart)
  - Spam
  - Banned Word
  - Virus Archive
  - FortiGuard
  - URL Filter
  - Fragmented
  - DLP

From the Application Monitor you can get an idea of which applications are being used over your network and who is using them by looking at the charts:

- Top Application by Bandwidth
- Top Applications by Session Count
- Top IP/User for...

From the Intrusion Monitor you can determine what are the Top Attacks against your network. The report will list:

- Attack Name
- Last time the attack was detected
- A count of how many times the attack was used

From the Email Monitor

- Total Emails (pie chart)
- Blocked Emails, broken down by
  - Protocol used
  - Reason/technique used to block

From the Archive & Data Leak Monitor you can see what is the:

- Top DLP usage by policy
- Total Dropped Archives

From the FortiGuard Quota you can monitor the status of quotas by seeing which ones are in effect listing:

- User name
- Webfilter Profile
- Used Quota

The second way to look at the logs of the Security Profiles activity is to look at the individual logs. This is useful for trouble shooting and verification of what is being tracked and how because individual log display more information about what happened to the traffic in question.

To look at the logs go to Log and Report > Traffic Log > Forward Traffic and search for individual logged events. In order to see just the Security Profiles based events you may have to first display a column that relates to Security Profiles such as Security Action, Security Event or Security Sub type. Once the appropriate column is displayed in the log window you can then filter based on the criteria that you are searching on. For instance if you were looking for examples of where your DLP profile stopped some traffic from getting out you could go to the Security event column and then filter for the event “dlp”. The log page will now only display dlp events. You could not further refine your filter until you were only looking at the logs that relate to the events they you are trying to track.

## Configuring packet logging options

You can use a number of CLI commands to further configure packet logging.

### Limiting memory use

When logging to memory, you can define the maximum amount of memory used to store logged packets.

```
config ips settings
 set packet-log-memory 256
end
```

The acceptable range is from 64 to 8192 kilobytes. This command affects only logging to memory.

### Limiting disk use

When logging to the FortiGate unit internal hard disk, you can define the maximum amount of space used to store logged packets.

```
config ips settings
 set ips-packet-quota 256
end
```

The acceptable range is from 0 to 4294967295 megabytes. This command affects only logging to disk.

### Configuring how many packets are captured

Since the packet containing the signature is sometimes not sufficient to troubleshoot a problem, you can specify how many packets are captured before and after the packet containing the IPS signature match.

```
config ips settings
 packet-log-history
 packet-log-post-attack
end
```

The `packet-log-history` command specifies how many packets are captured before and including the one in which the IPS signature is detected. If the value is more than 1, the packet containing the signature is saved in the packet log, as well as those preceding it, with the total number of logged packets equalling the `packet-log-history` setting. For example, if `packet-log-history` is set to 7, the FortiGate unit will save the packet containing the IPS signature match and the six before it.

The acceptable range for `packet-log-history` is from 1 to 255. The default is 1.



Setting `packet-log-history` to a value larger than 1 can affect the performance of the FortiGate unit because network traffic must be buffered. The performance penalty depends on the model, the setting, and the traffic load.

---

The `packet-log-post-attack` command specifies how many packets are logged after the one in which the IPS signature is detected. For example, if `packet-log-post-attack` is set to 10, the FortiGate unit will save the ten packets following the one containing the IPS signature match.

The acceptable range for `packet-log-post-attack` is from 0 to 255. The default is 0.

## Using wildcards and Perl regular expressions

Many Security Profiles feature list entries can include wildcards or Perl regular expressions.

For more information about using Perl regular expressions, see <http://perldoc.perl.org/perlretut.html>.

### Regular expression vs. wildcard match pattern

A wildcard character is a special character that represents one or more other characters. The most commonly used wildcard characters are the asterisk (\*), which typically represents zero or more characters in a string of characters, and the question mark (?), which typically represents any one character.

In Perl regular expressions, the '.' character refers to any single character. It is similar to the '?' character in wildcard match pattern. As a result:

- `example.com` not only matches `example.com` but also `examplea.com`, `exampleb.com`, `examplec.com`, and so on.



To add a question mark (?) character to a regular expression from the FortiGate CLI, enter Ctrl+V followed by ?. To add a single backslash character (\) to a regular expression from the CLI you must add precede it with another backslash character. For example, `example\\.com`.

To match a special character such as '.' and '\*' use the escape character '\\'. For example:

- To match `example.com`, the regular expression should be: `example\\.com`

In Perl regular expressions, '\*' means match 0 or more times of the character before it, not 0 or more times of any character. For example:

- `exam*.com` matches `exammmm.com` but does not match `example.com`

To match any character 0 or more times, use '.\*' where '.' means any character and the '\*' means 0 or more times. For example, the wildcard match pattern `exam*.com` should therefore be `exam\\.*.com`.

## Word boundary

In Perl regular expressions, the pattern does not have an implicit word boundary. For example, the regular expression “test” not only matches the word “test” but also any word that contains “test” such as “atest”, “mytest”, “testimony”, “atestb”. The notation “\b” specifies the word boundary. To match exactly the word “test”, the expression should be \btest\b.

## Case sensitivity

Regular expression pattern matching is case sensitive in the web and Email Filter filters. To make a word or phrase case insensitive, use the regular expression /i. For example, /bad language/i will block all instances of “bad language”, regardless of case.

## Perl regular expression formats

Table 100 lists and describes some example Perl regular expressions.

**Table 100:**Perl regular expression formats

Expression	Matches
abc	“abc” (the exact character sequence, but anywhere in the string)
^abc	“abc” at the beginning of the string
abc\$	“abc” at the end of the string
a b	Either “a” or “b”
^abc abc\$	The string “abc” at the beginning or at the end of the string
ab{2,4}c	“a” followed by two, three or four “b”s followed by a “c”
ab{2,}c	“a” followed by at least two “b”s followed by a “c”
ab*c	“a” followed by any number (zero or more) of “b”s followed by a “c”
ab+c	“a” followed by one or more b's followed by a c
ab?c	“a” followed by an optional “b” followed by a “c”; that is, either “abc” or “ac”
a.c	“a” followed by any single character (not newline) followed by a “c”
a\.c	“a.c” exactly
[abc]	Any one of “a”, “b” and “c”
[Aa]bc	Either of “Abc” and “abc”
[abc]+	Any (nonempty) string of “a”s, “b”s and “c”s (such as “a”, “abba”, “acbabcacaa”)
[^abc]+	Any (nonempty) string which does not contain any of “a”, “b”, and “c” (such as “defg”)
\d\d	Any two decimal digits, such as 42; same as \d{2}
/i	Makes the pattern case insensitive. For example, /bad language/i blocks any instance of bad language regardless of case.

**Table 100:**Perl regular expression formats (continued)

<code>\w+</code>	A “word”: A nonempty sequence of alphanumeric characters and low lines (underscores), such as <code>foo</code> and <code>12bar8</code> and <code>foo_1</code>
<code>100\s*mk</code>	The strings “100” and “mk” optionally separated by any amount of white space (spaces, tabs, newlines)
<code>abc\b</code>	“abc” when followed by a word boundary (for example, in “abc!” but not in “abcd”)
<code>perl\B</code>	“perl” when not followed by a word boundary (for example, in “perlert” but not in “perl stuff”)
<code>\x</code>	Tells the regular expression parser to ignore white space that is neither preceded by a backslash character nor within a character class. Use this to break up a regular expression into (slightly) more readable parts.
<code>/x</code>	Used to add regular expressions within other text. If the first character in a pattern is forward slash ‘/’, the ‘/’ is treated as the delimiter. The pattern must contain a second ‘/’. The pattern between ‘/’ will be taken as a regular expressions, and anything after the second ‘/’ will be parsed as a list of regular expression options (‘i’, ‘x’, etc). An error occurs if the second ‘/’ is missing. In regular expressions, the leading and trailing space is treated as part of the regular expression.

### Examples of regular expressions

Block any word in a phrase

```
/block|any|word/
```

Block purposely misspelled words

Spammers often insert other characters between the letters of a word to fool spam blocking software.

```
/^.*v.*i.*a.*g.*r.*o.*$/i
```

```
/cr[eéèëë][\+|-|=<>\\.\\,;!\\?%&~#\$@\\^\\°\\$£€\\{\\}()\\[\\]_01]dit/i
```

Block common spam phrases

The following phrases are some examples of common phrases found in spam messages.

```
/try it for free/i
```

```
/student loans/i
```

```
/you’re already approved/i
```

```
/special[\\+|-|=<>\\.\\,;!\\?%&~#\$@\\^\\°\\$£€\\{\\}()\\[\\]_1]offer/i
```

## Monitor interface reference

The Monitor submenus allow you to view the Security Profiles activity occurring on your network. You must have Security Profiles and sensors applied to firewall policies, as well as logging enabled for the profiles and sensors, for the monitors to display any information regarding this activity.

This topic contains the following:

- [AV Monitor](#)
- [Intrusion Monitor](#)
- [Web Monitor](#)
- [Email Monitor](#)
- [Archive & Data Leak Monitor](#)
- [Application Monitor](#)

## AV Monitor

The AV Monitor submenu allows you to view statistical information regarding viruses that were detected on your unit from *Security Profiles > Monitor > AV Monitor*. The information displays in a bar chart as well as in a table below the bar chart. The table contains detailed information.



You must have antivirus logging enabled for this within the profile itself, as well as within log settings and an antivirus profile is applied to a firewall policy.

---

### **AV Monitor page**

Displays monitored information about viruses that were detected by the unit.

**Tip:** To view information about a specific virus, select a bar within the chart; the virus FortiGuard definition displays.

---

<b>Refresh</b>	Select to refresh the information on the page.
<b>Reset</b>	Select to reset the information to clear the current information from the page. New information is included on the page.
<b>Top Viruses (all policies) since &lt;yyyy-mm-dd hh:mm:ss&gt;</b>	The top viruses detected by the unit using all firewall policies.
<b>#</b>	The order that the viruses are listed in the table.
<b>Virus Name</b>	The name of the virus.
<b>Last Detected</b>	The last time that the virus was detected.
<b>Count</b>	The number of times the virus has been detected.

---

## Intrusion Monitor

The Intrusion Monitor submenu allows you to view statistical information regarding attacks that were detected on your unit from *Security Profiles > Monitor > Intrusion Monitor*. The information



displays in a bar chart as well as in a table below the bar chart. The table contains detailed information.

---

### ***Intrusion Monitor page***

Displays monitored information about attacks that were detected by the unit.

**Tip:** To view information about a specific attack, select a bar within the chart; the attack FortiGuard definition displays.

---

<b>Refresh</b>	Select to refresh the information on the page.
<b>Reset</b>	Select to reset the information to clear the current information from the page. New information is included on the page.
<b>Top Attacks (all policies) since &lt;yyyy-mm-dd hh:mm:ss&gt;</b>	A bar chart displaying the top attacks detected by the unit.
<b>#</b>	The order that the attacks are listed in the table.
<b>Attack Name</b>	The name of the attack.
<b>Last Detected</b>	The last time that the attack was detected.
<b>Count</b>	The number of times the attack has been detected.

---

## **Web Monitor**

The Web Monitor submenu allows you to view statistical information regarding the web activity from *Security Profiles > Monitor > Web Monitor*. The information displays in both a pie chart and a bar chart .

---

### **Web Monitor page**

Displays monitored information about web activity detected by the unit.

---

<b>Refresh</b>	Select to refresh the information on the page.
<b>Reset</b>	Select to reset the information to clear the current information from the page. New information is included on the page.
<b>Report By</b>	Select whether to view the web filter monitored information by web filter technique or by FortiGuard web filter category. If you choose FortiGuard web filter category, you are viewing the information that was gathered from the category settings for FortiGuard web filter from the web filter profile.
<b>Web Monitor since &lt;yyyy-mm-dd hh:mm:ss&gt;</b>	
<b>Total Requests (HTTP)</b>	A pie chart representing the total requests detected.

---

---

<b>Blocked Requests (HTTP)</b>	A bar chart representing the total blocked requests detected. The information is broken down to spam, banned words, file filter, viruses, archives, FortiGuard, URL filter, and fragmented.
--------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

---

<b>Total Web Requests (HTTP):</b> <number>	The total number of web requests over HTTP that occurred.
-----------------------------------------------	-----------------------------------------------------------

---

## Email Monitor

The Email Monitor submenu allows you to view statistical information regarding email filtering from *Security Profiles > Monitor > Email Monitor*. The information displays in both a pie chart and bar chart.

---

### **Email Monitor page**

Displays monitored information about email filter activity detected by the unit.

---

<b>Refresh</b>	Select to refresh the information on the page.
----------------	------------------------------------------------

---

<b>Reset</b>	Select to reset the information to clear the current information from the page. New information is included on the page.
--------------	--------------------------------------------------------------------------------------------------------------------------

---

<b>Total Emails</b>	A pie chart representing the total number of emails scanned by the unit.
---------------------	--------------------------------------------------------------------------

---

<b>Blocked Emails</b>	A bar chart representing the total number of blocked emails, broken down by protocol. The colors indicate the type of scanning that occurred.
-----------------------	-----------------------------------------------------------------------------------------------------------------------------------------------

---

<b>Total Emails:</b> <number>	The total number of email messages detected by the unit.
----------------------------------	----------------------------------------------------------

---

## Archive & Data Leak Monitor

The Archive & Data Leak Monitor submenu allows you to view statistical information regarding log archives, as well as DLP usage. This page displays the information in a bar chart in *Security Profiles > Monitor > Archive & Data Leak Monitor*.

---

### **Archive & Data Leak Monitor page**

Displays monitored information about archive and DLP activity detected by the unit.

---

<b>Refresh</b>	Select to refresh the information on the page.
----------------	------------------------------------------------

---

<b>Reset</b>	Select to reset the information to clear the current information from the page. New information is included on the page.
--------------	--------------------------------------------------------------------------------------------------------------------------

---

<b>Report By:</b>	Select what type of DLP information you want to view. You can view DLP usage by DLP sensor, firewall policy usage, or by protocol.
-------------------	------------------------------------------------------------------------------------------------------------------------------------

---

<b>Top DLP Usage by DLP Sensor</b> <yyyy-mm-dd hh:mm:ss>	The bar chart that displays DLP usage monitored using DLP sensor information.
----------------------------------------------------------------	-------------------------------------------------------------------------------

---

<b>Top DLP Usage by Policy</b> <yyyy-mm-dd hh:mm:ss>	The bar chart that displays DLP usage monitored using firewall policy traffic information.
<b>Top DLP Usage by Protocol</b> <yyyy-mm-dd hh:mm:ss>	The bar chart that displays DLP usage monitored using protocol information.
<b>Total Dropped Archives:</b> <number>	The total number of dropped DLP archives.

## Application Monitor

The Application Monitor submenu allows you to view statistical information regarding application usage in *Security Profiles > Monitor > Application Monitor*.

### **Application Monitor page**

Displays monitored information about the application usage detected by the unit.

**Tip:** To view top source IP addresses for a specific application, select a bar in the chart to view that application's source IP addresses.

<b>Refresh</b>	Select to refresh the information on the page.
<b>Reset</b>	Select to reset the information to clear the current information from the page. New information is included on the page.
<b>Top Application Usage by</b> <yyyy-mm-dd hh:mm:ss>	The bar chart that displays the top applications being used detected by the unit.
<b>Resolve Host Name</b>	Appears after selecting a bar for a specific application, for example SSL. Select to resolve the host name.  <b>Tip:</b> Hover your mouse over the bar to view the address and total MB (or KB) used for that application.
<b>Report By:</b>	Appears after selecting a bar for a specific application, for example, SSL. Select to view the detailed information by destination address, or source address.
<b>Display User Name</b>	Appears after selecting <i>Source Address</i> from the drop-down list beside <i>Report By</i> .  Select to display user names.

## FortiGuard Quota

The FortiGuard Quota submenu allows you to view statistical information regarding quota usage by users in *Security Profiles > Monitor > FortiGuard Quota*.

---

**FortiGuard Quota page**

Lists the users and the amount of quota that they have used.

---

<b>Page Controls</b>	Use to navigate through the list.
<b>User Name</b>	The user name of the user that has FortiGuard quota enabled for them.
<b>Webfilter Profile</b>	The web filter profile that was used for detecting users' FortiGuard quota usage.
<b>Used Quota</b>	The amount of used quota by a user.

---

## Endpoint Monitor

You can view monitored endpoints in *Security Profiles > Monitor > Endpoint Monitor*. An endpoint is added to the list when it uses a security policy that has *Endpoint Security* enabled.

---

**Endpoint Monitor page**

Provides information about endpoints, such as endpoint traffic.

Note: The pie chart displays information in percent and indicates which is non-compliant and which is compliant.

---

<b>Refresh</b>	Updates the list, providing current endpoints that are being monitored.
<b>Report By</b>	Select to view endpoint information by traffic, status or application usage.  When you select <i>Status</i> , a pie chart appears along with information about the total endpoints ( <i>Total Endpoints</i> ). When you select <i>Traffic</i> or <i>Application usage</i> , a bar chart appears; select a bar to view detailed information.

---

# Chapter 16 SSL VPN for FortiOS 5.0

- [Introduction to SSL VPN](#) provides useful general information about VPN and SSL, how the FortiGate unit implements them, and gives guidance on how to choose between SSL and IPsec.
- [Basic Configuration](#) explains how to configure the FortiGate unit and the web portal. Along with these configuration details, this chapter also explains how to grant unique access permissions, configure the SSL virtual interface (`ssl.root`), and describes the SSL VPN OS Patch Check feature that allows a client with a specific OS patch to access SSL VPN services.
- [The SSL VPN client](#) provides an overview of the FortiClient software required for tunnel mode, where to obtain the software, install it and the configuration information required for remote users to connect to the internal network.
- [Setup examples](#) explores several configuration scenarios with step-by-step instructions. While the information provided is enough to set up the described SSL VPN configurations, these scenarios are not the only possible SSL VPN setups.

# Introduction to SSL VPN

Over the past several years, as organizations have grown and become more complex, secure remote access to network resources has become critical for day-to-day operations. In addition, businesses are expected to provide clients with efficient, convenient services including knowledge bases and customer portals. Employees travelling across the country or around the world require timely and comprehensive access to network resources. As a result of the growing need for providing remote/mobile clients with easy, cost-effective and secure access to a multitude of resources, the concept of a Virtual Private Network was developed.

SSL VPNs establish connectivity using SSL, which functions at Levels 4 - 5 (Transport and Session). Information is encapsulated at Levels 6 - 7 (Presentation and Application), and SSL VPNs communicate at the highest levels in the OSI model. SSL is not strictly a Virtual Private Network (VPN) technology allows clients to connect to remote networks in a secure way. A VPN is a secure logical network created from physically separate networks. VPNs use encryption and other security methods to ensure that only authorized users can access the network. VPNs also ensure that the data transmitted between computers cannot be intercepted by unauthorized users. When data is encoded and transmitted over the Internet, the data is said to be sent through a "VPN tunnel". A VPN tunnel is a non-application oriented tunnel that allows the users and networks to exchange a wide range of traffic regardless of application or protocol.

The advantages of a VPN over an actual physical private network are two-fold. Rather than utilizing expensive leased lines or other infrastructure, you use the relatively inexpensive, high-bandwidth Internet. Perhaps more important though is the universal availability of the Internet - in most areas, access to the Internet is readily obtainable without any special arrangements or long wait times.

SSL (Secure Sockets Layer) as HTTPS is supported by most web browsers for exchanging sensitive information securely between a web server and a client. SSL establishes an encrypted link, ensuring that all data passed between the web server and the browser remains private and secure. SSL protection is initiated automatically when a user (client) connects to a web server that is SSL-enabled. Once the successful connection is established, the browser encrypts all the information before it leaves the computer. When the information reaches its destination, it is decrypted using a secret (private) key. Any data sent back is first encrypted, and is decrypted when it reaches the client.

FortiOS supports the SSL and TLS versions defined below.

<b>Version</b>	<b>RFC</b>
SSL 2.0	<a href="#">RFC 6176</a>
SSL 3.0	<a href="#">RFC 6101</a>
TLS 1.0	<a href="#">RFC 2246</a>
TLS 1.1	<a href="#">RFC 4346</a>
TLS 1.2	<a href="#">RFC 5246</a>

Table 101. SSL and TLS version support table.

## SSL VPN modes of operation

When a remote client connects to the FortiGate unit, the FortiGate unit authenticates the user based on user name, password, and authentication domain. A successful login determines the access rights of remote users according to user group. The user group settings specify whether the connection will operate in web-only mode or tunnel mode.

### Web-only mode

Web-only mode provides remote users with a fast and efficient way to access server applications from any thin client computer equipped with a web browser. Web-only mode offers true clientless network access using any web browser that has built-in SSL encryption and the Sun Java runtime environment.

Support for SSL VPN web-only mode is built into the FortiOS operating system. The feature comprises of an SSL daemon running on the FortiGate unit, and a web portal, which provides users with access to network services and resources including HTTP/HTTPS, telnet, FTP, SMB/CIFS, VNC, RDP and SSH.

In web-only mode, the FortiGate unit acts as a secure HTTP/HTTPS gateway and authenticates remote users as members of a user group. After successful authentication, the FortiGate unit redirects the web browser to the web portal home page and the user can access the server applications behind the FortiGate unit.

When the FortiGate unit provides services in web-only mode, a secure connection between the remote client and the FortiGate unit is established through the SSL VPN security in the FortiGate unit and the SSL security in the web browser. After the connection has been established, the FortiGate unit provides access to selected services and network resources through a web portal.

FortiGate SSL VPN web portals have a 1- or 2-column page layout and portal functionality is provided through small applets called widgets. Widget windows can be moved or minimized. The controls within each widget depend on its function. There are predefined web portals and the administrator can create additional portals.

Configuring the FortiGate unit involves selecting the appropriate web portal configuration in the user group settings. These configuration settings determine which server applications can be accessed. SSL encryption is used to ensure traffic confidentiality.

For information about client operating system and browser requirements, see the Release Notes for your FortiGate firmware.

### Tunnel mode

Tunnel mode offers remote users the freedom to connect to the internal network using the traditional means of web-based access from laptop computers, as well as from airport kiosks, hotel business centers, and Internet cafés. If the applications on the client computers used by your user community vary greatly, you can deploy a dedicated SSL VPN client to any remote client through its web browser. The SSL VPN client encrypts all traffic from the remote client computer and sends it to the FortiGate unit through an SSL VPN tunnel over the HTTPS link between the web browser and the FortiGate unit. Another option is split tunneling, which ensures that only the traffic for the private network is sent to the SSL VPN gateway. Internet traffic is sent through the usual unencrypted route. This conserves bandwidth and alleviates bottlenecks.

In tunnel mode, remote clients connect to the FortiGate unit and the web portal login page using Microsoft Internet Explorer, Firefox, Chrome, Mac OS, or Linux. The FortiGate unit acts as a secure HTTP/HTTPS gateway and authenticates remote users as members of a user group. After successful authentication, the FortiGate unit redirects the web browser to the web portal

home page dictated by the user group authentication settings. If the user does not have the SSL VPN client installed, they will be prompted to download the SSL VPN client (an ActiveX or Java plugin) and install it using controls provided through the web portal. SSL VPN tunnel mode can also be initiated from a standalone application on Windows, Mac OS, and Linux.

When the user initiates a VPN connection with the FortiGate unit through the SSL VPN client, the FortiGate unit establishes a tunnel with the client and assigns the client a virtual IP address from a range of reserved addresses. The client uses the assigned IP address as its source address for the duration of the connection. After the tunnel has been established, the user can access the network behind the FortiGate unit.

Configuring the FortiGate unit to establish a tunnel with remote clients involves enabling the feature through SSL VPN configuration settings and selecting the appropriate web portal configuration for tunnel-mode access in the user group settings. The security policy and protection profiles on the FortiGate unit ensure that inbound traffic is screened and processed securely.



The user account used to install the SSL VPN client on the remote computer must have administrator privileges.



If you are using Windows Vista, you must disable UAC (User Account Control) before installing the SSL VPN tunnel client. This UAC setting must be disabled before the SSL VPN tunnel client is installed. IE7 in Windows Vista runs in Protected Mode by default. To install SSL VPN client ActiveX, you need to launch IE7 by using 'Run as administrator' (right-click the IE7 icon and select 'Run as administrator').

---

For information about client operating system requirements, see the Release Notes for your FortiGate firmware. For information on configuring tunnel mode, see [“Tunnel mode and split tunneling” on page 2203](#).

## Port forwarding mode

While tunnel mode provides a Layer 3 tunnel that users can run any application over it, the user needs to install the tunnel client, and have the required administrative rights to do so. In some situations, this may not be desirable, yet the simple web mode does not provide enough flexibility for application support. For example, using an email client that needs to communicate with a POP3 server. The port forward mode, or proxy mode, provides this middle ground between web mode and tunnel mode.

SSL VPN port forwarding listens on local ports on the user's computer. When it receives data from a client application, the port forward module encrypts and sends the data to the FortiGate unit, which then forwards the traffic to the application server.

The port forward module is implemented with a Java applet, which is downloaded and runs on the user's computer. The applet provides the up-to-date status information such as addressing and bytes sent and received.



On the user end, the user logs into the FortiGate SSL VPN portal, and selects a port forward bookmark configured for a specific application. The bookmark defines the server address and port as well as which port to listen to on the user's computer.



The user must configure the application on the PC to point to the local proxy instead of the application server. For information on this configuration change, see the application documentation.

This mode only supports client/server applications that are using a static TCP port. It will not support client/server applications using dynamic ports or traffic over UDP.

For information on configuring a port forward tunnel, see [“Port forward tunnel” on page 2203](#).

## Application support

With Citrix application servers, the server downloads an ICA configuration file to the user's PC. The client application uses this information to connect to the Citrix server. The FortiGate unit will read this file and append a SOCKS entry to set the SOCKS proxy to localhost. The Citrix client will then be able to connect to the SSL VPN port forward module to provide the connection. When configuring the port forwarding module, an selection is available for Citrix servers.

For Windows Remote Desktop Connections, when selecting the RDP option, the tunnel will launch the RDP client and connect to the local loopback address after the port forward module has been initiated.

## SSL VPN and IPv6

FortiOS supports SSL VPN using IPv6 addressing using IPv6 configurations for security policies and addressing including:

- Policy matching for IPv6 addresses
- Support for DNS resolving in SSL VPN
- Support IPv6 for ping
- FTP applications
- SMB
- Support IPV6 for all the java applets (Telnet, VNC, RDP and so on)

## Traveling and security

Because SSL VPN provides a means for “on-the-go” users to dial in to the network while away from the office, you need to ensure that wherever and however they choose to dial in is secure, and not potentially compromising the corporate network.

When setting up the portal, you can include two options to ensure corporate data is safe; a host check for antivirus software, and a cache cleaner.

### Host check

You can enable a host integrity checker to scan the remote client. The integrity checker probes the remote client computer to verify that it is safe before access is granted. Security attributes recorded on the client computer (for example, in the Windows registry, in specific files, or held in memory due to running processes) are examined and uploaded to the FortiGate unit.

For more information, see [“Host check” on page 2212](#).

## Cache cleaning

You can enable a cache cleaner to remove any sensitive data that would otherwise remain on the remote computer after the session ends. For example, all cache entries, browser history, cookies, encrypted information related to user authentication, and any temporary data generated during the session are removed from the remote computer. If the client's browser cannot install and run the cache cleaner, the user is not allowed to access the SSL-VPN portal.

For more information, see [“Configuring cache cleaning” on page 2214](#).

# Basic Configuration

Configuring SSL VPN involves a number of configurations within FortiOS that you need to complete to make it all come together. This chapter describes the components required, and how and where to configure them to set up the FortiGate unit as an SSL VPN server. The configurations and steps are high level, to show you the procedures needed, and where in FortiOS they are located. For real-world examples, see the chapter, [“Setup examples” on page 2221](#).

There are three or four key steps to configuring an SSL VPN tunnel. The first three in the points below are mandatory, while the other is optional. This chapter will outline these four key steps, as well as additional configuration you can do for tighter security and monitoring.

## The key steps are:

- Create user accounts and user groups for the remote clients.  
([“User accounts and groups” on page 2195](#))
- Create a web portal to define user access to network resources.  
([“Configuring SSL VPN web portals” on page 2199](#))
- Configure the security policies.  
([“Configuring security policies” on page 2204](#))
- For tunnel-mode operation, add routing to ensure that client tunnel-mode packets reach the SSL VPN interface.  
([“Routing in tunnel mode” on page 2211](#))
- Setup logging of SSL VPN activities.  
([“SSL VPN logs” on page 2216](#))

## User accounts and groups

The first step for an SSL VPN tunnel is to add the users and user groups that will access the tunnel. You may already have users defined for other authentication-based security policies. These users and groups are identified when creating the security policy when defining the authentication rules.

The user group is associated with the web portal that the user sees after logging in. You can use one policy for multiple groups, or multiple policies to handle differences between the groups such as access to different services, or different schedules.

### To create a user account

- in the web-based manager, go to *User & Device > User > User Definition*, and select *Create New*.
- in the CLI, use the commands in `config user local`.

All users accessing the SSL tunnel must be in a firewall user group. User names can be up to 64 characters long.

### To create user groups

- in the web-based manager, go to *User & Device > User > User Groups* and select *Create New*.
- in the CLI, use the commands in `config user group`.

## Authentication

Remote users must be authenticated before they can request services and/or access network resources through the web portal. The authentication process can use a password defined on the FortiGate unit or optionally use established external authentication mechanisms such as RADIUS or LDAP.

To authenticate users, you can use a plain text password on the FortiGate unit (Local domain), forward authentication requests to an external RADIUS, LDAP or TACACS+ server, or utilize PKI certificates.

For information about how to create RADIUS, LDAP, TACACS+ or PKI user accounts and certificates, see the [Authentication](#) chapter of [The Handbook](#).



FortiOS supports LDAP password renewal notification and updates through SSL VPN. Configuration is enabled using the CLI commands:

```
config user ldap
 edit <username>
 set password-expiry-warning enable
 set password-renewal enable
 end
```

---

## MAC host check

When a remote client attempts to log in to the portal, you can have the FortiGate unit check against the client's MAC address to ensure that only a specific computer or device is connecting to the tunnel. This can ensure better security should a password be compromised.

MAC addresses can be tied to specific portals and can be either the entire MAC address or a subset of it. MAC host checking is configured in the CLI using the commands:

```
conf vpn ssl web portal
 edit portal
 set mac-addr-check enable
 set mac-addr-action allow
 config mac-addr-check-rule
 edit "rule1"
 set mac-addr-list 01:01:01:01:01:01 08:00:27:d4:06:5d
 set mac-addr-mask 48
 end
 end
 end
```

## IP addresses for users

After the FortiGate unit authenticates a request for a tunnel-mode connection, the FortiGate unit assigns the SSL VPN client an IP address for the session. The address is assigned from an address range (IP Pool) which is a firewall address that defines an IP address range.



Take care to prevent overlapping IP addresses. Do not assign to clients any IP addresses that are already in use on the private network. As a precaution, consider assigning IP addresses from a network that is not commonly used (for example, 10.254.254.0/24).

---

### To set tunnel-mode client IP address range - web-based manager

1. Go to *Firewall Objects > Address > Addresses* and select *Create New*.
2. Enter an *Name*, for example, *SSL\_VPN\_tunnel\_range*.
3. Select a *Type of IP Range*.
4. In the *Subnet/IP Range* field, enter the starting and ending IP addresses that you want to assign to SSL VPN clients, for example *10.254.254.[80-100]*.
5. In *Interface*, select *Any*.
6. Select *OK*.

### To set tunnel-mode client IP address range - CLI

If your SSL VPN tunnel range is for example *10.254.254.80 - 10.254.254.100*, you could enter

```
config firewall address
 edit SSL_tunnel_users
 set type iprange
 set end-ip 10.254.254.100
 set start-ip 10.254.254.80
 end
end
```

You can select the tunnel-mode IP Pools in two places:

- The *VPN > SSL > Config* page *IP Pools* setting applies to all web portals that do not specify their own IP Pools.
- The web portal Tunnel Mode widget IP Pools setting, if used, applies only to the web portal and overrides the setting in *VPN > SSL > Config*. See [“Tunnel mode and split tunneling” on page 2203](#).

## Authentication of remote users

When remote users connect to the SSL VPN tunnel, they must perform authentication before being able to use the internal network resources. This can be as simple as assigning users with their own passwords, connecting to an LDAP server or using more secure options. FortiOS provides a number of options for authentication as well as security option for those connected users.

The web portal can include bookmarks to connect to internal network resources. A web (HTTP/HTTPS) bookmark can include login credentials so that the FortiGate unit automatically logs the user into the web site. This means that the user logs into the SSL VPN and then does not have to enter any more credentials to visit preconfigured web sites.

Both the administrator and the end user can configure bookmarks, including SSO bookmarks. To add bookmarks as a web portal user, see [“Adding bookmarks” on page 2202](#).

### Setting the client authentication timeout

The client authentication timeout controls how long an authenticated user will remain connected. When this time expires, the system forces the remote client to authenticate again. As with the idle timeout, a shorter period of time is more secure. The default value is 28800 seconds (8 hours). You can only modify this timeout value in the CLI.

For example, to change the authentication timeout to 18 000 seconds, enter the following commands:

```
config vpn ssl settings
 set auth-timeout 18000
end
```

You can also set the idle timeout for the client, to define how long the user does not access the remote resources before they are logged out. For information see [“SSL connection configuration” on page 2200](#).

### Allow one time login per user

You can set the SSL VPN tunnel such that each user can only log into the tunnel one time concurrently per user per login. That is, once logged into the portal, they cannot go to another system and log in with the same credentials again.

To do this, go to *VPN > SSL > Portal* and select to disable *Allow Multiple Concurrent Sessions for Each User*. It is enabled by default.

To configure in the CLI, enter the commands:

```
config vpn ssl web portal
 edit <portal_name>
 set limit-user-logins enable
 end
```

Once set, once the user has logged in, no other user can use the same login credentials.

### Strong authentication with security certificates

The FortiGate unit supports strong (two-factor) authentication through X.509 security certificates (version 1 or 3). The FortiGate unit can require clients to authenticate using a certificate. Similarly, the client can require the FortiGate unit to authenticate using a certificate.

For information about obtaining and installing certificates, see the [Authentication](#) chapter of *The Handbook*.

You can select the *Require Client Certificate* option in *SSL VPN config* so that clients must authenticate using certificates. The client browser must have a local certificate installed, and the FortiGate unit must have the corresponding CA certificate installed.

When the remote client initiates a connection, the FortiGate unit prompts the client browser for its client-side certificate as part of the authentication process.

#### To require client authentication by security certificates - web-based manager

1. Go to *VPN > SSL > Config*.
2. Select *Require Client Certificate*.
3. Select *Apply*.

#### To require client authentication by security certificates - CLI

```
config vpn ssl settings
 set reqclientcert enable
end
```

If your SSL VPN clients require strong authentication, the FortiGate unit must offer a CA certificate that the client browser has installed.

In the FortiGate unit SSL VPN settings, you can select which certificate the FortiGate offers to authenticate itself. By default, the FortiGate unit offers its factory installed (self-signed) certificate from Fortinet to remote clients when they connect.

#### To enable FortiGate unit authentication by certificate - web-based manager

1. Go to *VPN > SSL > Config*.
2. From the *Server Certificate* list, select the certificate that the FortiGate unit uses to identify itself to SSL VPN clients.
3. Select *Apply*.

#### To enable FortiGate unit authentication by certificate - CLI

For example, to use the `example_cert` certificate

```
config vpn ssl settings
 set servercert example_cert
end
```



FortiOS will check the server certificate to verify that the certificate is valid. Only valid server certificates should be used.

---

#### NSA Suite B cryptography support

FortiOS supports the use of ECDSA Local Certificates for SSL VPN Suite B. The National Security Agency (NSA) developed Suite B algorithms in 2005 to serve as a cryptographic base for both classified and unclassified information at an interoperable level.

FortiOS allows you to import, generate, and use ECDSA certificates defined by the Suite B cryptography set. To generate ECDSA certificates, use the following command in the CLI:

```
exec vpn certificate local generate ec
```

## Configuring SSL VPN web portals

The SSL VPN portal enables remote users to access internal network resources through a secure channel using a web browser. FortiGate administrators can configure log in privileges for system users and which network resources are available to the users.

This step in the configuration of the SSL VPN tunnel sets up the infrastructure; the addressing, encryption, and certificates needed to make the initial connection to the FortiGate unit. This step also is where you set up what the remote user sees when the connection is successful. The portal view defines what resources are available to the remote users and what functionality they have on the network.

## SSL connection configuration

To configure the basic SSL VPN settings for encryption and log in options, go to *VPN > SSL > Config*.

<b>IP Pools</b>	Select <i>Edit</i> to select the range or subnet firewall addresses that represent IP address ranges reserved for tunnel-mode SSL VPN clients.
<b>Server Certificate</b>	Select the signed server certificate to use for authentication. If you leave the default setting (Self-Signed), the FortiGate unit offers its factory installed certificate from Fortinet, to remote clients when they connect.
<b>Require Client Certificate</b>	Select to use group certificates for authenticating remote clients. When the remote client initiates a connection, the FortiGate unit prompts the client for its client-side certificate as part of the authentication process.  For information on using PKI to provide client certificate authentication, see the <a href="#">Authentication Guide</a> .
<b>Encryption Key Algorithm</b>	Select the algorithm for creating a secure SSL connection between the remote client web browser and the FortiGate unit. This will depend on what the web browser of the client can support.  The FortiGate unit supports a range of cryptographic cipher suites to match the capabilities of various web browsers. The web browser and the FortiGate unit negotiate a cipher suite before any information is transmitted over the SSL link.
<b>Idle Timeout</b>	Type the period of time (in seconds) that the connection can remain idle before the user must log in again. The range is from 10 to 28800 seconds. Setting the value to 0 will disable the idle connection timeout. This setting applies to the SSL VPN session. The interface does not time out when web application sessions or tunnels are up.  You can also set the authentication timeout for the client, to define how long the user can remain connected to the network. For information see <a href="#">“Setting the client authentication timeout” on page 2197</a> .
<b>Login Port</b>	Enter the port number for HTTPS access.
<b>Enable Endpoint Registration</b>	Select so that FortiClient registers with the FortiGate unit when connecting. If you configured a registration key by going to <i>System &gt; Config &gt; Advanced</i> , the remote user is prompted to enter the key. This only occurs on the first connection to the FortiGate unit.
<b>Advanced (DNS and WINS Servers)</b>	Enter up to two DNS servers and/or two WINS servers to be provided for the use of clients.

## Portal configuration

The portal configuration determines what the remote user sees when they log in to the portal. Both the system administrator and the user have the ability to customize the SSL VPN portal.

To view the portals settings page, go to *VPN > SSL > Portal*.



There are three pre-defined default web portal configurations available:

- *full-access*
- *tunnel-access*
- *web-access*

Each web portal type include similar configuration options. Select between the different portals by selecting one from the drop-down list in the upper right corner of the window. You can also create a custom portal by selecting the plus sign next to the portal drop-down list.

<b>Name</b>	The name for the portal
<b>Portal Message</b>	This is a text header that appears on the top of the web portal.
<b>Theme</b>	A color styling for the web portal.
<b>Page Layout</b>	Select one or two column layouts for the widgets that appear on the web portal page.
<b>Enable Tunnel Mode</b>	If your web portal provides tunnel mode access, you need to configure the <i>Tunnel Mode</i> widget. These settings determine how tunnel mode clients are assigned IP addresses.
<b>Enable Split Tunneling</b>	Select so that the VPN carries only the traffic for the networks behind the FortiGate unit. The user's other traffic follows its normal route.
<b>IP Pools</b>	Select an IP Pool for users to acquire an IP address when connecting to the portal. There is always a default pool available if you do not create your own.
<b>Client Options</b>	<p>These options affect how the FortiClient application behaves when connected to the FortiGate VPN tunnel. When enabled, a check box for the corresponding option appears on the VPN login screen in FortiClient, and is not enabled by default.</p> <p><b>Save Password</b> - When enabled, if the user selects this option, their password is stored on the user's computer and will automatically populate each time they connect to the VPN.</p> <p><b>Auto Connect</b> - When enabled, if the user selects this option, when the FortiClient application is launched, for example after a reboot or system startup, FortiClient will automatically attempt to connect to the VPN tunnel.</p> <p><b>Always Up (Keep Alive)</b> - When enabled, if the user selects this option, the FortiClient connection will not shut down. When not selected, during periods of inactivity, FortiClient will attempt to stay connected every three minutes for a maximum of 10 minutes.</p>
<b>Enable Web Mode</b>	Select to enable web mode access.
<b>Applications</b>	Select the applications the user can access when connected over the VPN portal.
<b>Include Session Info</b>	Select to display the Session Information widget on the portal page. The <i>Session Information</i> widget displays the login name of the user, the amount of time the user has been logged in and the inbound and outbound traffic statistics.

<b>Include Connection Tool</b>	Select to display the Connection Tool widget on the portal page. Use the <i>Connection Tool</i> widget to connect to a internal network resource without adding a bookmark to the bookmark list. You select the type of resource and specify the URL or IP address of the host computer.
<b>Include Bookmarks</b>	Select to include bookmarks on the web portal. Bookmarks are used as links to internal network resources. When a bookmark is selected from a bookmark list, a pop-up window appears with the web page. Telnet, VNC, and RDP require a browser plug-in. FTP and Samba replace the bookmarks page with an HTML file-browser.
<b>Prompt Mobile Users to Download FortiClient App</b>	If a remote user is using web browser to connects to the SSL VPN in web mode they are prompted to download the FortiClient Application. The remote user can accept or reject the notification. If the user accepts, they are redirected to the FortiClient web site.
<b>Allow Multiple Concurrent Sessions for Each User</b>	You can set the SSL VPN tunnel such that each user can only log into the tunnel one time concurrently per user per login. That is, once logged into the portal, they cannot go to another system and log in with the same credentials again. To prevent multiple logins, clear the check box.

## Adding bookmarks

A web bookmark can include login credentials to automatically log the SSL VPN user into the web site. When the administrator configures bookmarks, the web site credentials must be the same as the user's SSL VPN credentials. Users configuring their own bookmarks can specify alternative credentials for the web site.

### To add a bookmark

1. On the *VPN > SSL > Portal* page, ensure *Include Bookmarks* is enabled.
2. Select *Create New* and enter the following information:

<b>Category</b>	Select a category, or group, to include the bookmark. If this is the first bookmark added, you will be prompted to add a category. Otherwise, select <i>Create</i> from the drop-down list.
<b>Name</b>	Enter a name for the bookmark.
<b>Type</b>	Select the type of link from the drop-down list. Telnet, VNC, and RDP require a browser plug-in. FTP and Samba replace the bookmarks page with an HTML file-browser.
<b>Location</b>	Enter the IP address source.
<b>SSO</b>	Select if you wish to use single sign-on for any links that require authentication.  When including a link using SSO, ensure to use the entire url. For example, <code>http://10.10.1.0/login</code> , rather than just the IP address.
<b>Description</b>	Enter a brief description of the link.

Select *OK*.

For more configuration options, see [“Additional configuration options”](#) on page 2210.

## Personal bookmarks

The administrator has the ability to view bookmarks the remote client has added to their SSL VPN login in the bookmarks widget. This enables the administrator to monitor and, if needed, remove unwanted bookmarks that do not meet with corporate policy.

To view and maintain remote client bookmarks, go to *VPN > SSL > Personal Bookmarks*.

On mid-range and high end FortiGate units, this feature is enabled by default. On low-end FortiGate units, it must be enabled.

### To enable personal bookmarks

1. Go to *System > Config > Features*.
2. In the *Display Options on GUI* section, select *SSVPN Personal Bookmark Management*.
3. Select *Apply*.

## Custom login screen

You can create a custom log in for your remote SSL VPN users. When configured with a security policy, when the user connects to the SSL VPN portal, a custom log in screen appears. With this screen, you can define the address, customize the look and define how many users are can connect at any one time to the portal.

When adding the URL Path, you only need to enter the subdirectory or site. The FortiGate unit will complete the remainder of the address. For example, if the sub site is *corpusers*, only enter *corpusers*. The final URL that appears is *http://172.20.120.230/corpusers*. The login port is separately configured by going to *VPN > SSL > Config*.

When configuring with the security policy, when you create *SSL VPN Authentication Rules*, you can select the specific portal login screen.

## Tunnel mode and split tunneling

If you want your web portal to have tunnel mode access, select *Tunnel Mode* when creating a new portal. Enable *Split Tunneling* so that the VPN carries only the traffic for the networks behind the FortiGate unit. The user's other traffic follows its normal route.

### Port forward tunnel

Port forwarding provides a method of connecting to application servers without configuring a tunnel mode connection, and requiring the installation of tunnel mode client. Set up the portal as described at [“Configuring SSL VPN web portals”](#) on page 2199. To configure the application, create a bookmark with the *Type* of *PortForward*.

Ensure that *Port Forward* is enabled in the *Applications* list.

## The Connection tool widget

The *Connection Tool* widget enables a user to connect to resources when isn't a bookmark. Ensure that what you want remote users to connect to is enabled in the *Applications* list of the *General* settings, by selecting the *Settings* button in the portal configuration window.

## To configure the Connection Tool widget - CLI

To change, for example, the full-access portal Connection Tool widget to allow all application types except Telnet, you would enter:

```
config vpn ssl web portal
 edit full-access
 config widget
 edit 3
 set allow-apps ftp rdp smb ssh vnc web
 end
 end
 end
end
```

## Configuring security policies

You will need at least one SSL VPN security policy. This is an identity-based policy that authenticates users and enables them to access the SSL VPN web portal. The SSL VPN user groups named in the policy determine who can authenticate and which web portal they will use. From the web portal, users can access protected resources or download the SSL VPN tunnel client application.

This section contains the procedures needed to configure security policies for web-only mode operation and tunnel-mode operation. These procedures assume that you have already completed the procedures outlined in [“User accounts and groups” on page 2195](#).

If you will provide tunnel mode access, you will need a second security policy — an ACCEPT tunnel mode policy to permit traffic to flow between the SSL VPN tunnel and the protected networks.

## Firewall addresses

Before you can create security policies, you need to define the firewall addresses you will use in those policies. For both web-only and tunnel mode operation, you need to create firewall addresses for all of the destination networks and servers to which the SSL VPN client will be able to connect.

For tunnel mode, you will already have defined firewall addresses for the IP address ranges that the FortiGate unit will assign to SSL VPN clients.

The source address for your SSL VPN security policies will be the predefined “all” address. Both the address and the netmask are 0 . 0 . 0 . 0 . The “all” address is used because VPN clients will be connecting from various addresses, not just one or two known networks. For improved security, if clients will be connecting from one or two known locations you should configure firewall addresses for those locations, instead of using the “all” address.

To create a firewall address, in the web-based manager, go to *Firewall Objects > Address > Address*, and select *Create New*.

## Create an SSL VPN security policy

At minimum, you need one SSL VPN security policy to authenticate users and provide access to the protected networks. You will need additional security policies only if you have multiple web portals that provide access to different resources. You can use one policy for multiple groups, or multiple policies to handle differences between the groups such as access to different services, or different schedules.

The SSL VPN security policy specifies:

- the remote address that corresponds to the IP address of the remote user.
- the local protected subnet address that corresponds to the IP address or addresses that remote clients need to access.

The local protected subnet address may correspond to an entire private network, a range of private IP addresses, or the private IP address of a server or host.

- the level of SSL encryption to use and the authentication method.
- which SSL VPN user groups can use the security policy.
- the times (schedule) and types of services that users can access.
- the UTM features and logging that are applied to the connection.



Do not use ALL as the destination address. If you do, you will see the “Destination address of Split Tunneling policy is invalid” error when you enable Split Tunneling

### To create an SSL-VPN security policy - web-based manager

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Select the *Policy Type* as *SSL*.
3. Enter the following information:

<b>Incoming Interface</b>	Select the name of the FortiGate network interface to that connects to the Internet.
<b>Remote Address</b>	Select <i>all</i> .
<b>Local Interface</b>	Select the FortiGate network interface that connects to the protected network.
<b>Local Protected Subnet</b>	Select the firewall address you created that represents the networks and servers to which the SSL VPN clients will connect.  If you want to associate multiple firewall addresses or address groups with the <i>Destination Interface/Zone</i> , from <i>Destination Address</i> , select the plus symbol. In the dialog box, move the firewall addresses or address groups from the <i>Available Addresses</i> section to the <i>Members</i> section, then select <i>OK</i> .
<b>SSL Client Certificate Restrictive</b>	Select to allow access only to holders of a (shared) group certificate. The holders of the group certificate must be members of an SSL VPN user group, and the name of that user group must be present in the Allowed field. See <a href="#">“Strong authentication with security certificates” on page 2198</a> .
<b>Cipher Strength</b>	Select the bit level of SSL encryption. The web browser on the remote client must be capable of matching the level that you select.

4. Under *Configure SSL-VPN Authentication Rules*, select *Create New*.

Add a user group to the policy. The New SSL VPN Authentication Rule window opens on top of the security policy. Enter the following information and then select OK. You can select Add again to add more groups.

<b>Group(s)</b>	Select user groups that can connect to the SSL VPN tunnel.
<b>User(s)</b>	Select individual users that can connect to the SSL VPN tunnel.
<b>Schedule</b>	Select always.
<b>SSL-VPN Portal</b>	Select the portal the users connect to.
<b>Custom Login</b>	Select to choose a configured login screen. For more information, see <a href="#">“Custom login screen” on page 2203</a> .

Your identity-based policies are listed in the security policy table. The FortiGate unit searches the table from the top down to find a policy to match the client’s user group. Using the move icon in each row, you can change the order of the policies in the table to ensure the best policy will be matched first. You can also use the icons to edit or delete policies.

### To create an SSL VPN security policy - CLI

To create the security policy by entering the following CLI commands.

```
config firewall policy
 edit 0
 set srcintf port1
 set dstintf port2
 set srcaddr all
 set dstaddr OfficeLAN
 set action ssl-vpn
 set nat enable
 config identity-based-policy
 edit 0
 set groups SSL-VPN
 set schedule always
 set service ALL
 set sslvpn-poprtal <portal_name>
 end
 end
end
```

## Create a tunnel mode security policy

If your SSL VPN will provide tunnel mode operation, you need to create a security policy to enable traffic to pass between the SSL VPN virtual interface and the protected networks. This is in addition to the SSL VPN security policy that you created in the preceding section.

The SSL VPN virtual interface is the FortiGate unit end of the SSL tunnel that connects to the remote client. It is named `ssl.<vdom_name>`. In the root VDOM, for example, it is named `ssl.root`. If VDOMs are not enabled on your FortiGate unit, the SSL VPN virtual interface is also named `ssl.root`.

### To configure the tunnel mode security policy - web-based manager

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.

3. Enter the following information and select *OK*.

<b>Incoming Interface</b>	Select the virtual SSL VPN interface, such as <i>ssl.root</i> .
<b>Source Address</b>	Select the firewall address you created that represents the IP address range assigned to SSL VPN clients, such as <i>SSL_VPN_tunnel_users</i> .
<b>Outgoing Interface</b>	Select the interface that connects to the protected network.
<b>Destination Address</b>	Select the firewall address that represents the networks and servers the SSL VPN clients will connect to.  To select multiple firewall addresses or address groups, select the plus sign next to the drop-down list.
<b>Service</b>	Select service in the left list and use the right arrow button to move them to the right list. Select the ALL service to allow the user group access to all services.
<b>Action</b>	Select <i>Accept</i> .
<b>Enable NAT</b>	Select <i>Enable NAT</i> . (Optional)

### To configure the tunnel mode security policy - CLI

```
config firewall policy
 edit <id>
 set srcintf ssl.root
 set dstintf <dst_interface_name>
 set srcaddr <tunnel_ip_address>
 set dstaddr <protected_network_address_name>
 set schedule always
 set service ALL
 set nat enable
 end
```

This policy enables the SSL VPN client to initiate communication with hosts on the protected network. If you want to enable hosts on the protected network to initiate communication with the SSL VPN client, you should create another Accept policy like the preceding one but with the source and destination settings reversed.

You must also add a static route for tunnel mode operation.

### Routing for tunnel mode

If your SSL VPN operates in tunnel mode, you must add a static route so that replies from the protected network can reach the remote SSL VPN client.

#### To add the tunnel mode route - web-based manager

1. Go to *Router > Static > Static Routes* and select *Create New*.  
For low-end FortiGate units, go to *System > Network > Routing* and select *Create New*.
2. Enter the *Destination IP/Mask* of the tunnel IP address that you assigned to the users of the web portal.
3. Select the SSL VPN virtual interface for the *Device*.
4. Select *OK*.

### To add the tunnel mode route - CLI

If you assigned 10.11.254.0/24 as the tunnel IP range, you would enter:

```
config router static
 edit <id>
 set device ssl.root
 set dst 10.11.254.0/24
 set gateway <gateway_IP>
 end
```

## Split tunnel Internet browsing policy

With split tunneling disabled, all of the SSL VPN client's requests are sent through the SSL VPN tunnel. But the tunnel mode security policy provides access only to the protected networks behind the FortiGate unit. Clients will receive no response if they attempt to access Internet resources. You can enable clients to connect to the Internet through the FortiGate unit.

### To add an Internet browsing policy

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter the following information and select *OK*.

<b>Incoming Interface</b>	Select the virtual SSL VPN interface, <i>ssl.root</i> , for example.
<b>Source Address</b>	Select the firewall address you created that represents the IP address range assigned to SSL VPN clients.
<b>Outgoing Interface</b>	Select the FortiGate network interface that connects to the Internet.
<b>Destination Address</b>	Select <i>all</i> .
<b>Action</b>	Select <i>Accept</i> .
<b>Enable NAT</b>	Select <i>Enable</i> .

### To configure the Internet browsing security policy - CLI

To enable browsing the Internet through port1, you would enter:

```
config firewall policy
 edit 0
 set srcintf ssl.root
 set dstintf port1
 set srcaddr SSL_tunne_users
 set dstaddr all
 set schedule always
 set service ALL
 set nat enable
 end
```



## Enabling a connection to an IPsec VPN

You might want to provide your SSL VPN clients access to another network, such as a branch office, that is connected by an IPsec VPN. To do this, you need only to add the appropriate security policy. For information about route-based and policy-based IPsec VPNs, see the [IPsec VPN Guide](#).

### Route-based connection

#### To configure interconnection with a route-based IPsec VPN - web-based manager

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter the following information and select *OK*.

<b>Incoming Interface</b>	Select the virtual SSL VPN interface, <i>ssl.root</i> , for example.
<b>Source Address</b>	Select the firewall address that represents the IP address range assigned to SSL VPN clients.
<b>Outgoing Interface</b>	Select the virtual IPsec interface for your IPsec VPN.
<b>Destination Address</b>	Select the address of the IPsec VPN remote protected subnet.
<b>Action</b>	Select <i>ACCEPT</i> .
<b>Enable NAT</b>	Enable.

#### To configure interconnection with a route-based IPsec VPN - CLI

If, for example, you want to enable SSL VPN users to connect to the private network (address name *OfficeAnet*) through the *toOfficeA* IPsec VPN, you would enter:

```
config firewall policy
 edit 0
 set srcintf ssl.root
 set dstintf toOfficeA
 set srcaddr SSL_tunnel_users
 set dstaddr OfficeAnet
 set action accept
 set nat enable
 set schedule always
 set service ALL
 end
```

### Policy-based connection

#### To configure interconnection with a policy-based IPsec VPN - web-based manager

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Select the *Policy Type* as *VPN* and leave the *Policy Subtype* as *IPsec*.

3. Enter the following information and select *OK*.

<b>Local Interface</b>	Select the virtual SSL VPN interface, <i>ssl.root</i> , for example.
<b>Local Protected Subnet</b>	Select the firewall address that represents the IP address range assigned to SSL VPN clients.
<b>Outgoing VPN Interface</b>	Select the FortiGate network interface that connects to the Internet.
<b>Remote Protected Subnet</b>	Select the address of the IPsec VPN remote protected subnet.
<b>VPN tunnel</b>	Select the Phase 1 configuration name of your IPsec VPN.
<b>Allow traffic to be initiated from the remote site</b>	Enable
<b>NAT inbound</b>	Enable

4. Configure inbound NAT from the CLI:

```
config firewall policy
 edit 0
 set natinbound enable
 end
```

#### To configure interconnection with a policy-based IPsec VPN - CLI

If, for example, you want to enable SSL VPN users to connect to the private network (address name OfficeAnet) through the OfficeA IPsec VPN, you would enter:

```
config firewall policy
 edit 0
 set srcintf ssl.root
 set dstintf port1
 set srcaddr SSL_tunnel_users
 set dstaddr OfficeAnet
 set action ipsec
 set schedule always
 set service ALL
 set inbound enable
 set outbound enable
 set natinbound enable
 set vpntunnel toOfficeA
 end
```

In this example, port1 is connected to the Internet.

## Additional configuration options

Beyond the basics of setting up the SSL VPN, you can configure a number of other options that can help to ensure your internal network is secure and limit the possibility of attacks and viruses entering the network from an outside source.

## Routing in tunnel mode

If are creating a SSL VPN connection in tunnel mode, you need to add a static route so that replies from the protected network can reach the remote SSL VPN client.

### To add the tunnel mode route - web-based manager

1. Go to *Router > Static > Static Routes* and select *Create New*.  
For low-end FortiGate units, go to *System > Network > Routing* and select *Create New*.
2. Enter the *Destination IP/Mask* of the tunnel IP address that you assigned to the users of the web portal.
3. Select the SSL VPN virtual interface for the *Device*.
4. Select *OK*.

### To add the tunnel mode route - CLI

If you assigned 10.11.254.0/24 as the tunnel IP range, you would enter:

```
config router static
 edit <id>
 set device ssl.root
 set dst 10.11.254.0/24
 set gateway <gateway_IP>
 end
```

## Changing the port number for web portal connections

You can specify a different TCP port number for users to access the web portal login page through the HTTPS link. By default, the port number is 443 and users can access the web portal login page using the following default URL:

```
https://<FortiGate_IP_address>:443/remote/login
```

where <FortiGate\_IP\_address> is the IP address of the FortiGate interface that accepts connections from remote users.

### To change the SSL VPN port - web-based manager

1. If *Current VDOM* appears at the bottom left of the screen, select *Global* from the list of VDOMs.
2. Go to *VPN > SSL > Config*.
3. Type an unused port number in *Login Port*, and select *Apply*.

### To change the SSL VPN port - CLI

This is a global setting. For example, to set the SSL VPN port to 10443, enter:

```
config global
 config system global
 set sslvpn-sport 10443
 end
```

## SSL offloading

Configuring SSL offloading that allows or denies client renegotiation, is configured in the CLI. This helps to resolve the issues that affect all SSL and TLS servers that support renegotiation, identified by the Common Vulnerabilities and Exposures system in CVE-2009-3555. The IETF is currently working on a TLS protocol change that will permanently resolve the issue. The SSL

offloading renegotiation feature is considered a workaround until the IETF permanently resolves the issue.

The CLI command is `ssl-client-renegotiation` and is found in `config firewall vip` command.

## Customizing the web portal login page

The default web portal login page shows only the *Name* and *Password* fields and the Login button, centred in the web browser window. You can customize the page with your company name or other information.

The login page is a form of replacement message, in HTML format. You can modify the content to display a customized message. Note that there are specific fields that must remain in the code to ensure the page appears correctly in the user's browser.



Before you begin, copy the default web portal login page text to a separate text file for safe-keeping. Afterward, if needed you can restore the text to the original version.

---

### To configure the SSL VPN login page - web-based manager

1. If you want to edit the global login page and *Current VDOM* appears at the bottom left of the screen, select *Global* from the list of VDOMs.
2. Go to *System > Config > Replacement Messages*.
3. Expand the *SSL VPN* row and select *SSL VPN login page*.
4. Edit the HTML text. Note the following content that must remain on the page:
  - The login page must contain a form with `ACTION="%%SSL_ACT%%"` and `METHOD="%%SSL_METHOD%%"`
  - The form must contain the `%%SSL_LOGIN%%` tag to provide the login form.
  - The form must contain the `%%SSL_HIDDEN%%` tag.

## Host check

When you enable AV, FW, or AV-FW host checking in the web portal Security Control settings, each client is checked for security software that is recognized by the Windows Security Center. As an alternative, you can create a custom host check that looks for security software selected from the Host Check list. For more information, see [“Portal configuration” on page 2200](#).

The Host Check list includes default entries for many security software products.



Host integrity checking is only possible with client computers running Microsoft Windows platforms.

## To configure host checking - CLI

To configure the full-access portal to check for AV and firewall software on client Windows computers, you would enter the following:

```
config vpn ssl web portal
 edit full-access
 set host-check av-fw
 end
```

To configure the full-access portal to perform a custom host check for FortiClient Host Security AV and firewall software, you would enter the following:

```
config vpn ssl web portal
 edit full-access
 set host-check custom
 set host-check-policy FortiClient-AV FortiClient-FW
 end
```

## Creating a custom host check list

You can add your own software requirements to the host check list using the CLI. Host integrity checking is only possible with client computers running Microsoft Windows platforms. Enter the following commands:

```
config vpn ssl web host-check-software
 edit <software_name>
 set guid <guid_value>
 set type <av | fw>
 set version <version_number>
 end
```

Enter the Globally Unique Identifier (GUID) for the host check application, if known. Windows uses GUIDs to identify applications in the Windows Registry. The GUID can be found in the Windows registry in the HKEY\_CLASSES\_ROOT section.

To get the exact versioning, in Windows right-click on the .EXE file of the application and select Properties. Select the *Version* tab.

## Windows OS check

The Windows patch check enables you to define the minimum Windows version and patch level allowed when connecting to the SSL VPN portal. When the user attempts to connect to the web portal, FortiOS performs a query on the version of Windows the user has installed. If it does not match the minimum requirement, the connection is denied. The Windows patch check is configured in the CLI.

The following example shows how you would add an OS check to the g1portal web portal. This OS check accepts all Windows XP users and Windows 2000 users running patch level 3.

To specify the acceptable patch level, you set the `latest-patch-level` and the `tolerance`. The lowest acceptable patch level is `latest-patch-level` minus `tolerance`. In this case, `latest-patch-level` is 3 and `tolerance` is 1, so 2 is the lowest acceptable patch level.

```
config vpn ssl web portal
 edit glportal
 set os-check enable
 config os-check-list windows-2000
 set action check-up-to-date
 set latest-patch-level 3
 set tolerance 1
 end
 config os-check-list windows-xp
 set action allow
 end
end
```

## Configuring cache cleaning

When the SSL VPN session ends, the client browser cache may retain some information. To enhance security, cache cleaning clears this information just before the SSL VPN session ends.



The cache cleaner is effective only if the session terminates normally. The cache is not cleaned if the session ends due to a malfunction, such as a power failure.

### To enable cache cleaning

To enable cache cleaning on the full-access portal, you would enter:

```
config vpn ssl web portal
 edit full-access
 set cache-cleaner enable
 end
```

Cache cleaning requires a browser plug-in. If the user does not have the plug-in, it is automatically downloaded to the client computer.

## Configuring virtual desktop

Available for Windows XP, Windows Vista, and Windows 7 client PCs, the virtual desktop feature completely isolates the SSL VPN session from the client computer's desktop environment. All data is encrypted, including cached user credentials, browser history, cookies, temporary files, and user files created during the session. When the SSL VPN session ends normally, the files are deleted. If the session ends due to a malfunction, files might remain, but they are encrypted, so the information is protected.

When the user starts an SSL VPN session which has virtual desktop enabled, the virtual desktop replaces the user's normal desktop. When the virtual desktop exits, the user's normal desktop is restored.

Virtual desktop requires the Fortinet cache cleaner plug in. If the plug in is not present, it is automatically downloaded to the client computer.

### To enable virtual desktop

To enable virtual desktop on the full-access portal and apply the application control list List1, for example, you would enter:

```
config vpn ssl web portal
 edit full-access
 set virtual-desktop enable
 set virtual-desktop-app-list List1
 end
```

### Configuring virtual desktop application control

You can control which applications users can run on their virtual desktop. To do this, you create an Application Control List of either allowed or blocked applications. When you configure the web portal, you select the list to use. Configure the application control list in the CLI.

#### To create an Application Control List - CLI

If you want to add BannedApp to List1, a list of blocked applications, you would enter:

```
config vpn ssl web virtual-desktop-app-list
 edit "List1"
 set action block
 config apps
 edit "BannedApp"
 set md5s "06321103A343B04DF9283B80D1E00F6B"
 end
 end
 end
```

### Configuring client OS Check

The SSLVPN client OS Check feature can determine if clients are running the Windows 2000, Windows XP, Windows Vista or Windows 7 operating system. You can configure the OS Check to do any of the following:

- allow the client access
- allow the client access only if the operating system has been updated to a specified patch (service pack) version
- deny the client access

The OS Check has no effect on clients running other operating systems.

#### To configure OS Check

OS Check is configurable only in the CLI.

```
config vpn ssl web portal
 edit <portal_name>
 set os-check enable
 config os-check-list {windows-2000 | windows-xp
 | windows-vista | windows-7}
 set action {allow | check-up-to-date | deny}
 set latest-patch-level {disable | 0 - 255}
 set tolerance {tolerance_num}
 end
 end
```

## Adding WINS and DNS services for clients

You can specify the WINS or DNS servers that are made available to SSL-VPN clients.

DNS servers provide the IP addresses that browsers need to access web sites. For Internet sites, you can specify the DNS server that your FortiGate unit uses. If SSL VPN users will access intranet sites using URLs, you need to provide them access to the intranet's DNS server. You specify a primary and a secondary DNS server.

A WINS server provides IP addresses for named servers in a Windows domain. If SSL VPN users will access a Windows network, you need to provide them access to the domain WINS server. You specify a primary and a secondary WINS server.

### To specify WINS and DNS services for clients - web-based manager

1. Go to *VPN > SSL > Config*.
2. Select the *Expand Arrow* to display the *Advanced* section.
3. Enter the IP addresses of DNS servers in the *DNS Server* fields as needed.
4. Enter the IP addresses of WINS servers in the *WINS Server* fields as needed.
5. Select *Apply*.

### To specify WINS and DNS services for clients - CLI

```
config vpn ssl settings
 set dns-server1 <address_ipv4>
 set dns-server2 <address_ipv4>
 set wins-server1 <address_ipv4>
 set wins-server2 <address_ipv4>
end
```

## Setting the idle timeout setting

The idle timeout setting controls how long the connection can remain idle before the system forces the remote user to log in again. For security, keep the default value of 300 seconds (5 minutes) or less.

### To set the idle timeout - web-based manager

1. Go to *VPN > SSL > Config*.
2. In the *Idle Timeout* field, enter the timeout value.  
The valid range is from 10 to 28800 seconds.
3. Select *Apply*.

### To set the idle timeout - CLI

```
config vpn ssl settings
 set idle-timeout <seconds_int>
end
```

## SSL VPN logs

Logging is available for SSP VPN traffic so you can monitor users connected to the FortiGate unit and their activity. For more information on configuring logs on the FortiGate unit, see the [Logging and Reporting](#) chapter of *The Handbook*.

### To enable logging of SSL VPN events - web-based manager

1. Go to *Log&Report > Log Config > Log Settings*.



2. Select *Enable*, and select *VPN activity event*.
3. Select *Apply*.

To view the SSL VPN log data, in the web-based manager, go to *Log&Report > Log & Archive Access* and select either the *Event Log* or *Traffic Log*.

In event log entries, look for the sub-types “sslvpn-session” and “sslvpn-user”.

For information about how to interpret log messages, see the [FortiGate Log Message Reference](#).

## Monitoring active SSL VPN sessions

You can go to *User & Device > Monitor* to view a list of active SSL VPN sessions. The list displays the user name of the remote user, the IP address of the remote client, and the time the connection was made. You can also see which services are being provided, and delete an active web session from the FortiGate unit.

### To monitor SSL VPNs - web-based manager

To view the list of active SSL VPN sessions, go to *VPN > SSL-VPN > Monitor*.

When a tunnel-mode user is connected, the *Description* field displays the IP address that the FortiGate unit assigned to the remote host.

If required, you can end a session/connection by selecting its check box and then selecting the *Delete* icon.

## Troubleshooting

Here is a list of common SSL VPN problems and the likely solutions.

<b>No response from SSL VPN URL</b>	Check SSL VPN port assignment (default 10443). Verify the SSL VPN security policy.
<b>Error: “The web page cannot be found.”</b>	Check URL: <code>https://&lt;FortiGate_IP&gt;:&lt;SSLVPN_port&gt;/remote/login</code>
<b>Tunnel connects, but there is no communication.</b>	Check that there is a static route to direct packets destined for the tunnel users to the SSL VPN interface. See <a href="#">“Routing for tunnel mode” on page 2207</a> .
<b>Tunnel-mode connection shuts down after a few seconds</b>	This issue occurs when there are multiple interfaces connected to the Internet, for example, a dual WAN configuration. Upgrade to the latest firmware then use the following CLI command: <pre>config vpn ssl settings     set route-source-interface enable end</pre>

---

**Error: “Destination address of Split Tunneling policy is invalid.”**

The SSL VPN security policy uses the ALL address as its destination. Specify the address of the protected network instead.

---

**When trying to connect using FortiClient the error message “Unable to logon to the server. Your user name or password may not be configured properly for this connection. (-12)” appears. When trying to login to the web portal, login and password are entered and login page will be sent back.**

Cookies must be enabled for SSL VPN to function in Web portal or with FortiClient.

Access to the web portal or tunnel will fail if Internet Explorer has the privacy Internet Options set to High. If set to High, Internet Explorer will:

Block cookies that do not have a compact privacy policy.

Block cookies that use personally identifiable information without your explicit consent.

---

# The SSL VPN client

The remote client connects to the SSL VPN tunnel in various ways, depending on the VPN configuration.

- Web mode requires nothing more than a web browser. Microsoft Internet Explorer, Firefox, and Apple Safari browsers are supported. For detailed information about supported browsers see the Release Notes for your FortiOS firmware.
- Tunnel mode establishes a connection to the remote protected network that any application can use. This requires FortiClient SSL VPN application that sends and receives data through the SSL VPN tunnel.

If the client computer runs Microsoft Windows, they can download the tunnel mode client from the web portal Tunnel Mode widget. After installing the client, they can start and stop tunnel operation from the Tunnel Mode widget, or open the tunnel mode client as a standalone application. The tunnel mode client is available on the Start menu at *All Programs > FortiClient > FortiClient SSL VPN*.

If the client computer runs Linux or Mac OS X, the user needs to download the tunnel mode client application from the Fortinet Support web site. See the Release Notes for your FortiOS firmware for the specific operating system versions that are supported. On Linux and Mac OS X platforms, tunnel mode operation cannot be initiated from the web portal Tunnel Mode widget. The remote user must use the standalone tunnel client application.

- The virtual desktop application creates a virtual desktop on a user's PC and monitors the data read/write activity of the web browser running inside the virtual desktop. When the application starts, it presents a 'virtual desktop' to the user. The user starts the web browser from within the virtual desktop and connects to the SSL VPN web portal. The browser file/directory operation is redirected to a new location, and the data is encrypted before it is written to the local disk. When the virtual desktop application exits normally, all the data written to the disk is removed. If the session terminates abnormally (power loss, system failure), the data left behind is encrypted and unusable to the user. The next time you start the virtual desktop, the encrypted data is removed.

## FortiClient

Remote users can use FortiClient software to initiate an SSL VPN tunnel to connect to the internal network. FortiClient uses local port TCP 1024 to initiate an SSL encrypted connection to the FortiGate unit, on port TCP 443. When connection using FortiClient, the FortiGate unit authenticates the FortiClient SSL VPN request based on the user group options. The FortiGate unit establishes a tunnel with the client and assigns a virtual IP address to the client PC. Once the tunnel has been established, the user can access the network behind the FortiGate unit.

FortiClient software is available for download at [www.forticlient.com](http://www.forticlient.com) and is available for Windows, Mac OS X, Apple iOS and Android.

## Tunnel mode client configuration

The FortiClient SSL VPN tunnel client requires basic configuration by the remote user to connect to the SSL VPN tunnel. When distributing the FortiClient software, provide the following information for the remote user to enter once the client software has been started. Once entered, they can select *Connect* to begin a SSL VPN session.

---

<b>Connection Name</b>	If you have pre-configured the connection settings, select the connection from the list and then select <i>Connect</i> . Otherwise, enter the settings in the fields below.
<b>Remote Gateway</b>	Enter the IP address or FQDN of the FortiGate unit that hosts the SSL VPN.
<b>Username</b>	Enter your user name.
<b>Client Certificate</b>	Use this field if the SSL VPN requires a certificate for authentication. Select the required certificate from the drop-down list. The certificate must be installed in the Internet Explorer certificate store.

---

# Setup examples


The examples in this chapter demonstrate the basic configurations needed for common connections to the SSL VPN tunnel and portals, applying the steps outlined in the chapter “Basic Configuration” on page 2195.

The example included are:

- [Secure internet browsing](#)
- [Split Tunnel](#)
- [Multiple user groups with different access permissions example](#)

## Secure internet browsing

This example sets up an SSL VPN tunnel to provide remote users the ability to access the Internet while travelling, and ensure that they are not subjected to malware and other dangers, by using the corporate firewall to filter all of their Internet traffic. Essentially, the remote user will connect to the corporate FortiGate unit to surf the Internet.



Using SSL VPN and FortiClient SSL VPN software, you create a means to use the corporate FortiGate to browse the web safely.

### Creating an SSL VPN IP pool and SSL VPN web portal

1. Go to *VPN > SSL > Config* and for *IP Pools* select *SSLVPN\_TUNNEL\_ADDR1*.
2. Create the SSL VPN portal to by going to *VPN > SSL > Portal* and selecting *tunnel-access* in the upper right-hand corner drop-down list box.
3. Select *OK*.

### Creating the SSL VPN user and user group

Create the SSL VPN user and add the user to a user group configured for SSL VPN use.

1. Go to *User & Device > User > User Definition* and select *Create New* to add the user:

---

<b>User Name</b>	twhite
------------------	--------

---

<b>Password</b>	password
-----------------	----------

---

2. Select *OK*.
3. Go to *User & Device > User > User Groups* and select *Create New* to add *twhite* to a group called *SSL VPN*:

---

<b>Name</b>	SSL Group
-------------	-----------

---

<b>Type</b>	Firewall
-------------	----------

---

4. Move *twhite* to the *Members* list.
5. Select *OK*.

## Creating a static route for the remote SSL VPN user

Create a static route to direct traffic destined for tunnel users to the SSL VPN tunnel.

1. Go to *Router > Static > Static* and select *Create New* to add the static route.  
For low-end FortiGate units, go to *System > Network > Routing* and select *Create New*.

---

<b>Destination IP/Mask</b>	10.212.134.0/255.255.255.0
----------------------------	----------------------------

---

<b>Device</b>	ssl.root
---------------	----------

---



The *Destination IP/Mask* matches the network address of the remote SSL VPN user.

2. Select *OK*.

## Creating security policies

Create an SSL VPN security policy with SSL VPN user authentication to allow SSL VPN traffic to enter the FortiGate unit. Create a normal security policy from *ssl.root* to *wan1* to allow SSL VPN traffic to connect to the Internet.

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Select the *Policy Type* as *VPN* and the *Policy Subtype* as *SSL-VPN*.
3. to add the SSL VPN security policy:

---

<b>Incoming Interface</b>	wan1
---------------------------	------

---

<b>Remote Address</b>	all
-----------------------	-----

---

<b>Local Interface</b>	ssl.root
------------------------	----------

---

<b>Local Protected Subnet</b>	all
-------------------------------	-----

---

4. Select *Create New* for Configure SSL-VPN Authentication Rules and add an authentication rule for the remote user:

<b>Selected User Groups</b>	Tunnel
<b>Selected Services</b>	All
<b>Schedule</b>	always
<b>SSL-VPN Portal</b>	tunnel-access

5. Select *OK*.
6. Select *Create New* to add a security policy that allows remote SSL VPN users to connect to the Internet
7. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.

<b>Incoming Interface</b>	ssl.root
<b>Source Address</b>	all
<b>Outgoing Interface</b>	wan1
<b>Destination Address</b>	all
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT

8. Select *OK*.

## Results


Using FortiClient SSLVPN application, log into the VPN using the address `https://172.20.120.136:443/` and log in as `twhite`. Once connected, you can browse the Internet.

From the FortiGate web-based manager go to *VPN > Monitor > SSL-VPN Monitor* to view the list of users connected using SSL VPN. The *Subsession* entry indicates the split tunnel which redirects to the Internet.

## Split Tunnel

For this example, the remote users are configured to be able to securely access head office internal network servers, and browse the Internet through the head office firewall. This will enable the remote user to use the FortiGate security to connect to the internal network and the web.

This solution describes how to configure FortiGate SSL VPN split tunnelling using the FortiClient SSL VPN software, available from the Fortinet Support site.



Using split tunneling, all communication from remote SSL VPN users to the head office internal network and to the Internet uses an SSL VPN tunnel between the user's PC and the head office FortiGate unit. Connections to the Internet are routed back out the head office FortiGate unit to the Internet. Replies come back into the head office FortiGate unit before being routed back through the SSL VPN tunnel to the remote user.

## Creating a firewall address for the head office server

1. Go to *Firewall Objects > Address > Addresses* and select *Create New* and add the head office server address:

<b>Name</b>	Head office server
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	192.168.1.12
<b>Interface</b>	Internal

2. Select *OK*.

## Creating an SSL VPN IP pool and SSL VPN web portal

1. Go to *VPN > SSL > Config*.
2. For *IP Pools* select *SSLVPN\_TUNNEL\_ADDR1*.
3. Create the SSL VPN portal to by going to *VPN > SSL > Portal* and select the plus sign in the upper right of the window.
4. Enter the following:

<b>Name</b>	Connect to head office server
<b>IP Pools</b>	SSLVPN_TUNNEL_ADDR1
<b>Enable Tunnel Mode</b>	Enable
<b>Split Tunneling</b>	Enable



5. Select *OK*.

### Creating the SSL VPN user and user group

Create the SSL VPN user and add the user to a user group.

1. Go to *User & Device > User > User Definition*, select *Create New* and add the user:

<b>User Name</b>	twhite
<b>Password</b>	password

2. Select *OK*.
3. Go to *User & Device > User > User Groups* and select *Create New* to add **twhite** to the SSL VPN user group:

<b>Name</b>	Tunnel
<b>Type</b>	Firewall

4. Move **twhite** to the *Members* list.
5. Select *OK*.

### Creating a static route for the remote SSL VPN user

Create a static route to direct traffic destined for tunnel users to the SSL VPN tunnel.

1. Go to *Router > Static > Static* and select *Create New*
2. For low-end FortiGate units, go to *System > Network > Routing* and select *Create New*:

<b>Destination IP/Mask</b>	10.212.134.0/255.255.255.0
<b>Device</b>	ssl.root

3. Select *OK*.

### Creating security policies

Create an SSL VPN security policy with SSL VPN user authentication to allow SSL VPN traffic to enter the FortiGate unit. Create a normal security policy from ssl.root to wan1 to allow SSL VPN traffic to connect to the Internet.

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Select the *Policy Type* of VPN and the *Policy Subtype* as SSL-VPN.
3. Complete the following:

<b>Incoming Interface</b>	wan1
<b>Remote Address</b>	all
<b>Local Interface</b>	internal
<b>Local Protected Subnet</b>	Head office server

- Under *Configure SSL-VPN Authentication Rules* select *Create New* to add an authentication rule for the remote user:

<b>Groups(s)</b>	Tunnel
<b>Service</b>	ALL
<b>Schedule</b>	always

- Select *OK*.  
Add a security policy that allows remote SSL VPN users to connect to the Internet.
- Select *Create New*.
- Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
- Complete the following and select *OK*:

<b>Incoming Interface</b>	ssl.root
<b>Source Address</b>	all
<b>Outgoing Interface</b>	wan1
<b>Destination Address</b>	all
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT

## Results


Using the FortiClient SSL VPN application on the remote PC, connect to the VPN using the address `https://172.20.120.136:443/` and log in with the `twhite` user account. Once connected, you can connect to the head office server or browse to web sites on the Internet.

From the web-based manager go to *VPN > Monitor > SSL-VPN Monitor* to view the list of users connected using SSL VPN. The *Subsession* entry indicates the split tunnel which redirects SSL VPN sessions to the Internet.

## Multiple user groups with different access permissions example

You might need to provide access to several user groups with different access permissions. Consider the following example topology in which users on the Internet have controlled access to servers and workstations on private networks behind a FortiGate unit.

**Figure 313:**SSL VPN configuration for different access permissions by user group



In this example configuration, there are two users:

- user1 can access the servers on Subnet\_1
- user2 can access the workstation PCs on Subnet\_2

You could easily add more users to either user group to provide them access to the user group's assigned web portal.

## General configuration steps

1. Create firewall addresses for
  - the destination networks
  - two non-overlapping tunnel IP address ranges that the FortiGate unit will assign to tunnel clients in the two user groups
2. Create two web portals.
3. Create two user accounts, user1 and user2.
4. Create two user groups. For each group, add a user as a member and select a web portal. In this example, user1 will belong to group1, which will be assigned to portal1.
5. Create security policies:
  - two SSL VPN security policies, one to each destination
  - two tunnel-mode policies to allow each group of users to reach its permitted destination network
6. Create the static route to direct packets for the users to the tunnel.

## Creating the firewall addresses

Security policies do not accept direct entry of IP addresses and address ranges. You must define firewall addresses in advance.

## Creating the destination addresses

SSL VPN users in this example can access either Subnet\_1 or Subnet\_2.

### To define destination addresses - web-based manager

1. Go to *Firewall Objects > Address > Addresses*.
2. Select *Create New*, enter the following information, and select *OK*:

<b>Name</b>	Subnet_1
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	10.11.101.0/24
<b>Interface</b>	port2

3. Select *Create New*, enter the following information, and select *OK*:

<b>Name</b>	Subnet_2
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	10.11.201.0/24
<b>Interface</b>	port3

### Creating the tunnel client range addresses

To accommodate the two groups of users, split an otherwise unused subnet into two ranges. The tunnel client addresses must not conflict with each other or with other addresses.

### To define tunnel client addresses - web-based manager

1. Go to *Firewall Objects > Address > Addresses*.
2. Select *Create New*, enter the following information, and select *OK*:

<b>Name</b>	Tunnel_group1
<b>Type</b>	IP Range
<b>Subnet / IP Range</b>	10.11.254.[1-50]
<b>Interface</b>	Any

3. Select *Create New*, enter the following information, and select *OK*.

<b>Name</b>	Tunnel_group2
<b>Type</b>	IP Range
<b>Subnet / IP Range</b>	10.11.254.[51-100]
<b>Interface</b>	Any

### Creating the web portals

To accommodate two different sets of access permissions, you need to create two web portals, portal1 and portal2, for example. Later, you will create two SSL VPN user groups, one to assign to portal1 and the other to assign to portal2.

### To create the portal1 web portal

1. Go to *VPN > SSL > Portal* and select the plus icon in the upper right corner.
2. Enter `portal1` in the *Name* field.
3. In *Applications*, select all of the application types that the users can access.
4. In *IP Pools*, select *Tunnel\_group1*.
5. Select *OK*.

### To create the portal2 web portal

1. Go to *VPN > SSL > Portal* and select the plus icon in the upper right corner.
2. Enter `portal2` in the *Name* field and select *OK*.
3. In *Applications*, select all of the application types that the users can access.
4. In *IP Pools*, select *Tunnel\_group2*.
5. Select *OK*.

Later, you can configure these portals with bookmarks and enable connection tool capabilities for the convenience of your users.

## Creating the user accounts and user groups

After enabling SSL VPN and creating the web portals that you need, you need to create the user accounts and then the user groups that require SSL VPN access.

Go to *User & Device > User > User Definition* and create `user1` and `user2` with password authentication. After you create the users, create the SSL VPN user groups.

### To create the user groups - web-based manager

1. Go to *User & Device > User > User Groups*.
2. Select *Create New* and enter the following information:

<b>Name</b>	group1
<b>Type</b>	Firewall

3. From the *Available* list, select `user1` and move it to the *Members* list by selecting the right arrow button.
4. Select *OK*.
5. Repeat steps 2 through 4 to create `group2`, assigned to `portal2`, with `user2` as its only member.

## Creating the security policies

You need to define security policies to permit your SSL VPN clients, web-mode or tunnel-mode, to connect to the protected networks behind the FortiGate unit. Before you create the security policies, you must define the source and destination addresses to include in the policy. See [“Creating the firewall addresses” on page 2227](#).

Two types of security policy are required:

- An SSL VPN policy enables clients to authenticate and permits a web-mode connection to the destination network. In this example, there are two destination networks, so there will be

two SSL VPN policies. The authentication, ensures that only authorized users access the destination network.

- A tunnel-mode policy is a regular ACCEPT security policy that enables traffic to flow between the SSL VPN tunnel interface and the protected network. Tunnel-mode policies are required if you want to provide tunnel-mode connections for your clients. In this example, there are two destination networks, so there will be two tunnel-mode policies.

**To create the SSL VPN security policies - web-based manager**

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Select the *Policy Type* of *VPN* and the *Policy Subtype* as *SSL-VPN*.
3. Enter the following information:

<b>Incoming Interface</b>	port1
<b>Remote Address</b>	All
<b>Local Interface</b>	port2
<b>Local Protected Interface</b>	Subnet_1

4. Under *Configure SSL-VPN Authentication Rules*, select *Create New* and enter the following information:

<b>Group(s)</b>	group1
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>SSL-VPN Portal</b>	portal1

5. Select *OK*, and then select *OK* again.
6. Select *Create New*.
7. Select the *Policy Type* of *VPN* and the *Policy Subtype* as *SSL-VPN*.
8. Enter the following information:

<b>Incoming Interface</b>	port1
<b>Remote Address</b>	All
<b>Local Interface</b>	port3
<b>Local Protected Interface</b>	Subnet_2

9. Under *Configure SSL-VPN Authentication Rules*, select *Create New* and enter the following information:

<b>Group(s)</b>	group2
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>SSL-VPN Portal</b>	portal1

10. Select *OK*, and then select *OK* again.

**To create the tunnel-mode security policies - web-based manager**

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	sslvpn tunnel interface (ssl.root)
<b>Source Address</b>	Tunnel_group1
<b>Outgoing Interface</b>	port2
<b>Destination Address</b>	Subnet_1
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Enable

4. Select *Create New*.
5. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
6. Enter the following information, and select *OK*:

<b>Incoming Interface</b>	sslvpn tunnel interface (ssl.root)
<b>Source Address</b>	Tunnel_group2
<b>Outgoing Interface</b>	port3
<b>Destination Address</b>	Subnet_2
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	Enable

## Create the static route to tunnel mode clients

Reply packets destined for tunnel mode clients must pass through the SSL VPN tunnel. You need to define a static route to accomplish this.

**To add a route to SSL VPN tunnel mode clients - web-based manager**

1. Go to *Router > Static > Static Routes* and select *Create New*.  
For low-end FortiGate units, go to *System > Network > Routing* and select *Create New*.

2. Enter the following information and select *OK*.

---

**Destination IP/Mask** 10.11.254.0/24

This IP address range covers both ranges that you assigned to SSL VPN tunnel-mode users. See [“Creating the tunnel client range addresses”](#) on page 2228.

---

**Device** Select the SSL VPN virtual interface, *ssl.root* for example.

---



In this example, the *IP Pools* field on the *VPN > SSL > Config* page is not used because each web portal specifies its own tunnel IP address range

---



# Chapter 17 Traffic Shaping for FortiOS

## 5.0

With the ever-increasing demands on network systems for a number of protocols, including email, HTTP traffic both internally and externally to the internet, voice over IP, FTP, and more, slow traffic is becoming a reality. Important traffic may even be dropped or slowed to an unusable speed. Web traffic delays can result in a loss of revenue for businesses.

Traffic shaping attempts to normalize traffic peaks and bursts to prioritize certain flows over others. There is a physical limitation to the amount of data which can be buffered and to the length of time it can be buffered.

FortiGate units provide Quality of Service (QoS) by applying bandwidth limits and prioritization. Using traffic shaping, you can adjust how your FortiGate unit allocates resources to different traffic types to improve performance and stability of latency sensitive or bandwidth intensive network applications.

This document describes Quality of Service (QoS), traffic shaping, FortiGate traffic shaping algorithms, and includes configuration procedures for traffic shaping on FortiGate units.

This FortiOS Handbook chapter contains the following sections:

[The purpose of traffic shaping](#) describes traffic shaping theories and quality of service.

[Traffic shaping methods](#) lists different methods of applying traffic shaping within FortiOS, and explains how to use TOS and Differentiated Services.

[Examples](#) provides basic application scenarios for shapers.

[Troubleshooting traffic shaping](#) lists diagnose commands to use for determining if traffic shapers are working correctly.

# The purpose of traffic shaping

Traffic shaping, or traffic management, controls the bandwidth available and sets the priority of traffic processed by the policy to control the volume of traffic for a specific period (bandwidth throttling) or rate the traffic is sent (rate limiting).

Traffic shaping attempts to normalize traffic peaks and bursts to prioritize certain flows over others. But there is a physical limitation to the amount of data which can be buffered and to the length of time. Once these thresholds have been surpassed, frames and packets will be dropped, and sessions will be affected in other ways.

A basic traffic shaping approach is to prioritize certain traffic flows over other traffic whose potential loss is less disadvantageous. This would mean that you accept certain sacrifices in performance and stability on low-priority traffic, to increase or guarantee performance and stability to high-priority traffic.

If, for example, you are applying bandwidth limitations to certain flows, you must accept the fact that these sessions can be limited and therefore negatively impacted.

Note that traffic shaping is effective for normal IP traffic at normal traffic rates. Traffic shaping is not effective during periods when traffic exceeds the capacity of the FortiGate unit. Because packets must be received by the FortiGate unit before they are subject to traffic shaping, if the FortiGate unit cannot process all of the traffic it receives, then dropped packets, delays, and latency are likely to occur.

To ensure that traffic shaping is working at its best, make sure that the interface Ethernet statistics show no errors, collisions or buffer overruns.

Accelerated interfaces (NPx network processors and CE) affect traffic shaping. For more information, see the [FortiGate Hardware](#) Guide.

## Quality of Service

Quality of Service (QoS) is the capability to adjust some quality aspects of your overall network traffic. This can include such techniques as priority-based queuing and traffic policing. Because bandwidth is finite and because some types of traffic are slow, jitter or packet loss sensitive, bandwidth intensive, or operation critical, QoS can be a useful tool for optimizing the performance of the various applications on your network.

Before implementing QoS, organizations should first identify the types of traffic that are important to the organization, the types of traffic that use high amounts of bandwidth, and the types of traffic that are sensitive to latency or packet loss.

For example, a company might want to guarantee sufficient bandwidth for revenue producing e-commerce traffic. They need to ensure that transactions can be completed and that clients do not experience service delays and interruptions. At the same time, the company may need to ensure low latency for voice over IP (VoIP) traffic used by sales and customer support, while traffic latency and bursts may be less critical to the success of other network applications such as long term, resumable file transfers. Many organizations discover that QoS is especially important for managing their voice and streaming multi-media traffic. These types of traffic can rapidly consume bandwidth and are sensitive to latency.

Discovering the needs and relative importance of each traffic type on your network will help you to design an appropriate overall approach, including how you will configure each available QoS component technique. Some organizations discover that they only need to configure bandwidth limits for some services. Other organizations determine that they need to fully configure

interface and security policy bandwidth limits for all services, and prioritize queuing of critical services relative to traffic rate.

You can implement QoS on FortiGate units using the following techniques:

---

<b>Traffic policing</b>	Drops packets that do not conform to bandwidth limitations.
<b>Traffic shaping</b>	Ensures that the traffic may consume bandwidth at least at the guaranteed rate by assigning a greater priority queue if the guarantee is not being met. Also ensures that the traffic cannot consume bandwidth greater than the maximum at any given instant in time. Flows greater than the maximum rate are subject to traffic policing.
<b>Queuing</b>	Transmits packets in order of their assigned priority queue for that physical interface. All traffic in a higher priority traffic queue must be completely transmitted before traffic in lower priority queues will be transmitted.

---

When deciding how to configure QoS techniques, it can be helpful to know when FortiGate units employ each technique in the overall traffic processing flow, and the considerations that arise from those mechanisms.

## Traffic policing

The FortiGate unit begins to process traffic as it arrives (ingress) and departs (egress) on an interface. In later phases of the network processing, such as enforcing maximum bandwidth use on sessions handled by a security policy, if the current rate for the destination interface or traffic regulated by that security policy is too high, the FortiGate unit may drop the packet. Time spent on prior processing, such as web filtering, decryption or IPS, is often wasted on packets that are not forwarded. This applies to VLAN interfaces and physical interfaces.

You can prevent this wasted effort on ingress by configuring the FortiGate unit to preemptively drop excess packets when they are received at the source interface, before most other traffic processing is performed:

```
config system interface
 edit <interface_name>
 set inbandwidth <rate_int>
 next
end
```

where <rate\_int> is the bandwidth limit in Kb/s. Excess packets will be dropped. If inbandwidth is 0, the rate is not limited.

A similar command is available that can be performed on egress as well using the CLI commands:

```
config system interface
 edit <interface_name>
 set outbandwidth <rate_int>
 next
end
```

As with ingress, setting the rate to 0 (zero) sets the rate to unlimited.

Rate limiting traffic accepted by the interface enables you to restrict incoming traffic to rates that, while no longer the full capacity of the interface, at the traffic shaping point in the processing are more likely to result in acceptable rates of outgoing traffic per destination

interface or all security policies. This conserves FortiGate processing resources for those packets that are more likely to be viable completely to the point of egress.

Excessive traffic policing can degrade network performance rather than improve it. For details on factors you may want to consider when configuring traffic policing, see [“Important considerations” on page 2241](#).

## Bandwidth guarantee, limit, and priority interactions

After packet acceptance, the FortiGate unit classifies traffic and may apply traffic policing at additional points during processing. It may also apply additional QoS techniques, such as prioritization and traffic shaping. Traffic shaping consists of a mixture of traffic policing to enforce bandwidth limits, and priority queue adjustment to assist packets in achieving the guaranteed rate.

If you have configured prioritization, the FortiGate unit prioritizes egressing packets by distributing them among FIFO (first in, first out) queues associated with each possible priority number. Each physical interface has six priority queues. Virtual interfaces do not have their own queues, and instead use the priority queues of the physical interface to which they are bound.

Each physical interface's six queues are queue 0 to queue 5, where queue 0 is the highest priority queue. However, for the reasons described below, you may observe that your traffic uses only a subset of those six queues. Some traffic may always use a certain queue number. Some queuing may vary by the packet rate or mixture of services. Some queue numbers may be used only by through traffic for which you have configured traffic shaping in the security policy that applies to that traffic session. For example:

- Administrative access traffic will always use queue 0.
- Traffic matching security policies **without** traffic shaping may use queue 0, queue 1, or queue 2. Which queue will be used depends on the priority value you have configured for packets with that ToS (type of service) bit value, if you have configured ToS-based priorities.
- Traffic matching security policies **with** traffic shaping may use any queue. Which queue will be used depends on whether the packet rate is currently below the guaranteed bandwidth (queue 0), or above the guaranteed bandwidth. Packets at rates greater than the maximum bandwidth limit are dropped.
- If the global tos-based-priority is low (3), the priority in a traffic-shaper is medium (2) and a packet flows through a policy that refers to the shaper, the packet will be assigned the priority defined by the shaper, in this case medium (2).

Prioritization and traffic shaping behavior varies by your configuration, the service types and traffic volumes, and by whether the traffic is through traffic, or the traffic originates from or terminates at the FortiGate unit itself.

### FortiGate traffic

Administrative access to the FortiGate through HTTPS or SSH, or IPsec tunnel negotiations, security policies do not apply, and therefore FortiGate units do not apply traffic shaping. Such traffic also uses the highest priority queue, queue 0. In other words:

packet priority = 0

Exceptions to this rule include traffic types that are connections related to a session governed by a security policy.

For example, if you have enabled scanning by FortiGuard antivirus, traffic from the sender technically terminates at the FortiGate proxy that scans that traffic type; the FortiGate unit initiates a second connection that transmits scanned content to its destination. Because the second connection's traffic is technically originating from the FortiGate proxy and therefore the

FortiGate unit itself, it uses the highest priority queue, queue 0. However, this connection is logically associated with through traffic, and is therefore subject to possible bandwidth enforcement and guarantees in its governing security policy. In this way, it behaves partly like other through traffic.

## Through traffic

For traffic passing through the FortiGate unit, the method a FortiGate unit uses to determine the priority queue varies by whether you have enabled Traffic Shaping. Packets may or may not use a priority queue directly or indirectly derived from the type of service (ToS) bit — sometimes used instead with differentiated services — in the packet's IP header.

If Traffic Shaping is not enabled in the security policy, the FortiGate unit neither limits nor guarantees bandwidth, and traffic for that session uses the priority queue determined directly by matching the ToS bit in its header with your configured values:

```
config system global
 set tos-based-priority {high | low | medium}
end
```

or, if you have configured a priority specifically for that TOS bit value:

```
config system tos-based-priority
 edit <id_int>
 set tos [0-15]
 set priority {high | low | medium}
 next
end
```

where `tos` is the value of the ToS bit in the packet's IP header, and `high` has a priority value of 0 and `low` is 2. Priority values configured in the second location will override the global ToS-based priority. In other words:

`packet priority = ToS-based priority`

For example, you might specify that packets with a ToS bit value of 2 should use queue 0, the highest priority queue:


```
config system tos-based-priority
 edit 15
 set tos 2
 set priority high
 next
end
```

If Traffic Shaping is enabled in the security policy using shared traffic shapers, the FortiGate unit may instead or also subject packets to traffic policing, or priority queue increase in an effort to meet bandwidth guarantees configured in the shaper:

```
config firewall shaper traffic-shaper
 edit <shaper_name>
 ...
 set priority {high | medium | low}
 set maximum-bandwidth <rate>
 set guaranteed-bandwidth <rate>
 end
```

where `high` has a priority value of 1 and `low` is 3, and `<rate>` is the bandwidth limit in kilobits per second.

**Figure 314:**Traffic queuing as packet rate increases



- If the current packet rate is less than Guaranteed Bandwidth, packets use priority queue 0. In other words:  
packet priority = 0
- If the current packet rate is greater than Guaranteed Bandwidth but less than Maximum Bandwidth, the FortiGate unit assigns a priority queue by adding the numerical value of the security policy-based priority, where the value of High is 1, and Low is 3, with the numerical value of the ToS-based priority, where high has a priority value of 0 and low is 2. Because the two values are added, depending on the your configured ToS-based priorities, packets in this category could use queues from queue 1 to queue 5. In other words:  
packet priority = ToS-based priority + security policy-based priority  
For example, if you have enabled Traffic Shaping in the security policy, and the security policy's Traffic Priority is Low (value 3), and the priority normally applied to packets with that ToS bit is medium (value 1), then packets have a total packet priority of 4, and use priority queue 4.
- If the current packet rate exceeds Maximum Bandwidth, excess packets are dropped.

### Calculation and regulation of packet rates

Packet rates specified for Maximum Bandwidth or Guaranteed Bandwidth are:

$$\text{rate} = \text{amount} / \text{time}$$

where rate is expressed in kilobits per second (Kb/s).

Burst size at any given instant cannot exceed the amount configured in Maximum Bandwidth. Packets in excess are dropped. Packets deduct from the amount of bandwidth available to subsequent packets and available bandwidth regenerates at a fixed rate. As a result, bandwidth available to a given packet may be less than the configured rate, down to a minimum of 0 Kb/s.

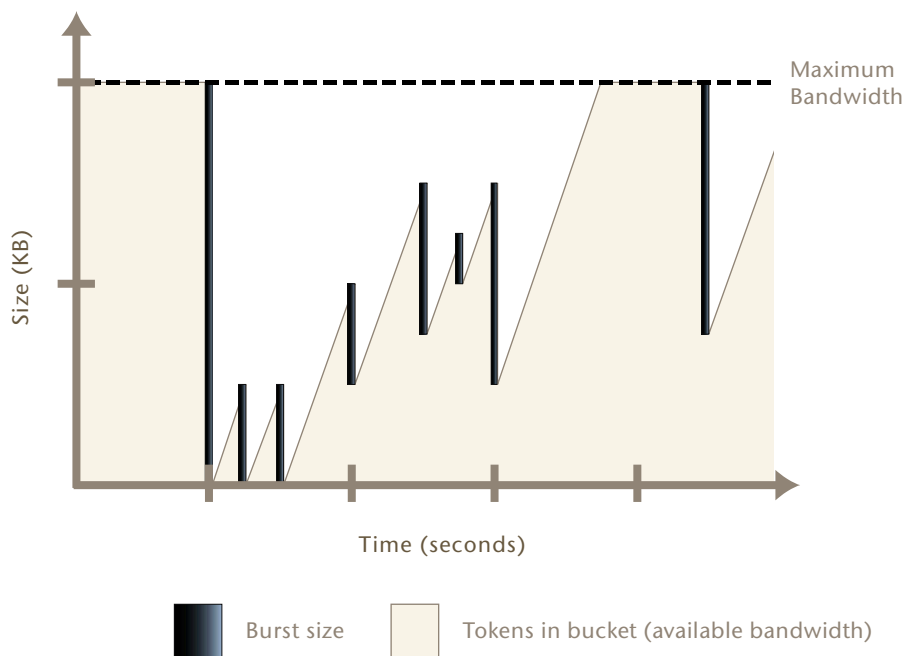
Rate calculation and behavior can alternatively be described using the token bucket metaphor, where:

- a traffic flow has an associated bucket, which represents burst size bounds, and is the size of your configured bandwidth limit
- the bucket receives tokens, which represent available bandwidth, at the fixed configured rate
- as time passes, tokens are added to the bucket, up to the capacity of the bucket; excess tokens are discarded
- when a packet arrives, the packet must deduct bandwidth tokens from the bucket equal to its packet size in order to egress
- packets cannot egress if there are insufficient tokens to pay for its egress; these nonconforming packets are dropped

Bursts are not redistributed over a longer interval, so bursts are propagated rather than smoothed, although their peak size is limited.

Maximum burst size is the capacity of the bucket (the configured bandwidth limit); actual size varies by the current number of tokens in the bucket, which may be less than bucket capacity, due to deductions from previous packets and the fixed rate at which tokens accumulate. A depleted bucket refills at the rate of your configured bandwidth limit. Bursts cannot borrow tokens from other time intervals. This behavior is illustrated in [Figure 315 on page 2239](#).

**Figure 315:**Bursts and bandwidth limits over time



By limiting traffic peaks and token regeneration in this way, the available bandwidth at a given moment may be less than bucket capacity, but your limit on the total amount per time interval is ensured. That is, total bandwidth use during each interval of 1 second is at most the integral of your configured rate.

You may observe that external clients, such as FTP or BitTorrent clients, initially report rates between Maximum Bandwidth and twice that of Maximum Bandwidth, depending on the size of their initial burst. This is notably so when a connection is initiated following a period of no network activity. The apparent discrepancy in rates is caused by a difference of perspective when delimiting time intervals. A burst from the client may initially consume all tokens in the

bucket, and before the end of 1 second, as the bucket regenerates, be allowed to consume almost another bucket's worth of bandwidth. From the perspective of the client, this constitutes one time interval. From the perspective of the FortiGate unit, however, the bucket cannot accumulate tokens while full; therefore, the time interval for token regeneration begins **after** the initial burst, and does not contain the burst. These different points of reference result in an initial discrepancy equal to the size of the burst — the client's rate contains it, but the FortiGate unit's rate does not. If the connection is sustained to its limit and time progresses over an increasing number of intervals, however, this discrepancy decreases in importance relative to the bandwidth total, and the client's reported rate will eventually approach that of the FortiGate unit's configured rate limit.

For example, your Maximum Bandwidth might be 50 Kb/s and there has been no network activity for one or more seconds. The bucket is full. A burst from an FTP client immediately consumes 50 Kb. Because the bucket completely regenerates over 1 second, by the time almost another 1 second has elapsed from the initial burst, traffic can consume another 49.999 Kb, for a total of 99.999 Kb between the two points in time. From the vantage point of an external FTP client regulated by this bandwidth limit, it therefore initially appears that the bandwidth limit is 99.999 Kb/s, almost twice the configured limit of 50 Kb/s. However, bucket capacity only regenerates at your configured rate of 50 Kb/s, and so the connection can only consume a maximum of 50 Kb during each second thereafter. The result is that as bandwidth consumption is averaged over an increasing number of time intervals, each of which are limited to 50 Kb/s, the effects of the first interval's doubled bandwidth size diminishes proportionately, and the client's reported rate eventually approach your configured rate limit. This effect is illustrated in [Table 102 on page 2240](#).

**Table 102:**Effects of a 50 Kb/s limit on client reported rates

Total size transferred (Kb)	Time (s)	Rate reported by client (Kb/s)
99.999 (50 + 49.999)	1	99.999
149.999	2	74.999
199.999	3	66.666
249.999	4	62.499
299.999	5	59.998
349.999	6	58.333
...	...	...

Guaranteed Bandwidth can also be described using a token bucket metaphor. However, because this feature attempts to achieve or exceed a rate rather than limit it, the FortiGate unit does not discard non-conforming packets, as it does for Maximum Bandwidth; instead, when the flow does not achieve the rate, the FortiGate unit increases the packets' priority queue, in an effort to increase the rate.

Guaranteed and maximum bandwidth rates apply to the bidirectional total for all sessions controlled by the security policy. For example, an FTP connection may entail two separate connections for the data and control portion of the session; some packets may be reply traffic rather than initiating traffic. All packets for both connections are counted when calculating the packet rate for comparison with the guaranteed and maximum bandwidth rate.



## Important considerations

In essence, by implementing QoS, you trade some performance and/or stability from traffic X by discarding packets or introducing latency in order to improve performance and stability of traffic Y. The best traffic shaping configuration for your network will appropriately balance the needs of each traffic flow by considering not only the needs of your particular organization, but also the resiliency and other characteristics of each particular service.

For example, you may find that web browsing traffic is both more resistant to interruptions or latency and less business critical than UDP or VoIP traffic, and so you might implement less restrictive QoS measures on UDP or VoIP traffic than on HTTP traffic.

An appropriate QoS configuration will also take into account the physical limits of your network devices, and the interactions of the aforementioned QoS mechanisms, described in [“Bandwidth guarantee, limit, and priority interactions” on page 2236](#).

You may choose to configure QoS differently based upon the hardware limits of your network and FortiGate unit. Traffic shaping may be less beneficial in extremely high-volume situations where traffic exceeds a network interface’s or your FortiGate model’s overall physical capacity. A FortiGate unit must have sufficient resources, such as memory and processing power, to process all traffic it receives, and to process it at the required rate; if it does not have this capacity, then dropped packets and increased latency are likely to occur. For example, if the total amount of memory available for queuing on a physical interface is frequently exceeded by your network’s typical packet rates, frames and packets must be dropped. In such a situation, you might choose to implement QoS using a higher model FortiGate unit, or to configure an incoming bandwidth limit on each interface.

Incorrect traffic shaping configurations can actually further degrade certain network flows, because excessive discarding of packets or increased latency beyond points that can be gracefully handled by that protocol can create additional overhead at upper layers of the network, which may be attempting to recover from these errors. For example, a configuration might be too restrictive on the bandwidth accepted by an interface, and may therefore drop too many packets, resulting in the inability to complete or maintain a SIP call.

To optimize traffic shaping performance, first ensure that the network interface’s Ethernet statistics are clean of errors, collisions, or buffer overruns. To check the interface, enter the following diagnose command to see the traffic statistics:

```
diagnose hardware deviceinfo nic <port_name>
```

If these are not clean, adjust FortiGate unit and settings of routers or other network devices that are connected to the FortiGate unit. For additional information, see [“Troubleshooting traffic shaping” on page 2266](#).

Once Ethernet statistics are clean, you may want to use only some of the available FortiGate QoS techniques, or configure them differently, based upon the nature of FortiGate QoS mechanisms described in [“Bandwidth guarantee, limit, and priority interactions” on page 2236](#). Configuration considerations include:

- For maximum bandwidth limits, ensure that bandwidth limits at the source interface and/or the security policy are not too low, which can cause the FortiGate unit to discard an excessive number of packets.
- For prioritization, consider the ratios of how packets are distributed between available queues, and which queue is used by which types of services. If you assign most packets to the same priority queue, it negates the effects of configuring prioritization. If you assign many high bandwidth services to high priority queues, lower priority queues may be starved for bandwidth and experience increased or indefinite latency. For example, you may want to prioritize a latency-sensitive service such as SIP over a bandwidth-intensive service such as

FTP. Consider also that bandwidth guarantees can affect the queue distribution, assigning packets to queue 0 instead of their typical queue in high-volume situations.

- You may or may not want to guarantee bandwidth, because it causes the FortiGate unit to assign packets to queue 0 if the guaranteed packet rate is not currently being met. Comparing queuing behavior for lower-bandwidth and higher-bandwidth situations, this would mean that effects of prioritization only become visible as traffic volumes rise and exceed their guarantees. Because of this, you might want only some services to use bandwidth guarantees, to avoid the possibility that in high-volume situations all traffic uses the same queue, thereby negating the effects of configuring prioritization.
- For prioritization, configure prioritization for all through traffic. You may want to configure prioritization by either ToS-based priority or security policy priority, but not both. This simplifies analysis and troubleshooting.

Traffic subject to both security policy and ToS-based priorities will use a combined priority from both of those parts of the configuration, while traffic subject to only one of the prioritization methods will use only that priority. If you configure both methods, or if you configure either method for only a subset of your traffic, packets for which a combined priority applies will frequently receive a lower priority queue than packets for which you have only configured one priority method, or for which you have not configured prioritization.

For example, if both ToS-based priority and security policy priority both dictate that a packet should receive a “medium” priority, in the absence of bandwidth guarantees, a packet will use queue 3, while if only ToS-based priority had been configured, the packet would have used queue 1, and if only security policy-based priority had been configured, the packet would have used queue 2. If no prioritization had been configured at all, the packet would have used queue 0.

For example alternative QoS implementations that illustrate these considerations, see [“Examples” on page 2259](#)

# Traffic shaping methods

In FortiOS, there are three types of traffic shaping configuration. Each has a specific function, and all can be used together in varying configurations. Policy shaping enables you to define the maximum bandwidth and guaranteed bandwidth set for a security policy. Per-IP shaping enables you to define traffic control on a more granular level. Application traffic shaping goes further, enabling traffic controls on specific applications or application groupings.

This chapter describes the types of traffic shapers and how to configure them in the web-based manager and the CLI.

## Traffic shaping options

When configuring traffic shaping for your network, there are three different methods to control the flow of network traffic to ensure that the desired traffic gets through while also limiting the bandwidth that users use for other less important or bandwidth consuming traffic. The three shaping options are:

- shared policy shaping - bandwidth management by security policies
- per-IP shaping - bandwidth management by user IP addresses
- application control shaping - bandwidth management by application

Shared policy shaping and per IP shaping are enabled within the security policy, while the application control shaping is configured in *Security Profiles > Application Control > Application Sensors*, and enabled in the security policy by enabling *Application Control* in the *Security Profiles* section.

The FortiGate unit offers three different traffic shaping options, all of which can be enabled at the same time within the same security policy. Generally speaking, the hierarchy for shapers in FortiOS is:

- Application Control shaper
- Security policy shaper
- Per-IP shaper

With this hierarchy, if an application control list has a traffic shaper defined, it will have precedence always over any other security policy shaper. For example, with the example above creating an application control for Facebook, the shaper defined for Facebook will supersede any security policy enabled traffic shapers. While the Facebook application may reach its maximum bandwidth, the user can still have the bandwidth room available from the shared shaper and, if enabled, the per-IP shaper.

Equally, any security policy shared shaper will have precedence over any per-IP shaper. However, traffic that exceeds any of these shapers will be dropped. For example, the policy shaper will take effect first, however, if the per-IP shaper limit is reached first, then traffic for that user will be dropped even if the shared shaper limit for the policy has not been exceeded.

## Shared policy shaping

Traffic shaping by security policy enables you to control the maximum and/or guaranteed throughput for a selected security policy. When configuring a shaper, you can select to apply the bandwidth shaping per policy or for all policies. Depending on your selection, the FortiGate unit will apply the shaping rules differently.

### Per policy

When selecting a shaper to be *per policy*, the FortiGate unit will apply the shaping rules defined to each security policy individually.

For example, the shaper is set to be per policy with a maximum bandwidth of 1000 Kb/s. There are four security policies monitoring traffic through the FortiGate unit. Three of these have the shaper enabled. Each security policy has the same maximum bandwidth of 1000 Kb/s.

Per policy traffic shaping is compatible with client/server (active-passive) transparent mode WAN optimization rules. Traffic shaping is ignored for peer-to-peer WAN optimization and for client/server WAN optimization not operating in transparent mode.

### All policies

When selecting a shaper to be for all policies - *For All Policies Using This Shaper* - the FortiGate unit applies the shaping rules to all policies using the same shaper. For example, the shaper is set to be per policy with a maximum bandwidth of 1000 Kb/s. There are four security policies monitoring traffic through the FortiGate unit. All four have the shaper enabled. Each security policy must share the defined 1000 Kb/s, and is set on a first come, first served basis. For example, if policy 1 uses 800 Kb/s, the remaining three must share 200 Kb/s. As policy 1 uses less bandwidth, it is opened up to the other policies to use as required. Once used, any other policies will encounter latency until free bandwidth opens from a policy currently in use.

## Maximum and guaranteed bandwidth

The maximum bandwidth instructs the security policy what the largest amount of traffic allowed using the policy. Depending on the service or the users included for the security policy, this number can provide a larger or smaller throughput depending on the priority you set for the shaper.

The *Maximum Bandwidth* can be set to a value of between 1 and 16776000 kbit/s. The Web-Based Manager gives an error if any value outside of this range is used, but in the CLI a value of 0 can be entered. Setting `maximum-bandwidth` to 0 (zero) prevents any traffic from going through the policy.

The guaranteed bandwidth ensures there is a consistent reserved bandwidth available for a given service or user. When setting the guaranteed bandwidth, ensure that the value is significantly less than the bandwidth capacity of the interface, otherwise no other traffic will pass through the interface or very little and potentially causing unwanted latency.

### Traffic priority

Select a Traffic Priority of high, medium or low, so the FortiGate unit manages the relative priorities of different types of traffic. For example, a policy for connecting to a secure web server needed to support e-commerce traffic should be assigned a high traffic priority. Less important services should be assigned a low priority. The firewall provides bandwidth to low-priority connections only when bandwidth is not needed for high-priority connections.

Be sure to enable traffic shaping on all security policies. If you do not apply any traffic shaping rule to a policy, the policy is set to high priority by default. Distribute security policies over all three priority queues.

## VLAN, VDOM and virtual interfaces

Policy-based traffic shaping does not use queues directly. It shapes the traffic and if the packet is allowed by the security policy, then a priority is assigned. That priority controls what queue the packet will be put in upon egress. VLANs, VDOMs, aggregate ports and other virtual devices do not have queues and as such, traffic is sent directly to the underlying physical device where it is queued and affected by the physical ports. This is also the case with IPsec connections.

## Shared traffic shaper configuration settings

To configure a shared traffic shaper go to *Firewall Objects > Traffic Shaper > Shared* and select *Create New*.

<b>Name</b>	Enter a name for the traffic shaper.
<b>Apply Shaper</b>	<p>When selecting a shaper to be <i>Per Policy</i>, the FortiGate unit will apply the shaping rules defined to each security policy individually. For example, the shaper is set to be per policy with a maximum bandwidth of 1000 Kb/s. There are four security policies monitoring traffic through the FortiGate unit. Three of these have the shaper enabled. Each security policy has the same maximum bandwidth of 1000 Kb/s.</p> <p>Per policy traffic shaping is compatible with client/server (active-passive) transparent mode WAN optimization rules. Traffic shaping is ignored for peer-to-peer WAN optimization and for client/server WAN optimization not operating in transparent mode.</p> <p>When selecting a shaper to be for all policies - <i>For All Policies Using This Shaper</i> - the FortiGate unit applies the shaping rules to all policies using the same shaper. For example, the shaper is set to be per policy with a maximum bandwidth of 1000 Kb/s. There are four security policies monitoring traffic through the FortiGate unit. All four have the shaper enabled. Each security policy must share the defined 1000 Kb/s, and is set on a first come, first served basis. For example, if policy 1 uses 800 Kb/s, the remaining three must share 200 Kb/s. As policy 1 uses less bandwidth, it is opened up to the other policies to use as required. Once used, any other policies will encounter latency until free bandwidth opens from a policy currently in use.</p>
<b>Traffic Priority</b>	<p>Select level of importance <i>Priority</i> so the FortiGate unit manages the relative priorities of different types of traffic. For example, a policy for connecting to a secure web server needed to support e-commerce traffic should be assigned a high traffic priority. Less important services should be assigned a low priority.</p> <p>If you do not apply any traffic shaping priority, the priority is set to high priority by default.</p>

<b>Maximum Bandwidth</b>	<p>The maximum bandwidth instructs the security policy what the largest amount of traffic allowed using the policy. Depending on the service or the users included for the security policy, this number can provide a larger or smaller throughput depending on the priority you set for the shaper.</p> <p>Setting <i>Maximum Bandwidth</i> to 0 (zero) provides unlimited bandwidth.</p>
<b>Guaranteed Bandwidth</b>	<p>The guaranteed bandwidth ensures there is a consistent reserved bandwidth available for a given service or user. When setting the guaranteed bandwidth, ensure that the value is significantly less than the bandwidth capacity of the interface, otherwise no other traffic will pass through the interface or very little and potentially causing unwanted latency.</p> <p>Setting <i>Guaranteed Bandwidth</i> to 0 (zero) provides unlimited bandwidth.</p>
<b>DSCP</b>	<p>Enter the number for the DSCP value. You can use the FortiGate Differentiated Services feature to change the DSCP (Differentiated Services Code Point) value for all packets accepted by a policy. The network can use these DSCP values to classify, mark, shape, and police traffic, and to perform intelligent queuing. DSCP features are applied to traffic by configuring the routers on your network to apply different service levels to packets depending on the DSCP value of the packet. For more information, see <a href="#">“Differentiated Services”</a>.</p>

## Example

The following steps create a Per Policy traffic shaper called “Throughput” with a maximum traffic amount of 720,000 Kb/s, and a guaranteed traffic of 150,000 Kb/s with a high traffic priority.

### To create the shared shaper - web-based manager

1. Go to *Firewall Objects > Traffic Shaper > Shared* and select *Create New*.
2. Enter the *Name* Throughput.
3. Select *Per Policy*.
4. Select the *Maximum Bandwidth* check box and enter the value 120000.
5. Select the *Guaranteed Bandwidth* check box and enter the value 150000.
6. Set the *Traffic Priority* to *High*.
7. Select *OK*.

### To create the shared shaper - CLI

```
config firewall shaper traffic-shaper
 edit Throughput
 set per-policy enable
 set maximum-bandwidth 720000
 set guaranteed-bandwidth 150000
 set priority high
 end
```

## Per-IP shaping

Traffic shaping by IP enables you to apply traffic shaping to all source IP addresses in the security policy. As well as controlling the maximum bandwidth users of a selected policy, you can also define the maximum number of concurrent sessions.

Per-IP traffic shaping enables you limit the behavior of every member of a policy to avoid one user from using all the available bandwidth - it now is shared within a group equally. Using a per-IP shaper avoids having to create multiple policies for every user you want to apply a shaper. Per-IP traffic shaping is not supported over NP2 interfaces.

### Per-IP traffic shaping configuration settings

To configure per-IP traffic shaping go to *Firewall Objects > Traffic Shaper > Per-IP*. and select *Create New*.

<b>Name</b>	Enter a name for the per-IP traffic shaper.
<b>Maximum Bandwidth</b>	<p>The maximum bandwidth instructs the security policy what the largest amount of traffic allowed using the policy. Depending on the service or the users included for the security policy, this number can provide a larger or smaller throughput depending on the priority you set for the shaper.</p> <p>Setting <i>Maximum Bandwidth</i> to 0 (zero) provides unlimited bandwidth.</p>
<b>Maximum Concurrent Connections</b>	Enter the maximum allowed concurrent connection.
<b>Forward DSCP Reverse DSCP</b>	Enter the number for the DSCP value. You can use the FortiGate Differentiated Services feature to change the DSCP (Differentiated Services Code Point) value for all packets accepted by a policy. The network can use these DSCP values to classify, mark, shape, and police traffic, and to perform intelligent queuing. DSCP features are applied to traffic by configuring the routers on your network to apply different service levels to packets depending on the DSCP value of the packet. For more information, see <a href="#">“Differentiated Services”</a> .

### Example

The following steps create a Per-IP traffic shaper called “Accounting” with a maximum traffic amount of 720,000 Kb/s, and the number of concurrent sessions of 200.

#### To create the shared shaper - web-based manager

1. Go to *Firewall Objects > Traffic Shaper > Per-IP*.
2. Select *Create New*.
3. Enter the *Name Accounting*.
4. Select the *Maximum Bandwidth* check box and enter the value 720000.
5. Select the *Maximum Concurrent Sessions* check box and enter the value 200.
6. Select *OK*.

### To create the shared shaper - CLI

```
config firewall shaper per-ip-shaper
 edit Accounting
 set max-bandwidth 720000
 set max-concurrent-sessions 200
 end
```

## Adding Per-IP traffic shapers to a security policy

Per-IP traffic shaping is supported by IPv6 security policies. You can add any Per-IP traffic shaper to an IPv6 security policy in the CLI.5.020142014

### To add a Per-IP traffic shaper to an IPv6 security policy - CLI

```
config firewall policy6
 edit 0
 set per-ip-shaper 'new-perip-shaper'
 end
```

## Application control shaping

Traffic shaping is also possible for specific applications for both shared and per IP shaping. Through the *Security Profiles > Application Control* feature, you can configure a specific application's maximum bandwidth. When configuring the application control features, if the application is set to pass, you can set the traffic shaping options. The shapers available are those set up in the *Firewall Objects > Traffic Shaping* menu.

For more information on configuring application control shapers, see the [Security Profiles Guide](#).

### Example

This example sets the traffic shaping definition for Facebook to a medium priority, a default traffic shaper.

#### To add traffic shaping for Facebook - web-based manager

1. Go to *Security Profiles > Application Control > Application Sensors*.
2. Select the *Create New* "Plus" icon in the upper right corner of the screen to create a new application group, and enter the name *Web*.
3. Select *OK*.
4. Select *Create New*.
5. Deselect the check box for *Category* to unselect all categories and select *Social.Networking*.
6. Select *Traffic Shaping* and select *medium-priority* from the drop-down list.
7. Select *OK*.



### To add traffic shaping for Facebook - CLI

```
config application list
 edit web
 config entries
 edit 1
 set category 23
 set application 17735
 set action pass
 set shaper medium-priority
 end
 end
 end
end
```

## Enabling in the security policy

All traffic shapers are enabled within a security policy, including the Application Control shapers. As such, the shapers are in effect after any DoS detection policies, and before any routing or packet scanning occurs.

Traffic shaping is also supported for IPv6 policies.

### To enable traffic shaping - web-based manager

1. Go to *Policy > Policy > Policy*.
2. Select *Create New* or select an existing policy and select *Edit*.
3. Select *Traffic Shaping*.
4. Select the shaping option and select the shaper from the drop-down list.
5. Select *OK*.

Shapers applied in the security policy affect outbound or traffic to a destination. To affect inbound, or download, traffic, select *Shared Traffic Shaper Reverse Direction*. For more information, see [“Reverse direction traffic shaping” on page 2249](#).

### To enable traffic shaping - CLI

```
config firewall policy
 edit <policy_number>
 ...
 set traffic-shaper <shaper_name>
 set per-ip-shaper <shaper_name>
 end
```

## Reverse direction traffic shaping

The shaper you select for the security policy (shared shaper) will affect the traffic in the direction defined in the policy. For example, if the source port is port 1 and the destination is port 3, the shaping affects the flow in this direction only, that is, the upload or outbound direction. By selecting *Shared Traffic Shaper Reverse Direction*, you can define the traffic shaper for the policy in the opposite direction, that is, the download or inbound direction. In this example, from port 3 to port 1.

### To add a reverse shaper

1. Go to *Policy > Policy > Policy*.

2. Select *Traffic Shaping*.
3. Select *Shared Traffic Shaper Reverse Direction* and select the shaper from the list.
4. Select *OK*.

## Setting the reverse direction only

There may be instances where you only need to have the traffic shaping for incoming connections. That is, the “reverse” direction to the typical traffic shaper.

### To add a reverse shaper - web-based manager

1. Go to *Policy > Policy > Policy*.
2. Select *Traffic Shaping*.
3. Select *Shared Traffic Shaper Reverse Direction* and select the shaper from the list.
4. Select *OK*.

### To configure a reverse-only shaper - CLI

```
config firewall policy
 edit <policy_number>
 ...
 set traffic-shaper-reverse <shaper_name>
 end
```

## Application control shaper

Application control shapers are in effect within the application control profile. Within the security policy options, select *Application Control* and select the application from the list.

## Type of Service priority

Type of service (ToS) is an 8-bit field in the IP header that enables you to determine how the IP datagram should be delivered, using criteria of Delay, Throughput, Priority, Reliability, and Cost. Each quality helps gateways determine the best way to route datagrams. A router maintains a ToS value for each route in its routing table. The lowest priority ToS is 0, the highest is 7 when bits 3, 4, and 5 are all set to 1. There are 4 other bits that are seldom used or reserved that are not included here.

Together these bits are the tos variable of the `tos-based-priority` command. The router tries to match the ToS of the datagram to the ToS on one of the possible routes to the destination. If there is no match, the datagram is sent over a zero ToS route. Using increased quality may increase the cost of delivery because better performance may consume limited network resources.

Each bit represents the priority as per RFC 1349:

- 1000 - minimize delay
- 0100 - maximize throughput
- 0010 - maximize reliability
- 0001 - minimize monetary cost

The TOS value is set in the CLI using the commands:

```
config system tos-based-priority
 edit <sequence_number>
 set tos [0-15]
 set priority [high | medium | low]
 end
```

Where `tos` is the value of the type of service bit in the IP datagram header with a value between 0 and 15, and `priority` is the priority of this type of service priority. These priority levels conform to the firewall traffic shaping priorities, as defined in [RFC 1349](#).

For example, if you want to configure the FortiGate unit so that reliability is the first priority, set the `tos` value to 4.

```
config system tos-based-priority
 edit 1
 set tos 4
 set priority high
 end
```

For a list of ToS values and their DSCP equivalents see [“Tos and DSCP mapping” on page 2257](#).

### Example

```
config system tos-based-priority
 edit 1
 set tos 1
 set priority low
 next
 edit 4
 set tos 4
 set priority medium
 next
 edit 6
 set tos 6
 set priority high
 next
end
```

## TOS in FortiOS

Traffic shaping and TOS follow the following sequence:

- The CLI command `tos-based-priority` acts as a `tos-to-priority` mapping. FortiOS maps the TOS to a priority when it receives a packet.
- Traffic shaping settings adjust the packet's priority according the traffic.
- Deliver the packet based on its priority.

## Differentiated Services

Differentiated Services describes a set of end-to-end Quality of Service (QoS) capabilities. End-to-end QoS is the ability of a network to deliver service required by specific network traffic

from one end of the network to another. By configuring differentiated services, you configure your network to deliver particular levels of service for different packets based on the QoS specified by each packet.

Differentiated Services (also called DiffServ) is defined by RFC 2474 and 2475 as enhancements to IP networking to enable scalable service discrimination in the IP network without the need for per-flow state and signaling at every hop. Routers that can understand differentiated services sort IP traffic into classes by inspecting the DS field in IPv4 header or the Traffic Class field in the IPv6 header.

You can use the FortiGate Differentiated Services feature to change the DSCP (Differentiated Services Code Point) value for all packets accepted by a policy. The network can use these DSCP values to classify, mark, shape, and police traffic, and to perform intelligent queuing. DSCP features are applied to traffic by configuring the routers on your network to apply different service levels to packets depending on the DSCP value of the packet.

If the differentiated services feature is not enabled, the FortiGate unit treats traffic as if the DSCP value is set to the default (00), and will not change IP packets' DSCP field. DSCP values are also not applied to traffic if the traffic originates from a FortiGate unit itself.

The FortiGate unit applies the DSCP value and IPsec encryption to the differentiated services (formerly TOS) field in the first word of the IP header. The typical first word of an IP header, with the default DSCP value, is 4500:

- 4 for IPv4
- 5 for a length of five words
- 00 for the default DSCP value

You can change the packet's DSCP field for traffic initiating a session (forward) or for reply traffic (reverse) and enable each direction separately and configure it in the security policy.

Changes to DSCP values in a security policy effect new sessions. If traffic must use the new DSCP values immediately, clear all existing sessions.

DSCP is enabled using the CLI command:


```
config firewall policy
 edit <policy_number>
 ...
 set diffserv-forward enable
 set diffservcode-forward <binary_integer>
 set diffserv-reverse enable
 set diffservcode-rev <binary_integer>
 end
```

For more information on the different DSCP commands, see the examples below and the [CLI Reference](#). If you only set `diffserv-forward` and `diffserv-reverse` without setting the corresponding `diffservcode` values, the FortiGate unit will reset the bits to zero.

For a list of DSCP values and their ToS equivalents see [“Tos and DSCP mapping” on page 2257](#). DSCP values can also be defined within a shared shaper as a single value, and per-IP shaper for forward and reverse directions.

## DSCP examples

For all the following DSCP examples, the FortiGate and client PC configuration is the following diagram and used firewall-based DSCP configurations.



### Example

In this example, an ICMP ping is executed between User 1 and FortiGate B, through a FortiGate unit. DSCP is disabled on FortiGate B, and FortiGate A contains the following configuration:

```
config firewall policy
 edit 2
 set srcintf port6
 set dstintf port3
 set src addr all
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
 set diffserv-forward enable
 set diffservcode-forward 101110
 end
```


As a result, FortiGate A changes the DSCP field for outgoing traffic, but not to its reply traffic. The binary DSCP values used map to the following hexadecimal

TOS field values, which are observable by a sniffer (also known as a packet tracer):

- DSCP 000000 is TOS field 0x00
- DSCP 101110 is TOS field 0xb8, the recommended DSCP value for expedited forwarding (EF)

If you performed an ICMP ping between User 1 and User 2, the following output illustrates the IP headers for the request and the reply by sniffers on each of FortiGate unit's network

interfaces. The right-most two digits of each IP header are the TOS field, which contains the DSCP value.



## Example


In this example, an ICMP ping is executed between User 1 and FortiGate B, through FortiGate A. A. DSCP is disabled on FortiGate B, and FortiGate A contains the following configuration:

```
config firewall policy
 edit 2
 set srcintf port6
 set dstintf port3
 set src addr all
 set dstaddr all
 set action accept
 set schedule always
 set service ANY"
 set diffserv-forward enable
 set diffserv-rev enable
 set diffservcode-forward 101110
 set diffservcode-rev 101111
 end
```

As a result, FortiGate A changes the DSCP field for both outgoing traffic and its reply traffic. The binary DSCP values in map to the following hexadecimal TOS field values, which are observable by a sniffer (also known as a packet tracer):

- DSCP 000000 is TOS field 0x00
- DSCP 101110 is TOS field 0xb8, the recommended DSCP value for expedited forwarding (EF)
- DSCP 101111 is TOS field 0xbc

If you performed an ICMP ping between User 1 and User 2, the output below illustrates the IP headers observed for the request and the reply by sniffers on each of FortiGate A's and FortiGate B's network interfaces. The right-most two digits of each IP header are the TOS field, which contains the DSCP value.



## Example

In this example, an ICMP ping is executed between User 1 and FortiGate B, through FortiGate A. DSCP is enabled for both traffic directions on FortiGate A, and enabled only for reply traffic on FortiGate B. FortiGate A contains the following configuration:

```
config firewall policy
 edit 2
 set srcintf port6
 set dstintf port3
 set src addr all
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
 set diffserv-forward enable
 set diffserv-rev enable
 set diffservcode-forward 101110
 set diffservcode-rev 101111
 end
```


FortiGate B contains the following configuration:

```
config firewall policy
 edit 2
 set srcintf wan2
 set dstintf internal
 set src addr all
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
 set diffserv-rev enable
 set diffservcode-rev 101101
 end
```

As a result, FortiGate A changes the DSCP field for both outgoing traffic and its reply traffic, and FortiGate B changes the DSCP field only for reply traffic. The binary DSCP values in this configuration map to the following hexadecimal TOS field values:

- DSCP 000000 is TOS field **0x00**
- DSCP 101101 is TOS field **0xb4**
- DSCP 101110 is TOS field **0xb8**, the recommended DSCP value for expedited forwarding (EF)
- DSCP 101111 is TOS field **0xbc**

If you performed an ICMP ping between User 1 and User 2, the output below illustrates the IP headers observed for the request and the reply by sniffers on each of FortiGate A's and FortiGate B's network interfaces. The right-most two digits of each IP header are the TOS field, which contains the DSCP value.



## Example

In this example, HTTPS and DNS traffic is sent from User 1 to FortiGate B, through FortiGate A. DSCP is enabled for both traffic directions on FortiGate A, and enabled only for reply traffic on FortiGate B. FortiGate A contains the following configuration:

```
config firewall policy
 edit 2
 set srcintf port6
 set dstintf port3
 set src addr all
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
 set diffserv-forward enable
 set diffserv-rev enable
 set diffservcode-forward 101110
 set diffservcode-rev 101111
 end
```

FortiGate B contains the following configuration:

```
config firewall policy
 edit 2
 set srcintf wan2
 set dstintf internal
 set src addr all
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
 set diffserv-rev enable
 set diffservcode-rev 101101
 end
```


As a result, FortiGate A changes the DSCP field for both outgoing traffic and its reply traffic, but FortiGate B changes the DSCP field only for reply traffic which passes through its internal interface. Since the example traffic does not pass through the internal interface, FortiGate B does not mark the packets. The binary DSCP values in this configuration map to the following hexadecimal TOS field values:

- DSCP 000000 is TOS field **0x00**
- DSCP 101101 is TOS field **0xb4**, which is configured on FortiGate B but not observed by the sniffer because the example traffic originates from the FortiGate unit itself, and therefore does not match that security policy.
- DSCP 101110 is TOS field **0xb8**, the recommended DSCP value for expedited forwarding (EF)
- DSCP 101111 is TOS field **0xbc**

If you sent HTTPS or DNS traffic from User 1 to FortiGate B, the following would illustrate the IP headers observed for the request and the reply by sniffers on each of FortiGate A's and



FortiGate B's network interfaces. The right-most two digits of each IP header are the TOS field, which contains the DSCP value.



## Tos and DSCP mapping

The table below lists the mapping of DSCP and ToS hexadecimal values for QoS.

**Table 103:**ToS to DSCP mappings

Service Class	DSCP Bits	DSCP Value	ToS Value	ToS Hexidecimal
<b>Network Control</b>	111000	56-63	224	0xE0
<b>Internetwork Control</b>	110000	48-55	192	0xC0
<b>Critical - Voice Data (RTP)</b>	101110	46	184	0xB8
	101000	40	160	0xA0
<b>Flash Override Video Data</b>	100010	34	136	0x88
	100100	36	144	0x90
	100110	38	152	0x98
	100000	32	128	0x80
<b>Flash Voice Control</b>	011010	26	104	0x68
	011100	28	112	0x70
	011110	30	120	0x78
	011000	24	96	0x60
<b>Immediate Deterministic (SNA)</b>	010010	18	72	0x48
	010100	20	80	0x50
	010110	22	88	0x58
	010000	16	64	0x40
<b>Priority Controlled Load</b>	001010	10	40	0x28
	001100	12	48	0x30
	001110	14	56	0x38
	001000	8	32	0x20
<b>Routine - Best Effort</b>	000000	0	0	0x00
<b>Routine - Penalty Box</b>	000010	2	8	0x08

## Traffic Shaper Monitor

You can view statistical information about traffic shapers and their bandwidth from *Firewall Objects > Monitor > Traffic Shaper Monitor*.

<b>Refresh</b>	Select to refresh the information on the page.
<b>Reset</b>	Select to reset the information to clear the current information from the page. New information is included on the page.
<b>Report By</b>	Select to display dropped packets or current bandwidth. The chart changes its name so you know what current information is displayed.
<b>Traffic Shaper Usage Dropped packets</b>	The bar chart displays the packets that were dropped by traffic shaper.
<b>Traffic Shaper Usage Current Bandwidth</b>	The bar chart displays the current bandwidth of traffic shapers.

# Examples

While it is possible to configure QoS using a combination of security policies and in ToS-based priorities, and to distribute traffic over all six of the possible queues for each physical interface, the results of those configurations can be more difficult to analyze due to their complexity. In those cases, prioritization behavior can vary by several factors, including traffic volume, ToS (type of service) or differentiated services markings, and correlation of session to a security policy.

The following simple examples illustrate QoS configurations using either prioritization by security policy, or prioritization by ToS bit, but not both. The examples also assume you are not configuring traffic shaping for interfaces that receive hardware acceleration from network processing units (NPU).

## QoS using priority from security policies

Configurations implementing QoS using the priority values defined in security policies are capable of applying bandwidth limits and guarantees.

In addition to configuring traffic shaping, you may also choose to limit bandwidth accepted by each interface. This can be useful in scenarios where bandwidth being received on source interfaces frequently exceeds the maximum bandwidth limit defined in the security policy. In this case, rather than wasting processing power on packets that will only be dropped later in the processing to enforce those limits, you may choose to preemptively police the traffic.

Note that if you implement QoS using security policies rather than ToS bit, the FortiGate unit applies QoS to all packets controlled by the policy. Control is less granular than prioritization by ToS bit, but has the benefits of correlating quality of service to a security policy, enabling you to distribute traffic over up to four of the possible 6 priority queues (queue 0 to queue 3), not requiring other devices in your network to set or respect the ToS bit, and of enabling you to configure bandwidth limits and guarantees.

In this example, we limit the bandwidth accepted by each source interface, limit the bandwidth used by sessions controlled by the security policy, and then configure prioritized queuing on the destination interface based upon the priority in the security policy, subject to alternative assignment to queue 0 when necessary to achieve the guaranteed packet rate.

### To limit bandwidth accepted by an interface

In the CLI, enter the following commands:

```
config system interface
 edit <name_str>
 set inbandwidth <rate_int>
 next
end
```

where <rate\_int> is the bandwidth limit in Kb/s. Excess packets will be dropped.

### To configure bandwidth guarantees, limits, and priorities

1. Go to *Firewall Objects > Traffic Shaper > Shared*, and select *Create New*.
2. Enter a name for the shaper.

3. Enter the *Guaranteed Bandwidth*, if any.  
Bandwidth guarantees affect prioritization. While packet rates are less than this rate, they use priority queue 0. If this is not the effect you intend, consider entering a small guaranteed rate, or enter 0 to effectively disable bandwidth guarantees.
4. Enter a *Maximum Bandwidth*.  
Packets greater than this rate will be discarded.
5. Select the *Traffic Priority*.  
High has a priority value of 1, while Low is 3. While the current packet rate is below Guaranteed Bandwidth, the FortiGate unit will disregard this setting, and instead use priority queue 0.
6. Select *OK*.

## Sample configuration

This sample configuration limits ingress bandwidth to 500 Kb/s. It also applies separate traffic shapers to FTP and HTTP traffic. In addition to the interface bandwidth limit, HTTP traffic is subject to a security policy bandwidth limit of 200 Kb/s.

All egressing FTP traffic greater than 10 Kb/s is subject to a low priority queue (queue 3), while all egressing HTTP traffic greater than 100 Kb/s is subject to a medium priority queue (queue 2). That is, unless FTP traffic rates are lower than their guaranteed rate, and web traffic rates are greater than their guaranteed rate, FTP traffic is lower priority than web traffic.

Traffic less than these guaranteed bandwidth rates use the highest priority queue (queue 0).

Set the inbandwidth limits. This setting is only available in the CLI:

```
config system interface
 edit wan1
 set inbandwidth 500
 next
end
```

Create the traffic shapers for FTP and HTTP.

### To configure the shapers - web-based manager

1. Go to *Firewall Objects > Traffic Shaper > Shared*, and select *Create New*.
2. Enter *FTP* for the name of the shaper.
3. Enter the *Guaranteed Bandwidth*, of 10 Kbps.
4. Enter a *Maximum Bandwidth* of 500 Kbps.
5. Select the *Traffic Priority* of *Low*.
6. Select *OK*.
7. Select *Create New*.
8. Enter *HTTP* for the name of the shaper.
9. Enter the *Guaranteed Bandwidth*, of 100 Kbps.
10. Enter a *Maximum Bandwidth* of 200 Kbps.
11. Select the *Traffic Priority* of *Medium*.
12. Select *OK*.

### To configure the shapers - CLI

```
config firewall shaper traffic-shaper
 edit FTP
 set maximum-bandwidth 500
 set guaranteed-bandwidth 10
 set per-policy enable
 set priority low
 end
next
edit HTTP
 set maximum-bandwidth 200
 set guaranteed-bandwidth 100
 set per-policy enable
 set priority medium
end
```

## QoS using priority from ToS or differentiated services

Configurations implementing QoS using the priority values defined in either global or specific ToS bit values are not capable of applying bandwidth limits and guarantees, but are capable of prioritizing traffic at per-packet levels, rather than uniformly to all services matched by the security policy.

In addition to configuring traffic prioritization, you may also choose to limit bandwidth being received by each interface. This can sometimes be useful in scenarios where you want to limit traffic levels, but do not want to configure traffic shaping within a security policy. This has the benefit of policing traffic at a point before the FortiGate unit performs most processing.

Note that if you implement QoS using ToS octet rather than security policies, the FortiGate unit applies QoS on a packet by packet basis, and priorities may be different for packets and services controlled by the same security policy. This is more granular control than prioritization by security policies, but has the drawbacks that quality of service is may not be uniform for multiple services controlled by the same security policy, packets will only use up to three of the six possible queues (queue 0 to queue 2), and bandwidth cannot be guaranteed. Other devices in your network must also be able to set or preserve ToS bits.

In this example, we limit the bandwidth accepted by each source interface, and then configure prioritized queuing on the destination interface based upon the value of the ToS bit located in the IP header of each accepted packet.

To limit bandwidth accepted by an interface, in the CLI, enter the following commands:

```
config system interface
 edit <name_str>
 set inbandwidth <rate_int>
 next
end
```

where <rate\_int> is the bandwidth limit in Kb/s. Excess packets will be dropped.

To configure priorities, in the CLI, configure the global priority value using the following commands:

```
config system global
 set tos-based-priority {high | low | medium}
end
```

where `high` has a priority value of 0 and `low` is 2.

If you want to prioritize some ToS bit values differently than the global ToS-based priority, configure the priority for packets with that ToS bit value using the following commands:

```
config system tos-based-priority
 edit <id_int>
 set tos [0-15]
 set priority {high | low | medium}
 next
end
```

where `and tos` is the value of the ToS bit in the packet's IP header, and `high` has a priority value of 0 and `low` is 2. Priority values configured in this location will override the global ToS-based priority.

## Sample configuration

This sample configuration limits ingress bandwidth to 500 Kb/s. It also queues egress traffic based upon the ToS bit in the IP header of ingress packets.

Unless specified for the packet's ToS bit value, packets use the low priority queue (queue 2). For ToS bit values 4 and 15, the priorities are specified as medium (value 1) and high (value 0), respectively.

```
config system interface
 edit wan1
 set inbandwidth 500
 next
end
config system global
 set tos-based-priority low
end
config system tos-based-priority
 edit 4
 set tos 4
 set priority medium
 next
 edit 15
 set tos 15
 set priority high
 next
end
```

## Example setup for VoIP

In this example, there are three traffic shaping requirements for a network:

- Voice over IP (VoIP) requires a guaranteed, high-priority for bandwidth for telephone communications.
- FTP bursts must be contained so as not to consume any available bandwidth. As such this traffic needs to be throttled to a smaller amount.
- A consistent bandwidth requirement is needed for all other email and web-based traffic.

To enable this requirement, you need to create three separate shapers and three security policies for each traffic type.

For this example, the actual values are not actual values, they are used for the simplicity of the example.

### Creating the traffic shapers

First create the traffic shapers that define the maximum and guaranteed bandwidth. The shared shapers will be used, some with per-policy and some all policies as shown in the table, to better control traffic.

#### VoIP shaper

The VoIP functionary is a key component to the business as a communication tool and as such requires a guaranteed bandwidth.

##### To create a VoIP shaper - web-based manager

1. Go to *Firewall Objects > Traffic Shaping > Shared*.
2. Enter the *Name* `voip`.
3. Select *Per Policy*.
4. Enter the *Maximum Bandwidth* of 1000 Kb/s
5. Enter the *Guaranteed Bandwidth* of 800 Kb/s.
6. Select a *Traffic Priority* of *High*.
7. Select *OK*.

##### To create a VoIP shaper - CLI

```
config firewall shaper traffic-shaper
 edit voip
 set maximum-bandwidth 1000
 set guaranteed-bandwidth 800
 set per-policy enable
 set priority high
 end
```

This ensures that whatever number of policies use this shaper, the defined bandwidth will always be the same. At the same time, the bandwidth is continually guaranteed at 800 Kb/s but if available can be as much as 1000 Kb/s. Setting the priority to high ensures that the FortiGate unit always considers VoIP traffic as the most important.

#### FTP shaper

The FTP shaper sets the maximum bandwidth to use to avoid sudden spikes by sudden uploading or downloading of large files, and interfering with other more important traffic.

### To create a FTP shaper - web-based manager

1. Go to *Firewall Objects > Traffic Shaping > Shared*.
2. Enter the *Name* ftp.
3. Select *For all Policies Using This Shaper*.
4. Enter the *Maximum Bandwidth* of 200 Kb/s
5. Enter the *Guaranteed Bandwidth* of 200 Kb/s.
6. Select a *Traffic Priority* of Low.
7. Select OK.

### To create a FTP shaper - CLI

```
config firewall shaper traffic-shaper
 edit ftp
 set maximum-bandwidth 200
 set guaranteed-bandwidth 200
 set priority low
 end
```

For this shaper, the maximum and guaranteed bandwidth are set low and to the same value. In this case, the bandwidth is restricted to a specific amount. By also setting the traffic priority low ensures more important traffic will be able to pass before FTP traffic.

## Regular traffic shaper

The regular shaper sets the maximum bandwidth and guaranteed bandwidth for everyday business traffic such as web and email traffic.

### To create a regular shaper - web-based manager

1. Go to *Firewall Objects > Traffic Shaping > Shared*.
2. Enter the *Name* daily\_traffic.
3. Select *Per Policy*.
4. Enter the *Maximum Bandwidth* of 600 Kb/s
5. Enter the *Guaranteed Bandwidth* of 600 Kb/s.
6. Select a *Traffic Priority* of Medium.
7. Select OK.

### To create a regular shaper - CLI

```
config firewall shaper traffic-shaper
 edit daily_traffic
 set maximum-bandwidth 600
 set guaranteed-bandwidth 600
 set per-policy enable
 set priority medium
 end
```

For this shaper, the maximum and guaranteed bandwidth are set to a moderate value of 600 Kb/s. It is also set for per policy, which ensures each security policy for day-to-day business traffic has the same distribution of bandwidth.



## Creating security policies

To employ the shaper, create security policies that use the shapers within the policies. Create a separate policy for each service and enable traffic shaping. For example, a policy for FTP traffic, a policy for SIP and so on.

For the following steps the VoIP traffic shaper is enabled as well as the reverse direction option. This ensures that return traffic for a VoIP call has the same guaranteed bandwidth as the outgoing call.

### To enable traffic shaping in the security policy - web-based manager

1. Go to *Policy > Policy > Policy* and select *Create New*.
2. Leave the *Policy Type* as *Firewall* and leave the *Policy Subtype* as *Address*.
3. Enter the following:

<b>Incoming interface</b>	Internal
<b>Source address</b>	All
<b>Outgoing interface</b>	WAN1
<b>Destination address</b>	All
<b>Schedule</b>	always
<b>Service</b>	SIP
<b>Action</b>	ACCEPT

4. Select *Traffic Shaping*.
5. From the drop-down list, select the voip shaper created in the previous steps.
6. Select *Reverse Direction Traffic Shaping*.
7. Select *OK*.

### To enable traffic shaping in the security policy - CLI

```
config firewall policy
 edit 6
 set srcintf internal
 set scraddr all
 set dstintf wan1
 set dstaddr all
 set action accept
 set schedule always
 set service sip
 set traffic-shaper voip
 set reverse-traffic-shaper voip
 end
```

# Troubleshooting traffic shaping

This chapter outlines some troubleshooting tips and steps to diagnose the shapers and whether they are working correctly. These diagnose commands include:

- `diagnose system tos-based-priority`
- `diagnose firewall shaper traffic-shaper`
- `diagnose firewall per-ip-shaper`
- `diagnose debug flow`

## Interface diagnosis

To optimize traffic shaping performance, first ensure that the network interface's Ethernet statistics are clean of errors, collisions, or buffer overruns. To check the interface, enter the following diagnose command to see the traffic statistics:

```
diagnose hardware deviceinfo nic <port_name>
```

## Shaper diagnose commands

There are specific diagnose commands you can use to verify the configuration and flow of traffic, including packet loss due to the employed shaper.

All of these diagnose troubleshooting commands are supported in both IPv4 and IPv6.

### TOS command

Use the following command to list command to view information of the TOS lists and traffic.

```
diagnose system tos-based-priority
```

This example displays the priority value currently correlated with each possible TOS bit value. Priority values are displayed in order of their corresponding TOS bit values, which can range between 0 and 15, from lowest TOS bit value to highest.

For example, if you have not configured TOS-based priorities, the following appears...

```
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
```

...reflecting that all packets are currently using the same default priority, high (value 0).

If you have configured a TOS-based priority of `low` (value 2) for packets with a ToS bit value of 3, the following appears...

```
0 0 0 2 0 0 0 0 0 0 0 0 0 0 0 0
```

...reflecting that most packets are using the default priority value, except those with a ToS bit value of 3.

## Shared shaper

To view information for the shared traffic shaper for security policies enter the command

```
diagnose firewall shaper traffic-shaper list
```

The resultant output displays the information on all available shapers. The more shapers available the longer the list. For example:

```
name Throughput
maximum-bandwidth 1200000 Kb/sec
guaranteed-bandwidth 50000 Kb/sec
current-bandwidth 0 B/sec
priority 1
packets dropped 0
```

Additional commands include:

`diagnose firewall shaper traffic-shaper state` - provides the total number of traffic shapers on the FortiGate unit.

`diagnose firewall shaper traffic-shaper stats` - provides summary statistics on the shapers. Sample output looks like the following:

```
shapers 9 ipv4 0 ipv6 0 drops 0
```

## Per-IP shaper

To view information for the per-IP shaper for security policies enter the command

```
diagnose firewall shaper per-ip-shaper list
```

The resultant output displays the information on all available per-IP shapers. The more shapers available the longer the list. For example:

```
name accounting_group
maximum-bandwidth 200000 Kb/sec
maximum-concurrent-session 55
packet dropped 0
```

Additional commands include:

`diagnose firewall shaper per-ip-shaper state` - provides the total number of per-ip shapers on the FortiGate unit.

`diagnose firewall shaper per-ip-shaper stats` - provides summary statistics on the shapers. Sample output looks like the following:

```
memory allocated 3 packet dropped: 0
```

You can also clear the per-ip statistical data to begin a fresh diagnoses using:

```
diagnose firewall shaper per-ip-shaper clear
```

## Packet loss with statistics on shapers

For each shaper there are counters that allow to verify if packets have been discarded. To view this information, in the CLI, enter the command `diagnose firewall shaper`. The results will look similar to the following output:

```
diagnose firewall shaper traffic-shaper list
name limit_GB_25_MB_50_LQ
maximum-bandwidth 50 Kb/sec
```

```
guaranteed-bandwidth 25 Kb/sec
current-bandwidth 51 Kb/sec
priority 3
dropped 1291985
```

The diagnose command output is different if the shapers are configured either per-policy or shared between policies.

For per-IP the output would be:

```
diagnose firewall shaper per-ip-shaper list

name accounting_group
maximum-bandwidth 200000 Kb/sec
maximum-concurrent-session 55
packet dropped 3264220
```

## Packet lost with the debug flow

When using the debug flow diagnostic command, there is a specific message information that a packet has exceeded the shaper limits and therefore discarded:

```
diagnose debug flow show console enable
diagnose debug flow filter addr 10.143.0.5
diagnose debug flow trace start 1000

id=20085 trace_id=11 msg="vd-root received a packet(proto=17,
 10.141.0.11:3735->10.143.0.5:5001) from port5."
id=20085 trace_id=11 msg="Find an existing session, id=0000eabc,
 original direction"
id=20085 trace_id=11 msg="exceeded shaper limit, drop"
```

## Session list details with dual traffic shaper

When a Security Policy has a different traffic shaper for each direction, it is reflected in the session list output from the CLI:

```
diagnose system session list

session info: proto=6 proto_state=02 expire=115 timeout=3600
 flags=00000000 sock
flag=00000000 sockport=0 av_idx=0 use=4
origin-shaper=Limit_25Mbps prio=1 guarantee 25600/sec max 204800/sec
 traffic 48/sec
reply-shaper=Limit_100Mbps prio=1 guarantee 102400/sec max 204800/sec
 traffic 0/sec
ha_id=0 hakey=44020
policy_dir=0 tunnel=/
state=may_dirty rem os rs
statistic(bits/packets/allow_err): org=96/2/1 reply=0/0/0 tuples=2
origin->sink: org pre->post, reply pre->post dev=2->3/3->2
 gwy=10.160.0.1/0.0.0.0
```

```
hook=pre dir=org act=dnat
 192.168.171.243:2538->192.168.182.110:80(10.160.0.1:80)
hook=post dir=reply act=snat
 10.160.0.1:80->192.168.171.243:2538(192.168.182.110:80)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=2 auth_info=0 chk_client_info=0 vd=0 serial=00011e81
 tos=ff/ff app=0 dd_type=0 dd_rule_id=0
```

## Additional Information

- Packets discarded by the shaper impact flow-control mechanisms like TCP. For more accurate testing results prefer UDP protocol.
- Traffic shaping accuracy is optimum for security policies without a protection profile where no FortiGate content inspection is processed.
- Do not oversubscribe an outbandwidth throughput. For example,  $\text{sum}[\text{guaranteed BW}] < \text{outbandwidth}$ . For accuracy in bandwidth calculation, it is required to set the “outbandwidth” parameter on the interfaces. For more information see [“Bandwidth guarantee, limit, and priority interactions” on page 2236](#).
- The FortiGate unit is not prioritizing traffic based on the DSCP marking configured in the security policy. However, TOS based prioritizing can be made at ingress. For more information see [“Differentiated Services” on page 2251](#).

# Chapter 18 Troubleshooting

This handbook chapter describes concepts of troubleshooting and solving issues that may occur with FortiGate units.

This FortiOS Handbook chapter contains the following chapters:

[Life of a Packet](#) explains the different layers and modules a packet goes through in FortiOS, including the order of operations.

[Verifying FortiGate admin access security](#) explains how to verify and configure administrative access.

[Troubleshooting resources](#) walks you through Fortinet's resources for troubleshooting.

[Troubleshooting tools](#) describes some of the basic commands and parts of FortiOS that can help you with troubleshooting.

[Troubleshooting methodologies](#) walks you through best practice concepts of FortiOS troubleshooting.

[Technical Support Organization Overview](#) describes how Fortinet Support operates, what they will need from you if you contact them, and what you can expect in general.

# Life of a Packet

Directed by security policies, a FortiGate unit screens network traffic from the IP layer up through the application layer of the TCP/IP stack. This chapter provides a general, high-level description of what happens to a packet as it travels through a FortiGate security system.

The FortiGate unit performs three types of security inspection:

- stateful inspection, that provides individual packet-based security within a basic session state
- flow-based inspection, that buffers packets and uses pattern matching to identify security threats
- proxy-based inspection, that reconstructs content passing through the FortiGate unit and inspects the content for security threats.

Each inspection component plays a role in the processing of a packet as it traverses the FortiGate unit in route to its destination. To understand these inspections is the first step to understanding the flow of the packet.

This section contains the following topics:

- [Stateful inspection](#)
- [Flow inspection](#)
- [Proxy inspection](#)
- [Comparison of inspection layers](#)
- [FortiOS functions and security layers](#)
- [Packet flow](#)
- [Example 1: client/server connection](#)
- [Example 2: Routing table update](#)
- [Example 3: Dialup IPsec VPN with application control](#)

## Stateful inspection

With stateful inspection, the FortiGate unit looks at the first packet of a session to make a security decision. Common fields inspected include TCP SYN and FIN flags to identify the start and end of a session, the source/destination IP, source/destination port and protocol. Other checks are also performed on the packet payload and sequence numbers to verify it as a valid communication and that the data is not corrupted or poorly formed.

What makes it stateful is that one or both ends must save information about the session history in order to communicate. In stateless communication, only independent requests and responses are used, that do not depend on previous data. For example, UDP is stateless by nature because it has no provision for reliability, ordering, or data integrity.

The FortiGate unit makes the decision to drop, pass or log a session based on what is found in the first packet of the session. If the FortiGate unit decides to drop or block the first packet of a session, then all subsequent packets in the same session are also dropped or blocked without being inspected. If the FortiGate unit accepts the first packet of a session, then all subsequent packets in the same session are also accepted without being inspected.

## Connections over connectionless

A connection is established when two end points use a protocol to establish connection through use of various methods such as segment numbering to ensure data delivery, and handshaking to establish the initial connection. Connections can be stateful because they record information about the state of the connection. Persistent connections reduce request latency because the end points do not need to re-negotiate the connection multiple times, but instead just send the information without the extra overhead. By contrast, connectionless communication does not keep any information about the data being sent or the state. It is based on an autonomous response/reply that is independent of other responses/replies that may have gone before. One example of connectionless communication is IP.

Benefits of connections over connectionless include being able to split data up over multiple packets, the data allows for a best-effort approach, and once the connection is established subsequent packets are not required to contain the full addressing information which saves on bandwidth. Connections are often reliable network services since acknowledgements can be sent when data is received.

## What is a session?

A session is established on an existing connection, for a defined period of time, using a determined type of communication or protocol. Sessions can have specific bandwidth, and time to live (TTL) parameters.

You can compare a session to a conversation. A session is established when one end point initiates a request by establishing a TCP connection on a particular port, the receiving end is listening on that port, and replies. You could telnet to port 80 even though telnet normally uses port 23, because at this level, the application being used cannot be determined.

However, the strong points of sessions and stateful protocols can also be their weak points. Denial of service (DoS) attacks involve creating so many sessions that the connection state information tables are full and the unit will not accept additional sessions.

## Differences between connections and sessions


In almost all cases, established sessions are stateful and all involve connections. However, some types of connections, such as UDP, are stateless, and are not sessions.

This means that not all traffic can be inspected by stateful inspection, because some of it is stateless. For example IP packets are stateless. Communications using HTTP are stateless, but HTTP often uses cookies to store persistent data in a way that approaches stateful.

Stateful inspection of sessions has the benefit of being able to apply the initial connection information to the packets that follow — the end points of the session will remain the same as will the protocol for example. That information can be examined for the first packet of the session and if it is malicious or not appropriate, the whole session can be dropped without committing significant resources.




**Figure 316:**Stateful inspection of packets through the FortiGate unit



## Flow inspection

With flow inspection (also called flow-based inspection), the FortiGate unit samples multiple packets in a session and multiple sessions, and uses a pattern matching engine to determine the kind of activity that the session is performing and to identify possible attacks or viruses. For example, if application control is operating, flow inspection can sample network traffic and identify the application that is generating the activity. Flow inspection using IPS samples network traffic and determines if the traffic constitutes an attack. Flow inspection can also be used for antivirus protection, web filtering, and data leak protection (DLP). Flow inspection occurs as the data is passing from its source to its destination. Flow inspection identifies and blocks security threats in real time as they are identified.

**Figure 317:**Flow inspection of packets through the FortiGate unit




Flow inspection typically requires less processing than proxy inspection, and therefore flow antivirus, web filtering, and DLP inspection performance can be better than proxy inspection performance. However, some threats can only be detected when a complete copy of the payload (for example a complete email attachment) is obtained so, proxy inspection tends to be more accurate and complete than flow inspection.

## Proxy inspection

Proxy inspection examines the content contained in content protocol sessions for security threats. Content protocols include HTTP, FTP, and email protocols. Security threats can be found in files and other content downloaded using these protocols. With proxy inspection, the FortiGate unit downloads the entire payload of a content protocol session and re-constructs it. For example, proxy inspection can reconstruct an email message and its attachments. After a satisfactory inspection the FortiGate unit passes the content on to the client. If the proxy inspection detects a security threat in the content, the content is removed from the communication stream before it reaches its destination. For example, if proxy inspection detects a virus in an email attachment, the attachment is removed from the email message before its sent to the client. Proxy inspection is the most thorough inspection of all, although it requires more processing power, and this may result in lower performance.

**Figure 318:**Proxy inspection of packets through the FortiGate unit



## Comparison of inspection layers

The three inspection methods each have their own strengths and weaknesses. The following table looks at all three methods side-by-side.

**Table 104:** Inspection methods comparison

Feature	Stateful	Flow	Proxy
Inspection unit per session	first packet	selected packets	complete content
Memory, CPU required	low	medium	high
Level of threat protection	good	better	best
Authentication	yes		
IPsec and SSL VPN	yes		
Antivirus protection		yes	yes
Web Filtering		yes	yes
Data Leak Protection (DLP)		yes	yes

**Table 104:** Inspection methods comparison

Feature	Stateful	Flow	Proxy
Application control		yes	
IPS		yes	
Delay in traffic		no	small
Reconstruct entire content		no	yes

## FortiOS functions and security layers

Within these security inspection types, FortiOS functions map to different inspections. The table below outlines when actions are taken as a packet progresses through its life within a FortiGate unit.

**Table 105: FortiOS security functions and security layers**

Security Function	Stateful	Flow	Proxy
Firewall	yes		
IPsec VPN	yes		
Traffic Shaping	yes		
User Authentication	yes		
Management Traffic	yes		
SSL VPN	yes		
Intrusion Prevention		yes	
Antivirus		yes	yes
Application Control		yes	
Web filtering		yes	yes
DLP			yes
Email Filtering		yes	yes
VoIP inspection			yes

## Packet flow

After the FortiGate unit's external interface receives a packet, the packet proceeds through a number of steps on its way to the internal interface, traversing each of the inspection types, depending on the security policy and security profile configuration. The diagram in [Figure 319 on page 2276](#) is a high level view of the packet's journey.

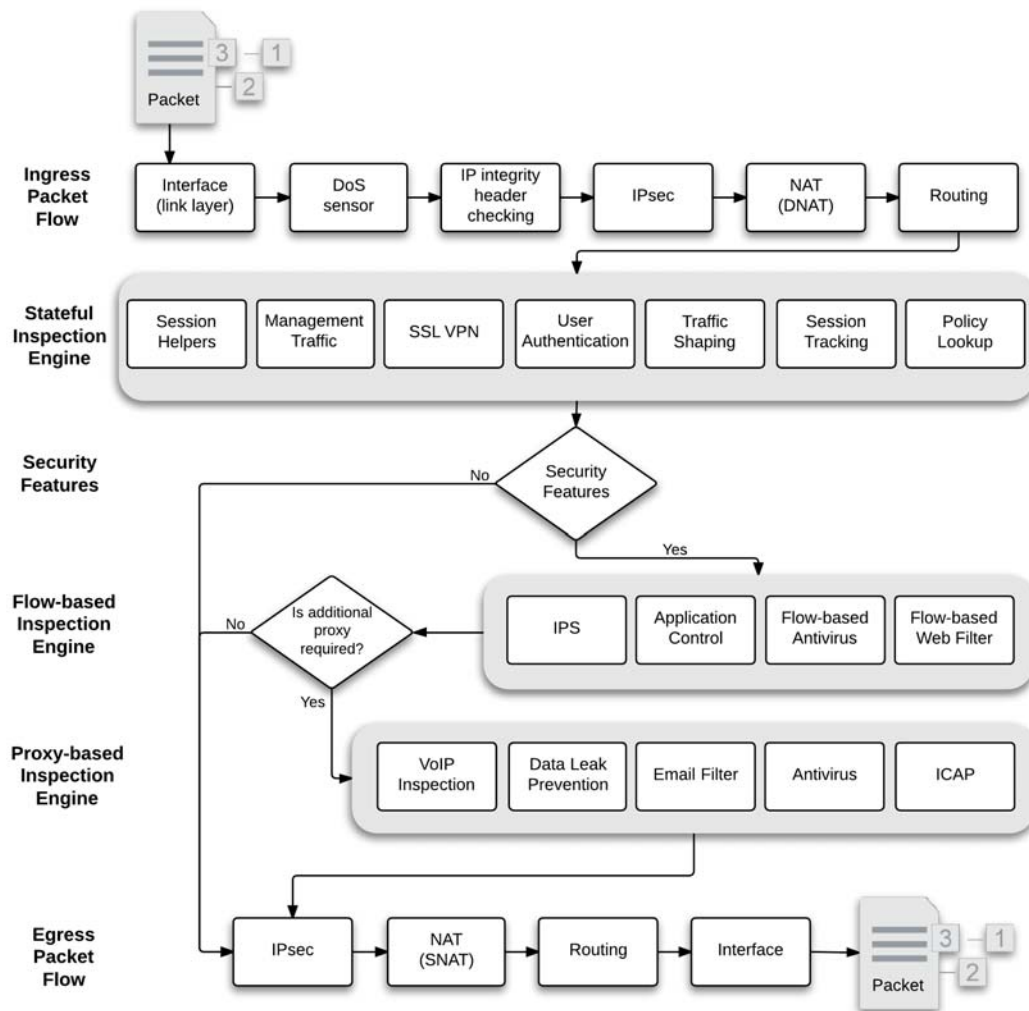
The description following is a high-level description of these steps as a packet enters the FortiGate unit towards its destination on the internal network. Similar steps occur for outbound traffic.

## Packet inspection (Ingress)

In [Figure 319 on page 2276](#), in the first set of steps (ingress), a number of header checks take place to ensure the packet is valid and contains the necessary information to reach its destination. This includes:

- Packet verification - during the IP integrity stage, verification is performed to ensure that the layer 4 protocol header is the correct length. If not, the packet is dropped.
- Session creation - the FortiGate unit attempts to create a session for the incoming data
- IP stack validation for routing - the firewall performs IP header length, version and checksum verifications in preparation for routing the packet.
- Verifications of IP options - the FortiGate unit validates the routing information

**Figure 319:**Packet flow process



## Interface

Ingress packets are received by a FortiGate interface. The packet enters the system, and the interface network device driver passes the packet to the Denial of Service (DoS) sensors, if enabled, to determine whether this is a valid information request or not.

## DoS sensor

DoS scans are handled very early in the life of the packet to determine whether the traffic is valid or is part of a DoS attack. Unlike signature-based IPS which inspects all the packets within a certain traffic flow, the DoS module inspects all traffic flows but only tracks packets that can be used for DoS attacks (for example TCP SYN packets), to ensure they are within the permitted parameters. Suspected DoS attacks are blocked, other packets are allowed.

## IP integrity header checking

The FortiGate unit reads the packet headers to verify if the packet is a valid TCP, UDP, ICMP, SCTP or GRE packet. The only verification that is done at this step to ensure that the protocol header is the correct length. If it is, the packet is allowed to carry on to the next step. If not, the packet is dropped.

## IPsec

If the packet is an IPsec packet, the IPsec engine attempts to decrypt it. The IPsec engine applies the correct encryption keys to the IPsec packet and sends the unencrypted packet to the next step. IPsec is bypassed when for non-IPsec traffic and for IPsec traffic that cannot be decrypted by the FortiGate unit.

## Destination NAT (DNAT)

The FortiGate unit checks the NAT table and determines the destination IP address for the traffic. This step determines whether a route to the destination address actually exists.

For example, if a user's browser on the internal network at IP address 192.168.1.1 visited the web site www.example.com using NAT, after passing through the FortiGate unit the source IP address becomes NATed to the FortiGate unit external interface IP address. The destination address of the reply back from www.example.com is the IP address of the FortiGate unit internal interface. For this reply packet to be returned to the user, the destination IP address must be destination NATed to 192.168.1.1.

DNAT must take place before routing so that the FortiGate unit can route packets to the correct destination.

## Routing

The routing step determines the outgoing interface to be used by the packet as it leaves the FortiGate unit. In the previous step, the FortiGate unit determined the real destination address, so it can now refer to its routing table and decide where the packet must go next.

Routing also distinguishes between local traffic and forwarded traffic and selects the source and destination interfaces used by the security policy engine to accept or deny the packet.

## Policy lookup

The policy look up is where the FortiGate unit reviews the list of security policies which govern the flow of network traffic, from the first entry to the last, to find a match for the source and

destination IP addresses and port numbers. The decision to accept or deny a packet, after being verified as a valid request within the stateful inspection, occurs here. A denied packet is discarded. An accepted packet will have further actions taken. If IPS is enabled, the packet will go to [Flow-based inspection engine](#), otherwise it will go to the [Proxy-based inspection engine](#).

If no other security options are enabled, then the session was only subject to stateful inspection. If the action is accept, the packet will go to Source NAT to be ready to leave the FortiGate unit.

## Session tracking

Part of the stateful inspection engine, session tracking maintains session tables that maintain information about sessions that the stateful inspection module uses for maintaining sessions, NAT, and other session related functions.

## User authentication

User authentication added to security policies is handled by the stateful inspection engine, which is why Firewall authentication is based on IP address. Authentication takes place after policy lookup selects a security policy that includes authentication. This is also known as identify-based policies. Authentication also takes place before security features are applied to the packet.

## Management traffic

This local traffic is delivered to the FortiGate unit TCP/IP stack and includes communication with the web-based manager, the CLI, the FortiGuard network, log messages sent to FortiAnalyzer or a remote syslog server, and so on. Management traffic is processed by applications such as the web server which displays the FortiOS web-based manager, the SSH server for the CLI or the FortiGuard server to handle local FortiGuard database updates or FortiGuard Web Filtering URL lookups.

## SSL VPN traffic

For local SSL VPN traffic, the internal packets are decrypted and are routed to a special interface. This interface is typically called `ssl.root` for decryption. Once decrypted, the packets go to policy lookup.

## ICAP traffic

If you enable ICAP in a security policy, HTTP (and optionally HTTPS) traffic intercepted by the policy is transferred to ICAP servers in the ICAP profile added to the policy. The FortiGate unit is the surrogate, or “middle-man”, and carries the ICAP responses from the ICAP server to the ICAP client; the ICAP client then responds back, and the FortiGate unit determines the action that should be taken with these ICAP responses and requests.

## Session helpers

Some protocols include information in the packet body (or payload) that must be analyzed to successfully process sessions for this protocol. For example, the SIP VoIP protocol uses TCP control packets with a standard destination port to set up SIP calls. To successfully process SIP VoIP calls, FortiOS must be able to extract information from the body of the SIP packet and use this information to allow the voice-carrying packets through the firewall.

FortiOS uses session helpers to analyze the data in the packet bodies of some protocols and adjust the firewall to allow those protocols to send packets through the firewall.

## Flow-based inspection engine

Flow-based inspection is responsible for IPS, application control, flow-based antivirus scanning and VoIP inspection. Packets are sent to flow-based inspection if the security policy that accepts the packets includes one or more of these security features.



Flow-based antivirus scanning is only available on some FortiGate models.

---

Once the packet has passed the flow-based engine, it can be sent to the proxy inspection engine or egress.

## Proxy-based inspection engine

The proxy inspection engine is responsible for carrying out antivirus protection, email filtering (antispam), web filtering and data leak prevention. The proxy engine will process multiple packets to generate content before it is able to make a decision for a specific packet.

## IPsec

If the packet is transmitted through an IPsec tunnel, it is at this stage the encryption and required encapsulation is performed. For non-IPsec traffic (TCP/UDP) this step is bypassed.

## Source NAT (SNAT)

When preparing the packet to leave the FortiGate unit, it needs to NAT the source address of the packet to the external interface IP address of the FortiGate unit. For example, a packet from a user at 192.168.1.1 accessing www.example.com is now using a valid external IP address as its source address.

## Routing

The final routing step determines the outgoing interface to be used by the packet as it leaves the FortiGate unit.

## Egress

Upon completion of the scanning at the IP level, the packet exits the FortiGate unit.

## Example 1: client/server connection

The following example illustrates the flow of a packet of a client/web server connection with authentication and FortiGuard URL and antivirus filtering.

This example includes the following steps:

### Initiating connection from client to web server

1. Client sends packet to web server.

2. Packet intercepted by FortiGate unit interface.
  - 2.1 Link level CRC and packet size checking. If the size is correct, the packet continues, otherwise it is dropped.
3. DoS sensor - checks are done to ensure the sender is valid and not attempting a denial of service attack.
4. IP integrity header checking, verifying the IP header length, version and checksums.
5. Next hop route
6. Policy lookup
7. User authentication
8. Proxy inspection
  - 8.1 Web Filtering
  - 8.2 FortiGuard Web Filtering URL lookup
  - 8.3 Antivirus scanning
9. Source NAT
10. Routing
11. Interface transmission to network
12. Packet forwarded to web server

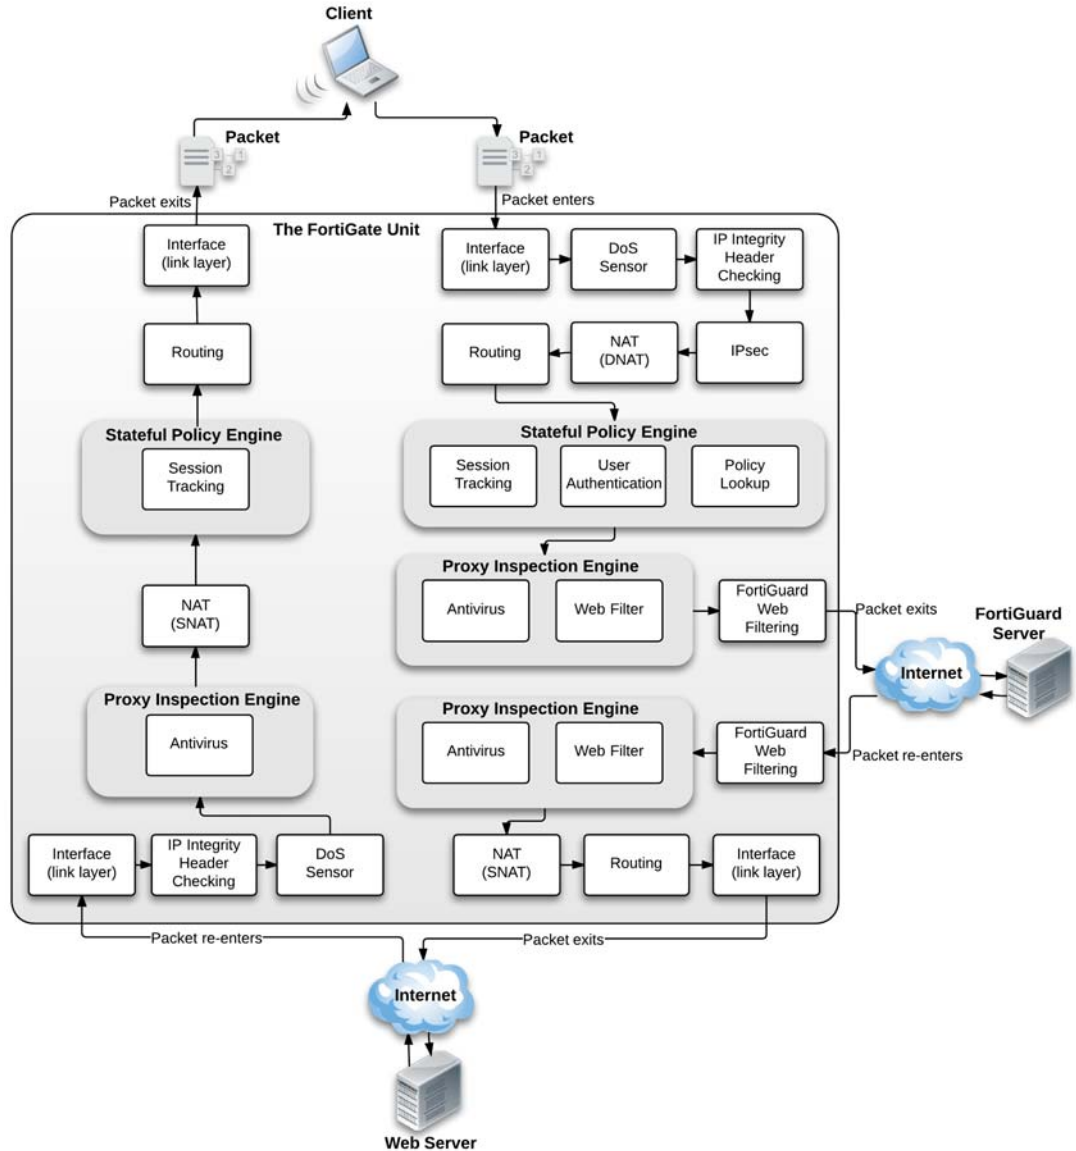
#### **Response from web server**

1. Web Server sends response packet to client.
2. Packet intercepted by FortiGate unit interface
  - 2.1 Link level CRC and packet size checking.
3. IP integrity header checking.
4. DoS sensor.
5. Proxy inspection
  - 5.1 Antivirus scanning.
6. Source NAT.
7. Stateful Policy Engine
  - 7.1 Session Tracking
8. Next hop route
9. Interface transmission to network
10. Packet returns to client

This process is illustrated in [Figure 320](#).



**Figure 320:**Client/server connection



## Example 2: Routing table update

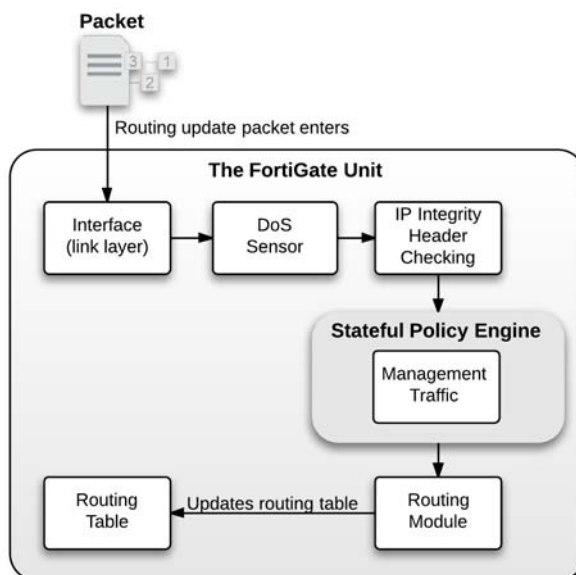
The following example illustrates the flow of a packet when there is a routing table update. As this is low level, there is no security involved. This example includes the following steps:

1. FortiGate unit receives routing update packet
2. Packet intercepted by FortiGate unit interface
  - 2.1 Link level CRC and packet size checking. If the size is correct, the packet continues, otherwise it is dropped.
3. DoS sensor - checks are done to ensure the sender is valid and not attempting a denial of service attack.
4. IP integrity header checking, verifying the IP header length, version and checksums.
5. Stateful policy engine
  - 5.1 Management traffic (local traffic)

6. Routing module
  - 6.1 Update routing table

Figure 321 illustrates the process steps.

**Figure 321:**Routing table update



### Example 3: Dialup IPsec VPN with application control

This example includes the following steps:

1. FortiGate unit receives IPsec packet from Internet
2. Packet intercepted by FortiGate unit interface
  - 2.1 Link level CRC and packet size checking. If the size is correct, the packet continues, otherwise it is dropped.
3. DoS sensor - checks are done to ensure the sender is valid and not attempting a denial of service attack.
4. IP integrity header checking, verifying the IP header length, version and checksums.
5. IPsec
  - 5.1 Determines that packet matched IPsec phase 1 configuration
  - 5.2 Unencrypted packet
6. Next hop route
7. Stateful policy engine
  - 7.1 Session tracking
8. Flow inspection engine
  - 8.1 IPS
  - 8.2 Application control
9. Source NAT
10. Routing

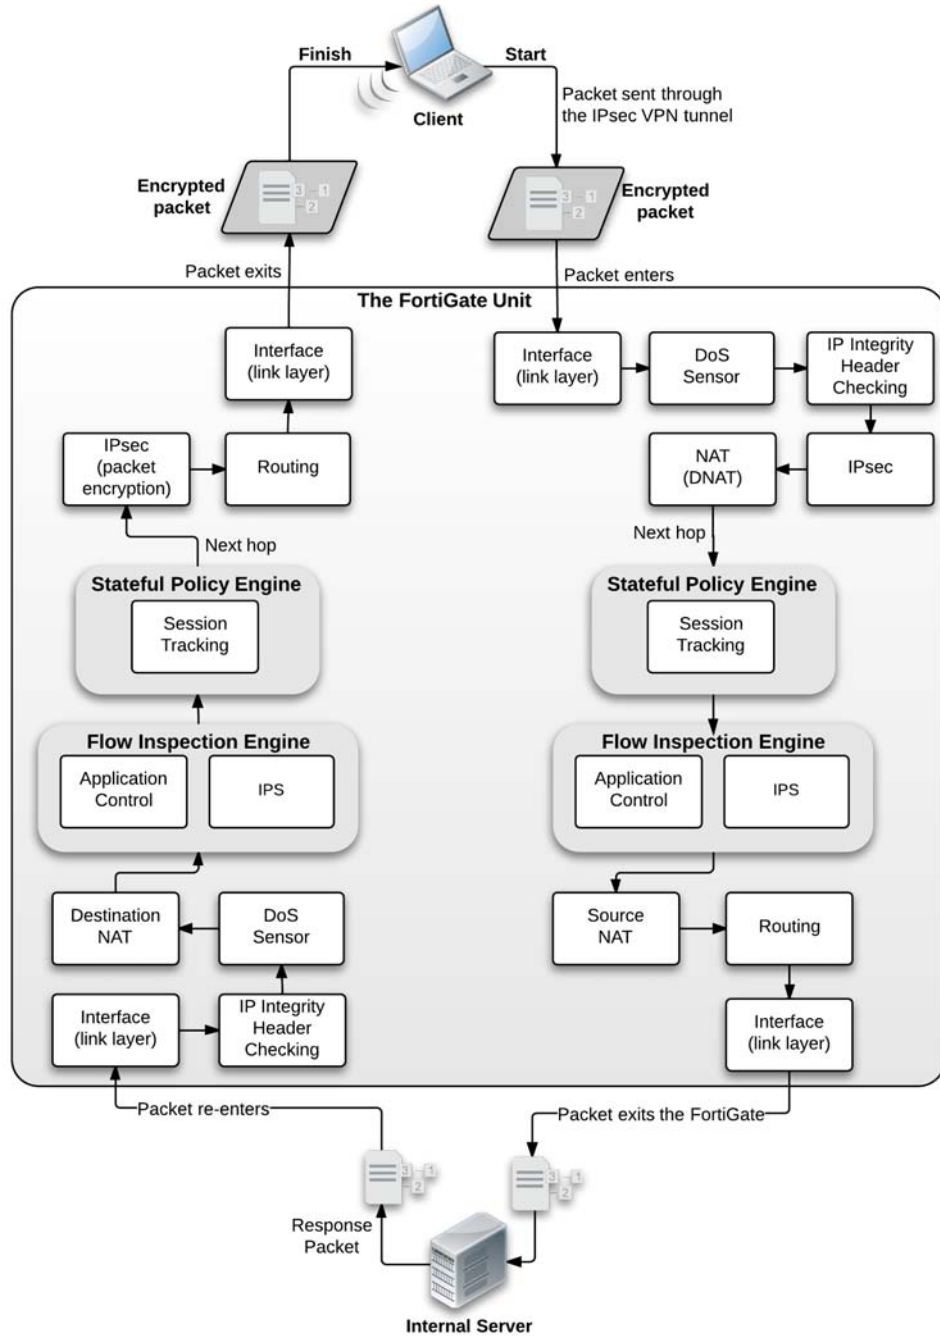
11. Interface transmission to network
12. Packet forwarded to internal server

**Response from server**

1. Server sends response packet
2. Packet intercepted by FortiGate unit interface
  - 2.1 Link level CRC and packet size checking
3. IP integrity header checking.
4. DoS sensor
5. Flow inspection engine
  - 5.1 IPS
  - 5.2 Application control
6. Stateful policy engine
  - 6.1 Session tracking
7. Next hop route
8. IPsec
  - 8.1 Encrypts packet
9. Routing
10. Interface transmission to network
11. Encrypted Packet returns to internet

Figure 322 illustrates the process.

**Figure 322:** Dialup IPsec with application control



# Verifying FortiGate admin access security

FortiOS provides a number of methods that help to enhance FortiGate administrative access security. This section describes FortiGate administrative access security best practices.

- Install the FortiGate unit in a physically secure location
- Add new administrator accounts
- Change the admin account name and limit access to this account
- Only allow administrative access to the external interface when needed
- When enabling remote access, configure Trusted Hosts and Two-factor Authentication
- Change the default administrative port to a non-standard port
- Enable Password Policy
- Maintain short login timeouts
- Modify administrator account Lockout Duration and Threshold values
- Disable auto installation via USB
- Auditing and Logging

## Install the FortiGate unit in a physically secure location

A good place to start with is physical security. Install the FortiGate unit in a secure location, such as a locked room or a room with restricted access. This way unauthorized users can't get physical access to the device.

If unauthorized users have physical access they can disrupt your entire network by disconnecting your FortiGate unit (either by accident or on purpose). They could also connect a console cable and attempt to log into the CLI. Also, when a FortiGate unit reboots, a person with physical access can interrupt the boot process and install different firmware.

## Add new administrator accounts

Rather than allowing all administrators to access the FortiGate unit with the admin administrator account you should create administrator accounts for each person that requires administrative access. That way you can track who has made configuration changes and performed other administrative activities. Keep the number of administrative accounts to a minimum to keep better control on who can access the device.

To add administrators go to *System > Admin > Administrators* and select *Create New*.

If you want administrators to have access to all FortiGate configuration options, their accounts should have the *prof\_admin* admin profile. Administrators with this profile can do anything except add new administrator accounts.

At least one account should always have the *super\_admin* profile as this profile is required to add and remove administrators. To improve security only a very few administrators (usually one) should be able to add new administrators.

If you want some administrator accounts to have limited access to the FortiGate configuration you can create custom admin profiles that only allow access to selected parts of the configuration. To add custom admin profiles, go to *System > Admin > Admin Profiles* and select *Create New*.

For example, if you want to add an admin profile that does not allow changing firewall policies, when you configure the admin profile set *Firewall Configuration* to *None* or *Read Only*.

## Change the admin account name and limit access to this account

The default super\_admin administrator account, admin, is a well known administrator name so if this account is available it could be easier for attackers to access the FortiGate unit because they know they can log in with this name, only having to determine the password. You can improve security by changing this name to one more difficult for an attacker to guess. To do this, create a new administrator account with the super\_admin admin profile and log in as that administrator. Then go to *System > Admin > Administrators* and edit the admin administrator and change the Administrator name.

Once the account has been renamed you could delete the super\_admin account that you just added. Consider also only using the super-admin account for adding or changing administrators. The less this account is used to less likely that it could be compromised. You could also store the account name and password for this account in a secure location in case for some reason the account name or password is forgotten.

## Only allow administrative access to the external interface when needed

When possible, don't allow administration access on the external interface and use internal access methods such as IPsec VPN or SSL VPN.

To disable administrative access on the external interface, go to *System > Network > Interfaces*, edit the external interface and disable HTTPS, PING, HTTP, SSH, and TELNET under *Administrative Access*.

This can also be done with CLI using following commands:

```
config system interface
 edit <external_interface_name>
 unset allowaccess
 end
```

Please note that this will disable all services on the external interface including CAPWAP, FMG-Access, SNMP, and FCT-Access.

If you need some of these services enabled on your external interface, for example CAPWAP and FMG-Access to ensure connectivity between FortiGate unit and respectively FortiAP and FortiManager, then you need to use following CLI command:

```
config system interface
 edit <external_interface_name>
 set allowaccess capwap fgfm
 end
```

## When enabling remote access, configure Trusted Hosts and Two-factor Authentication

If you have to have remote access and can't use IPsec or SSL VPN then you should only allow HTTPS and SSH and use secure access methods such as trusted hosts and Two-factor authentication.

### Configuring Trusted Hosts

Setting trusted hosts for administrators limits what computers an administrator can log in the FortiGate unit from. When you identify a trusted host, the FortiGate unit will only accept the administrator's login from the configured IP address or subnet. Any attempt to log in with the same credentials from any other IP address or any other subnet will be dropped. To ensure the administrator has access from different locations, you can enter up to ten IP addresses or subnets. Ideally, this should be kept to a minimum. For higher security, use an IP address with a net mask of 255.255.255.255, and enter an IP address (non-zero) in each of the three default trusted host fields.

Trusted hosts are configured when adding a new administrator by going to *System > Admin > Administrators* in the web-based manager and selecting *Restrict this Admin Login from Trusted Hosts Only*, or `config system admin` in the CLI.

The trusted hosts apply to the web-based manager, ping, snmp and the CLI when accessed through SSH. CLI access through the console port is not affected.

Also ensure all entries contain actual IP addresses, not the default 0.0.0.0.

### Configuring Two-factor Authentication

FortiOS 5.0 provides support for FortiToken and FortiToken Mobile. FortiToken Mobile is a Fortinet application that enables you to generate One Time Passwords (OTPs) on a mobile device for FortiGate two-factor authentication. The user's mobile device and the FortiGate unit must be connected to the Internet to activate FortiToken mobile. Once activated, users can generate OTPs on their mobile device without having network access. FortiToken Mobile is available for iOS and Android devices from their respective Application stores. No cellular network is required for activation.

The latest FortiToken Mobile documentation is available from the [FortiToken](#) page of the [Fortinet Technical Documentation](#) website.

Two free trial tokens are included with every registered FortiGate unit. Additional tokens can be purchased from your reseller or from Fortinet.

To assign a token to an administrator go to *System > Admin > Administrators* and either add a new or select an existing administrator to assign the token to. Configure the administrator as required, you need to enter your email address and phone number in order to receive the activation code for the FortiToken mobile. Select *Enable Two-factor Authentication*. Select the token to associate with the administrator. Select *OK* to assign the token to the administrator.

To configure your FortiGate unit to send email or SMS messages go to *System > Config > Messaging Servers*.

## Change the default administrative port to a non-standard port

Administration Settings under *System > Admin > Settings* or `config system global` in the CLI, enable you to change the default port configurations for administrative connections to the FortiGate unit for added security. When connecting to the FortiGate unit when the port has changed, the port must be included. For example, if you are connecting to the FortiGate unit using HTTPS over port 8081, the url would be `https://192.168.1.99:8081`

If you make a change to the default port number for HTTP, HTTPS, Telnet, or SSH, ensure that the port number is not used for other services.

## Enable Password Policy

Brute force password software can launch more than just dictionary attacks. It can discover common passwords where a letter is replaced by a number. For example, if “p4ssw0rd” is used as a password, it can be cracked.

Password policies, available by going to *System > Admin > Settings > Enable Password Policy*, enable you to create a password policy that any administrator who updates their passwords, must follow. Using the available options you can define the required length of the password, what it must contain (numbers, upper and lower case, and so on) and an expiry time frame. The FortiGate unit will warn of any password that is added and does not meet the criteria.

## Maintain short login timeouts

To avoid the possibility of an administrator walking away from the management computer and leaving it exposed to unauthorized personnel, you can add an idle time-out. That is, if the web-based manager is not used for a specified amount of time, the FortiGate unit will automatically log the administrator out. To continue their work, they must log in again.

The time-out can be set as high as 480 minutes, or eight hours, although this is not recommend.

To set the idle time out, go to *System > Admin > Settings* and enter the amount of time for the Idle Timeout. A best practice is to keep the default of 5 min.

When logging into the console using SSH, the default time of inactivity to successfully log into the FortiGate unit is 120 seconds (2 minutes). You can configure the time to be shorter by using the CLI to change the length of time the command prompt remains idle before the FortiGate unit will log the administrator out. The range can be between 10 and 3600 seconds. To set the logout time enter the following CLI commands:

```
config system global
 set admin-ssh-grace-time <number_of_seconds>
end
```

## Modify administrator account Lockout Duration and Threshold values

Account lockout policies control how and when accounts are locked out of the FortiGate unit. These policies are described and implemented as follows:



## Administrator account Lockout Duration

If someone violates the lockout controls by entering an incorrect user name and/or password, account lockout duration sets the length of time the account is locked. The lockout duration can be set to a specific length of time using a value between 1 and 4294967295 seconds. The default value is 60 seconds.

When it's required use the CLI to modify the lockout duration as follow:

```
config system global
 set admin-lockout-duration <integer>
end
```

## Administrator account Lockout Threshold

The lockout threshold sets the number of invalid logon attempts that are allowed before an account is locked out. You may set a value that balances the need to prevent account cracking against the needs of an administrator who may have difficulty accessing their account.

Its normal for an administrator to sometimes take a few attempts to logon with the right password.

The lockout threshold can be set to any value from 1 to 10. The Default value is 3, which is normally a good setting. However, to improve security you could reduce it to 1 or 2 as long as administrators know to take extra care when entering their passwords.

Use the following CLI command to modify the lockout threshold:

```
config system global
 set admin-lockout-threshold <integer>
end
```

Keep in mind that the higher the lockout value, the higher the risk that someone may be able to break into the FortiGate unit.

## Disable auto installation via USB

An attacker with a physical access to the device could load a new configuration or firmware on the FortiGate using the USB port, reinitializing the device through a power cut. To avoid this, execute the following CLI commands:

```
config system auto-install
 set auto-install-config disable
 set auto-install-image disable
end
```

## Auditing and Logging

Audit web facing administration interfaces. By default, FortiGate logs all deny action, you can check these actions by going to *Log & Report > Event Log > System*. This default behavior should not be changed. Also secure log files in a central location such as FortiCloud and configure alert email which provides an efficient and direct method of notifying an administrator of events. You can configure log settings by going to *Log & Report > Log Config*.

An auditing schedule should be established to routinely inspect logs for signs of intrusion and probing.

# Troubleshooting resources

Before you begin troubleshooting, you need to know Fortinet's troubleshooting resources. Doing so will shorten the time to solve your issue. Indeed, an administrator can save time and effort during the troubleshooting process by first checking if the issue has been experienced before. Several self-help resources are available to provide valuable information about FortiOS technical issues, including:

## Technical Documentation

Installation Guides, Administration Guides, Quick Start Guides, and other technical documents are available online at the following URL:

<http://docs.fortinet.com>

## Fortinet Video Library

The Fortinet Video Library hosts a collection of video which provide valuable information about Fortinet products.

<http://video.fortinet.com>

## Release Notes

Issues that are uncovered after the technical documentation has been published will often be listed in the Release Notes that accompany the device.

## Knowledge Base

The Fortinet Knowledge Base provides access to a variety of articles, white papers, and other documentation providing technical insight into a range of Fortinet products. The Knowledge Base is available online at the following URL:

<http://kb.fortinet.com>

## Fortinet Technical Discussion Forums

An online technical forums allow administrators to contribute to discussions about issues related to their Fortinet products. Searching the forum can help the administrator identify if an issue has been experienced by another user. The support forums can be accessed at the following URL:

<http://support.fortinet.com/forum>

## Fortinet Training Services Online Campus

The Fortinet Training Services Online Campus hosts a collection of tutorials and training materials which can be used to increase knowledge of the Fortinet products.

<http://campus.training.fortinet.com>

## Fortinet Customer Support

You have defined your problem, researched a solution, put together a plan to find the solution, and executed that plan. At this point if the problem has not been solved, its time to contact Fortinet Customer Support for assistance.

<http://support.fortinet.com>

# Troubleshooting tools

FortiOS provides a number of tools that help with troubleshooting both hardware and software issues. These tools include diagnostics and ports; ports are used when you need to understand the traffic coming in or going out on a specific port, for example, UDP 53, which is used by the FortiGate unit for DNS lookup and RBL lookup.

This section also contains information about troubleshooting FortiGuard issues.

This section contains the following topics:

- [FortiOS diagnostics](#)
- [FortiOS ports](#)
- [FortiAnalyzer/FortiManager ports](#)
- [FortiGuard troubleshooting](#)

## FortiOS diagnostics

A collection of diagnostic commands are available in FortiOS for troubleshooting and performance monitoring. Within the CLI commands, the two main groups of diagnostic commands are `get` and `diagnose` commands. Both commands display information about system resources, connections, and settings that enable you to locate and fix problems, or to monitor system performance.

This topic includes diagnostics commands to help with:

- [Check date and time](#)
- [Resource usage](#)
- [Proxy operation](#)
- [Hardware NIC](#)
- [Traffic trace](#)
- [Session table](#)
- [Firewall session setup rate](#)
- [Finding object dependencies](#)
- [Flow trace](#)
- [Packet sniffing and packet capture](#)
- [FA2 and NP2 based interfaces](#)
- [Debug command](#)
- [The execute tac report command](#)
- [Other commands](#)

Additional diagnostic commands related to specific features are covered in the chapter for that specific feature. For example in-depth diagnostics for dynamic routing are covered in the dynamic routing chapter.

### Check date and time

The system date and time are important for FortiGuard services, when logging events, and when sending alerts. The wrong time will make the log entries confusing and difficult to use.

Use Network Time Protocol (NTP) to set the date and time if possible. This is an automatic method that does not require manual intervention. However, you must ensure the port is allowed through the firewalls on your network. FortiToken synchronization requires NTP in many situations.

### How to check the date and time - web-based manager

#### 1. Go to *System Information* > *System Time* on the dashboard.

Alternately, you can check the date and time using the CLI commands `execute date` and `execute time`.

#### 2. If required, select *Change to adjust the date and time settings*.

You can set the time zone, date and time, and select NTP usage. In the CLI, use the following commands to change the date and time:

```
config system global
 set timezone (use ? to get a list of IDs and descriptions of their
 timezone)
 set
config system ntp
 config ntpserver
 edit 1
 set server "ntp1.fortinet.net"
 next
 edit 2
 set server "ntp2.fortinet.net"
 next
 end
 set ntpsync enable
 set syncinterval 60
end
```

## Resource usage

Each program running on a computer has one or more processes associated with it. For example if you open a Telnet program, it will have an associated telnet process. The same is true in FortiOS. All the processes have to share the system resources in FortiOS including memory and CPU.

Use `get system performance status` command to show the FortiOS performance status.

Sample output:

```
FGT#get system performance status
CPU states: 0% user 0% system 0% nice 100% idle
CPU0 states: 0% user 0% system 0% nice 100% idle
CPU1 states: 0% user 0% system 0% nice 100% idle
CPU2 states: 0% user 0% system 0% nice 100% idle
CPU3 states: 0% user 0% system 0% nice 100% idle
Memory states: 25% used
Average network usage: 0 kbps in 1 minute, 0 kbps in 10 minutes, 0 kbps
in 30 minutes
Average sessions: 5 sessions in 1 minute, 5 sessions in 10 minutes, 4
sessions in 30 minutes
Average session setup rate: 0 sessions per second in last 1 minute, 0
sessions per second in last 10 minutes, 0 sessions per second in
last 30 minutes
Virus caught: 0 total in 1 minute
IPS attacks blocked: 0 total in 1 minute
Uptime: 0 days, 12 hours, 7 minutes
```

Monitor the CPU/memory usage of internal processes using the following command:

```
get system performance top <delay> <max_lines>
```

The data listed by the command includes the name of the daemon, the process ID, whether the process is sleeping or running, the CPU percentage being used, and the memory percentage being used.

Sample output:

```
FGT#get system performance top 10 100
Run Time: 0 days, 11 hours and 30 minutes
0U, 0S, 100I; 1977T, 1470F, 121KF
 pyfcgid 120 S 0.0 1.3
 pyfcgid 121 S 0.0 1.3
 pyfcgid 122 S 0.0 1.3
 pyfcgid 53 S 0.0 1.3
 ipsengine 75 S < 0.0 1.3
 ipsengine 66 S < 0.0 1.3
 ipsengine 73 S < 0.0 1.3
 ipsengine 74 S < 0.0 1.3
 ipsengine 79 S < 0.0 1.3
 ipsengine 80 S < 0.0 1.3
 cmdbsvr 43 S 0.0 1.0
 proxyworker 110 S 0.0 1.0
 proxyworker 111 S 0.0 1.0
 httpsd 125 S 0.0 0.8
 httpsd 52 S 0.0 0.8
 httpsd 124 S 0.0 0.8
 newcli 141 R 0.0 0.7
 newcli 128 S 0.0 0.7
 fgfmd 102 S 0.0 0.7
 iked 86 S 0.0 0.7
```

## Proxy operation

Monitor proxy operations using the following command:

```
diag test application <application> <option>
```

The <application> value can include the following:

acd	Aggregate Controller.
ddnscd	DDNS client daemon.
dhcp6c	DHCP6 client daemon.
dhcprelay	DHCP relay daemon.
dlpfingerprint	DLP fingerprint daemon.
dlpfpcache	DLP fingerprint cache daemon.
dnsproxy	DNS proxy.
dsd	DLP Statistics daemon.
forticldd	FortiCloud daemon.
forticron	FortiCron daemon.
fsd	FortiExplorer daemon.
ftpd	FTP proxy.
harelay	HA relay daemon.
http	HTTP proxy.
imap	IMAP proxy.
info-sslvpn	SSL-VPN info daemon.
ipldbd	IP load balancing daemon.
ipsengine	ips sensor
ipsmonitor	ips monitor
ipsufd	IPS urlfilter daemon.
l2tpcd	L2TP client daemon.
ltd	USB LTE daemon.
miglogd	Miglog logging daemon.
nat64d	NAT 64 daemon.
nntp	NNTP proxy.
pop3	POP3 proxy.

pptpcd	PPTP client.
proxyacceptor	Proxy acceptor.
proxyworker	Proxy worker.
quarantined	Quarantine daemon.
radiusd	RADIUS daemon.
reportd	Report daemon.
reputation	Client reputation daemon.
scanunit	Scanning unit.
sflowd	sFlow daemon.
smtp	SMTP proxy.
snmpd	SNMP daemon.
sqldb	SQL database daemon.
ssh	SSH proxy.
sslacceptor	SSL proxy.
sslworker	SSL proxy.
swctrl_authd	Switch controller authentication daemon.
uploadd	Upload daemon.
urlfilter	URL filter daemon.
wa_cs	WAN optimization cs server.
wa_dbd	WAN optimization storage server.
wad	WAN optimization proxy.
wad_diskd	WAN optimization disk access daemon.
wccpd	WCCP daemon.
wpad	WPA daemon.

The `<option>` value depends from the application value used in the command. Here are some examples:

- If the application is `http`, the CLI command will be `diag test application http <option>`

The `<option>` value can be one from the following:



2	Drop all connections
22	Drop max idle connections
222	Drop all idle connections
4	Display connection stat
44	Display info per connection
444	Display connections per state
4444	Display per-VDOM statistics
44444	Display information about idle connections
55	Display tcp info per connection
6	Display ICAP information
70	Disable ICAP 'Allow: 204' (default)
71	Enable ICAP 'Allow: 204'
72	Drop all ICAP server connections
11	Display the SSL session ID cache statistics
12	Clear the SSL session ID cache statistics
13	Display the SSL session ID cache
14	Clear the SSL session ID cache
80	Show Fortinet bar SSL-VPN bookmark info
81	Show Fortinet bar SSL-VPN bookmark cache
82	Show Fortinet bar SSL-VPN bookmark LRU list

- If the application is `ipsmonitor`, the CLI command will be `diag test application ipsmonitor <option>`

The `<option>` value can be one from the following:

1	Display IPS engine information
2	Toggle IPS engine enable/disable status
3	Display restart log
4	Clear restart log
5	Toggle bypass status

6	Submit attack characteristics now
10	IPS queue length
11	Clear IPS queue length
12	IPS L7 socket statistics
13	IPS session list
14	IPS NTurbo statistics
15	IPSA statistics
97	Start all IPS engines
98	Stop all IPS engines
99	Restart all IPS engines and monitor

## Hardware NIC

Monitor hardware network operations using the following command:

```
diag hardware deviceinfo nic <interface>
```

The information displayed by this command is important as errors at the interface are indicative of data link or physical layer issues which may impact the performance of the FortiGate unit.

The following is sample output when <interface> = internal:

```
System_Device_Name port5
Current_HWaddr 00:09:0f:68:35:60
Permanent_HWaddr 00:09:0f:68:35:60
Link up
Speed 100
Duplex full
[.....]
Rx_Packets=5685708
Tx_Packets=4107073
Rx_Bytes=617908014
Tx_Bytes=1269751248
Rx_Errors=0
Tx_Errors=0
Rx_Dropped=0
Tx_Dropped=0
[... .]
```

The `diag hardware deviceinfo nic` command displays a list of hardware related error names and values. The following table explains the items in the list and their meanings.

**Table 106:** Possible hardware errors and meanings

Field	Definition
Rx_Errors = rx error count	Bad frame was marked as error by PHY.
Rx_CRC_Errors + Rx_Length_Errors - Rx_Align_Errors	This error is only valid in 10/100M mode.
Rx_Dropped or Rx_No_Buffer_Count	Running out of buffer space.
Rx_Missed_Errors	Equals Rx_FIFO_Errors + CEXTERR (Carrier Extension Error Count). Only valid in 1000M mode, which is marked by PHY.
Tx_Errors = Tx_Aborted_Errors	ECOL (Excessive Collisions Count). Only valid in half-duplex mode.
Tx_Window_Errors	LATECOL (Late Collisions Count). Late collisions are collisions that occur after 64-byte time into the transmission of the packet while working in 10 to 100Mb/s data rate and 512-byte time into the transmission of the packet while working in the 1000Mb/s data rate. This register only increments if transmits are enabled and the device is in half-duplex mode.
Rx_Dropped	See Rx_Errors.
Tx_Dropped	Not defined.
Collisions	Total number of collisions experienced by the transmitter. Valid in half-duplex mode.
Rx_Length_Errors	Transmission length error.
Rx_Over_Errors	Not defined.
Rx_CRC_Errors	Frame CRC error.
Rx_Frame_Errors	Same as Rx_Align_Errors. This error is only valid in 10/100M mode.
Rx_FIFO_Errors	Same as Rx_Missed_Errors - a missed packet count.
Tx_Aborted_Errors	See Tx_Errors.
Tx_Carrier_Errors	The PHY should assert the internal carrier sense signal during every transmission. Failure to do so may indicate that the link has failed or the PHY has an incorrect link configuration. This register only increments if transmits are enabled. This register is not valid in internal SerDes 1 mode (TBI mode for the 82544GC/EI) and is only valid when the Ethernet controller is operating at full duplex.
Tx_FIFO_Errors	Not defined.
Tx_Heartbeat_Errors	Not defined.
Tx_Window_Errors	See LATECOL.
Tx_Single_Collision_Frames	Counts the number of times that a successfully transmitted packet encountered a single collision. The value only increments if transmits are enabled and the Ethernet controller is in half-duplex mode.
Tx_Multiple_Collision_Frames	A Multiple Collision Count which counts the number of times that a transmit encountered more than one collision but less than 16. The value only increments if transmits are enabled and the Ethernet controller is in half-duplex mode.

**Table 106:**Possible hardware errors and meanings

Field	Definition
Tx_Deferred	Counts defer events. A defer event occurs when the transmitter cannot immediately send a packet due to the medium being busy because another device is transmitting, the IPG timer has not expired, half-duplex deferral events are occurring, XOFF frames are being received, or the link is not up. This register only increments if transmits are enabled. This counter does not increment for streaming transmits that are deferred due to TX IPG.
Rx_Frame_Too_Longs	The Rx frame is over size.
Rx_Frame_Too_Shots	The Rx frame is too short.
Rx_Align_Errors	This error is only valid in 10/100M mode.
Symbol Error Count	Counts the number of symbol errors between reads - SYMERRS. The count increases for every bad symbol received, whether or not a packet is currently being received and whether or not the link is up. This register only increments in internal SerDes mode.

## Traffic trace

Traffic tracing allows a specific packet stream to be followed. This is useful to confirm packets are taking the route you expected on your network.

View the characteristics of a traffic session through specific security policies using:

```
diag sys session
```

Trace per-packet operations for flow tracing using:

```
diag debug flow
```

Trace per-Ethernet frame using:

```
diag sniffer packet
```

## Session table

A session is a communication channel between two devices or applications across the network. Sessions enable FortiOS to inspect and act on a sequential group of packets in a session all together instead of inspecting each packet individually. Each of these sessions has an entry in the session table that includes important information about the session.

### Use as a tool

Session tables are useful troubleshooting tools because they allow you to verify connections that you expect to see open. For example, if you have a web browser open to browse the Fortinet website, you would expect a session entry from your computer, on port 80, to the IP for the Fortinet website. Another troubleshooting method is if there are too many sessions for FortiOS to process, you can examine the session table for evidence why this is happening.

The FortiGate session table can be viewed from either the CLI or the web-based manager. The most useful troubleshooting data comes from the CLI. The session table in web-based manager also provides some useful summary information, particularly the current policy number that the session is using.

## Web-based manager session information

In the web-based manager there are actually two places to view session information — the policy session monitor, and the dashboard Top Sources, Top Destinations and Top Applications

### Top Sessions Dashboard

Top Sources Dashboard shows Top Sessions by source Address, Top Destinations shows Top sessions by Destination address, and Top Applications shows Top Sessions by applications. If there are not enough entries in the session table, try browsing to a different web site and re-examine the table. The *Policy ID* shows which security policy matches the session. The sessions that do not have a *Policy ID* entry originate from the FortiGate device

### Session monitor

The session monitor is the session table. It lists the protocol used, source and destination addresses, source and destination ports, what policy ID was matched (if any), how long until the session expires, and how long it has been established.

If there is no policy ID listed in the session entry, the traffic originated from the FortiGate unit. Otherwise all sessions must match a security policy to pass through the FortiGate unit. You can specify a filter to show Forward Traffic only. To do this, click on the Edit icon (it looks like a pencil)

As there are potentially many sessions active at one time, there are different methods you can use to filter unimportant sessions out of your search. The easiest filter is to display only IPv4 or IPv6 sessions. By default both are displayed.

#	Src	Src Port	Dst	Dst Port	Policy ID	Expiry (sec)	Duration (sec)	
1	192.168.1.200:53659	53659	157.55.56.147	40004	<a href="#">1</a>	26	153	
2	192.168.1.200:53659	53659	157.55.56.151	40020	<a href="#">1</a>	124	55	
3	192.168.1.200:53659	53659	157.56.52.38	40045	<a href="#">1</a>	88	91	
4	192.168.1.200:53659	53659	75.158.90.51	56715	<a href="#">1</a>	83	150	
5	192.168.1.200:61730	61730	70.78.76.207	49149	<a href="#">1</a>	3,584	28	
6	192.168.1.200:53659	53659	157.55.130.150	40020	<a href="#">1</a>	11	168	
7	192.168.1.200:61638	61638	111.253.246.196	45660	<a href="#">1</a>	3,594	398	
8	192.168.1.200:53659	53659	157.55.235.148	40004	<a href="#">1</a>	66	113	
9	192.168.1.200:53659	53659	157.55.235.160	40044	<a href="#">1</a>	128	52	
10	192.168.1.200:53875	53875	216.2.48.143	8888	<a href="#">1</a>	56	123	
11	192.168.1.200:53659	53659	65.55.223.27	40012	<a href="#">1</a>	105	74	
12	192.168.1.200:53659	53659	213.199.179.141	40021	<a href="#">1</a>	127	52	
13	192.168.1.200:53659	53659	213.199.179.145	40037	<a href="#">1</a>	10	169	
14	192.168.1.200:53659	53659	64.4.23.156	40032	<a href="#">1</a>	54	125	
15	192.168.1.200:53659	53659	157.55.235.161	40013	<a href="#">1</a>	88	91	
16	192.168.1.200:53659	53659	65.55.223.13	40021	<a href="#">1</a>	124	55	
17	192.168.1.200:53659	53659	65.55.223.38	40029	<a href="#">1</a>	103	76	
18	192.168.1.200:53876	53876	208.91.112.195	8888	<a href="#">1</a>	56	123	
19	192.168.1.200:53876	53876	208.91.112.197	8888	<a href="#">1</a>	56	123	

### How to find which security policy a specific connection is using

Every program and device on your network must have a communication channel, or session, open to pass information. The FortiGate unit manages these sessions with its many features from traffic shaping, to antivirus scanning, and even blocking known bad web sites. Each session has an entry in the session table. In the web, you can use the Session Monitor or Top Session Dashboard to view session information.

You may want to find information for a specific session, say a secure web browser session, for troubleshooting. For example if that web browser session is not working properly, you can check the session table to ensure the session is still active, and that it is going to the proper

address. It can also tell you the security policy number it matches, so you can check what is happening in that policy.

### 1. Know your connection information.

You need to be able to identify the session you want. For this you need the source IP address (usually your computer), the destination IP address if you have it, and the port number which is determined by the program being used. Some common ports are:

- port 80 (HTTP for web browsing),
- port 22 (SSH used for secure login and file transfers)
- port 23 (telnet for a text connection)
- port 443 (HTTPS for secure web browsing)

### 2. Find your session and policy ID.

Follow *System > Dashboard > Top Sources* to the session table monitor. Find your session by finding your source IP address, destination IP address if you have it, and port number. The policy ID is listed after the destination information. If the list of sessions is very long, you can filter the list to make it easier to find your session.

### 3. When there are many sessions, use a filter to help you find your session.

If there are multiple pages of sessions it is difficult to find a single session. To help you in your search you can use a filter to block out sessions that you don't want. Select the filter icon next to Src Address. In the window that pops up, enter your source IP address and select Apply. Now only sessions that originate from your IP address will be displayed in the session table. If the list is still too long, you can do the same for the Src port. That will make it easy to find your session and the security policy ID. When you are finished remember to clear the filters.

## CLI session information

The session table output from the CLI (`diag sys session list`) is very verbose. Even on a system with a small amount of traffic, displaying the session table will generate a large amount of output. For this reason, filters are used to display only the session data of interest.

You can filter a column in the web-based manager by clicking the funnel icon on the column heading or from the CLI by creating a filter.

An entry is placed in the session table for each traffic session passing through a security policy. The following command will list the information for a session in the table:

```
diag sys session list
```

### Sample Output:

```
FGT# diag sys session list
session info: proto=6 proto_state=05 expire=89 timeout=3600
 flags=00000000 av_idx=0 use=3
bandwidth=204800/sec guaranteed_bandwidth=102400/sec
 traffic=332/sec prio=0 logtype=session ha_id=0 hakey=4450
tunnel=/
state=log shape may_dirty
statistic(bytes/packets/err): org=3408/38/0 reply=3888/31/0 tuples=2
origin->sink: org pre->post, reply pre->post oif=3/5
 gwy=192.168.11.254/10.0.5.100
hook=post dir=org act=snat
 10.0.5.100:1251->192.168.11.254:22(192.168.11.105:1251)
hook=pre dir=reply act=dnat
 192.168.11.254:22->192.168.11.105:1251(10.0.5.100:1251)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 domain_info=0 auth_info=0 ftgd_info=0 ids=0x0 vd=0
 serial=00007c33 tos=ff/ff
```

Since output can be verbose, the filter option allows specific information to be displayed, for example:

```
diag sys session filter <option>
```

The <option> values available include the following:

clear	Clear session filter.
dintf	Destination interface.
dport	Destination port.
dst	Destination IP address.
duration	duration
expire	expire
negate	Inverse filter.
nport	NAT'd source port
nsrc	NAT'd source ip address
policy	Policy ID.
proto	Protocol number.
proto-state	Protocol state.
sintf	Source interface.
sport	Source port.
src	Source IP address.
vd	Index of virtual domain. -1 matches all.

Even though UDP is a sessionless protocol, the FortiGate unit still keeps track of the following two different states:

- UDP reply not seen with a value of 0
- UDP reply seen with a value of 1

The following illustrates FW session states from the session table:

**Table 107:**

State	Meaning
log	Session is being logged.
local	Session is originated from or destined for local stack.
ext	Session is created by a firewall session helper.
may_dirty	Session is created by a policy. For example, the session for <code>ftp control channel</code> will have this state but <code>ftp data channel</code> will not. This is also seen when NAT is enabled.
ndr	Session will be checked by IPS signature.
nds	Session will be checked by IPS anomaly.
br	Session is being bridged (TP) mode.

## Firewall session setup rate

The number of sessions that can be established in a set period of time is useful information. A session is an end-to-end TCP/IP connection for communication with a limited lifespan. If you record the setup rate during normal operation, when you experience problems you have that setup rate with the current number to see if its very different. While this will not solve your problems, it can be a useful step to help you define your problem.

A reduced firewall session setup rate could be the result of a number of things from a lack of system resources on the FortiGate unit, to reaching the limit of your session count for your VDOM.

### To view your session setup rate - web-based manager

1. Got to *System > Dashboard*.
2. Maximize *Top Sources*
3. Read the *New Sessions per Second* value displayed at the bottom.

If the *Top Sessions* widget is not visible on your dashboard, go to the + *Widget* button at the top of the window. When a window pops up, select *Top Sessions* for it to be added to the dashboard.



## To view your session setup rate method 1- CLI

```
FGT# get sys performance status
CPU states: 0% user 0% system 0% nice 100% idle
Memory states: 10% used
Average network usage: 0 kbps in 1 minute, 0 kbps in 10 minutes,
13 kbps in 30 minutes
Average sessions: 31 sessions in 1 minute, 30 sessions in 10
minutes, 31 sessions in 30 minutes
Average session setup rate: 0.5 sessions per second in last 1
minute, 0 sessions per second in last 10 minutes, 0 sessions per
second in last 30 minutes
Virus caught: 0 total in 1 minute
IPS attacks blocked: 0 total in 1 minute
Uptime: 44 days, 18 hours, 42 minutes
```

The information you are looking for is the Average sessions section, highlighted in the above output. In this example you can see there were 31 sessions in 1 minute, or an average of 0.5 sessions per second. The values for 10 minutes and 30 minutes allow you to take a longer average for a more reliable value if your FortiGate unit is working at maximum capacity. The smallest FortiGate unit can have 1 000 sessions established per second across the unit.

Remember that session setup rate is a global command. If you have multiple VDOMs configured with many sessions in each one, the session setup rate per VDOM will be slower than if there were no VDOMs configured.

## Finding object dependencies

An administrator may not be permitted to delete a configuration object if there are other configuration objects that depend on it. This command identifies other objects which depend on or *make reference to* the configuration object in question. If an error is displayed that an object is in use and cannot be deleted, this command can help identify the source of the problem.

Another use is if you have a virtual interface with objects that depend on it, you need to find and remove those dependencies before you delete that interface.

### CLI method

When running multiple VDOMs, this command is run in the Global configuration only and it searches for the named object both in the Global and VDOM configuration most recently used:

```
diag sys checkused <path.object.mkey>
```

For example, to verify which objects are referred to in a security policy with an ID of 1, enter the command as follows:

```
diag sys checkused firewall.policy.policyid 1
```

To check what is referred to by interface `port1`, enter the following command:

```
diag sys checkused system.interface.name port1
```

To show all the dependencies for an interface, enter the command as follows:

```
diag sys checkused system.interface.name <interface name>
```

### Sample Output:

```
entry used by table firewall.address:name '10.98.23.23_host'
entry used by table firewall.address:name 'NAS'
entry used by table firewall.address:name 'all'
entry used by table firewall.address:name 'fortinet.com'
entry used by table firewall.vip:name 'TORRENT_10.0.0.70:6883'
entry used by table firewall.policy:policyid '21'
entry used by table firewall.policy:policyid '14'
entry used by table firewall.policy:policyid '19'
```

In this example, the interface has dependent objects, including four address objects, one VIP, and three security policies.

### Web-based manager method

In the web-based manager, the object dependencies for an interface can be easily checked and removed.

#### To remove interface object dependencies - web-based manager

1. Go to *System > Interfaces*.  
The number in the *Ref.* column is the number of objects that refer to this interface.
2. Select the number in the *Ref.* column for the desired interface.  
A Window listing the dependencies will appear.
3. Use these detailed entries to locate and remove object references to this interface.  
The trash can icon will change from gray when all object dependencies have been removed.
4. Remove the interface by selecting the check box for the interface, and select *Delete*.

### Flow trace

To trace the flow of packets through the FortiGate unit, use the following command:

```
diag debug flow trace start
```

Follow packet flow by setting a flow filter using this command:

```
diag debug flow filter <option>
```

Filtering options include the following:

```
addr IP address
clear clear filter
daddr destination IP address
dport destination port
negate inverse filter
port port
proto protocol number
saddr source IP address
sport source port
vd index of virtual domain, -1 matches all
```

Enable the output to be displayed to the CLI console using the following command:

```
diag debug flow show console
```



diag debug flow output is recorded as event log messages and are sent to a FortiAnalyzer unit if connected. **Do not let this command run longer than necessary since it generates significant amounts of data.**

---

Start flow monitoring with a specific number of packets using this command:

```
diag debug flow trace start <N>
```

Stop flow tracing at any time using:

```
diag debug flow trace stop
```

The following is an example of the flow trace for the device at the following IP address:

```
203.160.224.97
```

```
diag debug enable
diag debug flow filter addr 203.160.224.97
diag debug flow show console enable
diag debug flow show function-name enable
diag debug flow trace start 100
```

### Flow trace output example - HTTP

Connect to the web site at the following address to observe the debug flow trace. The display may vary slightly :

```
http://www.fortinet.com
```

Comment: SYN packet received:

```
id=20085 trace_id=209 func=resolve_ip_tuple_fast
line=2700 msg="\vd-root received a packet(proto=6,
192.168.3.221:1487->203.160.224.97:80) from port5."
```

SYN sent and a new session is allocated:

```
id=20085 trace_id=209 func=resolve_ip_tuple line=2799
msg="allocate a new session-00000e90"
```

Lookup for next-hop gateway address:

```
id=20085 trace_id=209 func=vf_ip4_route_input line=1543
msg="find a route: gw-192.168.11.254 via port6"
```

Source NAT, lookup next available port:

```
id=20085 trace_id=209 func=get_new_addr line=1219
msg="find SNAT: IP-192.168.11.59, port-31925"
direction"
```

Matched security policy. Check to see which policy this session matches:

```
id=20085 trace_id=209 func=fw_forward_handler line=317
msg="Allowed by Policy-3: SNAT"
```

Apply source NAT:

```
id=20085 trace_id=209 func=__ip_session_run_tuple
line=1502 msg="SNAT 192.168.3.221->192.168.11.59:31925"
```

SYN ACK received:

```
id=20085 trace_id=210 func=resolve_ip_tuple_fast line=2700
msg="vd-root received a packet(proto=6, 203.160.224.97:80-
>192.168.11.59:31925) from port6."
```

Found existing session ID. Identified as the reply direction:

```
id=20085 trace_id=210 func=resolve_ip_tuple_fast line=2727
msg="Find an existing session, id-00000e90, reply
direction"
```

Apply destination NAT to inverse source NAT action:

```
id=20085 trace_id=210 func=__ip_session_run_tuple
line=1516 msg="DNAT 192.168.11.59:31925-
>192.168.3.221:1487"
```

Lookup for next-hop gateway address for reply traffic:

```
id=20085 trace_id=210 func=vf_ip4_route_input line=1543
msg="find a route: gw-192.168.3.221 via port5"
```

ACK received:

```
id=20085 trace_id=211 func=resolve_ip_tuple_fast line=2700
msg="vd-root received a packet(proto=6,
192.168.3.221:1487->203.160.224.97:80) from port5."
```

Match existing session in the original direction:

```
id=20085 trace_id=211 func=resolve_ip_tuple_fast line=2727
msg="Find an existing session, id-00000e90, original
direction"
```

Apply source NAT:

```
id=20085 trace_id=211 func=__ip_session_run_tuple
line=1502 msg="SNAT 192.168.3.221->192.168.11.59:31925"
```

Receive data from client:

```
id=20085 trace_id=212 func=resolve_ip_tuple_fast
line=2700 msg="vd-root received a packet (proto=6,
192.168.3.221:1487->203.160.224.97:80) from port5."
```

Match existing session in the original direction:

```
id=20085 trace_id=212 func=resolve_ip_tuple_fast
line=2727 msg="Find an existing session, id-00000e90,
original direction"
```

Apply source NAT:

```
id=20085 trace_id=212 func=__ip_session_run_tuple
line=1502 msg="SNAT 192.168.3.221->192.168.11.59:31925"
```

Receive data from server:

```
id=20085 trace_id=213 func=resolve_ip_tuple_fast
line=2700 msg="vd-root received a packet (proto=6,
203.160.224.97:80->192.168.11.59:31925) from port6."
```

Match existing session in reply direction:

```
id=20085 trace_id=213 func=resolve_ip_tuple_fast
line=2727 msg="Find an existing session, id-00000e90,
reply direction"
```

Apply destination NAT to inverse source NAT action:

```
id=20085 trace_id=213 func=__ip_session_run_tuple
line=1516 msg="DNAT 192.168.11.59:31925-
>192.168.3.221:1487"
```

### Flow trace output example - IPsec (policy-based)

```
id=20085 trace_id=1 msg="vd-root received a packet (proto=1,
10.72.55.240:1->10.71.55.10:8) from internal."
id=20085 trace_id=1 msg="allocate a new session-00001cd3"
id=20085 trace_id=1 msg="find a route: gw-66.236.56.230 via wan1"
id=20085 trace_id=1 msg="Allowed by Policy-2: encrypt"
id=20085 trace_id=1 msg="enter IPsec tunnel-RemotePhase1"
id=20085 trace_id=1 msg="encrypted, and send to 15.215.225.22 with
source 66.236.56.226"
id=20085 trace_id=1 msg="send to 66.236.56.230 via intf-wan1"
id=20085 trace_id=2 msg="vd-root received a packet (proto=1,
10.72.55.240:1-1071.55.10:8) from internal."
id=20085 trace_id=2 msg="Find an existing session, id-00001cd3,
original direction"
id=20085 trace_id=2 msg="enter IPsec ="encrypted, and send to
15.215.225.22 with source 66.236.56.226" tunnel-RemotePhase1"
id=20085 trace_id=2 msgid=20085 trace_id=2 msg="send to 66.236.56.230
via intf-wan1"
```

## Packet sniffing and packet capture

FortiOS devices can sniff packets using commands in the CLI or capture packets using the web-based manager. The differences between the two methods are not large.

Packet sniffing in the CLI is well suited for spot checking traffic from the CLI, but if you have complex filters to enter it can be a lot of work to enter them each time. You can also save the sniffing output; however, you must log to a file and then analyze the file later by hand.

Packet capture in the web-based manager makes it easy to set up multiple filters at once and just run one or two as you need them. You also have controls to start and stop capturing as you wish. Packet capture output is downloaded to your local computer as a \*.pcap file which requires a third party application to read the file, such as Wireshark. This method is useful to send Fortinet support information to help resolve an issue.

Features	Packet sniffing	Packet capture
Command location	CLI	web-based manager
Third party software required	puTTY to log plaintext output	Wireshark to read *.pcap files
Read output in plain text file	yes	no
Read output as *.pcap file using Wireshark	no	yes
Easily configure single quick and simple filter	yes	no
Record packet interface	yes	no
Configure complex sniffer filters on multiple interface	no	yes
sniff IPv6	hard	easy
sniff non-IP packets	no	yes
Filter packets by protocol and/or port	easy	easy
Filter packets by source and/or destination address	easy	easy

## Packet sniffing

Before you start sniffing packets on the CLI, you should be prepared to capture the output to a file — there can be huge amounts of data that you will not be able to see without saving it to a file. One method is to use a terminal program like puTTY to connect to the FortiGate unit's CLI. Then once the packet sniffing count is reached you can end the session and analyze the output in the file.

Details within packets passing through particular interfaces can be displayed using the packet sniffer with the following command:

```
diag sniffer packet <interface> <filter> <verbose> <count> <tsformat>
```

The <interface> value is required, with the rest being optional. If not included the default values will be "none".

For example the simplest valid sniffer command would be:

```
diag sniffer packet any
```

The <interface> value can be any physical or virtual interface name. Use *any* to sniff packets on all interfaces.

The `<filter>` value limits the display of packets using filters, including Berkeley Packet Filtering (BPF) syntax. The `<filter>` value must be enclosed in quotes.

```
'[[src|dst] host <host_name_or_IP1>] [[src|dst] host
 <host_name_or_IP2>] [[arp|ip|ip6|gre|esp|udp|tcp] [port_no]]
 [[arp|ip|ip6|gre|esp|udp|tcp] [port_no]]'
```

If a second host is specified in the filter, only the traffic between the two hosts will be displayed. Optionally, you can use logical OR to match only one of the hosts, or match one of multiple protocols or ports. When defining a port, there are up to two parts — protocol and port number.

For example, to display UDP 1812 traffic or TCP 8080 traffic, use the following:

```
'udp port 1812 or tcp port 8080'
```

To display all IP traffic that has a source of 192.168.1.2 and a destination of 192.168.2.3:

```
'ip src host 192.168.1.2 and dst host 192.168.2.3'
```

The `<verbose>` option allows different levels of information to be displayed. The verbose levels include:

- 1 Print header of packets
- 2 Print header and data from the IP header of the packets
- 3 Print header and data from the Ethernet header of the packets
- 4 Print header of packets with interface name
- 5 Print header and data from ip of packets with interface name
- 6 Print header and data from ethernet of packets with interface name

The `<count>` value indicates the number of packets to sniff before stopping. If this variable is not included, or is set to zero, the sniffer will run until you manually halt it with Ctrl-C.

The `<tsformat>` value define the format of timestamp. It can be:

a: absolute UTC time, yyyy-mm-dd hh:mm:ss.ms

l: absolute LOCAL time, yyyy-mm-dd hh:mm:ss.ms

otherwise: relative to the start of sniffing, ss.ms

## Packet capture

FortiOS 5 includes packet capture to the web-based manager. To configure packet capture filters, go to *System > Network > Packet Capture*.

When you add a packet capture filter, enter the following information and select *OK*.

<b>Interface</b>	Select the interface to sniff from the dropdown menu. You must select one interface. You cannot change the interface without deleting the filter and creating a new one, unlike the other fields.
<b>Max Packets to Capture</b>	Enter the number of packets to capture before the filter stops. This number cannot be zero. You can halt the capturing before this number is reached.
<b>Enable Filters</b>	Select this option to specify your filter fields
<b>Host(s)</b>	Enter one or more hosts IP address Separate multiple hosts with commas. Enter a range using a dash without spaces, for example 172.16.1.5-172.16.1.15 or enter a subnet.
<b>Port(s)</b>	Enter one or more ports to capture on the selected interface. Separate multiple ports with commas. Enter a range using a dash without spaces, for example 88-90
<b>VLAN(s)</b>	Enter one or more vlans (if there is any). Separate multiple vlans with commas.
<b>Protocol</b>	Enter one or more protocol. Separate multiple protocol with commas. Enter a range using a dash without spaces, for example 1-6, 17, 21-25
<b>Include IPv6 packets</b>	Select this option if you are troubleshooting IPv6 networking, or if your network uses IPv6. Otherwise, leave it disabled.
<b>Capture Non-IP packets</b>	The protocols available in the list are all IP based except for ICMP (ping). To capture non-IP based packets select this feature. Some examples of non-IP packets include IPsec, IGMP, ARP, and as mentioned ICMP.

If you select a filter and go back to edit it, you have the added option of starting and stopping packet capture in the edit window, or downloading the captured packets. You can also see the filter status and the number of packets captured.

You can also select the filter and select *Start* to start capturing packets. While the filter is running, you will see the number of captured packets increasing until it reaches the max packet count or you select *Stop*. While the filter is running you cannot download the output file.

When the packet capture is complete, you can select *Download* to send the packet capture filter captured packets to your local computer as a \*.pcap file. To read this file format, you will need to use Wireshark or a similar third party application. Using this tool you will have extensive analytics available to you and the full contents of the packets that were captured.



## FA2 and NP2 based interfaces

Many Fortinet products contain network processors. Some of these products contain FortiAccel (FA2) network processors while others contain NP2 network processors. Network processor features, and therefore offloading requirements, vary by network processor model.

When using the FA2- and NP2-based interfaces, only the initial session setup will be seen through the `diag debug flow` command. If the session is correctly programmed into the ASIC (fastpath), the debug flow command will no longer see the packets arriving at the CPU. If the NP2 functionality is disabled, the CPU will see all the packets, however, this should only be used for troubleshooting purposes.

First, obtain the NP2 and port numbers with the following command:

```
diag npu np2 list
```

### Sample output:

```
ID PORTS
-- -----
0 port1
0 port2
0 port3
0 port4
ID PORTS
-- -----
1 port5
1 port6
1 port7
1 port8
ID PORTS
-- -----
2 port9
2 port10
2 port11
2 port12
ID PORTS
-- -----
3 port13
3 port14
3 port15
3 port16
```

Run the following commands:

```
diag npu np2 fastpafth disable <dev_id>
```

(where `dev_id` is the NP2 number)

Then, run this command:

```
diag npu np2 fastpath-sniffer enable port1
```

### Sample output:

```
NP2 Fast Path Sniffer on port1 enabled
```

This will cause all traffic on *port1* of NP2 to be sent to the CPU meaning a standard sniffer trace can be taken and other diag commands should work if it was a standard CPU driven port.

These commands are only for the newer NP2 interfaces. FA2 interfaces are more limited as the sniffer will only capture the initial packets before the session is offloaded into HW (FA2). The same holds true for the `diag debug flow` command as only the session setup will be shown, however, this is usually enough for this command to be useful.

## Debug command

Debug output provides continuous, real-time event information. Debugging output continues until it is explicitly stopped or until the unit is rebooted. Debugging output can affect system performance and will be continually generated even though output might not be displayed in the CLI console.

Debug information displayed in the console will scroll in the console display and may prevent CLI commands from being entered, for example, the command to disable the debug display. To turn off debugging output as the display is scrolling by, press the `↑` key to recall the recent diag debug command, press backspace, and type "0", followed by `Enter`.

Debug output display is enabled using the following command:

```
diag debug enable
```

When finished examining the debug output, disable it using:

```
diag debug disable
```

Once enabled, indicate the debug information that is required using this command:

```
diag debug <option> <level>
```

Debug command options include the following:

<code>application</code>	application
<code>authd</code>	Authentication daemon.
<code>cli</code>	Debug CLI.
<code>cmdb-trace</code>	Trace CLI.
<code>config-error-log</code>	Configure error log info.
<code>console</code>	console
<code>crashlog</code>	Crash log info.
<code>disable</code>	Disable debug output.
<code>enable</code>	Enable debug output.
<code>flow</code>	Trace packet flow in kernel.
<code>fsso-polling</code>	FSSO active directory poll module.
<code>info</code>	Show active debug level settings.
<code>kernel</code>	kernel
<code>rating</code>	Display rating info.

report	Report for tech support.
reset	Reset all debug level to default.
rtmon	rtmon daemon
sql-log-error	SQL log database error info
urlfilter	urlfilter

The debug level can be set at the end of the command. Typical values are 2 and 3, for example:

```
diag debug application DHCPS 2
diag debug application spamfilter 2
```

Fortinet support will advise which debugging level to use.

Timestamps can be enabled to the debug output using the following command:

```
diag debug console timestamp enable
```

### Debug output example

This example shows the IKE negotiation for a secure logging connection from a FortiGate unit to a FortiAnalyzer system.

```
diag debug reset
diag vpn ike log-filter src-addr4 192.168.11.2
diag debug enable
```

### Sample Output:

```
FGh_FtiLog1: IPsec SA connect 0 192.168.11.2->192.168.10.201:500,
natt_mode=0 rekey=0 phase2=FGh_FtiLog1
FGh_FtiLog1: using existing connection, dpd_fail=0
FGh_FtiLog1: found phase2 FGh_FtiLog1
FGh_FtiLog1: IPsec SA connect 0 192.168.11.2 -> 192.168.10.201:500
negotiating
FGh_FtiLog1: overriding selector 225.30.5.8 with 192.168.11.2
FGh_FtiLog1: initiator quick-mode set pfs=1536...
FGh_FtiLog1: try to negotiate with 1800 life seconds.
FGh_FtiLog1: initiate an SA with selectors:
192.168.11.2/0.0.0.0->192.168.10.201, ports=0/0, protocol=0/0
Send IKE Packet(quick_outI1):192.168.11.2:500(if0) ->
192.168.10.201:500, len=348
Initiator: sent 192.168.10.201 quick mode message #1 (OK)
FGh_FtiLog1: set retransmit: st=168, timeout=6.
```

In this example:

```
192.168.11.2->192.168.10.201:500 Source and Destination gateway IP
address
```

```
dpd_fail=0
```

Found existing Phase 1

```
pfs=1536...
```

Create new Phase 2 tunnel

## The execute tac report command

`exec tac report` is an execute command that runs an exhaustive series of diagnostic commands. It runs commands that are only needed if you are using certain features like HA, VPN tunnels, or a modem. The report takes a few minutes to complete due to the amount of output generated. If you have your CLI output logged to a file, you can run this command to familiarize yourself with the CLI commands involved.

When you call Fortinet Customer Support, you will be asked to provide information about your unit and its current state using the output from this CLI command.

## Other commands

### ARP table

To view the ARP cache, use the following command:

```
get sys arp
```

To view the ARP cache in the system, use this command:

```
diag ip arp list
```

### Sample output:

```
index=14 ifname=internal 224.0.0.5 01:00:5e:00:00:05 state=00000040
 use=72203 confirm=78203 update=72203 ref=1
index=13 ifname=dmz 192.168.3.100 state=00000020 use=1843
 confirm=650179 update=644179 ref=2 ? VIP
index=13 ifname=dmz 192.168.3.109 02:09:0f:78:69:ff state=00000004
 use=71743 confirm=75743 update=75743 ref=1
index=14 ifname=internal 192.168.11.56 00:1c:23:10:f8:20
 state=00000004 use=10532 confirm=10532 update=12658 ref=4
```

To remove the ARP cache, use this command:

```
execute clear system arp table
```

To remove a single ARP entry, use:

```
diag ip arp delete <interface name> <IP address>
```

To remove all entries associated with a particular interface, use this command:

```
diag ip arp flush <interface name>
```

To add static ARP entries, use the following command:

```
config system arp-table
```

## Time and date settings

Check time and date settings for log message timestamp synchronization (the Fortinet support group may request this) and for certificates that have a time requirement to check for validity. Use the following commands:

```
execute time
current time is: 12:40:48
last ntp sync:Thu Mar 16 12:00:21 2006
execute date
current date is: 2006-03-16
```

To force synchronization with an NTP server, toggle the following command:

```
set ntpsync enable/disable
```

If all devices have the same time, it helps to correlate log entries from different devices.

## IP address

There may be times when you want to verify the IP addresses assigned to the FortiGate unit interfaces are what you expect them to be. This is easily accomplished from the CLI using the following command.

```
diag ip address list
```

The output from this command lists the IP address and mask if available, the index of the interface (a sort of ID number) and the devname is the name of the interface. While physical interface names are set, virtual interface names can vary. Listing all the virtual interface names is a good use of this command. For `vsys_ha` and `vsys_fgfm`, the IP addresses are the local host — these are internally used virtual interfaces.

```
diag ip address list
IP=10.31.101.100->10.31.101.100/255.255.255.0 index=3 devname=internal
IP=172.20.120.122->172.20.120.122/255.255.255.0 index=5 devname=wan1
IP=127.0.0.1->127.0.0.1/255.0.0.0 index=8 devname=root
IP=127.0.0.1->127.0.0.1/255.0.0.0 index=11 devname=vsys_ha
IP=127.0.0.1->127.0.0.1/255.0.0.0 index=13 devname=vsys_fgfm
```

Other related commands include flushing the IP addresses (`diag ip address flush`), which will force a reload of the IP addresses. This can be useful if you think an IP address is wrong and don't want to reboot the unit. You can add or delete a single IP address (`diag ip address add <ipv4_addr>` or `diag ip address delete <ipv4_addr>`).

## FortiOS ports

In the TCP and UDP stacks, there are 65 535 ports available for applications to use when communicating with each other. Many of these ports are commonly known to be associated with specific applications or protocols. These known ports can be useful when troubleshooting your network.

Use the following ports while troubleshooting the FortiGate device:

**Table 108:**

<b>Port(s)</b>	<b>Functionality</b>
UDP 53	DNS lookup, RBL lookup
UDP 53 or UDP 8888	FortiGuard Antispam or Web Filtering rating lookup
UDP 53 (default) or UDP 8888 and UDP 1027 or UDP 1031	FDN Server List - source and destination port numbers vary by originating or reply traffic. See the article "How do I troubleshoot performance issues when FortiGuard Web Filtering is enabled?" in the Knowledge Base.
UDP 123	NTP Synchronization
UDP 162	SNMP Traps
UDP 514	SYSLOG - All FortiOS versions can use syslog to send log messages to remote syslog servers. FortiOS v2.80 and v3.0 can also view logs stored remotely on a FortiAnalyzer unit.
TCP 22	Configuration backup to FortiManager unit or FortiGuard Analysis and Management Service.
TCP 25	SMTP alert email, encrypted virus sample auto-submit
TCP 389 or TCP 636	LDAP or PKI authentication
TCP 443	FortiGuard Antivirus or IPS update - When requesting updates from a FortiManager unit instead of directly from the FDN, this port must be reconfigured as TCP 8890.
TCP 443	FortiGuard Analysis and Management Service
TCP 514	FortiGuard Analysis and Management Service log transmission (OFTP)
TCP 541	SSL Management Tunnel to FortiGuard Analysis and Management Service (FortiOS v3.0 MR6 or later)
TCP 514	Quarantine, remote access to logs and reports on a FortiAnalyzer unit, device registration with FortiAnalyzer units (OFTP)
TCP 1812	RADIUS authentication
TCP 8000 and TCP 8002	FSSO
TCP 10151	FortiGuard Analysis and Management Service contract validation

## FortiAnalyzer/FortiManager ports

If you have a FortiAnalyzer unit or FortiManager unit on your network you may need to use the following ports for troubleshooting network traffic.

**Table 109:**

Functionality	Port(s)
DNS lookup	UDP 53
NTP synchronization	UDP 123
Windows share	UDP 137-138
SNMP traps	UDP 162
Syslog, log forwarding	UDP 514
Log and report upload	TCP 21 or TCP 22
SMTP alert email	TCP 25
User name LDAP queries for reports	TCP 389 or TCP 636
RVS update	TCP 443
RADIUS authentication	TCP 1812
Log aggregation client	TCP 3000

## FortiGuard troubleshooting

The FortiGuard service provides updates to Antivirus, IPsec, Webfiltering, and more. The FortiGuard Distribution System (FDS) involves a number of servers across the world that provide updates to your FortiGate unit. Problems can occur both with connection to FDS, and its configuration on your local FortiGate unit. Some of the more common troubleshooting methods are listed here including

- [Troubleshooting process for FortiGuard updates](#)
- [FortiGuard server settings](#)
- [FortiGuard URL rating](#)

### Troubleshooting process for FortiGuard updates

The following process are the logical steps to take when troubleshooting FortiGuard update problems. This includes antivirus (AV), intrusion protection services (IPS), antispam (AS), and web filtering (WB).

1. Does the device have a valid licence that includes these services?

Each device requires a valid FortiGuard license to access updates for some or all of these services. You can verify the support contract status for your devices at the Fortinet Support website — <https://support.fortinet.com/>.

2. If the device is part of an HA cluster, do all members of the cluster have the same level of support?  
As with the previous step, you can verify the support contract status for all the devices in your HA cluster at the Fortinet Support website.
3. Have services been enabled on the device?  
To see the FortiGuard information and status for a device, in the web-based manager go to *System > Config > FortiGuard*. On that page you can verify the status of each component, and if required enable each service. If there are problems, see the FortiGuard section of the FortiOS Handbook.
4. Is the device able to communicate with FortiGuard servers?  
At *System > Config > FortiGuard* you can also attempt to update AV and IPS, or test the availability of WF and AS default and alternate ports. If there are problems, see the FortiGuard section of the FortiOS Handbook.
5. Is there proper routing to reach the FortiGuard servers?  
Ensure there is a static or dynamic route that enables your FortiGate unit to reach the FortiGuard servers. Usually a generic default route to the internet is enough, but you may need to verify this if your network is complex.
6. Are there issues with DNS?  
An easy way to test this is to attempt a traceroute from behind the FortiGate unit to an external network using the FQDN for a location. If the traceroute FQDN name does not resolve, you have general DNS problems.
7. Is there anything upstream that might be blocking FortiGuard traffic, either on the network or ISP side?  
Many firewalls block all ports by default, and often ISPs block ports that are low. There may be a firewall between the FortiGate unit and the FortiGuard servers that is blocking the traffic. FortiGuard uses port 53 by default, so if it is being blocked you need to either open a hole for it, or change the port it is using.
8. Is there an issue with source ports?  
It is possible that ports used to contact FortiGuard are being changed before reaching FortiGuard or on the return trip before reaching your FortiGate unit. A possible solution for this is to use a fixed-port at NATd firewalls to ensure the port remains the same. Packet sniffing can be used to find more information on what is happening with ports.
9. Are there security policies that include antivirus?  
If no security policies include antivirus, the antivirus database will not be updated. If antivirus is included, only the database type used will be updated.

## FortiGuard server settings

Your local FortiGate unit connects to remote FortiGuard servers get updates to FortiGuard information such as new viruses that may have been found or other new threats. This section demonstrates ways to display information about FortiGuard server information on your FortiGate unit, and how to use that information and update it to fix potential problems. This includes

- [Displaying the server list](#)
- [Sorting the server list](#)
- [Calculating weight](#)

### Displaying the server list

The `get webfilter status` command shows the list of FDS servers the FortiGate unit is using to send web filtering requests. Rating requests are only sent to the server on the top of



# Troubleshooting methodologies

Before you begin troubleshooting anything but the most minor issues, you need to prepare. Doing so will shorten the time to solve your issue. This section helps to explain how you prepare before troubleshooting, as well as creating a troubleshooting plan and contacting support.

This section contains the following topics:

- [Establish a baseline](#)
- [Define the problem](#)
- [Gathering Facts](#)
- [Create a troubleshooting plan](#)
- [Obtain any required additional equipment](#)
- [Ensure you have administrator level access to required equipment](#)
- [Contact Fortinet customer support for assistance](#)

## Establish a baseline

FortiGate units operate at all layers of the OSI model. For this reason troubleshooting problems can become complex. If you establish a normal operation parameters, or baseline, for your system before the problem occurs it will help reduce the complexity when you are troubleshooting.

Many of the guiding questions in the following sections are some form of comparing the current problem situation to normal operation on your FortiGate unit. For this reason it is a best practice that you know what your normal operating status is, and have a record of it you can refer to. This can easily be accomplished by monitoring the system performance with logs, SNMP tools, or regularly running information gathering commands and saving the output. This regular operation data will show trends, and enable you to see when changes happen and there may be a problem.



Back up your FortiOS configuration on a regular basis. This is a good practice for everyday as well as when troubleshooting. You can restore the backed up configuration when needed and save the time and effort of re-creating it from the factory default settings.

---

Some fundamental CLI commands you can use to obtain normal operating data for your system:

---

<code>get system status</code>	Displays versions of firmware and FortiGuard engines, and other system information.
<code>get system performance status</code>	Displays CPU and memory states, average network usage, average sessions and session setup rate, virus caught, IPS attacks blocked, and uptime.
<code>get hardware memory</code>	Displays informations about memory

---

<code>get system session status</code>	Displays total number of sessions
<code>get router info routing-table all</code>	Displays all the routes in the routing table including their type, source, and other useful data.
<code>get ips session</code>	Displays memory used and max available to IPS as well and counts.
<code>get webfilter ftgd-statistics</code>	Displays list of FortiGuard related counts of status, errors, and other data.
<code>diagnose firewall statistic show</code>	Displays the amount of network traffic broken down into categories such as email, VoIP, TCP, UDP, IM, Gaming, P2P, and Streaming.
<code>diag system session list</code>	Displays current detailed sessions list
<code>show system dns</code>	Displays configured DNS servers
<code>diag sys ntp status</code>	Displays informations about ntp servers

These commands are just a sample. Feel free to include any extra information gathering commands that apply to your system. For example if you have active VPN connections, record information about them using the `get vpn *` series of commands.

For an extensive snapshot of your system, run the CLI command used by TAC to gather extensive information about a system — `exec tac report`. It runs many diagnostic commands that are for specific configurations. This means no matter what features you are using, this command will record their current state. Then if you need to perform troubleshooting at a later date, you can run the same command again and compare the differences to quickly locate suspicious output you can investigate.

## Define the problem

The following questions can help determine the scope of the problem and isolate it:

- What is the problem?  
Do not assume that the problem is being experienced is the actual problem. First determine that the problem does not lie elsewhere before starting to troubleshoot the FortiGate device.
- Has it ever worked before?  
If the device never worked from the first day, you may not want to spend time troubleshooting something that could well be defective. See “Troubleshooting bootup”
- Can the problem be reproduced at will or is it intermittent?  
If the problem is intermittent, it may be dependent on system load. Also an intermittent problem can be very difficult to troubleshoot due to the difficulty reproducing the issue.
- What has changed?  
Do not assume that nothing has changed in the network. Use the FortiGate event log to see if any configuration changes were made. The change could be in the operating environment, for example, a gradual increase in load as more sites are forwarded through the firewall.  
If something has changed, see what the affect is if the change is rolled back.
- Determine the scope of the problem - after you have isolated the problem what applications, users, devices, and operating systems does it effect?

Before you can solve a problem, you need to understand it. Often this step can be the longest in this process.

Ask questions such as:

- What is not working? Be specific.
- Is there more than one thing not working?
- Is it partly working? If so, what parts are working?
- Is it a connectivity issue for the whole device, or is there an application that isn't reaching the Internet?

Be as specific as possible with your answers, even if it takes awhile to find the answers.

These questions will help you define the problem. Once the problem is defined, you can search for a solution and then create a plan on how to solve it.

## Gathering Facts

Fact gathering is an important part of defining the problem. Record the following information as it applies to the problem:

- Where did the problem occur?
- When did the problem occur and to whom?
- What components are involved?
- What is the affected application?
- Can the problem be traced using a packet sniffer?
- Can the problem be traced in the session table or using system debugging?
- Can log files be obtained that indicate a failure has occurred?

Answers to these questions will help you narrow down the problem, and what you have to check during your troubleshooting. The more things you can eliminate, the fewer things you need to check during troubleshooting. For this reason, be as specific and accurate as you can while gathering facts.

## Create a troubleshooting plan

Once you have defined the problem, and searched for a solution you can create a plan to solve that problem. Even if your search didn't find a solution to your problem you may have found some additional things to check to further define your problem.

The plan should list all the possible causes of the problem that you can think of, and how to test for each possible cause.

Your troubleshooting plan will act as a checklist so that you know what you have tried and what is left to check. This is important to have if more than one person will be doing the troubleshooting. Without a written plan, people will become easily confused and steps will be skipped. Also if you have to hand over the problem to someone else, providing them with a detailed list of what data has been gathered and what solutions have been already tried demonstrates a good level of professionalism.

Be ready to add to your plan as needed. After you are part way through, you may discover that you forgot some tests or a test you performed discovered new information. This is normal.

Also if you contact support, they will require information about your problem as well as what you have already tried to fix the problem. This should all be part of your plan.

## Providing Supporting Elements

If the Fortinet Technology Assistance Center (TAC) needs to be contacted to help you with your issue, be prepared to provide the following information:

- The firmware build version (use the `get system status` command)
- A network topology diagram
- A recent configuration file
- Optionally, a recent debug log
- Tell the support team what troubleshooting steps have already been performed and the results.



Do not provide the output from `exec tac` report unless Support requests it. The output from that command is very large and is not required in many cases.

---

For additional information about contacting Fortinet Customer Support, see [“Technical Support Organization Overview” on page 2325](#).

All of this is your troubleshooting plan.

## Obtain any required additional equipment

You may require additional networking equipment, computers, or other equipment to test your solution.

Normally network administrators have additional networking equipment available either to loan you, or a lab where you can bring the FortiGate unit to test.

If you do not have access to equipment, check for shareware applications that can perform the same task. Often there are software solutions when hardware is too expensive.

## Ensure you have administrator level access to required equipment

Before troubleshooting your FortiGate unit, you will need administrator access to the equipment. If you are a client on a FortiGate unit with virtual domains enabled, often you can troubleshoot within your own VDOM. However, you should inform your FortiGate unit's super admin that you will be doing troubleshooting.

Also, you may need access to other networking equipment such as switches, routers, and servers to help you test. If you do not normally have access to this equipment, contact your network administrator for assistance.

## Contact Fortinet customer support for assistance

You have defined your problem, researched a solution, put together a plan to find the solution, and executed that plan. At this point if the problem has not been solved, it's time to contact Fortinet Customer Support for assistance.

For more information, see [“Technical Support Organization Overview” on page 2325](#).

# Technical Support Organization Overview

This section explains how Fortinet's technical support works, as well as how you can easily create an account to get technical support for when issues arise that you cannot solve yourself.

This section contains the following topics:

- [Fortinet Global Customer Services Organization](#)
- [Creating an account](#)
- [Registering a device](#)
- [Reporting problems](#)
- [Assisting technical support](#)
- [Support priority levels](#)
- [Return material authorization process](#)

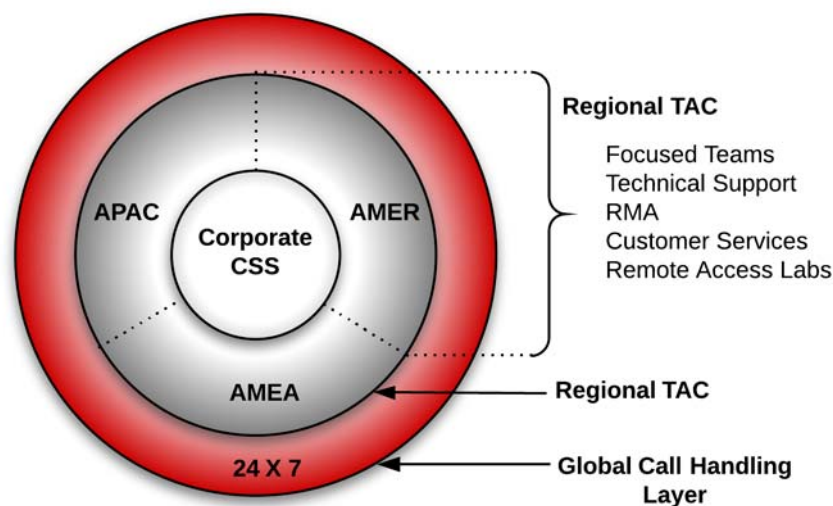
## Fortinet Global Customer Services Organization

The Fortinet Global Customer Services Organization is composed of three regional *Technical Assistance Centers* (TAC):

- The Americas (AMER)
- Europe, Middle East, and Africa (EMEA)
- Asia Pacific (APAC)

The regional TACs are contacted through a global call center. Incoming service requests are then routed to the appropriate TAC. Each regional TAC delivers technical support to the customers in its regions during its hours of operation. These TACs also combine to provide seamless, around-the-clock support for all customers.

**Figure 323:**Fortinet regions and TAC



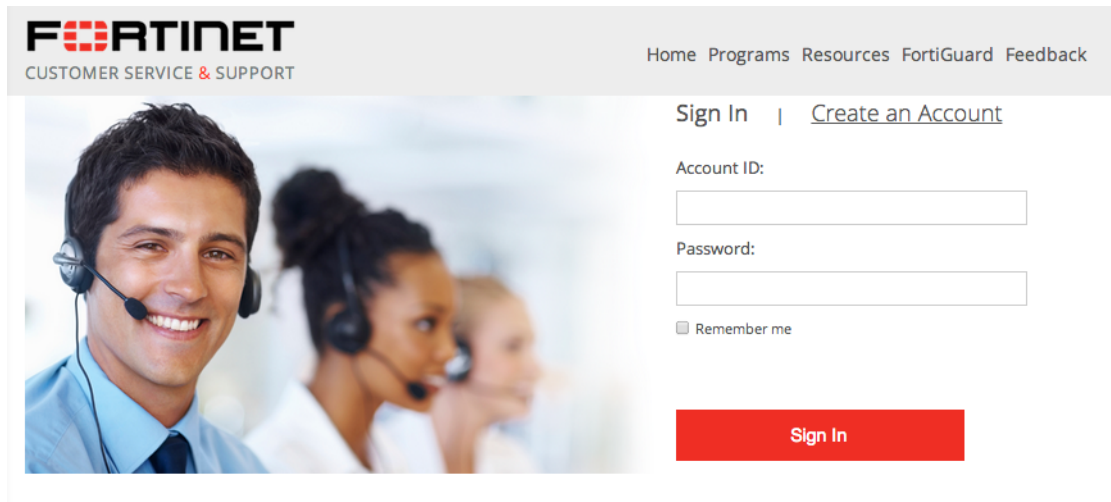
## Creating an account

To receive technical support and service updates, Fortinet products in the organization must be registered. The **Product Registration Form** on the support website will allow the registration to be completed online. Creating an account on the support website is the first step in registering products.

Go to the Fortinet support site shown below:

<https://support.fortinet.com/>

**Figure 324:**Customer service and support home page



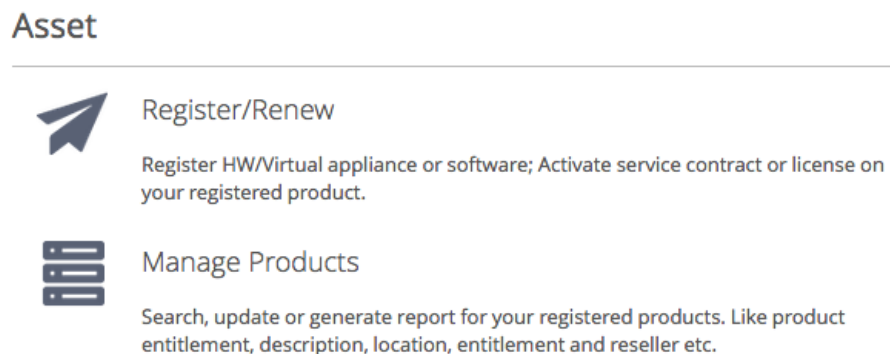
Once the support account has been created, product details can be provided by going to the *Product Register/Renew* and *Manage Product* buttons displayed on the home page. Alternately, the product registration can be completed at a later time.

## Registering a device

Complete the following steps when registering a device for support purposes:

1. Log in using the *Username* and *Password* defined when the account was created
2. Under the *Asset* section, select *Register/Renew* to go to the Registration Wizard. Alternatively, use the *Asset* menu at the top of the page.

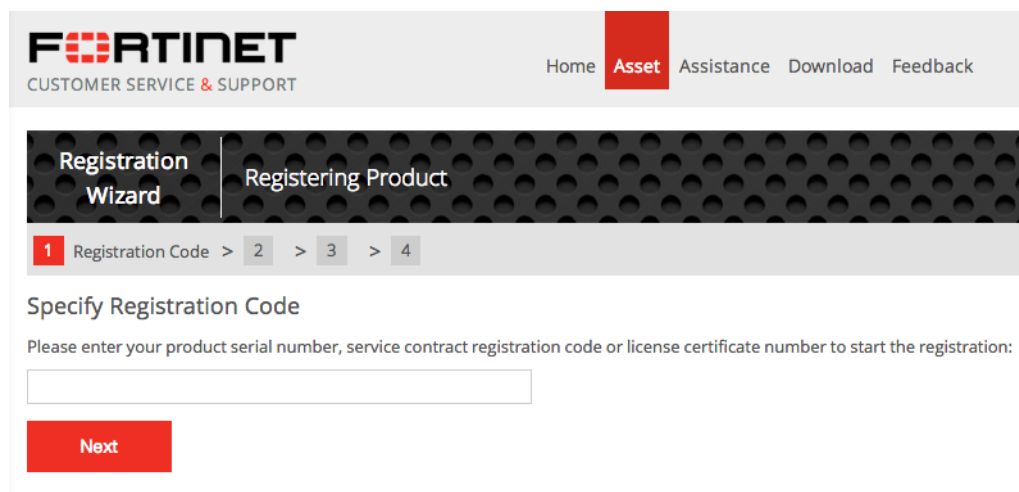
**Figure 325:**Register/Renew and Manage Products menu



3. Get a serial number from the back of the FortiGate unit or from the exterior of the FortiGate shipping box.

4. Enter the serial number, service contract registration code or license certificate number to start the product registration.

**Figure 326:**Adding a product to a support account



5. Enter your registration information.
6. Read and accept the license agreement.
7. Complete the verification process.
8. Select *Finish* to complete the registration process.

**Figure 327:**Registration wizard



## Reporting problems

Problems can be reported to a Fortinet Technical Assistance Center in the following ways:

- By logging an online ticket
- By phoning a technical support center

### Logging online tickets

Problem reporting methods differ depending on the type of customer.

#### Fortinet partners

Fortinet Partners are entitled to priority web-based technical support. This service is designed for partners who provide initial support to their customers and who need to open a support ticket with Fortinet on their behalf. We strongly encourage submission and follow up of support tickets using this service.

The support ticket can be submitted after logging into the partner website using one of the following links using FortiPartner account details:

<http://partners.fortinet.com>

This link will redirect to the general *Fortinet Partner Portal* extranet website. Click *Support > Online Support Ticket*.

<https://forticare.fortinet.com/customersupport/Login/CommonLogin.aspx>

## Fortinet customers

There are two methods to report a technical issue on the Fortinet Support website: creating a technical support ticket by product or creating any type of ticket with the Ticket Wizard for more options.

Fortinet customers should complete the following steps to create a support ticket by product:

1. Log in to the support website at the following address with the account credentials used when the account was created:  
<https://support.fortinet.com>
2. Navigate to the top menu, click *Asset* and select *Manage/View Products*.
3. In the product list, select the product that is causing the problem.
4. On the left side bar, go to the *Assistance* category, and select *Technical Request* to create a TA Ticket.
5. Complete the *Create TA Ticket* fields.
6. Click *View Products*.
7. In the *Products List*, select the product that is causing the problem.
8. Complete the *Create Support Ticket* fields.
9. Select *Finish* to complete the support ticket.

Fortinet customers who would like to submit a customer service ticket, DOA ticket, RMA ticket, or FortiGuard service ticket should use the Ticket Wizard and complete the following steps:

1. Log in to the support website at the following address with the account credentials used when the account was created:  
<https://support.fortinet.com>
2. Navigate to the top menu, click *Assistance* and select *Create a Ticket* from the drop down menu.
3. Select a ticket type and complete the remaining steps in the Ticket Wizard.
4. Select *Finish* to complete the ticket.

## Following up on online tickets

Perform the following steps to follow up on an existing issue.

Partners should log into the following web site:

<http://partners.fortinet.com>

Customers should log into the following site:

<http://support.fortinet.com>.

1. Log in with the account credentials used when the account was created.
2. Navigate to the top menu, click *Assistance*, and select *Manage Tickets*.
3. Use the search field on the View Tickets page to locate the tickets assigned to the account.
4. Select the appropriate ticket number. Closed tickets cannot be updated. A new ticket must be submitted if it concerns the same problem.
5. Add a *New Comment* or *Attachment*.



6. Click *Submit* when complete.



Every web ticket update triggers a notification to the ticket owner, or ticket queue supervisor.

---

## Telephoning a technical support center

The Fortinet Technical Assistance Centers can also be contacted by phone.

Call Fortinet Support Center at 1-408-486-7899 (international) or go to [http://www.fortinet.com/support/contact\\_support.html](http://www.fortinet.com/support/contact_support.html) and select your country from the drop-down list for local contact number.

## Assisting technical support

The more information that can be provided to Fortinet technical support, the better they can assist in resolving the issue. Every new support request should contain the following information:

- A valid contact name, phone number, and email address.
- A clear and accurate problem description.
- A detailed network diagram with complete IP address schema.
- The configuration file, software version, and build number of the Fortinet device.
- Additional log files such as *Antivirus* log, *Attack* log, *Event* log, *Debug* log or similar information to include in the ticket as an attachment. If a third-party product is involved, for example, email server, FTP server, router, or switch, please provide the information on its software revision version, configuration, and brand name.

## Support priority levels

Fortinet technical support assigns the following priority levels to support cases:

### Priority 1

This **Critical** priority is assigned to support cases in which:

- The network or system is down causing customers to experience a total loss of service.
- There are continuous or frequent instabilities affecting traffic-handling capability on a significant portion of the network.
- There is a loss of connectivity or isolation to a significant portion of the network.
- This issue has created a hazard or an emergency.

### Priority 2

This **Major** priority is assigned to support cases in which:

- The network or system event is causing intermittent impact to end customers.
- There is a loss of redundancy.
- There is a loss of routine administrative or diagnostic capability.

- There is an inability to deploy a key feature or function.
- There is a partial loss of service due to a failed hardware component.

### Priority 3

This **Medium** priority is assigned to support cases in which:

- The network event is causing only limited impact to end customers.
- Issues seen in a test or pre-production environment exist that would normally cause adverse impact to a production network.
- The customer is making time sensitive information requests.
- There is a successful workaround in place for a higher priority issue.

### Priority 4

This **Minor** priority is assigned to support cases in which:

- The customer is making information requests and asking standard questions about the configuration or functionality of equipment.

Customers must report Priority 1 and 2 issues by phone directly to the Fortinet EMEA Support Center.

For lower priority issues, you may submit an assistance request (ticket) via the web system.

The web ticket system also provides a global overview of all ongoing support requests.

## Return material authorization process

In some cases hardware issues are experienced and a replacement unit must be sent. This is referred to as a Return Material Authorization (RMA). In these cases or RMAs, the support contract must be moved to the new device. Customers can move the support contract from the failing production unit to the new device through the support web site.

### To move the support contract to a new device

1. Log in to the support web site with the credentials indicated when the account was created.
2. From *Manage Products*, locate the serial number of the defective unit from the list of devices displayed for the account. The *Product Info* for the selected device will be displayed.
3. In the left side bar under the *Assistance section*, select *RMA Transfer*.
4. Enter the *Original Serial Number* of the original device, enter the *New Serial Number*, and click *Replace* to complete the transfer.

This will transfer the support contract from the defective unit to the new unit with the serial number provided.

# Chapter 19 Virtual Domains

This FortiOS Handbook chapter contains the following sections:

[Virtual Domains](#) provides an overview of the VDOM technologies, and the basic concepts and rules for using them. We recommend that you begin with this chapter before attempting to configure VDOMs on your FortiGate unit.

[Virtual Domains in NAT/Route mode](#) provides detailed explanations and examples for configuring VDOM features in your FortiGate unit using the NAT/Route mode.

[Virtual Domains in Transparent mode](#) provides detailed explanations, as well as basic and advanced examples for configuring these features in your FortiGate unit using Transparent mode.

[Inter-VDOM routing](#) describes inter-VDOM routing concepts and scenarios, and gives examples that illustrate them.

[Troubleshooting Virtual Domains](#) provides diagnostic and troubleshooting information for some potential VDOM issues.

# Virtual Domains

Virtual domains (VDOMs) are a method of dividing a FortiGate unit into two or more virtual units that function as multiple independent units. VDOMs can provide separate firewall policies and, in NAT/Route mode, completely separate configurations for routing and VPN services for each connected network or organization.

This chapter will cover the basics of VDOMs, how they change your FortiGate unit, and how to work with VDOMs.

VDOMs let you split your physical FortiGate unit into multiple virtual units. The resulting benefits range from limiting Transparent mode ports to simplified administration, to reduced space and power requirements.

When VDOMs are disabled on any FortiGate unit, there is still one VDOM active: the root VDOM. It is always there in the background. When VDOMs are disabled, the root VDOM is not visible but it is still there.

The root VDOM must be there because the FortiGate unit needs a management VDOM for management traffic among other things. It is also why when you enable VDOMs, all your configuration is preserved in the root VDOM-because that is where you originally configured it.

This section includes:

- [Benefits of Virtual Domains](#)
- [Enabling and accessing Virtual Domains](#)
- [Configuring Virtual Domains](#)

## Benefits of Virtual Domains

VDOMs provide the following benefits:

- [Improving Transparent mode configuration](#)
- [Easier administration](#)
- [Continued security](#)
- [Savings in physical space and power](#)
- [More flexible MSSP configurations](#)

### Improving Transparent mode configuration

When VDOMs are not enabled and you put your FortiGate unit into Transparent mode, all the interfaces on your unit become broadcast interfaces. The problem with this is that there are no interfaces free to do anything else.

With multiple VDOMs you can have one of them configured in Transparent mode, and the rest in NAT/Route mode. In this configuration, you have an available transparent mode FortiGate unit you can drop into your network for troubleshooting, and you also have the standard NAT for networking.

### Easier administration

VDOMs provide separate security domains that allow separate zones, user authentication, firewall policies, routing, and VPN configurations. VDOMs separate security domains and

simplify administration of complex configurations—you do not have to manage as many settings at one time. For more information, see [“Global and per-VDOM settings” on page 2338](#).

By default, each FortiGate unit has a VDOM named root. This VDOM includes all of the unit’s physical interfaces, modem, VLAN subinterfaces, zones, firewall policies, routing settings, and VPN settings.

Also, you can optionally assign an administrator account restricted to one VDOM. If the VDOM is created to serve an organization, this feature enables the organization to manage its own configuration. For more information, see [“Administrators in Virtual Domains” on page 2357](#).

Each physical FortiGate unit requires a FortiGuard license to access security updates. VDOMs do not require any additional FortiGuard licenses, or updating — all the security updates for all the VDOMs are performed once per update at the global level. Combined this can be a potentially large money and time saving feature in your network.

Management systems such as SNMP, logging, alert email, FDN-based updates, and NTP-based time setting use addresses and routing in the management VDOM to communicate with the network. They can connect only to network resources that communicate with the management VDOM. Using a separate VDOM for management traffic enables easier management of the FortiGate unit global settings, and VDOM administrators can also manage their VDOMs more easily. For more information, see [“Changing the management virtual domain” on page 2361](#).

## Continued security

When a packet enters a VDOM, it is confined to that VDOM and is subject to any firewall policies for connections between VLAN subinterfaces or zones in that VDOM, just like those interfaces on a FortiGate unit without VDOMs enabled.

To travel between VDOMs, a packet must first pass through a firewall policy on a physical interface. The packet then arrives at another VDOM on that same FortiGate unit, but on a different interface, where it must pass through another firewall before entering. It doesn’t matter if the interface is physical or virtual — inter-VDOM packets still require the same security measures as when passing through physical interfaces.

VDOMs provide an additional level of security because regular administrator accounts are specific to one VDOM — an administrator restricted to one VDOM cannot change information on other VDOMs. Any configuration changes and potential errors will apply only to that VDOM and limit any potential down time. Using this concept, you can farther split settings so that the management domain is only accessible by the super\_admin and does not share any settings with the other VDOMs.

## Savings in physical space and power

To increase the number of physical FortiGate units, you need more rack space, cables, and power to install the new units. You also need to change your network configuration to accommodate the new physical units. In the future, if you need fewer physical units you are left with expensive hardware that is idle.

Increasing VDOMs involves no additional hardware, no additional cabling, and very few changes to existing networking configurations. VDOMs save physical space and power. You are limited only by the size of the VDOM license you buy and the physical resources on the FortiGate unit.

For example, if you are using one FortiGate 620B unit with 10 VDOMs instead of 10 physical units, over a year you will save an estimated 18,000 kWh. You could potentially save ten times that amount with a 100 VDOM license.

By default, FortiGate units support a maximum of 10 VDOMs in any combination of NAT/Route and Transparent modes. For FortiGate models numbered 3000 and higher, you can purchase a

license key to increase the maximum number of VDOMs beyond 10. For more information on VDOM licences, see [“Virtual Domain Licensing” on page 2351](#).

## More flexible MSSP configurations

If you are a managed security and service provider (MSSP), VDOMs are fundamental to your business. As a service provider you have multiple customers, each with their own needs and service plans. VDOMs allow you to have a separate configuration for each customer, or group of customers; with up to 500 VDOMs configured per FortiGate unit on high end models. See [“Virtual Domain Licensing” on page 2351](#).

Not only does this provide the exact level of service needed by each customer, but administration of the FortiGate unit is easier as well - you can provide uninterrupted service generally with immediate changes as required. Most importantly, it allows you to only use the resources that each customer needs. Inter-VDOM links allow you to customize the level of interaction you need between each of your customers and your administrators. See [“Inter-VDOM routing” on page 2400](#).

## Enabling and accessing Virtual Domains

While Virtual Domains are essentially the same as your regular FortiGate unit for menu configuration, CLI command structure, and general task flow, there are some small differences.

After first enabling VDOMs on your FortiGate unit, you should take the time to familiarize yourself with the interface. This section will help walk you through virtual domains.

This section includes:

- [Enabling Virtual Domains](#)
- [Viewing the VDOM list](#)
- [Global and per-VDOM settings](#)
- [Resource settings](#)
- [Virtual Domain Licensing](#)
- [Logging in to VDOMs](#)

### Enabling Virtual Domains

Using the default admin administration account, you can enable or disable VDOM operation on the FortiGate unit.

#### **To enable VDOM configuration - web-based manager**

1. Log in with a super\_admin account.
2. Go to *System > Dashboard > Status*.
3. Under *System Information > Virtual Domain*, select *Enable* and confirm your selection.

The FortiGate unit logs off all sessions. You can now log in again as admin. For more information, see [“Administrators in Virtual Domains” on page 2357](#).

**Figure 328:**System Information

Host Name	FG100D3G12801307 [Change]
Serial Number	FG100D3G12801307
HA Status	Standalone [Configure]
System Time	Tue Nov 6 15:50:05 2012 (FortiGuard) [Change]
Firmware Version	v5.0,build0128 (GA) [Update] [Details]
System Configuration	[Backup] [Restore] [Revisions]
Current Administrator	admin [Change Password] /1 in Total [Details]
Uptime	5 day(s) 7 hour(s) 27 min(s)
Virtual Domain	Enabled [Disable]

VDOMs are enabled

### To enable VDOM configuration - CLI

```
config system global
 set vdom-admin enable
end
```

### Changes to the web-based manager and CLI

When Virtual Domains are enabled, your FortiGate unit will change. The changes will be visible in both the web-based manager and CLI, just the web-based manager, or just the CLI.

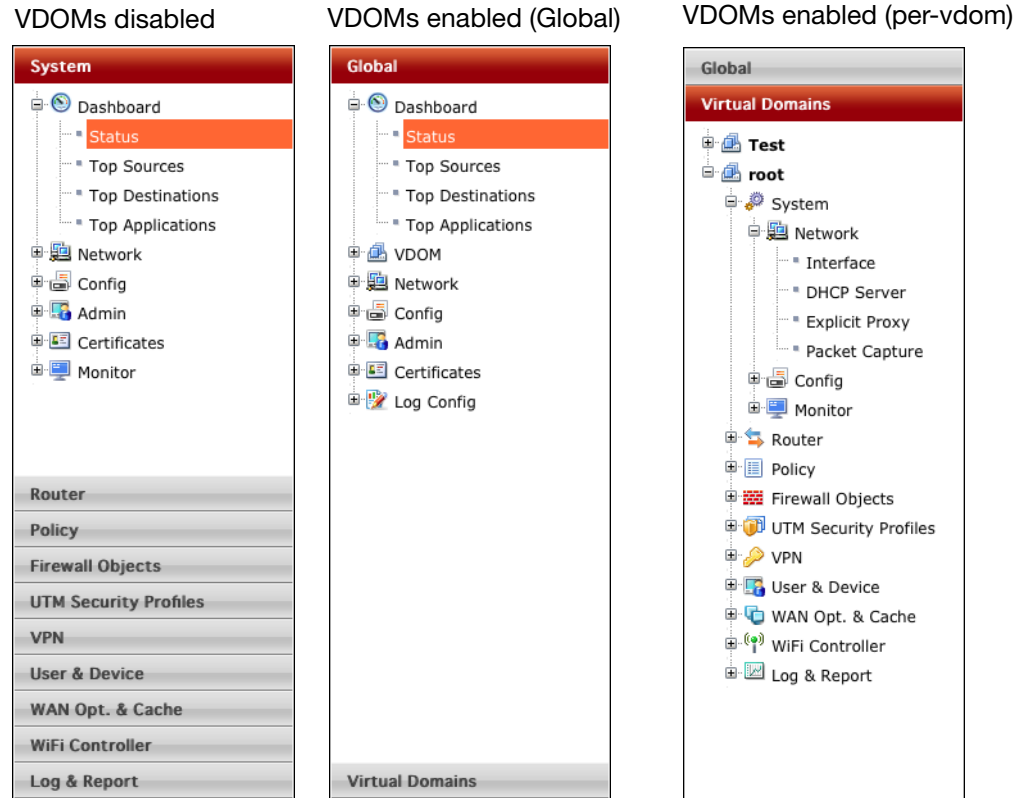
When enabling VDOMs, the web-based manager and the CLI are changed as follows:

- Global and per-VDOM configurations are separated. This is indicated in the Online Help by Global and VDOM icons. See [“Global and per-VDOM settings” on page 2338](#).
- Only admin accounts using the super\_admin profiles can view or configure global options. See [“Administrators in Virtual Domains” on page 2357](#).
- Admin accounts using the super\_admin profile can configure all VDOM configurations.
- All other administrator accounts can configure only the VDOM to which they are assigned.

The following changes are specific to the web-based manager:

- In the Global view, the System section of the left-hand menu is renamed to Global, and includes a VDOM sub-menu.
- The Log Config menu is moved from Log & Report into the new Global section.
- For admin accounts using the super\_admin profile, a new section called Virtual Domains is added at the bottom of the left-hand menu. It lists all the individual VDOMs as expandable menus, with all VDOM specific options in that menu, which allows you to easily select which VDOM to configure, including the root VDOM. See [Figure 329 on page 2336](#).

**Figure 329:**Menu with VDOMs disabled, at the global level, and VDOM level



In the CLI, admin accounts using the super\_admin profile must specify either the global or a VDOM-specific shell before entering commands:

- To change FortiGate unit system settings, from the top level you must first enter `config global` before entering commands.
- To change VDOM settings, from the top level you must first enter `config vdom` `edit <vdom_name>`

before entering your commands for that VDOM. For information on which commands are global and which are per-VDOM, see [“Global and per-VDOM settings” on page 2338](#).

### Changes to FortiGate unit settings

Settings configured outside of a VDOM are called global settings. These settings affect the entire FortiGate unit and include areas such as interfaces, HA, maintenance, some antivirus settings, and some logging settings. In general, any unit settings that should only be changed by the top level administrator are global settings.

Settings configured within a VDOM are called VDOM settings. These settings affect only that specific VDOM and include areas such as operating mode, routing, firewall, VPN, some antivirus, some logging, and reporting.

For more information, see [“Global and per-VDOM settings” on page 2338](#).



## Viewing the VDOM list

The VDOM list shows all virtual domains, their status, and which VDOM is the management VDOM. It is accessible if you are logged in on an administrator account with the super\_admin profile such as the “admin” administrator account.

In the VDOM list you can create or delete VDOMs, edit VDOMs, change the management VDOM, and enable or disable VDOMs.

You can access the VDOM list in *Global > VDOM > VDOM*.

---



The root domain cannot be disabled, even if it is not the management VDOM.

---

**Figure 330:**List of VDOMs

<input type="checkbox"/>	Name	Operation Mode	Interfaces	Enable	Comments	Ref.
<input type="checkbox"/>	root	NAT	example_wlan , modem , ssl.root , wan1 , wan2 , wlan , wlan_employee	✓		0
<input type="checkbox"/>	vdom1	NAT	dmz , internal , ssl.vdom1	✓		0
<input type="checkbox"/>	vdom2	NAT	ssl.vdom2	✗		0

**Create New** Select to add a new VDOM. See [“Creating a Virtual Domain” on page 2354.](#)

**Edit** Select to change an existing selected VDOM.

**Delete** Select to delete the selected VDOM. See [“Deleting a VDOM” on page 2356.](#)

**Switch Management** Select to switch the management VDOM. Also shows the current management VDOM.  
 You must select an active non-management VDOM before this option becomes available.  
 See [“Changing the management virtual domain” on page 2361.](#)

**Selected** When checked, this checkbox indicates this VDOM has been selected. Nearly all operations such as Edit, Delete, and Switch Management require a VDOM to first be selected.

**Name** The name of the VDOM. VDOMs are listed in alphabetical order.  
 When the VDOM is active, you can select the VDOM name to enter that VDOM. See [“Enabling and accessing Virtual Domains” on page 2334.](#)

**Operation Mode** Indicates the operation mode as either NAT (for NAT/Route mode) or TP (for Transparent mode).

**Interfaces** The interfaces associated with this VDOM. Each VDOM also includes an interface that starts with “ssl.” that is created by default.

**Enable** A green checkmark indicates this VDOM is active. See [“Disabling a Virtual Domain” on page 2355.](#)  
 A grey X indicated this VDOM is disabled. See [“Disabling a Virtual Domain” on page 2355.](#)

**Comments** Comments entered when the VDOM was created are displayed here.

**Ref.** The number of references to this VDOM in the configuration.

## Global and per-VDOM settings

Settings configured outside of a VDOM are called global settings. These settings affect the entire FortiGate unit and include areas such as interfaces, HA, maintenance, some antivirus, and some logging. In general, any unit settings that should only be changed by the top level administrator are global settings.

Settings configured within a VDOM are called VDOM settings. These settings affect only that specific VDOM and include areas such as operating mode, routing, firewall, VPN, some antivirus, some logging settings, and reporting.

When Virtual Domains are not enabled, the entire FortiGate unit is effectively a single VDOM. Per-VDOM limits apply. For some resource types, the global limit cannot be reached with only one VDOM.

Some FortiGate unit documentation indicates which parts of the web-based manager, or the CLI are global and which are per-VDOM using the icons shown below. These icons are also present in the Online Help, available on your FortiGate unit.

**Figure 331:**Global and VDOM icons



For more information on CLI commands, see the [FortiGate CLI Reference](#).

This section includes:

- [Global settings - web-based manager](#)
- [Per-VDOM settings - web-based manager](#)
- [Global settings - CLI](#)
- [Per-VDOM settings - CLI](#)

### Global settings - web-based manager

The following table lists commands in the web-based manager that are considered global settings when VDOMs are enabled.

The following configuration settings affect all virtual domains. When virtual domains are enabled, only accounts with the default super\_admin profile can access global settings.

**Table 106:**Global configuration settings

System	
	Dashboard > Status - Host name
	Dashboard > Status - HA Status
	Dashboard > Status - System Time
	Dashboard > Status - Firmware version
	Dashboard > Status - Configuration backup and restore
	VDOM > VDOM - list
	VDOM > VDOM - edit VDOM (mode and resources)
	VDOM > Global Resources
	Network > Interfaces
	Network > DNS - DNS and DDNS settings
	Config > HA
	Config > SNMP
	Config > Replacement Message - messages and images

**Table 106:**Global configuration settings (Continued)

	Config > FortiGuard - configuration Config > Advanced - scripts, USB Auto-install, debug log download Config > Messaging Servers Admin > Administrators Admin > Admin Profile Admin > Settings - web administration ports, password policy, display settings, timeouts, LCD panel Certificates - local, remote, and CA certificates, CRLs
<b>Log&amp;Report</b>	Log Config - Log Setting

### Per-VDOM settings - web-based manager

The following table lists commands in the web-based manager that are considered per-VDOM settings when VDOMs are enabled.

**Table 107:**VDOM configuration settings

<b>System</b>	Dashboard > Status - read-only except for administrator password Network > Interfaces (and zones) Network > DHCP Server Network > Explicit Proxy Network > Routing Table (Transparent mode only) Network > Packet Capture Config > Replacement Message (messages and images) Config > Advanced Monitor > DHCP Monitor
<b>Router</b>	All settings, including dead gateway detection
<b>Policy</b>	All settings
<b>Firewall Objects</b>	All settings
<b>Security Profiles</b>	All settings
<b>VPN</b>	All settings
<b>User &amp; Device</b>	All settings
<b>WiFi Controller</b>	All settings

**Table 107:**VDOM configuration settings (Continued)

<b>Log&amp;Report</b>	Traffic and Event Logs
	Reports
	FortiCloud
	Log Config > Log Setting and Alert E-mail
	Logging Monitor

### Global settings - CLI

The following table lists commands in the web-based manager that are considered global settings when VDOMs are enabled.

From a super\_admin profile account, use this command to configure features that apply to the complete FortiGate unit including all virtual domains. Virtual domain configuration (vdom-admin) must be enabled first.

This command syntax shows how you access the commands within config global. For information on these commands, refer to the relevant sections in this Reference. If there are

multiple versions of the same command with a “2” or “3” added, the additional commands are not listed but fall under the unnumbered command of the same name.

```
config global
 config antivirus heuristic
 config antivirus quarfilepattern
 config antivirus service
 config application name
 config dlp settings
 config endpoint-control app-detect
 config firewall ssl
 config gui console
 config ips decoder
 config ips global
 config ips rule
 config log fortianalyzer setting
 config log fortiguard setting
 config log memory global-setting
 config log syslogd filter
 config log syslogd setting
 config log webtrends ...
 config spamfilter fortishield
 config spamfilter options
 config system accprofile
 config system admin
 config system alertemail
 config system amc
 config system auto-install
 config system autoupdate ...
 config system aux
 config system bug-report
 config system central-management
 config system chassis-loadbalance
 config system console
 config system ddns
 config system dialinsvr
 config system dns
 config system dynamic-profile
 config system fips-cc
 config system fortiguard
 config system fortiguard-log
 config system global
 config system ha
 config system interface
 config system npu
 config system ntp
 config system password-policy
 config system replacemsg ...
 config system replacemsg-image
 config system resource-limits
```

```
config system session-helper
config system session-sync
config system sflow
config system snmp ...
config system switch-interface
config system tos-based-priority
config system vdom-link
config system vdom-property
config vpn certificate ...
config wanopt storage
config webfilter fortiguard
config wireless-controller global
config wireless-controller timers
config wireless-controller vap
execute backup
execute batch
execute central-mgmt
execute cfg reload
execute cfg save
execute cli check-template-status
execute cli status-msg-only
execute date
execute disconnect-admin-session
execute disk
execute enter
execute factoryreset
execute firmware-list
execute formatlogdisk
execute forticlient
execute fortiguard-log
execute ha disconnect
execute ha manage
execute ha synchronize
execute log ...
execute log-report
execute reboot
execute report-config
execute restore
execute revision
execute router ... (except clear)
execute scsi-dev
execute send-fds-statistics
execute set-next-reboot
execute sfp-mode-sgmii
execute shutdown
execute tac
execute time
execute update-ase
execute update-av
```

```
execute update-ips
execute update-netscan
execute update-now
execute upload
execute usb-disk
execute vpn certificate ...
execute wireless-controller ... (except reset-wtp)
get firewall vip ...
end
```

## Per-VDOM settings - CLI

The following table lists commands in the web-based manager that are considered VDOM-specific settings when VDOMs are enabled.

From the super\_admin account, you can use the commands below to add and configure virtual domains. The number of virtual domains you can add is dependent on the FortiGate model. Virtual domain configuration (vdom-admin) must be enabled.

Once you add a virtual domain you can configure it by adding zones, firewall policies, routing settings, and VPN settings. You can also move physical interfaces from the root virtual domain to other virtual domains and move VLAN subinterfaces from one virtual domain to another.

By default all physical interfaces are in the root virtual domain. You cannot remove an interface from a virtual domain if the interface is part of any of the following configurations:

- routing
- proxy arp
- DHCP server
- zone
- firewall policy
- redundant pair
- link aggregate (802.3ad) group

Delete these objects, or modify them, to be able to remove the interface.



This command syntax shows how you access the commands within a VDOM. Refer to the relevant sections in this Reference for information on these commands.

```
config vdom
 edit <vdom_name>
 config antivirus profile
 config antivirus quarantine
 config antivirus settings
 config application list
 config application rule-settings
 config dlp ... (except settings)
 config endpoint-control app-detect
 config endpoint-control profile
 config endpoint-control settings
 config firewall ... (except ssl)
 config ftp-proxy
 config icap
 config imp2p
 config ips DoS
 config ips custom
 config ips rule-settings
 config ips sensor
 config ips settings
 config log custom-field
 config log disk
 config log eventfilter
 config log fortianalyzer
 config log gui
 config log memory
 config log syslogd
 config log trafficfilter
 config log visibility
 config netscan
 config router
 config spamfilter ... (except fortishield and options)
 config system 3g-modem
 config system admin
 config system arp-table
 config system carrier-endpoint-translation
 config system dhcp ...
 config system dhcp6 ...
 config system dns-database
 config system dns-server
 config system gre-tunnel
 config system interface
 config system ipv6-tunnel
 config system modem
 config system monitors
 config system object-tag
 config system proxy-arp
```

```
config system replacemsg-group
config system session-ttl
config system settings
config system sit-tunnel
config system switch-interface
config system wccp
config system zone
config user ...
config voip
config vpn ...
config wanopt
config web-proxy
config webfilter (except fortiguard)
config wireless-controller (except global and timers)
execute backup
execute clear system arp table
execute cli check-template-status
execute cli status-msg-only
execute dhcp lease-clear
execute dhcp lease-list
execute dhcp6 lease-clear
execute dhcp6 lease-list
execute enter
execute fortitoken ...
execute fssso refresh
execute interface dhcpclient-renew
execute interface pppoe-reconnect
execute log ...
execute log-report ...
execute modem dial
execute modem hangup
execute modem trigger
execute mrouter clear
execute netscan ...
execute ping, ping6
execute ping-options, ping6-options
execute restore
execute revision
execute router clear bgp
execute router clear ospf process
execute router restart
execute sfp-mode-sgmii
execute ssh
execute tac
execute telnet
execute traceroute
execute tracert6
execute upload
execute usb-disk
```

```
execute vpn ipsec tunnel
execute vpn sslvpn ...
execute wireless-controller reset-wtp
next
edit <another_vdom>
 config ...
 execute ...
end
end
```

For more information, see [“Global and per-VDOM settings”](#) on page 2338.

## Resource settings

Your FortiGate unit has a limited amount of hardware resources such as memory, disk storage, CPU operations. When Virtual Domains are disabled, this limit is not a major concern because all sessions, users, and other processes share all the resources equally.

When using Virtual Domains, hardware resources can be divided differently between Virtual Domains as they are needed. Minimum levels of resources can be specified for each VDOM, so that no Virtual Domain will suffer a complete lack of resources.

For example, if one VDOM has only a web server and logging server connected, and a second VDOM has an internal network of 20 users, these two VDOMs will require different levels of resources. The first VDOM will require many sessions but no user accounts. This compares to the second VDOM where user accounts and management resources are required, but fewer sessions.

Using the global and per-VDOM resource settings, you can customize the resources allocated to each VDOM to ensure the proper level of service is maintained on each VDOM.

This section includes:

- [Global resource settings](#)
- [Per-VDOM resource settings](#)

### Global resource settings

Global Resources apply to the whole FortiGate unit. They represent all of the hardware capabilities of your unit. By default the values are set to their maximum values. These values vary by your model due to each model having differing hardware capabilities.


It can be useful to change the maximum values for some resources to ensure there is enough memory available for other resources that may be more important to your configuration.

To use the earlier example, if your FortiGate unit is protecting a number of web servers and other publicly accessible servers you would want to maximize the available sessions and proxies while minimizing other settings that are unused such as user settings, VPNs, and dial-up tunnels.

Global Resources are only configurable at the global level, and only the admin account has access to these settings.

**Note** that global resources, such as the log disk quote resource, will only be visible if your FortiGate unit hardware supports those resources, such as having a hard disk to support the log disk resource.

**Figure 332:**Global Resources- web-based manager

	Resource	Configured Maximum	Default Maximum	Current Usage
<input type="checkbox"/>	Sessions	0	0	26
<input type="checkbox"/>	VPN IPsec Phase1 Tunnels	10000	10000	0
<input type="checkbox"/>	VPN IPsec Phase2 Tunnels	10000	10000	0
<input type="checkbox"/>	Dial-up Tunnels	0	0	0
<input type="checkbox"/>	Firewall Policies	100000	100000	3
<input type="checkbox"/>	Firewall Addresses	20000	20000	11
<input type="checkbox"/>	Firewall Address Groups	10000	10000	0
<input type="checkbox"/>	Firewall Custom Services	0	0	0
<input type="checkbox"/>	Firewall Service Groups	0	0	0
<input type="checkbox"/>	Firewall One-time Schedules	0	0	0
<input type="checkbox"/>	Firewall Recurring Schedules	0	0	5
<input type="checkbox"/>	Local Users	0	0	0
<input type="checkbox"/>	User Groups	0	0	0
<input type="checkbox"/>	SSL VPN	0	0	0
<input type="checkbox"/>	Concurrent web proxy users	2000	2000	0
<input type="checkbox"/>	log disk quota	0	0	0

**To view global resource settings - web-based manager**

Select *Global > VDOM > Global Resources*.

The following information is displayed:

<b>Edit</b>	Select to edit the <i>Configured Maximum</i> value for a single selected <i>Resource</i> . If multiple <i>Resources</i> are selected, <i>Edit</i> is not available.
<b>Reset to default value</b>	Select to return one or more selected <i>Resources</i> to factory default settings.
<b>Checkbox</b>	Select a <i>Resource</i> for editing or resetting to default values.
<b>Resource</b>	The name of the available global resources.
<b>Configured Maximum</b>	The currently configured maximum for this resource. This value can be changed by selecting the <i>Resource</i> and editing it.
<b>Default Maximum</b>	The factory configured maximum value for this resource. You cannot set the <i>Configured Maximum</i> higher than the <i>Default Maximum</i> .
<b>Current Usage</b>	The amount of this resource that is currently being used. This value is useful for determining when and if you may need to adjust <i>Configured Maximum</i> values for some resources on your FortiGate unit.

**To view global resource settings - CLI**

```
config global
 config system resource-limits
 get
```

When viewing the global resource limits in the CLI, the output appears similar to:

```
FGT1000A (global) # config system resource-limits
FGT1000A (resource-limits) # get

session : 0
ipsec-phase1 : 10000
ipsec-phase2 : 10000
dialup-tunnel : 0
firewall-policy : 100000
firewall-address : 20000
firewall-addrgrp : 10000
custom-service : 0
service-group : 0
onetime-schedule : 0
recurring-schedule : 0
user : 0
user-group : 0
sslvpn : 0
proxy : 2000
```



For explicit proxies, when configuring limits on the number of concurrent users, you need to allow for the number of users based on their authentication method. Otherwise you may run out of user resources prematurely.

- Each session-based authenticated user is counted as a single user using their authentication membership (RADIUS, LDAP, FSAE, local database etc.) to match users in other sessions. So one authenticated user in multiple sessions is still one user.
- For all other situations, the source IP address is used to determine a user. All sessions from a single source address are assumed to be from the same user.

---

### Per-VDOM resource settings

While Global resources apply to resources shared by the whole FortiGate unit, per-VDOM resources are specific to only one Virtual Domain.

By default all the per-VDOM resource settings are set to no limits. This means that any single VDOM can use up all the resources of the entire FortiGate unit if it needs to do so. This would starve the other VDOMs for resources to the point where they would be unable to function. For this reason, it is recommended that you set some maximums on resources that are most vital to your customers.

Each Virtual Domain has its own resource settings. These settings include both maximum, and minimum levels. The maximum level is the highest amount of that resource that this VDOM can use if it is available on the FortiGate unit. Minimum levels are a guaranteed level that this minimum level of the resource will always be available no matter what the other VDOMs may be using.

**Figure 333:**per-VDOM resources - web-based manager

Resource	Maximum	Guaranteed	Current
Sessions	0	0	24
VPN IPsec Phase1 Tunnels	0	0	0
VPN IPsec Phase2 Tunnels	0	0	0
Dial-up Tunnels	0	0	0
Firewall Policies	0	0	3
Firewall Addresses	0	0	3
Firewall Address Groups	0	0	0
Firewall Custom Services	0	0	0
Firewall Service Groups	0	0	0
Firewall One-time Schedules	0	0	0
Firewall Recurring Schedules	0	0	1
Local Users	0	0	0
User Groups	0	0	0
SSL VPN	0	0	0
Concurrent web proxy users	0	0	0
log disk quota	0	0	0

For example, consider a FortiGate unit that has ten VDOMs configured. vdom1 has a maximum of 5000 sessions and a minimum of 1000 sessions. If the FortiGate unit has a global maximum of 20,000 sessions, it is possible that vdom1 will not be able to reach its 5000 session upper limit. However, at all times vdom1 is guaranteed to have 1000 sessions available that it can use. On the other hand, if the remaining nine VDOMs use only 1000 sessions each, vdom1 will be able to reach its maximum of 5000.

**To view per-VDOM resource settings - web-based manager**

1. Select *Global > VDOM > VDOM*.
2. Select the `root` VDOM, and select *Edit*.
3. Adjust the settings in the *Resource Usage* section of the page.

<b>Resource</b>	Name of the resource. Includes dynamic and static resources.
<b>Maximum</b>	Override the global limit to reduce the amount of each resource available for this VDOM. The maximum must be the same as or lower than the global limit. The default value is 0, which means the maximum is the same as the global limit.  <b>Note:</b> If you set the maximum resource usage for a VDOM you cannot reduce the default maximum global limit for all VDOMs below this maximum.
<b>Guaranteed</b>	Enter the minimum amount of the resource available to this VDOM regardless of usage by other VDOMs. The default value is 0, which means that an amount of this resource is not guaranteed for this VDOM.
<b>Current</b>	The amount of the resource that this VDOM currently uses.

4. Select *OK*.

## To view per-VDOM resource settings - CLI

```
config global
 config system vdom-property
 edit root
 get
```

When viewing the per-VDOM resource limits in the CLI, the output appears similar to the following. Note that the first two lines are not part of the resource limits. In the CLI, the first number is the maximum value, and the second number is the guaranteed minimum.

```
FGT1KA3607500810 (vdom-property) # edit root
FGT1KA3607500810 (root) # get

name : root
description : property limits for vdom root
session : 0 0
ipsec-phase1 : 0 0
ipsec-phase2 : 0 0
dialup-tunnel : 0 0
firewall-policy : 0 0
firewall-address : 0 0
firewall-addrgrp : 0 0
custom-service : 0 0
service-group : 0 0
onetime-schedule : 0 0
recurring-schedule : 0 0
user : 0 0
user-group : 0 0
sslvpn : 0 0
proxy : 0 0
```

## Virtual Domain Licensing

For FortiGate models in the 1U category and higher, you can purchase a license key to increase the maximum number of VDOMs. Most Enterprise and Large Enterprise (2U) models can support up to 500 VDOMs. Chassis-based models can support over 500 VDOMs. For specific information, see the product data sheet.

Configuring 500 or more VDOMs will result in reduced system performance. See [“FortiGate unit running very slowly” on page 2440](#).



Your FortiGate unit has limited resources that are divided among all configured VDOMs. These resources include system memory and CPU. Running security features on many VDOMs at once can limit resources available for basic processing. If you require many VDOMs, all with active security features, it is recommended to upgrade to a more powerful FortiGate unit.



It is important to backup your configuration before upgrading the VDOM license on your FortiGate unit or units, especially with FortiGate units in HA mode.

---

### To obtain a VDOM license key

1. Log in with a super\_admin account.
2. Go to *System > Dashboard > Status*.
3. Record your FortiGate unit serial number as shown in “[System Information](#)” on page 2335.
4. Under *License Information > Virtual Domain*, select *Purchase More*.



If you do not see the *Purchase More* option on the System Dashboard, your FortiGate model does not support more than 10 VDOMs.

**Figure 334:**VDOM License Information

License Information		
<b>Support Contract</b>		
Registration	Unreachable	
<b>FortiGuard Services</b>		
AntiVirus	Unreachable	[Configure]
AV Definitions	9.00795 (Updated 2008-12-08)	[Update]
Extended set	0.00000 (Updated 2003-01-01)	
Intrusion Protection	Unreachable	[Configure]
IPS Definitions	2.00720 (Updated 2009-12-01)	[Update]
Vulnerability Compliance and Management	Unreachable	[Configure]
VCM Plugin	1.00098 (Updated 2010-02-11)	[Update]
Web Filtering	Unreachable	[Configure]
Email Filtering	Unreachable	[Configure]
Analysis & Management Service	Unreachable	
Services Account ID	[Change]	
<b>Virtual Domain</b>		
VDOMs Allowed	10	[Purchase More]
<b>Endpoint Security</b>		
FortiClient Software	Unreachable	
Application Signature Package	1.131 (Updated 2010-02-16)	

Purchase a larger VDOM license

5. You will be taken to the Fortinet customer support web site where you can log in and purchase a license key for 25, 50, 100, 250, 500, or more VDOMs.
6. When you receive your license key, go to the Dashboard and select *Upload License* under *License Information, Virtual Domains*.
7. In the *Input License Key* field, enter the 32-character license key you received from Fortinet customer support.
8. Select *Apply*.

To verify the new VDOM license, in global configuration go to *System > Dashboard*. Under *License Information, Virtual Domains* the maximum number of VDOMs allowed is shown.



VDOMs created on a registered FortiGate unit are recognized as real devices by any connected FortiAnalyzer unit. The FortiAnalyzer unit includes VDOMs in its total number of registered devices. For example, if three FortiGate units are registered on the FortiAnalyzer unit and they contain a total of four VDOMs, the total number of registered FortiGate units on the FortiAnalyzer unit is seven. For more information, see the [FortiAnalyzer Administration Guide](#).

## Logging in to VDOMs

Only super\_admin administrator accounts can access all global settings on the FortiGate unit and all of the VDOMs as well. Other administrator accounts can access and configure only their



single VDOM and they must connect to an interface that is part of that VDOM. For example, administratorB is the admin for vdomB. If he tries to log into vdomA, or an interface that is part of vdomA he will not be able to log on. For more information on administrators in VDOMs, see [“Administrators in Virtual Domains” on page 2357](#).

Management services communicate using the management VDOM, which is the root VDOM by default. For more information, see [“Changing the management virtual domain” on page 2361](#).



Management traffic requires an interface that has access to the Internet. If there is no interface assigned to the VDOM containing the management traffic, services including updates will not function. For more information, see [“Changing the management virtual domain” on page 2361](#).

---

### To access a VDOM with a super\_admin account - web-based manager

1. Log in with a super\_admin account.
2. In the *Virtual Domains* menu on the left-hand side, select the VDOM to configure.  
The menu will expand to show the various pages and settings for that VDOM.
3. When you have finished configuring the VDOM, you can
  - open the *Global* menu to return to global configuration
  - log out.

### To access a VDOM with a super\_admin account - CLI

With the super\_admin, logging into the CLI involves also logging into the specific VDOM. If you need a reminder, use `edit ?` to see a list of existing VDOMs before you editing a VDOM.



If you misspell a VDOM you are trying to switch to, you will create a new VDOM by that name. Any changes you make will be part of the new VDOM, and not the intended VDOM. If you are having problems where your changes aren't visible, back up to the top level and use `edit ?` to see a list of VDOMs to ensure this has not happened. If it has happened, see [“Deleting a VDOM” on page 2356](#).

```
config vdom
 edit ?
 edit <chosen_vdom>
 ..
 <enter vdom related commands>
 ..
end
exit
```

### To access a VDOM with a non super\_admin account - web-based manager

1. Connect to the FortiGate unit using an interface that belongs to the VDOM to be configured.
2. Log in using an administrator account that has access to the VDOM.  
The main web-based manager page opens. The interface is largely the same as if the device has VDOMs disabled. From here you can access VDOM-specific settings.

### To access a VDOM with a non-super\_admin account - CLI

A non-super\_admin account has access to only one VDOM and must log in through an interface that belongs to the same VDOM, but the process is the same as logging into a non-VDOM unit.

```
Login: regular_admin
Password: <password>
..
<enter vdom related commands>
..
exit
```

## Configuring Virtual Domains

Only a super\_admin administrator account such as the default “admin” account can create, disable, or delete VDOMs. That account can create additional administrators for each VDOM.

This section includes:

- [Creating a Virtual Domain](#)
- [Disabling a Virtual Domain](#)
- [Deleting a VDOM](#)
- [Administrators in Virtual Domains](#)

### Creating a Virtual Domain

Once you have enabled Virtual Domains on your FortiGate unit, you can create additional Virtual Domains beyond the default root Virtual Domain.

By default new Virtual Domains are set to NAT/Route operation mode. If you want a Virtual Domain to be in Transparent operation mode, you must manually change it. See [“Virtual Domains in Transparent mode” on page 2380](#).

You can name new Virtual Domains as you like with the following restrictions:

- only letters, numbers, “-”, and “\_” are allowed
- no more than 11 characters are allowed
- no spaces are allowed
- VDOMs cannot have the same names as interfaces, zones, switch interfaces, or other VDOMs.



When creating large numbers of VDOMs you should not enable advanced features such as proxies, web filtering, and antivirus due to limited FortiGate unit resources. Also when creating large numbers of VDOMs, you may experience reduced performance for the same reason.

---

### To create a VDOM - web-based manager

1. Log in with a super\_admin account.
2. Go to *System > Dashboard > Status* and ensure that Virtual Domains are enabled. If not, see [“Enabling and accessing Virtual Domains” on page 2334](#).
3. Select *System > VDOM > VDOM*.
4. Select *Create New*.
5. Enter a unique name for your new VDOM.

6. Enter a short and descriptive comment to identify this VDOM.
7. Select *OK*.  
Repeat Steps 4 through 7 to add additional VDOMs.

#### To create a VDOM - CLI

```
config vdom
 edit <new_vdom_name>
end
```



If you want to edit an existing Virtual Domain in the CLI, and mistype the name a new Virtual Domain will be created with this new misspelled name. If you notice expected configuration changes are not visible, this may be the reason. You should periodically check your VDOM list to ensure there are none of these misspelled VDOMs present.

---

## Disabling a Virtual Domain

The status of a VDOM can be Enabled, or Disabled.

Active status VDOMs can be configured. Active is the default status when a VDOM is created. The management VDOM must be an Active VDOM. For more information on the management VDOM, see [“Changing the management virtual domain” on page 2361](#).

Disabled status VDOMs are considered “offline”. The configuration remains, but you cannot use the VDOM, and only the `super_admin` administrator can view it. You cannot delete a disabled VDOM without first enabling it, and removing references to it like usual—there is no *Delete* icon for disabled status VDOMs. You can assign interfaces to a disabled VDOM. See [“Deleting a VDOM” on page 2356](#).

The following procedures show how to disable a VDOM called “test-vdom”.

#### To disable a VDOM - web-based manager

1. Go to *Global > VDOM > VDOM*.
2. Open the VDOM for editing.
3. Ensure *Enable* is not selected and then select *OK*.  
The VDOM’s Enable icon in the VDOM list is a grey X.

#### To disable a VDOM - CLI

```
config vdom
 edit test-vdom
 config system settings
 set status disable
 end
 end
```

#### To enable a VDOM - web-based manager

1. Go to *Global > VDOM > VDOM*.
2. Open the VDOM for editing.
3. Ensure *Enable* is selected and then select *OK*.  
The VDOM’s Enable icon in the VDOM list is a green checkmark.

### To enable a VDOM - CLI

```
config vdom
 edit test-vdom
 config system settings
 set status enable
 end
 end
end
```

## Deleting a VDOM

Deleting a VDOM removes it from the FortiGate unit configuration.

Before you can delete a VDOM, all references to it must be removed. This includes any objects listed in [“Per-VDOM settings - web-based manager” on page 2340](#). If there are any references to the VDOM remaining, you will see an error message and not be able to delete the VDOM.

The VDOM must also be enabled. A disabled VDOM cannot be deleted. You cannot delete the root VDOM or the management VDOM.



Before deleting a VDOM, a good practice is to reset any interface referencing that VDOM to its default configuration, with “root” selected as the Virtual Domain.

---

The following procedures show how to delete the `test-vdom` VDOM.

### To delete a VDOM - web-based manager

1. Go to *Global > VDOM > VDOM*.
2. Select the check box for the VDOM and then select the *Delete* icon.  
If the *Delete* icon is not active, there are still references to the VDOM that must first be removed. The *Delete* icon is available when all the references to this VDOM are removed.
3. Confirm the deletion.

### To delete a VDOM - CLI

```
config vdom
 delete test-vdom
end
```

## Removing references to a VDOM

When you are doing to delete a VDOM, all references to that VDOM must first be removed. It can be difficult to find all the references to the VDOM. This section provides a list of common objects that must be removed before a VDOM can be deleted, and a CLI command to help list the dependencies.

Interfaces are an important part of VDOMs. If you can move all the interfaces out of a VDOM, generally you will be able to delete that VDOM.

## Common objects that refer to VDOMs

When you are getting ready to delete a VDOM check for, and remove the following objects that refer to that VDOM or its components:

- Routing - both static and dynamic routes
- Firewall addresses, policies, groups, or other settings
- Security Features/Profiles
- VPN configuration
- Users or user groups
- Logging
- DHCP servers
- Network interfaces, zones, custom DNS servers
- VDOM Administrators

## Administrators in Virtual Domains

When Virtual Domains are enabled, permissions change for administrators. Administrators are now divided into per-VDOM administrators, and `super_admin` administrators. Only `super_admin` administrator accounts can create other administrator accounts and assign them to a VDOM.

This section includes:

- [Administrator VDOM permissions](#)
- [Creating administrators for Virtual Domains](#)
- [Virtual Domain administrator dashboard display](#)

### Administrator VDOM permissions

Different types of administrator accounts have different permissions within VDOMs. For example, if you are using a `super_admin` profile account, you can perform all tasks. However, if you are using a regular admin account, the tasks available to you depend on whether you have read only or read/write permissions. The following table shows what tasks can be performed by which administrators.

**Table 108:**Administrator VDOM permissions

Tasks	Regular administrator account		Super_admin profile administrator account
	Read only permission	Read/write permission	
<b>View global settings</b>	yes	yes	yes
<b>Configure global settings</b>	no	no	yes
<b>Create or delete VDOMs</b>	no	no	yes
<b>Configure multiple VDOMs</b>	no	no	yes
<b>Assign interfaces to a VDOM</b>	no	no	yes
<b>Revision Control Backup and Restore</b>	no	no	yes

**Table 108:**Administrator VDOM permissions

<b>Create VLANs</b>	no	yes - for 1 VDOM	yes - for all VDOMs
<b>Assign an administrator to a VDOM</b>	no	no	yes
<b>Create additional admin accounts</b>	no	yes - for 1 VDOM	yes - for all VDOMs
<b>Create and edit protection profiles</b>	no	yes - for 1 VDOM	yes - for all VDOMs

The only difference in admin accounts when VDOMs are enabled is selecting which VDOM the admin account belongs to. Otherwise, by default the administration accounts are the same as when VDOMs are disabled and closely resemble the `super_admin` account in their privileges.

### Creating administrators for Virtual Domains

Using the admin administrator account, you can create additional administrator accounts and assign them to VDOMs.



The newly-created administrator can access the FortiGate unit only through network interfaces that belong to their assigned VDOM or through the console interface. The network interface must be configured to allow management access, such as HTTPS and SSH. Without these in place, the new administrator will not be able to access the FortiGate unit and will have to contact the `super_admin` administrator for access.

The following procedure creates a new Local administrator account called `admin_sales` with a password of `fortinet` in the `sales` VDOM using the `admin_prof` default profile.

#### To create an administrator for a VDOM - web-based manager

1. Log in with a `super_admin` account.
2. Go to *System > Admin > Administrators*.
3. Select *Create New*.
4. Select *Regular* for Type, as you are creating a Local administrator account.
5. Enter the necessary information about the administrator: email, password, etc.
6. If this admin will be accessing the VDOM from a particular IP address or subnet, enable *Restrict this Admin Login from Trusted Hosts Only* and enter the IP in *Trusted Host #1*. See [“Using trusted hosts” on page 2359](#).
7. Select `prof_admin` for the *Admin Profile*.
8. Select `sales` from the list of *Virtual Domains*.
9. Select *OK*.

## To create administrators for VDOMs - CLI

```
config global
 config system admin
 edit <new_admin_name>
 set vdom <vdom_for_this_account>
 set password <pwd>
 set accprofile <an_admin_profile>
 ...
 end
```

## Using trusted hosts

Setting trusted hosts for all of your administrators increases the security of your network by further restricting administrative access. In addition to knowing the password, an administrator must connect only through the subnet or subnets you specify. You can even restrict an administrator to a single IP address if you define only one trusted host IP address with a netmask of 255.255.255.255.

When you set trusted hosts for all administrators, the FortiGate unit does not respond to administrative access attempts from any other hosts. This provides the highest security. If you leave even one administrator unrestricted, the unit accepts administrative access attempts on any interface that has administrative access enabled, potentially exposing the unit to attempts to gain unauthorized access.

The trusted hosts you define apply both to the web-based manager and to the CLI when accessed through Telnet or SSH. CLI access through the console is not affected.

The trusted host addresses all default to 0.0.0.0/0.0.0.0 for IPv4, or ::/0 for IPv6. If you set one of the zero addresses to a non-zero address, the other zero addresses will be ignored. The only way to use a wildcard entry is to leave the trusted hosts at 0.0.0.0/0.0.0.0 or ::0. However, this configuration is less secure.

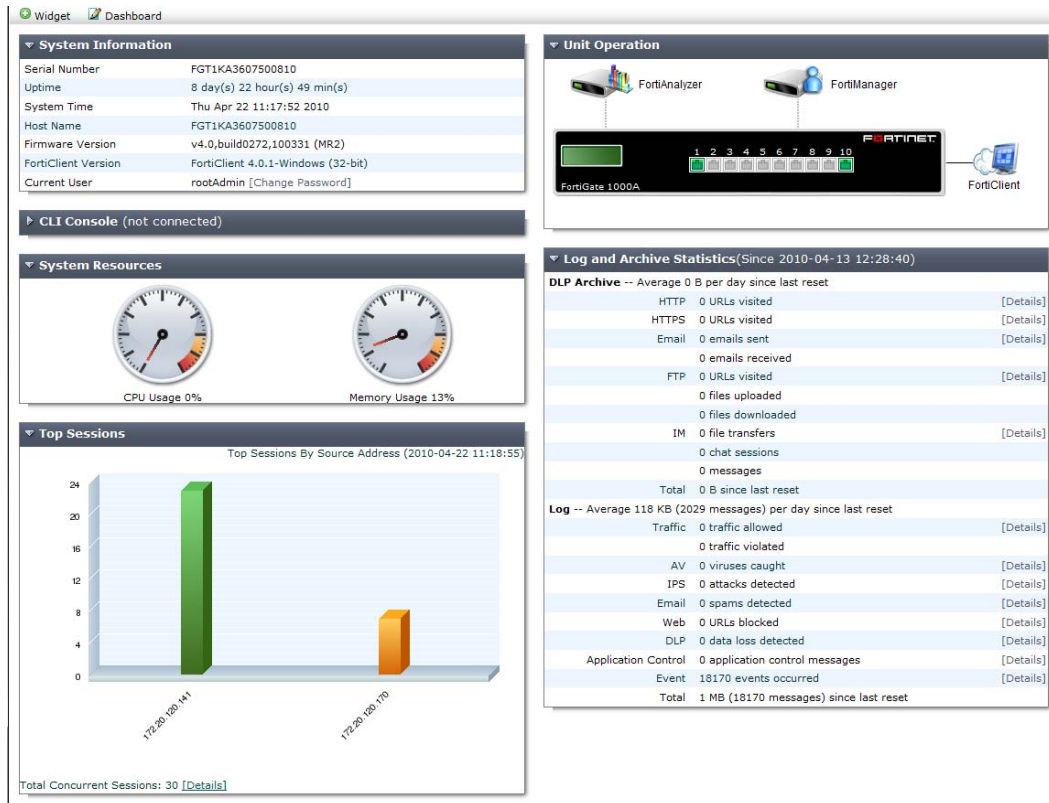
## Virtual Domain administrator dashboard display

When administrators logs into their virtual domain, they see a different dashboard than the global administrator will see. The VDOM dashboard displays information only relevant to that VDOM — no global or other VDOM information is displayed.

**Table 109:**

Information	per-VDOM	Global
System Information	read-only	yes
License Information	no	yes
CLI console	yes	yes
Unit Operation	read-only	yes
Alert Message Console	no	yes
Top Sessions	limited to VDOM sessions	yes
Traffic	limited to VDOM interfaces	yes
Statistics	yes	yes

**Figure 335:VDOM administrator dashboard**





# Virtual Domains in NAT/Route mode

Virtual domains (VDMs) are a method of dividing a FortiGate unit into two or more virtual units that each function as independent units. Each virtual domain has separate routing and security policies. A single FortiGate unit with virtual domains is flexible enough to serve multiple departments of an organization, separate organizations, or be the basis for a service provider's managed security service.



The examples in this chapter are intended to be followed in order as procedures build on previous procedures. If you do not complete the previous procedures, the procedure you are working on may not work properly. If this happens, consult previous procedures or FortiGate documentation.

This chapter contains the following sections:

- [Virtual domains in NAT/Route mode](#)
- [Example NAT/Route VDOM configuration](#)

## Virtual domains in NAT/Route mode

Once you have enabled virtual domains and created one or more VDOMs, you need to configure them. Configuring VDOMs on your FortiGate unit includes tasks such as the ones listed here; while you may not require all for your network topology, it is recommended that you perform them in the order given:

- [Changing the management virtual domain](#)
- [Configuring interfaces in a NAT/Route VDOM](#)
- [Configuring VDOM routing](#)
- [Configuring security policies for NAT/Route VDOMs](#)
- [Configuring security profiles for NAT/Route VDOMs](#)

### Changing the management virtual domain

The management virtual domain is the virtual domain where all the management traffic for the FortiGate unit originates. This management traffic needs access to remote servers, such as FortiGuard services and NTP, to perform its duties. It needs access to the Internet to send and receive this traffic.

Management traffic includes, but is not limited to:

- DNS lookups
- logging to FortiAnalyzer or syslog
- FortiGuard service
- sending alert emails
- Network time protocol traffic (NTP)
- Sending SNMP traps
- Quarantining suspicious files and email.

By default the management VDOM is the root domain. When other VDOMs are configured on your FortiGate unit, management traffic can be moved to one of these other VDOMs.

Reasons to move the management VDOM include selecting a non-root VDOM to be your administration VDOM, or the root VDOM not having an interface with a connection to the Internet.



You cannot change the management VDOM if any administrators are using RADIUS authentication.

---

The following procedure will change the management VDOM from the default `root` to a VDOM named `mgmt_vdom`. It is assumed that `mgmt_vdom` has already been created and has an interface that can access the Internet.

#### To change the management VDOM - web-based manager

1. Select *Global > VDOM > VDOM*.
2. Select the checkbox next to the required VDOM.
3. Select *Switch Management*.

The current management VDOM is shown in square brackets, “[root]” for example.

#### To change the management VDOM - CLI

```
config global
 config system global
 set management-vdom mgmt_vdom
 end
```

Management traffic will now originate from `mgmt_vdom`.

## Configuring interfaces in a NAT/Route VDOM

A VDOM must contain at least two interfaces to be useful. These can be physical interfaces or VLAN interfaces. By default, all physical interfaces are in the root VDOM. When you create a new VLAN, it is in the root VDOM by default.

When there are VDOMs on the FortiGate unit in both NAT and Transparent operation modes, some interface fields will be displayed as “-” on *System > Network > Interfaces*. Only someone with a `super_admin` account can view all the VDOMs.



When moving an interface to a different VDOM, firewall IP pools and virtual IPs for this interface are deleted. You should manually delete any routes that refer to this interface. Once the interface has been moved to the new VDOM, you can add these services to the interface again.

---



When configuring VDOMs on FortiGate units with accelerated interfaces you must assign both interfaces in the pair to the same VDOM for those interfaces to retain their acceleration. Otherwise they will become normal interfaces.

---

This section includes the following topics:

- [Adding a VLAN to a NAT/Route VDOM](#)
- [Moving an interface to a VDOM](#)
- [Deleting an interface](#)
- [Adding a zone to a VDOM](#)

### Adding a VLAN to a NAT/Route VDOM

The following example shows one way that multiple companies can maintain their security when they are using one FortiGate unit with VLANs that share interfaces on the unit.

This procedure will add a VLAN interface called `client1-v100` with a VLAN ID of 100 to an existing VDOM called `client1` using the physical interface called `port2`.



The physical interface does not need to belong to the VDOM that the VLAN belongs to.

#### To add a VLAN subinterface to a VDOM - web-based manager

1. Go to *Global > Network > Interfaces*.
2. Select *Create New*.
3. Enter the following information and select *OK*:

<b>Name</b>	client1-v100
<b>Interface</b>	port2
<b>VLAN ID</b>	100
<b>Virtual Domain</b>	Client1
<b>Addressing mode</b>	Manual
<b>IP/Netmask</b>	172.20.120.110/255.255.255.0
<b>Administrative Access</b>	HTTPS, SSH

You will see an expand arrow added to the `port2` interface. When the arrow is expanded, the interface shows the `client1-v100` VLAN subinterface.

#### To add a VLAN subinterface to a VDOM - CLI

```
config global
 config system interface
 edit client1-v100
 set type vlan
 set vlanid 100
 set vdom Client1
 set interface port2
 set ip 172.20.120.110 255.255.255.0
 set allowaccess https ssh
 end
```

## Moving an interface to a VDOM

Interfaces belong to the root VDOM by default. Moving an interface is the same procedure no matter if its moving from the root VDOM or a any other VDOM.

If you have an accelerated pair of physical interfaces both interfaces must be in the same VDOM or you will lose their acceleration.

The following procedure will move the port3 interface to the Client2 VDOM. This is a common action when configuring a VDOM. It is assumed that the Client2 VDOM has already been created. It is also assumed that your FortiGate unit has a port3 interface. If you are using a different model, your physical interfaces may not be named `port2`, `external` or `port3`.

### To move an existing interface to a different VDOM - web-based manager

1. Go to *Global > Network > Interfaces*.
2. Select *Edit* for the port3 interface.
3. Select `Client2` as the new *Virtual Domain*.
4. Select *OK*.

### To move an existing interface to a different VDOM - CLI

```
config global
 config system interface
 edit port3
 set vdom Client2
 end
```

## Deleting an interface

Before you can delete a virtual interface, or move an interface from one VDOM to another, all references to that interface must be removed. For a list of objects that can refer to an interface see [“Per-VDOM settings - web-based manager” on page 2340](#).

The easiest way to be sure an interface can be deleted is when the Delete icon is no longer greyed out. If it remains greyed out when an interface is selected, that interface still has objects referring to it, or it is a physical interface that cannot be deleted.

### To delete a virtual interface - web-based manager

1. Ensure all objects referring to this interface have been removed.
2. Select *Global > Network > Interfaces*.
3. Select the interface to delete.
4. Select the delete icon.

## Adding a zone to a VDOM

Grouping interfaces and VLAN subinterfaces into zones simplifies policy creation. You can configure policies for connections to and from a zone, but not between interfaces in a zone.

Zones are VDOM-specific. A zone cannot be moved to a different VDOM. Any interfaces in a zone cannot be used in another zone. To move a zone to a new VDOM requires deleting the current zone and re-creating a zone in the new VDOM.

The following procedure will create a zone called `accounting` in the `client2` VDOM. It will not allow intra-zone traffic, and both `port3` and `port2` interfaces belong to this zone. This is a method of grouping and isolating traffic over particular interfaces—it is useful for added security and control within a larger network.

### To add a zone to a VDOM - web-based manager

1. In *Virtual Domains*, select the client2 VDOM.
2. Go to *System > Network > Interfaces*.
3. Select *Create New > Zone*.
4. Enter the following information and select *OK*:

<b>Zone Name</b>	accounting
<b>Block intra-zone traffic</b>	Select
<b>Interface Members</b>	port3, port2

### To add a zone to a VDOM - CLI

```
config vdom
 edit client2
 config system zone
 edit accounting
 set interface port3 port2
 set intrazone deny
 end
 end
 end
```

## Configuring VDOM routing

Routing is VDOM-specific. Each VDOM should have a default static route configured as a minimum. Within a VDOM, routing is the same as routing on your FortiGate unit without VDOMs enabled.

When configuring dynamic routing on a VDOM, other VDOMs on the FortiGate unit can be neighbors. The following topics give a brief introduction to the routing protocols, and show specific examples of how to configure dynamic routing for VDOMs. Figures are included to show the FortiGate unit configuration after the successful completion of the routing example.

This section includes:

- [Default static route for a VDOM](#)
- [Dynamic Routing in VDOMs](#)

### Default static route for a VDOM

The routing you define applies only to network traffic entering non-ssl interfaces belonging to this VDOM. Set the administrative distance high enough, typically 20, so that automatically configured routes will be preferred to the default.

In the following procedure, it is assumed that a VDOM called “Client2” exists. The procedure will create a default static route for this VDOM. The route has a destination IP of 0.0.0.0, on the port3 interface. It has a gateway of 10.10.10.1, and an administrative distance of 20.

The values used in this procedure are very standard, and this procedure should be part of configuring all VDOMs.

### To add a default static route for a VDOM - web-based manager

1. In *Virtual Domains*, select the client2 VDOM.
2. Go to *Router > Static > Static Routes*.
3. Select *Create New*.

4. Enter the following information and select *OK*:

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Device</b>	port2
<b>Gateway</b>	10.10.10.1
<b>Distance</b>	20

#### To add a default static route for a VDOM - CLI

```
config vdom
 edit client2
 config router static
 edit 4
 set device port2
 set dst 0.0.0.0 0.0.0.0
 set gateway 10.10.10.1
 set distance 20
 end
 end
 end
```

### Dynamic Routing in VDOMs

Dynamic routing is VDOM-specific, like all other routing. Dynamic routing configuration is the same with VDOMs as with your FortiGate unit without VDOMs enabled, once you are at the routing menu. If you have multiple VDOMs configured, the dynamic routing configuration between them can become quite complex.

VDOMs provide some interesting changes to dynamic routing. Each VDOM can be a neighbor to the other VDOMs. This is useful in simulating a dynamic routing area or AS or network using only your FortiGate unit.

You can separate different types of routing to different VDOMs if required. This allows for easier troubleshooting. This is very useful if your FortiGate unit is on the border of a number of different routing domains.

For more information on dynamic routing in FortiOS, see [“Dynamic Routing Overview” on page 304](#).

Inter-VDOM links must have IP addresses assigned to them if they are part of a dynamic routing configuration. Inter-VDOM links may or may not have IP addresses assigned to them. Without IP addresses, you need to be careful how you configure routing. While the default static route can be assigned an address of 0.0.0.0 and rely instead on the interface, dynamic routing almost always requires an IP address.

#### RIP

The RIP dynamic routing protocol uses hop count to determine the best route, with a hop count of 1 being directly attached to the interface and a hop count of 16 being unreachable. For example if two VDOMs on the same FortiGate unit are RIP neighbors, they have a hop count of 1.

#### OSPF

OSPF communicates the status of its network links to adjacent neighbor routers instead of the complete routing table. When compared to RIP, OSPF is more suitable for large networks, it is

not limited by hop count, and is more complex to configure. For smaller OSPF configurations its easiest to just use the backbone area, instead of multiple areas.

## BGP

BGP is an Internet gateway protocol (IGP) used to connect autonomous systems (ASes) and is used by Internet service providers (ISPs). BGP stores the full path, or path vector, to a destination and its attributes which aid in proper routing.

## Configuring security policies for NAT/Route VDOMs

Security policies are VDOM-specific. This means that all firewall settings for a VDOM, such as firewall addresses and security policies, are configured within the VDOM.

In VDOMs, all firewall related objects are configured per-VDOM including addresses, service groups, security profiles, schedules, traffic shaping, and so on. If you want firewall addresses, you will have to create them on each VDOM separately. If you have many addresses, and VDOMs this can be tedious and time consuming. Consider using a FortiManager unit to manage your VDOM configuration — it can get firewall objects from a configured VDOM or FortiGate unit, and push those objects to many other VDOMs or FortiGate units. See the [FortiManager Administration Guide](#).



You can customize the *Policy* display by including some or all columns, and customize the column order onscreen. Due to this feature, security policy screenshots may not appear the same as on your screen.

## Configuring a security policy for a VDOM

Your security policies can involve only the interfaces, zones, and firewall addresses that are part of the current VDOM, and they are only visible when you are viewing the current VDOM. The security policies of this VDOM filter the network traffic on the interfaces and VLAN subinterfaces in this VDOM.

A firewall service group can be configured to group multiple services into one service group. When a descriptive name is used, service groups make it easier for an administrator to quickly determine what services are allowed by a security policy.

In the following procedure, it is assumed that a VDOM called `client2` exists. The procedure will configure an outgoing security policy. The security policy will allow all HTTPS and SSH traffic for the `SalesLocal` address group on `VLAN_200` going to all addresses on `port3`. This traffic will be scanned and logged.

### To configure a security policy for a VDOM - web-based manager

1. In *Virtual Domains*, select the `client2` VDOM.
2. Go to *Policy > Policy*.
3. Select *Create New*.
4. Enter the following information and select *OK*:

<b>Source Interface/Zone</b>	VLAN_200
<b>Source Address</b>	SalesLocal
<b>Destination Interface/Zone</b>	port3

<b>Destination Address</b>	any
<b>Schedule</b>	always
<b>Service</b>	Multiple - HTTPS, SSH
<b>Action</b>	ACCEPT
<b>Log Allowed Traffic</b>	enable

### To configure a security policy for a VDOM - CLI

```

config vdom
 edit Client2
 config firewall policy
 edit 12
 set srcintf VLAN_200
 set srcaddr SalesLocal
 set dstintf port3(dmz)
 set dstaddr any
 set schedule always
 set service HTTPS SSH
 set action accept
 set status enable
 set logtraffic enable
 end
 end
 end
end

```

## Configuring security profiles for NAT/Route VDOMs

In NAT/Route VDOMs, security profiles are exactly like regular FortiGate unit operation with one exception. In VDOMs, there are no default security profiles.

If you want security profiles in VDOMs, you must create them yourself. If you have many security profiles to create in each VDOM, you should consider using a FortiManager unit. It can get existing profiles from a VDOM or FortiGate unit, and push those profiles down to multiple other VDOMs or FortiGate units. See [FortiManager Administration Guide](#).

When VDOMs are enabled, you only need one FortiGuard license for the physical unit, and download FortiGuard updates once for the physical unit. This can result in a large time and money savings over multiple physical units if you have many VDOMs.

## Configuring VPNs for a VDOM

Virtual Private Networking (VPN) settings are VDOM-specific, and must be configured within each VDOM. Configurations for IPsec Tunnel, IPsec Interface, PPTP and SSL are VDOM-specific. However, certificates are shared by all VDOMs and are added and configured globally to the FortiGate unit.

## Example NAT/Route VDOM configuration

Company A and Company B each have their own internal networks and their own ISPs. They share a FortiGate unit that is configured with two separate VDOMs, with each VDOM running in



NAT/Route mode enabling separate configuration of network protection profiles. Each ISP is connected to a different interface on the FortiGate unit.

This network example was chosen to illustrate one of the most typical VDOM configurations.

This example has the following sections:

- [Network topology and assumptions](#)
- [General configuration steps](#)
- [Creating the VDOMs](#)
- [Configuring the FortiGate interfaces](#)
- [Configuring the vdomA VDOM](#)
- [Configuring the vdomB VDOM](#)
- [Testing the configuration](#)

## Network topology and assumptions

Both companies have their own ISPs and their own internal interface, external interface, and VDOM on the FortiGate unit.

For easier configuration, the following IP addressing is used:

- all IP addresses on the FortiGate unit end in “.2” such as 10.11.101.2.
- all IP addresses for ISPs end in “.7”, such as 172.20.201.7.
- all internal networks are 10.\*.\* networks, and sample internal addresses end in “.55”.

The IP address matrix for this example is as follows.

Address	Company A	Company B
ISP	172.20.201.7	192.168.201.7
Internal network	10.11.101.0	10.012.101.0
FortiGate / VDOM	172.20.201.2 (port1) 10.11.101.2 (port4)	192.168.201.2 (port3) 10.012.101.2 (port2)

The Company A internal network is on the 10.11.101.0/255.255.255.0 subnet. The Company B internal network is on the 10.12.101.0/255.255.255.0 subnet.

There are no switches or routers required for this configuration.


There are no VLANs in this network topology.

The interfaces used in this example are port1 through port4. Different FortiGate models may have different interface labels. port1 and port3 are used as external interfaces. port2 and port4 are internal interfaces.

The administrator is a super\_admin account. If you are using a non-super\_admin account, refer to [“Global and per-VDOM settings” on page 2338](#) to see which parts a non-super\_admin account can also configure.

When configuring security policies in the CLI always choose a policy number that is higher than any existing policy numbers, select services before profile-status, and profile-status before profile. If these commands are not entered in that order, they may not be available to enter.

**Figure 336:**Example VDOM configuration



## General configuration steps

For best results in this configuration, follow the procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

1. [Creating the VDOMs](#)
2. [Configuring the FortiGate interfaces](#)
3. [Configuring the vdomA VDOM, and Configuring the vdomB VDOM](#)
4. [Testing the configuration](#)

## Creating the VDOMs

In this example, two new VDOMs are created — vdomA for Company A and vdomB for Company B. These VDOMs will keep the traffic for these two companies separate while enabling each company to access its own ISP.

### To create two VDOMs - web-based manager

1. Log in with a super\_admin account.
2. Go to *Global > VDOM > VDOM*, and select *Create New*.
3. Enter vdomA and select *OK*.
4. Select *OK* again to return to the VDOM list.
5. Select *Create New*.
6. Enter vdomB and select *OK*.

### To create two VDOMs - CLI

```
config vdom
 edit vdomA
 next
 edit vdomB
end
```

## Configuring the FortiGate interfaces

This section configures the interfaces that connect to the companies' internal networks, and to the companies' ISPs.

All interfaces on the FortiGate unit will be configured with an IP address ending in ".2" such as 10.11.101.2. This will simplify network administration both for the companies, and for the FortiGate unit global administrator. Also the internal addresses for each company differ in the second octet of their IP address - Company A is 10.11.\*, and Company B is 10.12.\*.

This section includes the following topics:

- [Configuring the vdomA interfaces](#)
- [Configuring the vdomB interfaces](#)



If you cannot change the VDOM of a network interface it is because something is referring to that interface that needs to be deleted. Once all the references are deleted the interface will be available to switch to a different VDOM. For example a common reference to the external interface is the default static route entry. See "[Configuring interfaces in a NAT/Route VDOM](#)" on [page 2362](#).

### Configuring the vdomA interfaces

The vdomA VDOM includes two FortiGate unit interfaces: port1 and external.

The port4 interface connects the Company A internal network to the FortiGate unit, and shares the internal network subnet of 10.11.101.0/255.255.255.0.

The external interface connects the FortiGate unit to ISP A and the Internet. It shares the ISP A subnet of 172.20.201.0/255.255.255.0.

#### To configure the vdomA interfaces - web-based manager

1. Go to *Global > Network > Interfaces*.
2. Select *Edit* on the port1 interface.
3. Enter the following information and select *OK*:

<b>Virtual Domain</b>	vdomA
<b>Addressing mode</b>	Manual
<b>IP/Netmask</b>	172.20.201.2/255.255.255.0

4. Select *Edit* on the port4 interface.
5. Enter the following information and select *OK*:

<b>Virtual Domain</b>	vdomA
<b>Addressing mode</b>	Manual
<b>IP/Netmask</b>	10.11.101.2/255.255.255.0

## To configure the vdomA interfaces - CLI

```
config global
 config system interface
 edit port1
 set vdom vdomA
 set mode static
 set ip 172.20.201.2 255.255.255.0
 next
 edit port4
 set vdom ABCdomain
 set mode static
 set ip 10.11.101.2 255.255.255.0
 end
 end
end
```

## Configuring the vdomB interfaces

The vdomB VDOM uses two FortiGate unit interfaces: port2 and port3.

The port2 interface connects the Company B internal network to the FortiGate unit, and shares the internal network subnet of 10.12.101.0/255.255.255.0.

The port3 interface connects the FortiGate unit to ISP B and the Internet. It shares the ISP B subnet of 192.168.201.0/255.255.255.0.

## To configure the vdomB interfaces - web-based manager

1. Go to *Global > Network > Interfaces*.
2. Select *Edit* on the port3 interface.
3. Enter the following information and select *OK*:

<b>Virtual domain</b>	vdomB
<b>Addressing mode</b>	Manual
<b>IP/Netmask</b>	192.168.201.2/255.255.255.0

4. Select *Edit* on the port2 interface.
5. Enter the following information and select *OK*:

<b>Virtual domain</b>	vdomB
<b>Addressing mode</b>	Manual
<b>IP/Netmask</b>	10.12.101.2/255.255.255.0

## To configure the vdomB interfaces - CLI

```
config global
 config system interface
 edit port3
 set vdom vdomB
 set mode static
 set ip 192.168.201.2 255.255.255.0
 next
 edit port2
 set vdom vdomB
 set mode static
 set ip 10.12.101.2 255.255.255.0
 end
```

## Configuring the vdomA VDOM

With the VDOMs created and the ISPs connected, the next step is to configure the vdomA VDOM.

Configuring the vdomA includes the following:

- [Adding vdomA firewall addresses](#)
- [Adding the vdomA security policy](#)
- [Adding the vdomA default route](#)

### Adding vdomA firewall addresses

You need to define the addresses used by Company A's internal network for use in security policies. This internal network is the 10.11.101.0/255.255.255.0 subnet.

The FortiGate unit provides one default address, "all", that you can use when a security policy applies to all addresses as the source or destination of a packet.

### To add the vdomA firewall addresses - web-based manager

1. In *Virtual Domains*, select *vdomA*.
2. Go to *Firewall Objects > Address > Address*.
3. Select *Create New*.
4. Enter the following information and select *OK*:

<b>Address Name</b>	Ainternal
<b>Type</b>	Subnet / IP Range
<b>Subnet / IP Range</b>	10.11.101.0/255.255.255.0
<b>Interface</b>	port4

### To add the ABCdomain VDOM firewall addresses - CLI

```
config vdom
 edit vdomA
 config firewall address
 edit Ainternal
 set type ipmask
 set subnet 10.11.101.0 255.255.255.0
 end
 end
 end
```

### Adding the vdomA security policy

You need to add the vdomA security policy to allow traffic from the internal network to reach the external network, and from the external network to internal as well. You need two policies for this domain.

#### To add the vdomA security policy - web-based manager

1. In *Virtual Domains*, select *vdomA*.
2. Go to *Policy > Policy*.
3. Select *Create New*.
4. Enter the following information and select *OK*:

<b>Source Interface/Zone</b>	port4
<b>Source Address</b>	Ainternal
<b>Destination Interface/Zone</b>	port1
<b>Destination Address</b>	all
<b>Schedule</b>	Always
<b>Service</b>	ANY
<b>Action</b>	ACCEPT

5. Select *Create New*.
6. Enter the following information and select *OK*:

<b>Source Interface/Zone</b>	port1
<b>Source Address</b>	all
<b>Destination Interface/Zone</b>	port4
<b>Destination Address</b>	Ainternal
<b>Schedule</b>	Always
<b>Service</b>	ANY
<b>Action</b>	ACCEPT

### To add the vdomA security policy - CLI

```
config vdom
 edit vdomA
 config firewall policy
 edit 1
 set srcintf port4
 set srcaddr Ainternal
 set dstintf port1
 set dstaddr all
 set schedule always
 set service ANY
 set action accept
 set status enable
 next
 edit 2
 set srcintf port1
 set srcaddr all
 set dstintf port4
 set dstaddr Ainternal
 set schedule always
 set service ANY
 set action accept
 set status enable
 end
 end
```

### Adding the vdomA default route

You also need to define a default route to direct packets from the Company A internal network to ISP A. Every VDOM needs a default static route, as a minimum, to handle traffic addressed to external networks such as the Internet.

The administrative distance should be set slightly higher than other routes. Lower admin distances will get checked first, and this default route will only be used as a last resort.

### To add a default route to the vdomA - web-based manager

1. For *Virtual Domains*, select *vdomA*
2. Go to *Router > Static > Static Routes*.
3. Select *Create New*.
4. Enter the following information and select *OK*:

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Device</b>	port1
<b>Gateway</b>	172.20.201.7
<b>Distance</b>	20

### To add a default route to the vdomA - CLI

```
config vdom
 edit vdomA
 config router static
 edit 1
 set device port1
 set gateway 172.20.201.7
 end
 end
```

## Configuring the vdomB VDOM

In this example, the vdomB VDOM is used for Company B. Firewall and routing settings are specific to a single VDOM.

vdomB includes the FortiGate port2 interface to connect to the Company B internal network, and the FortiGate port3 interface to connect to ISP B. Security policies are needed to allow traffic from port2 to external and from external to port2 interfaces.

This section includes the following topics:

- [Adding the vdomB firewall address](#)
- [Adding the vdomB security policy](#)
- [Adding a default route to the vdomB VDOM](#)

### Adding the vdomB firewall address

You need to define addresses for use in security policies. In this example, the vdomB VDOM needs an address for the port2 interface and the “all” address.

#### To add the vdomB firewall address - web-based manager

1. In *Virtual Domains*, select *vdomB*.
2. Go to *Firewall Objects > Address > Address*.
3. Select *Create New*.
4. Enter the following information and select *OK*:

<b>Address Name</b>	Binternal
<b>Type</b>	Subnet / IP Range
<b>Subnet / IP Range</b>	10.12.101.0/255.255.255.0
<b>Interface</b>	port2

#### To add the vdomB firewall address - CLI

```
config vdom
 edit vdomB
 config firewall address
 edit Binternal
 set type ipmask
 set subnet 10.12.101.0 255.255.255.0
 end
 end
end
```



## Adding the vdomB security policy

You also need a security policy for the Company B domain. In this example, the security policy allows all traffic.

### To add the vdomB security policy - web-based manager

1. Log in with a super\_admin account.
2. In *Virtual Domains*, select vdomB.
3. Go to *Policy > Policy*.
4. Select *Create New*.
5. Enter the following information and select *OK*:

<b>Source Interface/Zone</b>	port2
<b>Source Address</b>	Binternal
<b>Destination Interface/Zone</b>	port3
<b>Destination Address</b>	all
<b>Schedule</b>	Always
<b>Service</b>	ANY
<b>Action</b>	ACCEPT

6. Select *Create New*.
7. Enter the following information and select *OK*:

<b>Source Interface/Zone</b>	port3
<b>Source Address</b>	all
<b>Destination Interface/Zone</b>	port2
<b>Destination Address</b>	Binternal
<b>Schedule</b>	Always
<b>Service</b>	ANY
<b>Action</b>	ACCEPT

### To add the vdomB security policy - CLI

```
config vdom
 edit vdomB
 config firewall policy
 edit 1
 set srcintf port2
 set dstintf port3
 set srcaddr Binternal
 set dstaddr all
 set schedule always
 set service ANY
 set action accept
 set status enable
 edit 1
 set srcintf port3
 set dstintf port2
 set srcaddr all
 set dstaddr Binternal
 set schedule always
 set service ANY
 set action accept
 set status enable
 end
 end
```

### Adding a default route to the vdomB VDOM

You need to define a default route to direct packets to ISP B.

#### To add a default route to the vdomB VDOM - web-based manager

1. Log in as the super\_admin administrator.
2. In *Virtual Domains*, select vdomB.
3. Go to *Router > Static > Static Routes*.
4. Select *Create New*.
5. Enter the following information and select *OK*:

<b>Destination IP/Mask</b>	0.0.0.0/0.0.0.0
<b>Device</b>	port3
<b>Gateway</b>	192.168.201.7
<b>Distance</b>	20

## To add a default route to the vdomB VDOM - CLI

```
config vdom
 edit vdomB
 config router static
 edit 1
 set dst 0.0.0.0/0
 set device external
 set gateway 192.168.201.7
 end
 end
 end
```

## Testing the configuration

Once you have completed configuration for both company VDOMs, you can use diagnostic commands, such as `tracert` in Windows, to test traffic routed through the FortiGate unit. Alternately, you can use the `tracert` command on a Linux system with similar output.

Possible errors during the traceroute test are:

- “\*\*\*Request timed out” - the trace was not able to make the next connection towards the destination fast enough
- “Destination host unreachable” - after a number of timed-out responses the trace will give up

Possible reasons for these errors are bad connections or configuration errors.

For additional troubleshooting, see [“Troubleshooting Virtual Domains” on page 2440](#).

## Testing traffic from the internal network to the ISP

In this example, a route is traced from the Company A internal network to ISP A. The test was run on a Windows PC with an IP address of 10.11.101.55.

The output here indicates three hops between the source and destination, the IP address of each hop, and that the trace was successful.

From the Company A internal network, access a command prompt and enter this command:

```
C:\>tracert 172.20.201.7
Tracing route to 172.20.201.7 over a maximum of 30 hops:
 1 <10 ms <10 ms <10 ms 10.11.101.2
 2 <10 ms <10 ms <10 ms 172.20.201.2
 3 <10 ms <10 ms <10 ms 172.20.201.7

Trace complete.
```

# Virtual Domains in Transparent mode

In Transparent mode, the FortiGate unit behaves like a layer-2 bridge but can still provide services such as antivirus scanning, web filtering, spam filtering and intrusion protection to traffic. There are some limitations in Transparent mode in that you cannot use SSL VPN, PPTP/L2TP VPN, DHCP server, or easily perform NAT on traffic. The limits in Transparent mode apply to IEEE 802.1Q VLAN trunks passing through the unit.

VDOMs can each be configured to operate either in Transparent or NAT/Route operation mode, with each VDOM behaving like a separate FortiGate unit operating in the respective mode. VLANs configured on a VDOM in Transparent mode are the same as VLANs configured on the FortiGate unit when VDOMs are disabled.

This chapter includes the following sections:

- [Transparent operation mode](#)
- [Configuring VDOMs in Transparent mode](#)
- [Example of VDOMs in Transparent mode](#)

## Transparent operation mode

In transparent mode, the FortiGate unit becomes a layer-2 IP forwarding bridge. This means that Ethernet frames are forwarded based on destination MAC address, and no other routing is performed. All incoming traffic that is accepted by the firewall, is broadcast out on all interfaces.

In transparent mode the FortiGate unit is a forwarding bridge, not a switch. A switch can develop a port table and associated MAC addresses, so that it can bridge two ports to deliver the traffic instead of broadcasting to all ports. In transparent mode, the FortiGate unit does not follow this switch behavior, but instead is the forwarding bridge that broadcasts all packets out over all interfaces, subject to security policies.

Features such as broadcast domains, forwarding domains, and STP apply to both FortiGate units and VDOMs in Transparent mode.

### Broadcast domains

A broadcast domain is a network segment in which any network equipment can transmit data directly to another device without going through a routing device. All the devices share the same subnet. The subnets are separated by layer-3 devices, such as routers, that can forward traffic from one broadcast domain to the next.

Broadcast domains are important to transparent mode FortiGate units because the broadcast domain is the limit of where the FortiGate unit can forward packets when it is in transparent mode.

### Forwarding domains

Address Resolution Protocol (ARP) packets are vital to communication on a network, and ARP support is enabled on FortiGate unit interfaces by default. Normally you want ARP packets to pass through the FortiGate unit. However, in Transparent mode ARP packets arriving on one interface are sent to all other interfaces including VLANs giving the appearance of duplicates of the same MAC address on different interfaces. Some layer-2 switches become unstable when

they detect these duplicate MAC addresses. Unstable switches may become unreliable or reset and cause network traffic to slow down considerably.

When you are using VLANs in Transparent mode, the solution to the duplicate MAC address issue is to use the `forward-domain` CLI command. This command tags VLAN traffic as belonging to a particular collision group, and only VLANs tagged as part of that collision group receive that traffic—it is like an additional set of VLANs. By default, all interfaces and VLANs are part of forward-domain collision group 0.

To assign VLAN 200 to collision group 2, VLAN 300 to collision group 3, and all other interfaces to stay in the default collision group 0 enter the following CLI commands:

```
config system interface
 edit vlan200
 set vlanid 200
 set forward_domain 2
 next
 edit vlan300
 set vlanid 300
 set forward_domain 3
 next
end
```

When using forwarding domains, you may experience connection issues with layer-2 traffic, such as ping, if your network configuration has

- packets going through the FortiGate unit in Transparent mode multiple times,
- more than one forwarding domain (such as incoming on one forwarding domain and outgoing on another)
- IPS and AV enabled.

## Spanning Tree Protocol

VDOMs and FortiGate units do not participate in the Spanning Tree Protocol (STP). STP is an IEEE 802.1 protocol that ensures there are no layer-2 loops on the network. Loops are created when there is more than one route for traffic to take and that traffic is broadcast back to the original switch. This loop floods the network with traffic, quickly reducing available bandwidth to zero.

If you use your VDOM or FortiGate unit in a network topology that relies on STP for network loop protection, you need to make changes to your FortiGate configuration. Otherwise, STP recognizes your FortiGate unit as a blocked link and forwards the data to another path. By default, your FortiGate unit blocks STP as well as other non-IP protocol traffic. Using the CLI, you can enable forwarding of STP and other layer-2 protocols through the interface. In this example, layer-2 forwarding is enabled on the port2 interface:

```
config global
 config system interface
 edit port2
 set l2forward enable
 set stpforward enable
 next
 end
```

There are different CLI commands to allow other common layer-2 protocols such as IPX, PPTP or L2TP on the network. For more information, see the [FortiOS CLI Reference](#).

## Differences between NAT/Route and Transparent mode

The differences between NAT/Route mode and Transparent mode include:

**Table 110:**Differences between NAT/Route and Transparent modes

Features	NAT/Route mode	Transparent mode
Specific Management IP address required	No	Yes
Perform Network Address Translation (NAT)	Yes	Yes
Stateful packet inspection	Yes	Yes
Layer-2 forwarding	Yes	Yes
Layer-3 routing	Yes	No
Unicast Routing / Policy Based routing	Yes	No
DHCP server	Yes	No
IPsec VPN	Yes	Yes
PPTP/L2TP VPN	Yes	No
SSL VPN	Yes	No
Security features	Yes	Yes
VLAN support	Yes	Yes - limited to VLAN trunks.
Ping servers (dead gateway detection)	Yes	No

To provide administrative access to a FortiGate unit or VDOM in Transparent mode, you must define a management IP address and a gateway. This step is not required in NAT/Route mode where you can access the FortiGate unit through the assigned IP address of any interface where administrative access is permitted.

If you incorrectly set the Transparent mode management IP address for your FortiGate unit, you will be unable to access your unit through the web-based manager. In this situation, you will need to connect to the FortiGate unit using the console cable and change the settings so you can access the unit. Alternately, if your unit has an LCD panel, you can change the operation mode and interface information through the LCD panel.

## Operation mode differences in VDOMs

A VDOM, such as root, can have a maximum of 255 interfaces in Network Address Translation (NAT) mode or Transparent mode. This includes VLANs, other virtual interfaces, and physical interfaces. To have more than a total of 255 interfaces configured, you need multiple VDOMs with multiple interfaces on each.

In Transparent mode without VDOMs enabled, all interfaces on the FortiGate unit act as a bridge — all traffic coming in on one interface is sent back out on all the other interfaces. This effectively turns the FortiGate unit into a two interface unit no matter how many physical interfaces it has. When VDOMs are enabled, this allows you to determine how many interfaces to assign to a VDOM running in Transparent mode. If there are reasons for assigning more than

two interfaces based on your network topology, you are able to. However, the benefit of VDOMs in this case is that you have the functionality of Transparent mode, but you can use interfaces for NAT/Route traffic as well.

You can add more VDOMs to separate groups of VLAN subinterfaces. When using a FortiGate unit to serve multiple organizations, this configuration simplifies administration because you see only the security policies and settings for the VDOM you are configuring. For information on adding and configuring virtual domains, see [“Benefits of Virtual Domains” on page 2332](#).

One essential application of VDOMs is to prevent problems caused when a FortiGate unit is connected to a layer-2 switch that has a global MAC table. FortiGate units normally forward ARP requests to all interfaces, including VLAN subinterfaces. It is then possible for the switch to receive duplicate ARP packets on different VLANs. Some layer-2 switches reset when this happens. As ARP requests are only forwarded to interfaces in the same VDOM, you can solve this problem by creating a VDOM for each VLAN. For a configuration example, see [“Example of VDOMs in Transparent mode” on page 2384](#).

## Configuring VDOMs in Transparent mode

In Transparent mode, your FortiGate unit becomes a layer-2 bridge — any traffic coming in on one port is broadcast out on all the other ports. If your FortiGate unit has many interfaces, this is not the best use of those interfaces. VDOMs can limit Transparent mode to only a few interfaces while allowing the rest of the FortiGate unit to remain in NAT/Route mode.

The essential steps to configure your FortiGate unit to work with VLANs in Transparent mode are:

- [Switching to Transparent mode](#)
- [Adding VLAN subinterfaces](#)
- [Creating security policies](#).

You can also configure the security profiles that manage antivirus scanning, web filtering and spam filtering.

In Transparent mode, you can access the FortiGate web-based manager by connecting to an interface configured for administrative access and using HTTPS to access the management IP address. On the FortiGate unit used for examples in this guide, administrative access is enabled by default on the internal interface and the default management IP address is 10.11.0.1.

### Switching to Transparent mode

A VDOM is in NAT/Route mode by default when it is created. You must switch it to Transparent mode, and add a management IP address so you can access the VDOM from your management computer.



Before applying the change to Transparent mode, ensure the VDOM has administrative access on the selected interface, and that the selected management IP address is reachable on your network.

---

#### To switch the `tpVDOM` VDOM to Transparent mode - web-based manager

1. Go to *Global > VDOM > VDOM*.
2. Edit the *tpVDOM*.
3. Select *Transparent* for *Operation mode*.

4. Enter the management IP/Netmask.

The IP address must be accessible to the subnet where the management computer is located. For example 10.11.0.99/255.255.255.0 will be able to access the 10.11.0.0 subnet.

5. Select *Apply*.

When you select *Apply*, the FortiGate unit will log you out. When you log back in, the VDOM will be in Transparent mode.

**To switch the tpVDOM VDOM to Transparent mode - CLI**

```
config vdom
 edit tpVDOM
 config system settings
 set opmode transparent
 set mangeip 10.11.0.99 255.255.255.0
 end
 end
```

## Adding VLAN subinterfaces

There are a few differences when adding VLANs in Transparent mode compared to NAT/Route mode.

In Transparent mode, VLAN traffic is trunked across the VDOM. That means VLAN traffic cannot be routed, changed, or inspected. For this reason when you assign a VLAN to a Transparent mode VDOM, you will see the *Addressing Mode* section of the interface configuration disappear in from the web-based manager. It is because with no routing, inspection, or any activities able to be performed on VLAN traffic the VDOM simply re-broadcasts the VLAN traffic. This requires no addressing.

Also any routing related features such as dynamic routing or Virtual Router Redundancy Protocol (VRRP) are not available in Transparent mode for any interfaces.

## Creating security policies

Security policies permit communication between the FortiGate unit's network interfaces based on source and destination IP addresses. Typically you will also limit communication to desired times and services for additional security.

In Transparent mode, the FortiGate unit performs antivirus and antispam scanning on each packet as it passes through the unit. You need security policies to permit packets to pass from the VLAN interface where they enter the unit to the VLAN interface where they exit the unit. If there are no security policies configured, no packets will be allowed to pass from one interface to another. For more information, see the [FortiGate Administration Guide](#), or [FortiGate Fundamentals Guide](#).

## Example of VDOMs in Transparent mode

In this example, the FortiGate unit provides network protection to two organizations — Company A and Company B. Each company has different policies for incoming and outgoing traffic, requiring three different security policies and protection profiles.

VDOMs are not required for this configuration, but by using VDOMs the profiles and policies can be more easily managed on a per-VDOM basis either by one central administrator or separate administrators for each company. Also future expansion is simply a matter of adding additional VDOMs, whilst not disrupt the existing VDOMs.



For this example, firewalls are only included to deal with web traffic. This is to provide an example without making configuration unnecessarily complicated.

This example includes the following sections:

- [Network topology and assumptions](#)
- [General configuration steps](#)
- [Configuring common items](#)
- [Creating virtual domains](#)
- [Configuring the Company\\_A VDOM](#)
- [Configuring the Company\\_B VDOM](#)
- [Configuring the VLAN switch and router](#)
- [Testing the configuration](#)

## Network topology and assumptions

Each organization's internal network consists of a different range of IP addresses:


- 10.11.0.0/255.255.0.0 for Company A.
- 10.12.0.0/255.255.0.0 for Company B.

For the procedures in this section, it is assumed that you have enabled VDOM configuration on your FortiGate unit. For more information, see [“Enabling and accessing Virtual Domains” on page 2334](#).

The VDOM names are similar to the company names for easy recognition. The root VDOM cannot be renamed and is not used in this example.

Interfaces used in this example are port1 and port2. Some FortiGate models may not have interfaces with these names. port1 is an external interface. port2 is an internal interface.

**Figure 337:**VLAN and VDOM Transparent example network topology



## General configuration steps

The following steps summarize the configuration for this example. For best results, follow the procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

1. [Configuring common items](#)
2. [Creating virtual domains](#)
3. [Configuring the Company\\_A VDOM](#)
4. [Configuring the Company\\_B VDOM](#)
5. [Configuring the VLAN switch and router](#)
6. [Testing the configuration](#)

## Configuring common items

Both VDOMs require you configure security profiles. These will be configured the same way, but need to be configured in both VDOMs.

The relaxed profile allows users to surf websites they are not allowed to visit during normal business hours. Also a quota is in place to restrict users to one hour of access to these websites to ensure employees do not take long and unproductive lunches.

### To create a strict web filtering profile - web-based manager

1. Go to the proper VDOM, and select *Security Profiles > Web Filter > Profile*.
2. Select *Create New*.
3. Enter *strict* for the *Name*.
4. Expand FortiGuard Web Filtering, and select block for all Categories except Business Oriented, and Other.
5. Block all Classifications except Cached Content, and Image Search.
6. Ensure *FortiGuard Quota* for all Categories and Classifications is Disabled.
7. Select *OK*.

### To create a strict web filtering profile - CLI

```
config vdom
 edit <vdom_name>
 config webfilter profile
 edit strict
 config ftgd-wf
 set allow g07 g08 g21 g22 c01 c03
 set deny g01 g02 g03 g04 g05 g06 c02 c04 c05 c06 c07
 end
 set web-ftgd-err-log enable
 end
 end
 end
```

### To create a relaxed web filtering profile - web-based manager

1. Go to the proper VDOM, and select *Security Profiles > Web Filter > Profile*.
2. Select *Create New*.
3. Enter *relaxed* for the *Name*.
4. Expand FortiGuard Web Filtering, and select block for Potentially Security Violating Category, and Spam URL Classification.

5. Enable FortiGuard Quotas to allow 1 hour for all allowed Categories and Classifications.

## Creating virtual domains

The FortiGate unit supports 10 virtual domains. Root is the default VDOM. It cannot be deleted or renamed. The root VDOM is not used in this example. New VDOMs are created for Company A and Company B

### To create the virtual domains - web-based manager

1. With VDOMs enabled, select *System > VDOM > VDOM*.
2. Select *Create New*.
3. Enter *Company\_A* for Name, and select *OK*.
4. Select *Create New*.
5. Enter *Company\_B* for Name, and select *OK*.

### To create the virtual domains - CLI

```
config system vdom
 edit Company_A
 next
 edit Company_B
end
```

## Configuring the Company\_A VDOM

This section describes how to add VLAN subinterfaces and configure security policies for the Company\_A VDOM.

This section includes the following topics:

- [Adding VLAN subinterfaces](#)
- [Creating the Lunch schedule](#)
- [Configuring Company\\_A firewall addresses](#)
- [Creating Company\\_A security policies](#)

### Adding VLAN subinterfaces

You need to create a VLAN subinterface on the port2 interface and another one on the port1 interface, both with the same VLAN ID.

### To add VLAN subinterfaces - web-based manager

1. Go to *System > Network > Interfaces*.
2. Select *Create New*.
3. Enter the following information and select *OK*:

<b>Name</b>	VLAN_100_int
<b>Interface</b>	port2
<b>VLAN ID</b>	100
<b>Virtual Domain</b>	Company_A

4. Select *Create New*.

5. Enter the following information and select *OK*:

<b>Name</b>	VLAN_100_ext
<b>Interface</b>	port1
<b>VLAN ID</b>	100
<b>Virtual Domain</b>	Company_A

#### To add the VLAN subinterfaces - CLI

```
config system interface
 edit VLAN_100_int
 set interface port2
 set vlanid 100
 set vdom Company_A
 next
 edit VLAN_100_ext
 set interface port1
 set vlanid 100
 set vdom Company_A
end
```

### Creating the Lunch schedule

Both organizations have the same lunch schedule, but only Company A has relaxed its security policy to allow employees more freedom in accessing the Internet during lunch. Lunch schedule will be Monday to Friday from 11:45am to 2:00pm (14:00).

#### To create a recurring schedule for lunchtime - web-based manager

1. In *Company\_A* VDOM, go to *Firewall Objects > Schedule > Recurring*.
2. Select *Create New*.
3. Enter *Lunch* as the name for the schedule.
4. Select *Mon, Tues, Wed, Thu, and Fri*.
5. Set the *Start* time as 11:45 and set the *Stop* time as 14:00.
6. Select *OK*.

#### To create a recurring schedule for lunchtime - CLI

```
config vdom
 edit Company_A
 config firewall schedule recurring
 edit Lunch
 set day monday tuesday wednesday thursday friday
 set start 11:45
 set end 14:00
 end
 end
```

### Configuring Company\_A firewall addresses

For Company A, its networks are all on the 10.11.0.0 network, so restricting addresses to that domain provides added security.

### To configure Company\_A firewall addresses - web-based manager

1. In the Company\_A VDOM, go to *Firewall Objects > Address > Address*.
2. Select *Create New*.
3. Enter *CompanyA* in the *Address Name* field.
4. Type *10.11.0.0/255.255.0.0* in the *Subnet / IP Range* field.
5. Select *OK*.

### To configure vdomA firewall addresses - CLI

```
config firewall address
 edit CompanyA
 set type ipmask
 set subnet 10.11.0.0 255.255.0.0
 end
```

## Creating Company\_A security policies

A security policy can include varying levels of security feature protection. This example only deals with web filtering. The following security policies use the custom security `strict` and `relaxed` profiles configured earlier. See [“Configuring common items” on page 2386](#).

For these security policies, we assume that all protocols will be on their standard ports, such as port 80 for http traffic. If the ports are changed, such as using port 8080 for http traffic, you will have to create custom services for protocols with non-standard ports, and assign them different names.

The firewalls configured in this section are:

- internal to external — always deny all
- external to internal — always deny all
- internal to external — always allow all, security features - web filtering: strict
- internal to external — Lunch allow all, security features - web filtering:relaxed

Security policies allow packets to travel between the internal VLAN\_100 interface to the external interface subject to the restrictions of the protection profile. Entering the policies in this order means the last one configured is at the top of the policy list, and will be checked first. This is important because the policies are arranged so if one does not apply the next is checked until the end of the list.

### To configure Company\_A security policies - web-based manager

1. Go to *Policy > Policy*.
2. Select *Create New*.

3. Enter the following information and select *OK*:

<b>Source Interface/Zone</b>	VLAN_100_int
<b>Source Address</b>	CompanyA
<b>Destination Interface/Zone</b>	VLAN_100_ext
<b>Destination Address</b>	all
<b>Schedule</b>	always
<b>Service</b>	all
<b>Action</b>	DENY

This policy is a catch all for outgoing traffic to ensure that if it doesn't match any of the other policies, it will not be allowed. This is standard procedure.

4. Select *Create New*.  
5. Enter the following information and select *OK*:

<b>Source Interface/Zone</b>	VLAN_100_ext
<b>Source Address</b>	all
<b>Destination Interface/Zone</b>	VLAN_100_int
<b>Destination Address</b>	CompanyA
<b>Schedule</b>	always
<b>Service</b>	all
<b>Action</b>	DENY

This policy is a catch all for incoming traffic to ensure that if it doesn't match any of the other policies, it will not be allowed. This is standard procedure.

6. Select *Create New*.

7. Enter the following information and select *OK*:

<b>Source Interface/Zone</b>	VLAN_100_int
<b>Source Address</b>	CompanyA
<b>Destination Interface/Zone</b>	VLAN_100_ext
<b>Destination Address</b>	all
<b>Schedule</b>	always
<b>Service</b>	all
<b>Action</b>	ACCEPT
<b>Security Features</b>	Enable
<b>Web Filtering</b>	strict

This policy enforces strict scanning at all times, while allowing all traffic. It ensures company policies are met for network security.

8. Select *Create New*.  
 9. Enter the following information and select *OK*:

<b>Source Interface/Zone</b>	VLAN_100_int
<b>Source Address</b>	CompanyA
<b>Destination Interface/Zone</b>	VLAN_100_ext
<b>Destination Address</b>	all
<b>Schedule</b>	Lunch
<b>Service</b>	all
<b>Action</b>	ACCEPT
<b>Security Features</b>	enable
<b>Web Filtering</b>	relaxed

This policy provides relaxed protection during lunch hours — going from strict down to scan for protocol options and web filtering. AntiVirus and Email Filtering remain at strict for security — relaxing them would not provide employees additional access to the Internet and it would make the company vulnerable.

10. Verify that the policies entered appear in the list with the last policy (lunch) at the top, and the first policy (deny all) at the bottom. Otherwise traffic will not flow as expected.

## To configure Company\_A security policies - CLI

```
config vdom
 edit Company_A
 config firewall policy
 edit 1
 set srcintf VLAN_100_int
 set dstintf VLAN_100_ext
 set srcaddr all
 set dstaddr all
 set action accept
 set schedule Lunch
 set webfiltering relaxed
 next
 edit 3
 set srcintf VLAN_100_int
 set dstintf VLAN_100_ext
 set srcaddr all
 set dstaddr all
 set action accept
 set schedule BusinessDay
 set service HTTP
 set profile_status enable
 set profile BusinessOnly
 end
 end
```

## Configuring the Company\_B VDOM

This section describes how to add VLAN subinterfaces and configure security policies for the Company B VDOM.

This section includes the following topics:

- [Adding VLAN subinterfaces](#)
- [Creating Company\\_B service groups](#)
- [Configuring Company\\_B firewall addresses](#)
- [Configuring Company\\_B security policies](#)

### Adding VLAN subinterfaces

You need to create a VLAN subinterface on the internal interface and another one on the external interface, both with the same VLAN ID.

#### To add VLAN subinterfaces - web-based manager

1. Go to *System > Network > Interfaces*.
2. Select *Create New*.
3. Enter the following information and select *OK*:

<b>Name</b>	VLAN_200_int
<b>Interface</b>	port2



<b>VLAN ID</b>	200
<b>Virtual Domain</b>	Company_B

4. Select *Create New*.
5. Enter the following information and select *OK*:

<b>Name</b>	VLAN_200_ext
<b>Interface</b>	port1
<b>VLAN ID</b>	200
<b>Virtual Domain</b>	Company_B

### To add the VLAN subinterfaces - CLI

```
config system interface
 edit VLAN_200_int
 set interface internal
 set vlanid 200
 set vdom Company_B
 next
 edit VLAN_200_ext
 set interface external
 set vlanid 200
 set vdom Company_B
end
```

### Creating Company\_B service groups

Company\_B does not want its employees to use any online chat software except NetMeeting, which the company uses for net conferencing. To simplify the creation of a security policy for this purpose, you create a service group that contains all of the services you want to restrict. A security policy can manage only one service or one group.

#### To create a chat service group - web-based manager

1. Go to *Firewall Objects > Service > Groups*.
2. Select *Create New*.
3. Enter *Chat* in the *Group Name* field.
4. For each of IRC, AOL, SIP-MSNmessenger and TALK, select the service in the *Available Services* list and select the right arrow to add it to the *Members* list.

If a particular service does not appear in the *Available Services* list, see the list in *Firewall Objects > Service > Services*. Some services do not appear by default unless edited.

5. Select *OK*.

#### To create a games and chat service group - CLI

```
config firewall service group
 edit Chat
 set member IRC SIP-MSNmessenger AOL TALK
 end
```

## Configuring Company\_B firewall addresses

Company B's network is all in the 10.12.0.0 network. Security can be improved by only allowing traffic from IP addresses on that network.

### To configure Company\_B firewall address - web-based manager

1. In the Company\_B VDOM, go to *Firewall Objects > Address > Address*.
2. Select *Create New*.
3. Enter `new` in the *Address Name* field.
4. Type `10.12.0.0/255.255.0.0` in the *Subnet / IP Range* field.
5. Select *OK*.

### To configure Company\_B firewall addresses - CLI

```
config vdom
 edit Company_B
 config firewall address
 edit all
 set type ipmask
 set subnet 10.12.0.0 255.255.0.0
 end
```

## Configuring Company\_B security policies

Security policies allow packets to travel between the internal and external VLAN\_200 interfaces subject to the restrictions of the protection profile.

### To configure Company\_B security policies - web-based manager

1. Go to *Policy > Policy*.
2. Select *Create New*.
3. Enter the following information and select *OK*:

<b>Source Interface/Zone</b>	VLAN_200_int
<b>Source Address</b>	all
<b>Destination Interface/Zone</b>	VLAN_200_ext
<b>Destination Address</b>	all
<b>Schedule</b>	BusinessDay
<b>Service</b>	games-chat
<b>Action</b>	DENY

This policy prevents the use of network games or chat programs (except NetMeeting) during business hours.

4. Enter the following information and select *OK*:

<b>Source Interface/Zone</b>	VLAN_200_int
<b>Source Address</b>	all
<b>Destination Interface/Zone</b>	VLAN_200_ext
<b>Destination Address</b>	all
<b>Schedule</b>	Lunch
<b>Service</b>	HTTP
<b>Action</b>	ACCEPT
<b>Protection Profile</b>	Relaxed

This policy relaxes the web category filtering during lunch hour.

5. Select *Create New*.
6. Enter the following information and select *OK*:

<b>Source Interface/Zone</b>	VLAN_200_int
<b>Source Address</b>	all
<b>Destination Interface/Zone</b>	VLAN_200_ext
<b>Destination Address</b>	all
<b>Schedule</b>	BusinessDay
<b>Service</b>	HTTP
<b>Action</b>	ACCEPT
<b>Protection Profile</b>	BusinessOnly

This policy provides rather strict web category filtering during business hours.

7. Select *Create New*.

8. Enter the following information and select *OK*:

<b>Source Interface/Zone</b>	VLAN_200_int
<b>Source Address</b>	all
<b>Destination Interface/Zone</b>	VLAN_200_ext
<b>Destination Address</b>	all
<b>Schedule</b>	always
<b>Service</b>	ANY
<b>Action</b>	ACCEPT
<b>Protection Profile</b>	Relaxed

Because it is last in the list, this policy applies to the times and services not covered in preceding policies. This means that outside of regular business hours, the Relaxed protection profile applies to email and web browsing, and online chat and games are permitted. Company B needs this policy because its employees sometimes work overtime. The other companies in this example maintain fixed hours and do not want any after-hours Internet access.

## To configure Company\_B security policies - CLI

```
config firewall policy
 edit 1
 set srcintf VLAN_200_int
 set srcaddr all
 set dstintf VLAN_200_ext
 set dstaddr all
 set schedule BusinessDay
 set service Games
 set action deny
 next
 edit 2
 set srcintf VLAN_200_int
 set srcaddr all
 set dstintf VLAN_200_ext
 set dstaddr all
 set action accept
 set schedule Lunch
 set service HTTP
 set profile_status enable
 set profile Relaxed
 next
 edit 3
 set srcintf VLAN_200_int
 set srcaddr all
 set dstintf VLAN_200_ext
 set dstaddr all
 set action accept
 set schedule BusinessDay
 set service HTTP
 set profile_status enable
 set profile BusinessOnly
 next
 edit 4
 set srcintf VLAN_200_int
 set srcaddr all
 set dstintf VLAN_200_ext
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
 set profile_status enable
 set profile Relaxed
end
```

## Configuring the VLAN switch and router

The Cisco switch is the first VLAN device internal passes through, and the Cisco router is the last device before the Internet or ISP.

This section includes the following topics:

- [Configuring the Cisco switch](#)
- [Configuring the Cisco router](#)

## Configuring the Cisco switch

On the Cisco Catalyst 2900 ethernet switch, you need to define the VLANs 100, 200 and 300 in the VLAN database, and then add configuration files to define the VLAN subinterfaces and the 802.1Q trunk interface.

Add this file to Cisco VLAN switch:

```
!
interface FastEthernet0/1
 switchport access vlan 100
!
interface FastEthernet0/5
 switchport access vlan 300
!
interface FastEthernet0/6
 switchport trunk encapsulation dot1q
 switchport mode trunk
!
```

Switch 1 has the following configuration:

<b>Port 0/1</b>	VLAN ID 100
<b>Port 0/3</b>	VLAN ID 200
<b>Port 0/6</b>	802.1Q trunk

## Configuring the Cisco router

The configuration for the Cisco router in this example is the same as in the basic example, except we add VLAN\_300. Each of the three companies has its own subnet assigned to it.

The IP addresses assigned to each VLAN on the router are the gateway addresses for the VLANs. For example, devices on VLAN\_100 would have their gateway set to 10.11.0.1/255.255.0.0.

```
!
interface FastEthernet0/0
 switchport trunk encapsulation dot1q
 switchport mode trunk
!
interface FastEthernet0/0.1
 encapsulation dot1Q 100
 ip address 10.11.0.1 255.255.0.0
!
interface FastEthernet0/0.3
 encapsulation dot1Q 200
 ip address 10.12.0.1 255.255.0.0
!
```

The router has the following configuration:

<b>Port 0/0.1</b>	VLAN ID 100
<b>Port 0/0.3</b>	VLAN ID 200
<b>Port 0/0</b>	802.1Q trunk

## Testing the configuration

Use diagnostic commands, such as `tracert`, to test traffic routed through the network.

You should test traffic between the internal VLANs as well as from the internal VLANs to the Internet to ensure connectivity.

For additional troubleshooting, see [“Troubleshooting Virtual Domains” on page 2440](#).

This section includes the following topics:

- [Testing traffic from VLAN\\_100 to the Internet](#)
- [Testing traffic from VLAN\\_100 to VLAN\\_200](#)

### Testing traffic from VLAN\_100 to the Internet

In this example, a route is traced from VLANs to a host on the Internet. The route target is `www.example.com`.

From a host on VLAN\_100, access a command prompt and enter this command:

```
C:\>tracert www.example.com
Tracing route to www.example.com [208.77.188.166]
over a maximum of 30 hops:
 1 <10 ms <10 ms <10 ms 10.100.0.1
 ...
 14 172 ms 141 ms 140 ms 208.77.188.166
Trace complete.
```

The number of steps between the first and the last hop, as well as their IP addresses, will vary depending on your location and ISP. However, all successful `tracerts` to `www.example.com` will start and end with these lines.

Repeat the `tracert` for VLAN\_200.

The `tracert` for each VLAN will include the gateway for that VLAN as the first step. Otherwise, the `tracert` should be the same for each VLAN.

### Testing traffic from VLAN\_100 to VLAN\_200

In this example, a route is traced between two internal networks. The route target is a host on VLAN\_200. The Windows `tracert` command `tracert` is used.

From VLAN\_100, access a Windows command prompt and enter this command:

```
C:\>tracert 10.12.0.2
Tracing route to 10.12.0.2 over a maximum of 30 hops:
 1 <10 ms <10 ms <10 ms 10.100.0.1
 2 <10 ms <10 ms <10 ms 10.12.0.2
Trace complete.
```

You can repeat this for different routes in the topology. In each case the IP addresses will be the gateway for the starting VLAN, and the end point at the ending VLAN.

# Inter-VDOM routing

Inter-VDOM routing and Inter-VDOM links allow communication between VDOMs. VDOM links are virtual interfaces that connect VDOMs. A VDOM link contains a pair of interfaces with each one connected to a VDOM, and forming either end of the inter-VDOM connection.

This chapter contains the following sections:

- [Benefits of inter-VDOM routing](#)
- [Getting started with VDOM links](#)
- [Dynamic routing over inter-VDOM links](#)
- [HA virtual clusters and VDOM links](#)
- [Example of inter-VDOM routing](#)

## Benefits of inter-VDOM routing

Inter-VDOM routing has a number of advantages over independent VDOM routing. These benefits include:

- [Freed-up physical interfaces](#)
- [More speed than physical interfaces](#)
- [Continued support for secure firewall policies](#)
- [Configuration flexibility](#)

### Continued support for secure firewall policies

VDOMs help to separate traffic based on your needs. This is an important step in satisfying regulations that require proof of secure data handling. This is especially important to health, law, accounting, and other businesses that handle sensitive data every day.

By keeping things separate, traffic has to leave the FortiGate unit and re-enter to change VDOMs. This forces traffic to go through the firewall when leaving and enter through another firewall, keeping traffic secure.

With inter-VDOM routing, the need for the physical interfaces is greatly reduced. However, firewall policies still need to be in place for traffic to pass through any interface, physical or virtual, and thus provide the same level of security both internally and externally. Configuration of firewall policies is the same for inter-VDOM links as for any other interface, and your data will continue to have the high level of security.

### Configuration flexibility

A typical VDOM uses at least two interfaces, typically physical interfaces, one for internal and one for external traffic. Depending on the configuration, more interfaces may be required. This means that the maximum number of VDOMs configurable on a FortiGate unit using physical interfaces is the number of interfaces available divided by two. VLANs can increase the number by providing multiple virtual interfaces over a single physical interface, but VLANs have some limitations. Using physical interfaces for inter-VDOM communication therefore limits the number of possible configurations on your FortiGate unit.



To overcome this limitation, inter-VDOM links can be created within the FortiGate unit. Using virtual interfaces, inter-VDOM links free up the physical interfaces for external traffic. Using VDOM links on a FortiGate unit with 8 physical interfaces, you can have 4 VDOMs communicating with each other (meshed configuration) and continue to have 2 physical interfaces each for internal and external connections. This configuration would have required 20 physical interfaces without inter-VDOM routing. With inter-VDOM routing it only requires 8 physical interfaces, with the other 12 interfaces being internal VDOM links.

Inter-VDOM routing allows you to make use of [Standalone VDOM configuration](#), [Management VDOM configuration](#) and [Meshed VDOM configuration](#) without being limited by the number of physical interfaces on your FortiGate unit.

## Getting started with VDOM links

Once VDOMs are configured on your FortiGate unit, configuring inter-VDOM routing and VDOM-links is very much like creating a VLAN interface. VDOM-links are managed through the web-based manager or CLI. In the web-based manager, VDOM link interfaces are managed in the network interface list.

This section includes the following topics:

- [Viewing VDOM links](#)
- [Creating VDOM links](#)
- [Deleting VDOM links](#)
- [NAT to Transparent VDOM links](#)

### Viewing VDOM links

VDOM links are displayed on the network interface list in the web-based manager. You can view VDOM links only if you are using a super\_admin account and in global configuration.

To view the network interface list, in the Global menu go to *System > Network > Interfaces*.

**Figure 338:**Interface list displaying interface names and information

	Name	IP/Netmask	Type	Access	Administrative Status	Virtual Domain
<input type="checkbox"/>	port1 (External)	172.20.120.170 / 255.255.255.0	Physical	HTTPS,PING,SSH	<span style="color: green;">●</span>	root
<input type="checkbox"/>	port2 (Internal)	10.11.102.2 / 255.255.255.0	Physical	PING	<span style="color: green;">●</span>	root
<input type="checkbox"/>	VLAN_100_INT	10.11.101.2 / 255.255.255.0	VLAN	PING,SSH	<span style="color: green;">●</span>	Company_A
<input type="checkbox"/>	VLAN_200_INT	10.12.101.2 / 255.255.255.0	VLAN	PING,SSH	<span style="color: green;">●</span>	Company_B
<input type="checkbox"/>	port3 (DMZ)	172.20.130.10 / 255.255.255.0	Physical	PING	<span style="color: green;">●</span>	root
<input type="checkbox"/>	port4	0.0.0.0 / 0.0.0.0	Physical	PING	<span style="color: red;">●</span>	root
<input type="checkbox"/>	port5	0.0.0.0 / 0.0.0.0	Physical	HTTPS,PING,SSH	<span style="color: red;">●</span>	Company_B
<input type="checkbox"/>	port6	0.0.0.0 / 0.0.0.0	Physical	HTTPS,PING,SSH	<span style="color: red;">●</span>	Company_A
<input type="checkbox"/>	port7	0.0.0.0 / 0.0.0.0	Physical	PING	<span style="color: red;">●</span>	root
<input checked="" type="checkbox"/>	port8 (External)	-	Physical	HTTPS,PING,SSH	<span style="color: green;">●</span>	tpVDM
<input type="checkbox"/>	tpVLAN111	-	VLAN	PING	<span style="color: green;">●</span>	tpVDM
<input type="checkbox"/>	port9	0.0.0.0 / 0.0.0.0	Physical	PING	<span style="color: red;">●</span>	vdom2
<input type="checkbox"/>	port10	10.13.201.2 / 255.255.255.0	Physical	HTTPS,PING,SSH	<span style="color: green;">●</span>	root
<input type="checkbox"/>	testLink (VDOM Link)	-	VDOM Link		<span style="color: green;">●</span>	root, vdom2
<input type="checkbox"/>	testLink0	0.0.0.0 / 0.0.0.0	Pair	HTTPS,PING,SSH	<span style="color: green;">●</span>	root
<input type="checkbox"/>	testLink1	0.0.0.0 / 0.0.0.0	Pair	HTTPS,PING,SSH	<span style="color: green;">●</span>	vdom2

Annotations in the image:  
 - Arrow pointing to 'testLink': VDOM link interface  
 - Arrow pointing to 'testLink0' and 'testLink1': VDOM link pair  
 - Arrow pointing to 'tpVDM': VDOM  
 - Arrow pointing to the first column header: Description of interface

<b>Create New</b>	Select the arrow to create a new interface or VDOM link. Interface options include VLAN, Aggregate, Redundant, or loopback interfaces.  For more information, see <a href="#">“Creating VDOM links” on page 2403</a> .
<b>Edit</b>	Select to change interface configuration for the selected interface.  This option not available if no interfaces or multiple interfaces are selected.
<b>Delete</b>	Select to remove an interface from the list. One or more interfaces must be selected for this option to be available.  You cannot delete permanent physical interfaces, or any interfaces that have configuration referring to them. See <a href="#">“Deleting VDOM links” on page 2405</a> or <a href="#">“Deleting an interface” on page 2364</a> .
<b>Column Settings</b>	Select to change which information is displayed about the interfaces, and in which order the columns appear. Use to display VDOM, VLAN, and other information.
<b>Checkbox</b>	Select the checkbox for an interface to edit or delete that interface.  Select multiple interfaces to delete those interfaces.  Optionally select the check box at the top of the column to select or unselect all checkboxes.
<b>Name</b>	The name of the interface.  The name of the VDOM link ( <code>vlink1</code> ) has an expand arrow to display or hide the pair of VDOM link interfaces. For more information, see <a href="#">“Viewing VDOM links” on page 2401</a> .
<b>IP/Netmask</b>	The IP address and netmask assigned to this interface.
<b>Type</b>	The type of interface such as physical, VLAN, or VDOM link pair.
<b>Access</b>	The protocols allowed for administrators to connect to the FortiGate unit.

<b>Administrative Status</b>	The status of this interface, either set to up (active) or down (disabled).
<b>Virtual Domain</b>	The virtual domain this interface belongs to. For more information on VDOMs, see <a href="#">“Virtual Domains in NAT/Route mode” on page 2361</a> .

## Creating VDOM links

VDOM links connect VDOMs together to allow traffic to pass between VDOMs as per firewall policies. Inter-VDOM links are virtual interfaces that are very similar to VPN tunnel interfaces except inter-VDOM links do not require IP addresses. See [“IP addresses and inter-VDOM links” on page 2404](#).

To create a VDOM link, you first create the point-to-point interface, and then bind the two interface objects associated with it to the virtual domains.

In creating the point-to-point interface, you also create two additional interface objects by default. They are called `vlink10` and `vlink11` - the interface name you chose with a 1 or a 0 to designate the two ends of the link.

Once the interface objects are bound, they are treated like normal FortiGate interfaces and need to be configured just like regular interfaces.

The assumptions for this example are as follows:

- Your FortiGate unit has VDOMs enabled and you have 2 VDOMs called `customer1` and `customer2` already configured. For more information on configuring VDOMs see [“Only a super\\_admin administrator account such as the default “admin” account can create, disable, or delete VDOMs. That account can create additional administrators for each VDOM.” on page 2354](#).
- You are using a `super_admin` account

### To configure an inter-VDOM link - web-based manager

1. Go to *Global > Network > Interfaces*.
2. Select *Create New > VDOM link*, enter the following information, and select *OK*.

<b>Name</b>	vlink1 <small>(The name can be up to 11 characters long. Valid characters are letters, numbers, “-”, and “_”. No spaces are allowed.)</small>
<b>Interface #0</b>	
<b>Virtual Domain</b>	customer1
<b>IP/Netmask</b>	10.11.12.13/255.255.255.0
<b>Administrative Access</b>	HTTPS, SSL

Interface #1	
<b>Virtual Domain</b>	customer2
<b>IP/Netmask</b>	172.120.100.13/255.255.255.0
<b>Administrative Access</b>	HTTPS, SSL

### To configure an inter-VDOM link - CLI

```

config global
 config system vdom-link
 edit vlink1
 end
 config system interface
 edit vlink10
 set vdom customer1
 next
 edit vlink11
 set vdom customer2
 end

```

Once you have created and bound the interface ends to VDOMs, configure the appropriate firewall policies and other settings that you require. To confirm the inter-VDOM link was created, find the VDOM link pair and use the expand arrow to view the two VDOM link interfaces. You can select edit to change any information.

### IP addresses and inter-VDOM links

Besides being virtual interfaces, here is one main difference between inter-VDOM links and regular interfaces— default inter-VDOM links do not require IP addresses. IP addresses are not required by default because an inter-VDOM link is an internal connection that can be referred to by the interface name in firewall policies, and other system references. This introduces three possible situations with inter-VDOM links that are:

- **unnumbered** - an inter-VDOM link with no IP addresses for either end of the tunnel
- **half numbered** - an inter-VDOM link with one IP address for one end and none for the other end
- **full numbered** - an inter-VDOM link with two IP addresses, one for each end.

Not using an IP address in the configuration can speed up and simplify configuration for you. Also you will not use up all the IP addresses in your subnets if you have many inter-VDOM links.

Half or full numbered interfaces are required if you are doing NAT, either SNAT or DNAT as you need an IP number on both ends to translate between.

You can use unnumbered interfaces in static routing, by naming the interface and using 0.0.0.0 for the gateway. Running traceroute will not show the interface in the list of hops. However you can see the interface when you are sniffing packets, which is useful for troubleshooting.

## Deleting VDOM links

When you delete the VDOM link, the two link objects associated with it will also be deleted. You cannot delete the objects by themselves. The example uses a VDOM routing connection called “vlink1”. Removing vlink1 will also remove its two link objects vlink10 and vlink11.



Before deleting the VDOM link, ensure all policies, firewalls, and other configurations that include the VDOM link are deleted, removed, or changed to no longer include the VDOM link.

---

### To remove a VDOM link - web-based manager

1. Go to *Global > Network > Interfaces*.
2. Select *Delete* for the VDOM link *vlink1*.

### To remove a VDOM link - CLI

```
config global
 config system vdom-link
 delete vlink1
 end
```

For more information, see the [FortiGate CLI Reference](#).

## NAT to Transparent VDOM links

Inter-VDOM links can be created between VDOMs in NAT mode and VDOMs in Transparent mode, but it must be done through the CLI, as the VDOM link type must be changed from the default PPP to Ethernet for the two VDOMs to communicate. The below example assumes one vdom is in NAT mode and one is Transparent.



An IP address must be assigned to the NAT VDOM's interface, but no IP address should be assigned to the Transparent VDOM's interface.

---

### To configure a NAT to Transparent VDOM link - CLI

```
config global
 config system vdom-link
 edit vlink1
 set type ethernet
 end
 config system interface
 edit vlink10
 set vdom (interface 1 name)
 set ip (interface 1 ip)
 next
 edit vlink11
 set vdom (interface 2 name)
 end
```

Ethernet-type is not recommended for standard NAT to NAT inter-VDOM links, as the default PPP-type link does not require the VDOM links to have addresses, while Ethernet-type does. VDOM link addresses are explained in [“IP addresses and inter-VDOM links”](#) on page 2404.

## Inter-VDOM configurations

By using fewer physical interfaces to inter-connect VDOMs, inter-VDOM links provide you with more configuration options.

None of these configurations use VLANs to reduce the number of physical interfaces. It is generally assumed that an internal or client network will have its own internal interface and an external interface to connect to its ISP and the Internet.

These inter-VDOM configurations can use any FortiGate model with possible limitations based on the number of physical interfaces. VLANs can be used to work around these limitations.

In the following inter-VDOM diagrams, red indicates the physical FortiGate unit, grey indicate network connections external to the FortiGate unit, and black is used for inter-VDOM links and VDOMs.


This section includes the following topics:

- [Standalone VDOM configuration](#)
- [Independent VDOMs configuration](#)
- [Management VDOM configuration](#)
- [Meshed VDOM configuration](#)

## Standalone VDOM configuration

The standalone VDOM configuration uses a single VDOM on your FortiGate unit — the root VDOM that all FortiGate units have by default. This is the VDOM configuration you are likely familiar with. It is the default configuration for FortiGate units before you create additional VDOMs.

**Figure 339:**Standalone VDOM



The configuration shown in [Figure 339](#) has no VDOM inter-connections and requires no special configurations or settings.


The standalone VDOM configuration can be used for simple network configurations that only have one department or one company administering the connections, firewalls and other VDOM-dependent settings.

However, with this configuration, keeping client networks separate requires many interfaces, considerable firewall design and maintenance, and can quickly become time consuming and complex. Also, configuration errors for one client network can easily affect other client networks, causing unnecessary network downtime.

## Independent VDOMs configuration

The independent VDOMs configuration uses multiple VDOMs that are completely separate from each other. This is another common VDOM configuration.

**Figure 340:**Independent VDOMs



This configuration has no communication between VDOMs and apart from initially setting up each VDOM, it requires no special configurations or settings. Any communication between VDOMs is treated as if communication is between separate physical devices.

The independent inter-VDOM configuration can be used where more than one department or one company is sharing the FortiGate unit. Each can administer the connections, firewalls and other VDOM-dependent settings for only its own VDOM. To each company or department, it appears as if it has its own FortiGate unit. This configuration reduces the amount of firewall configuration and maintenance required by dividing up the work.


However, this configuration lacks a management VDOM for VDOMs 1, 2, and 3. This is illustrated in Figure 50. This management VDOM would enable an extra level of control for the FortiGate unit administrator, while still allowing each company or department to administer its own VDOM.



## Management VDOM configuration

In the management VDOM configuration, the root VDOM is the management VDOM. The other VDOMs are connected to the management VDOM with inter-VDOM links. There are no other inter-VDOM connections.

**Figure 341:**Management VDOM configuration



The inter-VDOM links connect the management VDOM to the other VDOMs. This does not require any physical interfaces, and the bandwidth of inter-VDOM links can be faster than physical interfaces, depending on the CPU workload.

Only the management VDOM is connected to the Internet. The other VDOMs are connected to internal networks. All external traffic is routed through the management VDOM using inter-VDOM links and firewall policies between the management VDOM and each VDOM. This ensures the management VDOM has full control over access to the Internet, including what types of traffic are allowed in both directions. There is no communication directly between the non-root VDOMs. Security is greatly increased with only one point of entry and exit. Only the management VDOM needs to be fully managed to ensure network security in this case. Each client network can manage its own configuration without compromising security or bringing down another client network.

The management VDOM configuration is ideally suited for a service provider business. The service provider administers the management VDOM with the other VDOMs as customers. These customers do not require a dedicated IT person to manage their network. The service provider controls the traffic and can prevent the customers from using banned services and prevent Internet connections from initiating those same banned services. One example of a banned service might be Instant Messaging (IM) at a company concerned about intellectual property. Another example could be to limit bandwidth used by file-sharing applications without banning that application completely. Firewall policies control the traffic between the customer VDOM and the management VDOM and can be customized for each customer.


The management VDOM configuration is limited in that the customer VDOMs have no inter-connections. In many situations this limitation is ideal because it maintains proper security. However, some configurations may require customers to communicate with each other, which would be easier if the customer VDOMs were inter-connected.

## Meshed VDOM configuration

The meshed VDOMs configuration, including partial and full mesh, has VDOMs inter-connected with other VDOMs. There is no special feature to accomplish this—they are just complex VDOM configurations.

Partial mesh means only some VDOMs are inter-connected. In a full mesh configuration, all VDOMs are inter-connected to all other VDOMs. This can be useful when you want to provide full access between VDOMs but handle traffic differently depending on which VDOM it originates from or is going to.

**Figure 342:** Meshed VDOMs



With full access between all VDOMs being possible, it is extra important to ensure proper security. You can achieve this level of security by establishing extensive firewall policies and ensuring secure account access for all administrators and users.

Meshed VDOM configurations can become complex very quickly, with full mesh VDOMs being the most complex. Ensure this is the proper solution for your situation before using this configuration. Generally, these configurations are seen as theoretical and are rarely deployed in the field.

## Dynamic routing over inter-VDOM links

BGP is supported over inter-VDOM links. Unless otherwise indicated, routing works as expected over inter-VDOM links.

If an inter-VDOM link has no assigned IP addresses to it, it may be difficult to use that interface in dynamic routing configurations. For example BGP requires an IP address to define any BGP router added to the network.

In OSPF, you can configure a router using a router ID and not its IP address. In fact, having no IP address avoids possible confusing between which value is the router ID and which is the IP address. However for that router to become adjacent with another OSPF router it will have to share the same subnet, which is technically impossible without an IP address. For this reason, while you can configure an OSPF router using an IP-less inter-VDOM link, it will likely be of limited value to you.

In RIP the metric used is hop count. If the inter-VDOM link can reach other nodes on the network, such as through a default route, then it may be possible to configure a RIP router on an inter-VDOM link. However, once again it may be of limited value due to limitations.

As stated earlier, BGP requires an IP address to define a router — an IP-less inter-VDOM link will not work with BGP.

In Multicast, you can configure an interface without using an IP address. However that interface will be unable to become an RP candidate. This limits the roles available to such an interface.

## HA virtual clusters and VDOM links

FortiGate HA is implemented by configuring two or more FortiGate units to operate as an HA cluster. To the network, the HA cluster appears to function as a single FortiGate unit, processing network traffic and providing normal security services such as firewall, VPN, IPS, virus scanning, web filtering, and spam filtering.

Virtual clustering extends HA features to provide failover protection and load balancing for a FortiGate unit operating with virtual domains. A virtual cluster consists of a cluster of two FortiGate units operating with virtual domains. Traffic on different virtual domains can be load balanced between the cluster units.

With virtual clusters (vclusters) configured, inter-VDOM links must be entirely within one vcluster. You cannot create links between vclusters, and you cannot move a VDOM that is linked into another virtual cluster. If your FortiGate units are operating in HA mode, with multiple vclusters when you create the vdom-link, the CLI command `config system vdom-link` includes an option to set which vcluster the link will be in.

### What is virtual clustering?

Virtual clustering is an extension of the FGCP for FortiGate units operating with multiple VDOMS enabled. Virtual clustering operates in active-passive mode to provide failover protection between two instances of a VDOM operating on two different cluster units. You can also operate virtual clustering in active-active mode to use HA load balancing to load balance sessions between cluster units. Alternatively, by distributing VDOM processing between the two cluster units you can also configure virtual clustering to provide load balancing by distributing sessions for different VDOMs to each cluster unit.

### Virtual clustering and failover protection

Virtual clustering operates on a cluster of two (and only two) FortiGate units with VDOMs enabled. Each VDOM creates a cluster between instances of the VDOMs on the two FortiGate units in the virtual cluster. All traffic to and from the VDOM stays within the VDOM and is processed by the VDOM. One cluster unit is the primary unit for each VDOM and one cluster unit is the subordinate unit for each VDOM. The primary unit processes all traffic for the VDOM. The subordinate unit does not process traffic for the VDOM. If a cluster unit fails, all traffic fails over to the cluster unit that is still operating.

### Virtual clustering and heartbeat interfaces

The HA heartbeat provides the same HA services in a virtual clustering configuration as in a standard HA configuration. One set of HA heartbeat interfaces provides HA heartbeat services for all of the VDOMs in the cluster. You do not have to add a heartbeat interface for each VDOM.

## Virtual clustering and HA override

For a virtual cluster configuration, override is enabled by default for both virtual clusters when you:

- Enable VDOM partitioning from the web-based manager by moving virtual domains to virtual cluster 2
- Enter `set vcluster2 enable` from the CLI config system ha command to enable virtual cluster 2.

Usually you would enable virtual cluster 2 and expect one cluster unit to be the primary unit for virtual cluster 1 and the other cluster unit to be the primary unit for virtual cluster 2. For this distribution to occur override must be enabled for both virtual clusters. Otherwise you will need to restart the cluster to force it to renegotiate.

## Virtual clustering and load balancing or VDOM partitioning

There are two ways to configure load balancing for virtual clustering. The first is to set the HA mode to active-active. The second is to configure VDOM partitioning. For virtual clustering, setting the HA Mode to active-active has the same result as active-active HA for a cluster without virtual domains. The primary unit receives all sessions and load balances them among the cluster units according to the load balancing schedule. All cluster units process traffic for all virtual domains.

Note: If override is enabled the cluster may renegotiate too often. You can choose to disable override at any time. If you decide to disable override, for best results, you should disable it for both cluster units.

In a VDOM partitioning virtual clustering configuration, the HA mode is set to active-passive. Even though virtual clustering operates in active-passive mode you can configure a form of load balancing by using VDOM partitioning to distribute traffic between both cluster units. To configure VDOM partitioning you set one cluster unit as the primary unit for some virtual domains and you set the other cluster unit as the primary unit for other virtual domains. All traffic for a virtual domain is processed by the primary unit for that virtual domain. You can control the distribution of traffic between the cluster units by adjusting which cluster unit is the primary unit for each virtual domain.

For example, you could have 4 VDOMs, two of which have a high traffic volume and two of which have a low traffic volume. You can configure each cluster unit to be the primary unit for one of the high volume VDOMs and one of the low volume VDOMs. As a result each cluster unit will be processing traffic for a high volume VDOM and a low volume VDOM, resulting in an even distribution of traffic between the cluster units. You can adjust the distribution at any time. For example, if a low volume VDOM becomes a high volume VDOM you can move it from one cluster unit to another until the best balance is achieved. From the web-based manager you configure VDOM partitioning by setting the HA mode to active-passive and distributing virtual domains between Virtual Cluster 1 and Virtual Cluster 2. You can also configure different device priorities, port monitoring, and remote link failover, for Virtual Cluster 1 and Virtual Cluster 2.

From the CLI you configure VDOM partitioning by setting the HA mode to a-p. Then you configure device priority, port monitoring, and remote link failover and specify the VDOMs to include in virtual cluster 1. You do the same for virtual cluster 2 by entering the config secondary-vcluster command.

Failover protection does not change. If one cluster unit fails, all sessions are processed by the remaining cluster unit. No traffic interruption occurs for the virtual domains for which the still functioning cluster unit was the primary unit. Traffic may be interrupted temporarily for virtual domains for which the failed unit was the primary unit while processing fails over to the still functioning cluster unit. If the failed cluster unit restarts and rejoins the virtual cluster, VDOM partitioning load balancing is restored.

## Example of inter-VDOM routing

This example shows how to configure a FortiGate unit to use inter-VDOM routing.

This section contains the follow topics:

- [Network topology and assumptions](#)
- [Creating the VDOMs](#)
- [Configuring the physical interfaces](#)
- [Configuring the VDOM links](#)
- [Configuring the firewall and Security Profile settings](#)
- [Testing the configuration](#)

### Network topology and assumptions

Two departments of a company, Accounting and Sales, are connected to one FortiGate-800 unit. To do its work, the Sales department receives a lot of email from advertising companies that would appear to be spam if the Accounting department received it. For this reason, each department has its own VDOM to keep firewall policies and other configurations separate. A management VDOM makes sense to ensure company policies are followed for traffic content.

The traffic between Accounting and Sales will be email and HTTPS only. It could use a VDOM link for a meshed configuration, but we will keep from getting too complex. With the configuration, inter-VDOM traffic will have a slightly longer path to follow than normal—from one department VDOM, through the management VDOM, and back to the other department VDOM. Since inter-VDOM links are faster than physical interfaces, this longer path should not be noticed.


Firewall policies will be in place. For added security, firewall policies will allow only valid office services such as email, web browsing, and FTP between either department and the Internet. Any additional services that are required can be added in the future.

The company uses a single ISP to connect to the Internet. The ISP uses DHCP to provide an IP address to the FortiGate unit. Both departments use the same ISP to reach the Internet.

Other assumptions for this example are as follows:

- Your FortiGate unit has interfaces labelled port1 through port4 and VDOMs are not enabled.
- You are using the super\_admin account.
- You have the FortiClient application installed.
- You are familiar with configuring interfaces, firewalls, and other common features on your FortiGate unit.

**Figure 343:**Management VDOM for two departments



## General configuration steps

This example includes the following general steps. For best results, follow the steps in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

1. [Creating the VDOMs](#)
2. [Configuring the physical interfaces](#)
3. [Configuring the VDOM links](#)
4. [Configuring the firewall and Security Profile settings](#)
5. [Testing the configuration](#)

## Creating the VDOMs

This procedure enables VDOMs and creates the Sales and Accounting VDOMs.

### To create the VDOMs - web-based manager

1. Log in as the `super_admin` administrator.
2. Go to *System > Dashboard > Status > System Information > Virtual Domain*, and select *Enable*.
3. Log in again.
4. Go to *System > VDOM > VDOM*.
5. Select *Create New*, enter `Accounting` for the VDOM Name, and select *OK*.
6. Select *Create New*, enter `Sales` for the VDOM Name, and select *OK*.

### To create the VDOMs - CLI

```
config system global
 set vdom enable
end

config system vdom
 edit Accounting
 next
 edit Sales
 next
end
```

## Configuring the physical interfaces

Next, the physical interfaces must be configured. This example uses three interfaces on the FortiGate unit - port2 (internal), port3(dmz), and port1(external). port2 and port3 interfaces each have a department's network connected. port1 is for all traffic to or from the Internet and will use DHCP to configure its IP address, which is common with many ISPs.

### To configure the physical interfaces - web-based manager

1. Go to *Global > Network > Interfaces*.
2. Select *Edit* for the port2 interface, enter the following information, and select *OK*.

<b>Alias</b>	AccountingLocal
<b>Virtual Domain</b>	Accounting
<b>Addressing mode</b>	Manual
<b>IP/Netmask</b>	172.100.1.1/255.255.0.0
<b>Administrative Access</b>	HTTPS, PING, SSH
<b>Description</b>	This is the accounting department internal interface.

3. Select *Edit* for the port3 interface, enter the following information, and select *OK*.

<b>Alias</b>	SalesLocal
<b>Virtual Domain</b>	Sales
<b>Addressing mode</b>	Manual
<b>IP/Netmask</b>	192.168.1.1/255.255.0.0
<b>Administrative Access</b>	HTTPS, PING, SSH
<b>Description</b>	This is the sales department internal interface.

4. Select *Edit* for the port1 interface, enter the following information, and select *OK*.

<b>Alias</b>	ManagementExternal
<b>Virtual Domain</b>	root

<b>Addressing Mode</b>	DHCP
<b>Distance</b>	5
<b>Retrieve default gateway from server</b>	Enable
<b>Override internal DNS</b>	Enable
<b>Administrative Access</b>	HTTPS, SSH, SNMP
<b>Description</b>	This is the accounting department internal interface.



When the mode is set to DHCP or PPOE on an interface you can set the distance field. This is the administrative distance for any routes learned through the gateway for this interface. The gateway is added to the static route table with these values. A lower distance indicates a preferred route.

### To configure the physical interfaces - CLI

```

config global
 config system interface
 edit port2
 set alias AccountingLocal
 set vdom Accounting
 set mode static
 set ip 172.100.1.1 255.255.0.0
 set allowaccess https ping ssh
 set description "The accounting dept internal interface"
 next
 edit port3
 set alias SalesLocal
 set vdom Sales
 set mode static
 set ip 192.168.1.1 255.255.0.0
 set allowaccess https ping ssh
 set description "The sales dept. internal interface"
 next
 edit port1
 set alias ManagementExternal
 set vdom root
 set mode DHCP
 set distance 5
 set gwdetect enable
 set dns-server-override enable
 set allowaccess https ssh snmp
 set description "The systemwide management interface."
 end
 end

```



## Configuring the VDOM links

To complete the connection between each VDOM and the management VDOM, you need to add the two VDOM links; one pair is the Accounting - management link and the other is for Sales - management link.

When configuring inter-VDOM links, you do not have to assign IP addresses to the links unless you are using advanced features such as dynamic routing that require them. Not assigning IP addresses results in faster configuration, and more available IP addresses on your networks.

If you require them, or if you simply want to assign IP addresses for clarity can do so.

### To configure the Accounting and management VDOM link - web-based manager

1. Go to *Global > Network > Interfaces*.
2. Select the expand arrow to select *Create New > VDOM link*.
3. Enter the following information, and select *OK*.

<b>Name</b>	AccountVlnk
<b>Interface #0</b>	
<b>Virtual Domain</b>	Accounting
<b>IP/Netmask</b>	0.0.0.0/0.0.0.0
<b>Administrative Access</b>	HTTPS, PING, SSH
<b>Description</b>	The Accounting VDOM side of the link.
<b>Interface #1</b>	
<b>Virtual Domain</b>	root
<b>IP/Netmask</b>	0.0.0.0/0.0.0.0
<b>Administrative Access</b>	HTTPS, PING, SSH
<b>Description</b>	The Management VDOM side of the link.

### To configure the Accounting and management VDOM link - CLI

```
config global
 config system vdom-link
 edit AccountVlnk
 next
 end
 config system interface
 edit AccountVlnk0
 set vdom Accounting
 set ip 0.0.0.0 0.0.0.0
 set allowaccess https ping ssh
 set description "Accounting side of the VDOM link"
 next
 edit AccountVlnk1
 set vdom root
 set ip 0.0.0.0 0.0.0.0
 set allowaccess https ping ssh
 set description "Management side of the VDOM link"
 end
 end
```

### To configure the Sales and management VDOM link - web-based manager

1. Go to *Global > Network > Interfaces*.
2. Select the expand arrow and select *Create New > VDOM link*.
3. Enter the following information, and select *OK*.

<b>Name</b>	SalesVlnk
<b>Interface #0</b>	
<b>Virtual Domain</b>	Sales
<b>IP/Netmask</b>	0.0.0.0/0.0.0.0
<b>Administrative Access</b>	HTTPS, PING, SSH
<b>Description</b>	The Sales VDOM side of the link.
<b>Interface #1</b>	
<b>Virtual Domain</b>	root
<b>IP/Netmask</b>	0.0.0.0/0.0.0.0
<b>Administrative Access</b>	HTTPS, PING, SSH
<b>Description</b>	The Management VDOM side of the link.

## To configure the Sales and management VDOM link - CLI

```
config global
 config system vdom-link
 edit SalesVlnk
 end
 config system interface
 edit SalesVlnk0
 set vdom Accounting
 set ip 0.0.0.0 0.0.0.0
 set allowaccess https ping ssh
 set description "Sales side of the VDOM link"
 next
 edit SalesVlnk1
 set vdom root
 set ip 0.0.0.0 0.0.0.0
 set allowaccess https ping ssh
 set description "Management side of the VDOM link"
 end
 end
end
```

## Configuring the firewall and Security Profile settings

With the VDOMs, physical interfaces, and VDOM links configured the firewall must now be configured to allow the proper traffic. Firewalls are configured per-VDOM, and firewall objects must be created for each VDOM separately.

For this example, the firewall group of services allowed between the internal networks and the Internet are the basic services for web browsing, file transfer, and email. These include: HTTP, HTTPS, SSL, FTP, DNS, NTP, POP3, and SMTP.

The only services allowed between Sales and Accounting are secure web browsing (HTTPS) and email (POP3 and SMTP)



The limited number of services ensures security between departments. The list of services can be expanded in the future if needed.

---

Security profile settings will block all non-essential business websites while logging all web traffic, scan and file filter all web and email protocols, and block game and peer-to-peer applications using application control.

For added security, FortiClient is required on internal computers with AntiVirus scanning configured. This is enforced by *Endpoint NAC* in firewall policies.

Using firewall addresses makes the firewall policies easier to read. Also if any changes need to be made in the future, you can simply update the addresses without changing the firewall policies. The addresses required are:

- `AccountingLocal` - all traffic from the internal accounting network
- `AccountingVlnk` - all traffic from the VDOM link between accounting and management VDOMs
- `SalesLocal` - all traffic from the internal sales network
- `SalesVlnk` - all traffic from the VDOM link between sales and management VDOM.

The Accounting VDOM requires `AccountingLocal`, `AccountingVlnk`, and `SalesLocal`. The Sales VDOM requires `SalesLocal`, `SalesVlnk`, and `AccountingLocal`.

The firewall policies required on the Accounting VDOM are:

- `AccountingLocal` to Internet
- Internet to `AccountingLocal`
- `SalesLocal` to `AccountingLocal`
- `AccountingLocal` to `SalesLocal`

The firewall policies required on the Sales VDOM are:

- `SalesLocal` to Internet
- Internet to `SalesLocal`
- `SalesLocal` to `AccountingLocal`
- `AccountingLocal` to `SalesLocal`

This section includes the following topics:

- [Configuring firewall service groups](#)
- [Configuring Security Profile settings for the Accounting VDOM](#)
- [Configuring firewall settings for the Accounting VDOM](#)
- [Configuring Security Profile settings for the Sales VDOM](#)
- [Configuring firewall settings for the Sales VDOM](#)
- [Configuring firewall settings between the Accounting and Sales VDOMs](#)

## Configuring firewall service groups

Service groups are an easy way to manage multiple services, especially if the same services are used on different networks.

The two service groups used here are intended for normal office traffic to the Internet, and for restricted traffic between departments. In both cases network traffic will be limited to the services listed to prevent any potential security risks or bandwidth-robbing applications.

These service groups can be changed as needed to either include additional valid services that are being used on the network, or to exclude services that are not required. Also, custom services can be created as needed for applications that are not listed.

### To configure two firewall service groups - web-based manager

1. Open the *Accounting* VDOM.
2. Go to *Firewall Objects > Service > Group*.
3. Select *Create New*, enter the following information, and select *OK*.

<b>Group Name</b>	OfficeServices
<b>Members</b>	HTTP, HTTPS, SSL, FTP, DNS, NTP, POP3, PING, SMTP

4. Select *Create New*, enter the following information, and select *OK*.

<b>Group Name</b>	AccountingSalesServices
<b>Members</b>	HTTPS, POP3, PING, SMTP

### To configure two firewall service groups - CLI

```
config vdom
 edit Accounting
 config firewall service group
 edit OfficeServices
 set member HTTP HTTPS SSL FTP DNS NTP POP3 PING SMTP
 next
 edit AccountingSalesServices
 set member HTTPS POP3 PING SMTP
 end
 end
 end
```

### Configuring Security Profile settings for the Accounting VDOM

Security Profile settings include web filtering, antivirus, application control, and other features. This example just uses those three features to ensure that

- the business environment is free from viruses
- employees do not surf grossly inappropriate websites, and
- employees do not use games or peer-to-peer applications at work.

### To configure web filtering for the Accounting VDOM - web-based manager

1. Open the *Accounting* VDOM.
2. Go to *Security Profiles > Web Filter > Profile*.
3. Select *Create New*.
4. Enter `webStrict` for the *Name*.
5. Select the arrow to expand the *FortiGuard Web Filtering* section.
6. Block all *Categories* except Business Oriented, Other, and Unrated.
7. Block all *Classifications* except Image Search..
8. Log all *Categories* and *Classifications*.
9. Select *OK*.

### To configure web filtering for the Accounting VDOM - CLI

```
config vdom
 edit Accounting
 config webfilter profile
 edit webStrict
 config ftgd-wf
 set allow g07 g08 g21 g22 c01 c03
 set deny g01 g02 g03 g04 g05 g06 c02 c04 c05 c06 c07
 end
 set web-ftgd-err-log enable
 end
 end
 end
```

### To configure AntiVirus for the Accounting VDOM - web-based manager

1. Open the *Accounting* VDOM.
2. Go to *Security Profiles > AntiVirus > Profile*.
3. Select *Create New*.

4. Enter `avStrict` for the *Name*.
5. Enable *Scan* for all protocols.
6. Enable *File filter* for all protocols, and select `built-in-patterns` for *Option*.
7. Enable logging for both *Scan* and *File Filter*.
8. Select *OK*.

#### To configure AntiVirus for the Accounting VDOM - CLI

```

config vdom
 edit Accounting
 config antivirus profile
 edit avStrict
 config http
 set options scan file-filter
 end
 config ftp
 set options scan file-filter
 end
 config imap
 set options scan file-filter
 end
 config pop3
 set options scan file-filter
 end
 config smtp
 set options scan file-filter
 end
 config nntp
 set options scan file-filter
 end
 config im
 set options scan file-filter
 end
 set filepattable 1
 set av-virus-log enable
 set av-block-log enable
 end
 end
 end

```

#### To configure application control for the Accounting VDOM - web-based manager

1. Open the *Accounting* VDOM.
2. Go to *Security Profiles > Application Control > Application Sensor*.
3. Select *Create New* (+ button at top right of page).
4. Enter `appStrict` for *Name* and select *OK*.
5. Select *Create New*.
6. In *Filters*, set *Category* to *game*.

7. In *Applications/Settings*, enter the following, and select *OK*.

<b>Action</b>	Block
<b>Packet Logging</b>	Enable

8. Select *Create New*.

9. In *Filters*, set *Category* to *p2p*.

10. In *Applications/Settings*, enter the following, and select *OK*.

<b>Action</b>	Block
<b>Packet Logging</b>	Enable

11. Select *Apply*.

### To configure application control for the Accounting VDOM - CLI

```
config vdom
 edit Accounting
 config application list
 edit appStrict
 config entries
 edit 1
 set category 2
 next
 edit 2
 set category 8
 end
 end
 end
 end
 end
```

### Configuring firewall settings for the Accounting VDOM

This configuration includes two firewall addresses and two firewall policies for the Accounting VDOM - one for the internal network, and one for the VDOM link with the management VDOM (root).

For added security, all traffic allowed will be scanned. Only valid office traffic will be allowed using the service group *OfficeServices*. The FortiClient application must be used to ensure additional protection for the sensitive accounting information.

All sales and accounting computers have the FortiClient application installed, so the firewall policies check that FortiClient is installed and that antivirus scanning is enabled.

Note the spelling of *AccountVlnk* which is due to the eleven character limit on VDOM link names.

### To configure firewall addresses - web-based manager

1. Open the *Accounting* VDOM.
2. Select *Firewall Objects > Address > Address*

3. Select *Create New*, enter the following information, and select *OK*.

<b>Address Name</b>	AccountingLocal
<b>Type</b>	Subnet/ IP Range
<b>Subnet / IP Range</b>	172.100.0.0
<b>Interface</b>	port1

4. Select *Create New*, enter the following information, and select *OK*.

<b>Address Name</b>	AccountManagement
<b>Type</b>	Subnet/ IP Range
<b>Subnet / IP Range</b>	10.0.1.0
<b>Interface</b>	AccountVlnk

#### To configure firewall addresses - CLI

```

config vdom
 edit Accounting
 config firewall address
 edit AccountingLocal
 set type iprange
 set subnet 172.100.0.0
 set associated-interface port1
 next
 edit AccountManagement
 set type iprange
 set subnet 10.0.1.0
 set associated-interface AccountVlnk
 end
 end
 end
end

```

#### To configure protocol options for Accounting VDOM - web-based manager

1. Open the *Accounting* VDOM.
2. Select *Policy > Policy > Protocol Options*.
3. Select *Create New*.
4. Enter `default` for the *Name*.
5. Select *OK*.

#### To configure the firewall policies from AccountingLocal to the Internet - web-based manager

1. Open the *Accounting* VDOM.
2. Go to *Policy > Policy*.
3. Select *Create New*, enter the following information, and then select *OK*.

<b>Source Interface/Zone</b>	port2
<b>Source Address</b>	AccountingLocal



<b>Destination Interface/Zone</b>	AccountVlnk
<b>Destination Address</b>	AccountManagement
<b>Schedule</b>	always
<b>Service</b>	OfficeServices
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	enable
<b>Security Features</b>	enabled
<b>Protocol Option</b>	default
<b>Web Filtering</b>	webStrict
<b>AntiVirus Filtering</b>	avStrict
<b>Application Control</b>	appStrict
<b>Enable Endpoint NAC</b>	Enforce_FortiClient_AV

4. Open the *root* VDOM.
5. Go to *Policy > Policy*.
6. Select *Create New*, enter the following information, and then select *OK*.

<b>Source Interface/Zone</b>	AccountVlnk
<b>Source Address</b>	AccountManagement
<b>Destination Interface/Zone</b>	port2
<b>Destination Address</b>	all
<b>Schedule</b>	always
<b>Service</b>	OfficeServices
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	enable
<b>Security Features</b>	enable
<b>Protocol Option</b>	default
<b>Web Filtering</b>	webStrict
<b>AntiVirus Filtering</b>	avStrict
<b>Application Control</b>	appStrict
<b>Enable Endpoint NAC</b>	disabled

## To configure the firewall policies from AccountingLocal to Internet - CLI

```
config vdom
 edit Accounting
 config firewall policy
 edit 1
 set srcintf "port2"
 set dstintf "AccountVlnk"
 set srcaddr "AccountingLocal"
 set dstaddr "AccountManagement"
 set action accept
 set schedule "always"
 set service "OfficeServices"
 set nat enable
 set av-profile avStrict
 set webfilter-profile webStrict
 set application-list appStrict
 set profile-protocol-options default
 set endpoint-check enable
 set endpoint-profile "FortiClient_installed"
 end
 end
 end

config vdom
 edit root
 config firewall policy
 edit 2
 set srcintf AccountVlnk
 set dstintf port1
 set srcaddr AccountManagement
 set dstaddr all
 set action accept
 set schedule always
 set service OfficeServices
 set nat enable
 set av-profile "scan"
 set webfilter-profile "scan"
 set application-list "AppControlList"
 set profile-protocol-options default
 set endpoint-check disable
 end
 end
 end
```

## To configure the firewall policies from Internet to AccountingLocal - web-based manager

1. Open the *root* VDOM.
2. Go to *Policy > Policy*.

3. Select *Create New*, enter the following information, and select *OK*.

<b>Source Interface/Zone</b>	port1
<b>Source Address</b>	all
<b>Destination Interface/Zone</b>	AccountVlnk
<b>Destination Address</b>	AccountManagement
<b>Schedule</b>	always
<b>Service</b>	OfficeServices
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	enable
<b>Security Features</b>	enable
<b>Protocol Option</b>	default
<b>Web Filtering</b>	webStrict
<b>AntiVirus Filtering</b>	avStrict
<b>Application Control</b>	appStrict
<b>Enable Endpoint NAC</b>	disabled

4. Open the *Accounting* VDOM.
5. Go to *Policy > Policy*.
6. Select *Create New*, enter the following information, and select *OK*.

<b>Source Interface/Zone</b>	AccountVlnk
<b>Source Address</b>	AccountManagement
<b>Destination Interface/Zone</b>	port2
<b>Destination Address</b>	AccountingLocal
<b>Schedule</b>	always
<b>Service</b>	OfficeServices
<b>Action</b>	ACCEPT
<b>Enable NAT</b>	enable
<b>Security Features</b>	enable
<b>Protocol Option</b>	default
<b>Web Filtering</b>	webStrict
<b>AntiVirus Filtering</b>	avStrict

<b>Application Control</b>	appStrict
<b>Enable Endpoint NAC</b>	disabled

### To configure the firewall policies from Internet to AccountingLocal - CLI

```

config vdom
 edit root
 config firewall policy
 edit 3
 set srcintf port1
 set dstintf AccountVlnk
 set srcaddr all
 set dstaddr AccountManagement
 set action accept
 set schedule always
 set service OfficeServices
 set nat enable
 set av-profile avStrict
 set webfilter-profile webStrict
 set application-list appstrict
 set profile-protocol-options default
 set endpoint-check disable
 end
 end
 config vdom
 edit Accounting
 config firewall policy
 edit 4
 set srcintf AccountVlnk
 set dstintf port2
 set srcaddr AccountManagement
 set dstaddr AccountingLocal
 set action accept
 set schedule always
 set service OfficeServices
 set nat enable
 set av-profile avStrict
 set webfilter-profile webStrict
 set application-list appstrict
 set profile-protocol-options default
 set endpoint-check disable
 end
 end
 end
 end

```

## Configuring Security Profile settings for the Sales VDOM

Security profile settings include web filtering, antivirus, application control, and other features. This example just uses those three features to ensure that

- the business environment is free from viruses
- employees do not surf grossly inappropriate websites, and
- employees do not use games or peer-to-peer applications at work.

Note that Sales web traffic is different from Accounting, and web filtering is different to account for this.

### To configure web filtering for the Sales VDOM - web-based manager

1. Open the *Sales VDOM*.
2. Go to *Security Profiles > Web Filter > Profile*.
3. Select *Create New*.
4. Enter `webStrict` for the *Name*.
5. In *FortiGuard Categories*, select all of the categories except *Bandwidth Consuming, General Interest - Business* and *Unrated*.
6. In *Change Action for Selected Categories* select *Block*.
7. Select *Apply*.

### To configure web filtering for the Sales VDOM - CLI

```
config vdom
 edit Sales
 config webfilter profile
 edit webStrict
 config ftgd-wf
 set allow g07 g08 g21 g22 c01 c03
 set deny g01 g02 g03 g04 g05 g06 c02 c04 c05 c06 c07
 end
 set web-ftgd-err-log enable
 end
 end
 end
```

### To configure AntiVirus for the Sales VDOM - web-based manager

1. Open the *Sales VDOM*.
2. Go to *Security Profiles > AntiVirus > Profile*.
3. Select *Create New*.
4. Enter `avStrict` for the *Name*.
5. Enable virus scan for all protocols.
6. Select *Apply*.

## To configure AntiVirus for the Sales VDOM - CLI

```
config vdom
 edit Sales
 config antivirus profile
 edit "avStrict"
 config http
 set options scan file-filter
 end
 config ftp
 set options scan file-filter
 end
 config imap
 set options scan file-filter
 end
 config pop3
 set options scan file-filter
 end
 config smtp
 set options scan file-filter
 end
 config nntp
 set options scan file-filter
 end
 config im
 set options scan file-filter
 end
 set filepattable 1
 set av-virus-log enable
 set av-block-log enable
 end
 end
 end
```

## To configure application control for the Sales VDOM - web-based manager

1. Open the *Accounting* VDOM.
2. Go to *Security Profiles > Application Control > Application Sensor*.
3. Select *Create New* (+ button at top right of page).
4. Enter *appStrict* for *Name* and select *OK*.
5. Select *Create New*.
6. In *Filters*, set *Category* to *game*.
7. In *Applications/Settings*, enter the following, and select *OK*.

<b>Action</b>	Block
<b>Packet Logging</b>	Enable

8. Select *Create New*.
9. In *Filters*, set *Category* to *p2p*.

10. In *Applications/Settings*, enter the following, and select *OK*.

<b>Action</b>	Block
<b>Packet Logging</b>	Enable

11. Select *Apply*.

### To configure application control for the Sales VDOM - CLI

```
config vdom
 edit Sales
 config application list
 edit "appStrict"
 config entries
 edit 1
 set category 2
 next
 edit 2
 set category 8
 end
 end
 end
 end
 end
```

### Configuring firewall settings for the Sales VDOM

Like the Accounting firewall settings, this configuration includes two firewall addresses and two firewall policies for the sales VDOM: one for the internal network, and one for the VDOM link with the management VDOM.

When entering the CLI commands, the number of the firewall policies must be high enough to be a new policy. Depending on the number of firewall policies on your FortiGate unit, this may require starting at a higher number than the 6 required for the default configuration. This number is added automatically when you configure firewall policies using the web manager interface.

The FortiClient application must be used on Sales network computers to ensure additional protection for the sensitive information and for protection against spam.

### To configure firewall addresses - web-based manager

1. Open the *Sales VDOM*.
2. Go to *Firewall Objects > Address > Address*.
3. Select *Create New*, enter the following information, and select *OK*.

<b>Address Name</b>	SalesLocal
<b>Type</b>	Subnet / IP Range
<b>Subnet / IP Range</b>	172.100.0.0
<b>Interface</b>	port3

4. Go to *Firewall Objects > Addresses*.

5. Select *Create New*, enter the following information, and select OK.

<b>Address Name</b>	SalesManagement
<b>Type</b>	Subnet / IP Range
<b>Subnet / IP Range</b>	10.0.1.0
<b>Interface</b>	SalesVlnk

#### To configure the firewall addresses - CLI

```

config vdom
 edit Sales
 config fireall address
 edit SalesLocal
 set type iprange
 set subnet 172.100.0.0
 set associated-interface port2
 next
 edit SalesManagement
 set type iprange
 set subnet 10.0.1.0
 set associated-interface SalesVlnk
 end
 end
 end

```

#### To configure the firewall policies from SalesLocal to the Internet - web-based manager

1. Open the *Sales* VDOM.
2. Go to *Policy > Policy*.
3. Select *Create New*, enter the following information, and select OK.

<b>Source Interface/Zone</b>	port3
<b>Source Address</b>	SalesLocal
<b>Destination Interface/Zone</b>	SalesVlnk
<b>Destination Address</b>	SalesManagement
<b>Schedule</b>	always
<b>Service</b>	OfficeServices
<b>Action</b>	ACCEPT
<b>Log Allowed Traffic</b>	enabled
<b>Enable Endpoint Control Check</b>	disabled
<b>Redirect Non-conforming Clients to Download Portal</b>	enabled

4. Open the *root* VDOM.
5. Go to *Policy > Policy*.



6. Select *Create New*, enter the following information, and select *OK*.

<b>Source Interface/Zone</b>	SalesVInk
<b>Source Address</b>	SalesManagement
<b>Destination Interface/Zone</b>	external
<b>Destination Address</b>	all
<b>Schedule</b>	always
<b>Service</b>	OfficeServices
<b>Action</b>	ACCEPT
<b>Protection Profile</b>	scan
<b>Log Allowed Traffic</b>	enabled
<b>Enable Endpoint Control Check</b>	disabled

## To configure the firewall policies from SalesLocal to the Internet - CLI

```
config vdom
 edit root
 config firewall policy
 edit 6
 set srcintf port2
 set srcaddr SalesLocal
 set dstintf SalesVlnk
 set dstaddr SalesManagement
 set schedule always
 set service OfficeServices
 set action accept
 set profile-status enable
 set profile scan
 set logtraffic enable
 set endpoint-check enable
 set endpoint-redir-portal enable
 end
 end
 end

config vdom
 edit Sales
 config firewall policy
 edit 7
 set srcintf SalesVlnk
 set srcaddr SalesManagement
 set dstintf external
 set dstaddr all
 set schedule always
 set service OfficeServices
 set action accept
 set profile-status enable
 set profile scan
 set logtraffic enable
 set endpoint-check enable
 end
 end
 end
```

## To configure the firewall policies from the Internet to SalesLocal - web-based manager

1. Open the *root* VDOM.
2. Go to *Policy > Policy*.
3. Select *Create New*, enter the following information, and select *OK*.

<b>Source Interface/Zone</b>	external
<b>Source Address</b>	all
<b>Destination Interface/Zone</b>	SalesVlnk
<b>Destination Address</b>	SalesManagement

<b>Schedule</b>	always
<b>Service</b>	OfficeServices
<b>Action</b>	ACCEPT
<b>Protection Profile</b>	scan
<b>Log Allowed Traffic</b>	enabled
<b>Enable Endpoint Control Check</b>	disabled

4. Open the *Sales* VDOM.
5. Go to *Policy > Policy*.
6. Select *Create New*, enter the following information, and select *OK*.

<b>Source Interface/Zone</b>	SalesVlnk
<b>Source Address</b>	SalesManagement
<b>Destination Interface/Zone</b>	port2
<b>Destination Address</b>	SalesLocal
<b>Schedule</b>	always
<b>Service</b>	OfficeServices
<b>Action</b>	ACCEPT
<b>Protection Profile</b>	scan
<b>Log Allowed Traffic</b>	enabled
<b>Enable Endpoint Control Check</b>	disabled
<b>Redirect Non-conforming Clients to Download Portal</b>	enabled

## To configure the firewall policies from the Internet to SalesLocal - CLI

```
config vdom
 edit root
 config firewall policy
 edit 8
 set srcintf external
 set srcaddr all
 set dstintf SalesVlnk
 set dstaddr SalesManagement
 set schedule always
 set service OfficeServices
 set action accept
 set profile-status enable
 set profile scan
 set logtraffic enable
 set endpoint-check enable
 set endpoint-redir-portal enable
 end
 end
 end

config vdom
 edit Sales
 config firewall policy
 edit 9
 set srcintf SalesVlnk
 set srcaddr SalesManagement
 set dstintf port2
 set dstaddr SalesLocal
 set schedule always
 set service OfficeServices
 set action accept
 set profile-status enable
 set profile scan
 set logtraffic enable
 set endpoint-check enable
 set endpoint-redir-portal enable
 end
 end
 end
```

## Configuring firewall settings between the Accounting and Sales VDOMs

Firewall policies are required for any communication between each internal network and the Internet. Policies are also required for the two internal networks to communicate with each other through the management VDOM.

The more limited AccountingSalesServices group of services will be used between Sales and Accounting to ensure the traffic is necessary business traffic only. These policies will result in a partially meshed VDOM configuration. The FortiClient application must be used to ensure additional protection for the sensitive accounting information.

Two firewall policies are required to allow traffic in both directions between Sales and Accounting.

**To configure the firewall policy between Sales and Accounting on the management VDOM - web-based manager**

1. Open the *root* VDOM.
2. Go to *Policy > Policy*.
3. Select *Create New*, enter the following information, and select *OK*.

<b>Source Interface/Zone</b>	SalesVInk
<b>Source Address</b>	SalesManagement
<b>Destination Interface/Zone</b>	AccountVInk
<b>Destination Address</b>	AccountingManagement
<b>Schedule</b>	always
<b>Service</b>	AccountingSalesServices
<b>Action</b>	ACCEPT
<b>Protection Profile</b>	scan
<b>Log Allowed Traffic</b>	enabled
<b>Enable Endpoint Control Check</b>	disabled
<b>Redirect Non-conforming Clients to Download Portal</b>	enabled

4. Go to *Policy > Policy*.
5. Select *Create New*, enter the following information, and select *OK*.

<b>Source Interface/Zone</b>	AccountVInk
<b>Source Address</b>	AccountingManagement
<b>Destination Interface/Zone</b>	SalesVInk
<b>Destination Address</b>	SalesManagement
<b>Schedule</b>	always
<b>Service</b>	AccountingSalesServices
<b>Action</b>	ACCEPT
<b>Protection Profile</b>	scan
<b>Log Allowed Traffic</b>	enabled
<b>Enable Endpoint Control Check</b>	disabled
<b>Redirect Non-conforming Clients to Download Portal</b>	enabled

## To configure the firewall policy between Sales and Accounting on the management VDOM - CLI

```
config vdom
 edit root
 config system firewall policy
 edit 9
 set srcintf SalesVlnk
 set srcaddr SalesManagement
 set dstintf AccountVlnk
 set dstaddr AccountManagement
 set schedule always
 set service AccountingSalesServices
 set action accept
 set profile-status enable
 set profile scan
 set logtraffic enable
 set endpoint-check enable
 set endpoint-redir-portal enable
 next
 edit 10
 set srcintf AccountVlnk
 set srcaddr AccountManagement
 set dstintf SalesVlnk
 set dstaddr SalesManagement
 set schedule always
 set service AccountingSalesServices
 set action accept
 set profile-status enable
 set profile scan
 set logtraffic enable
 set endpoint-check enable
 set endpoint-redir-portal enable
 end
 end
 end
```

### Testing the configuration

Once the inter-VDOM routing has been configured, tests must be conducted to confirm proper operation. If there are any problems, use the troubleshooting tips to resolve them.

This section includes the following topics:

- [Testing connectivity](#)
- [Troubleshooting Tips](#)

### Testing connectivity

Testing connectivity ensures that physical networking connections as well as FortiGate unit interface configurations, including firewall policies, are properly configured.

The easiest way to test connectivity is to use the `ping` and `tracert` commands to confirm the connectivity of different routes on the network. Include testing:

- from AccountingLocal to Internet
- from Internet to AccountingLocal
- from SalesLocal to Internet
- from Internet to SalesLocal
- from AccountingLocal to SalesLocal.

When using the commands on a Windows computer, go to a command line prompt and enter either `ping <IP address>` or `tracert <IP address>`.

When using the commands on a FortiGate unit, go to the CLI and enter either `exec ping <IP address>` or `exec traceroute <IP address>`.

### Troubleshooting Tips

When there are problems with connectivity, the following troubleshooting tips will help resolve the issues.

- If a multiple hop test, such as `tracert`, is not successful then reduce it to a single hop to simplify the test. Test each link of the path to see which hop is down. If all hops are up, check the FortiGate unit policies to ensure they allow basic traffic to flow as expected.
- If ping does not work, confirm that the FortiGate unit interfaces have Ping enabled and also ensure Ping is enabled in the firewall policies. Otherwise the Ping traffic will be blocked.
- If one protocol does not work but others do work, check the FortiGate unit firewall policies for that one protocol to ensure it is allowed.
- If there are unexplained connectivity problems, check the local computer to ensure it does not have a software firewall running that may be blocking traffic. MS Windows computers have a firewall running by default that can cause problems.

For additional troubleshooting, see [“Troubleshooting Virtual Domains” on page 2440](#).

# Troubleshooting Virtual Domains

When you are configuring VDOMs you may run into some issues, with your VDOM configuration, your network configuration, or your device setup. This section addresses common problems and specific concerns that an administrator of a VDOM network may have.

This section includes:

- [VDOM admin having problems gaining access](#)
- [FortiGate unit running very slowly](#)
- [General VDOM tips and troubleshooting](#)

## VDOM admin having problems gaining access

With VDOMs configured, administrators have an extra layer of permissions and may have problems accessing their information.

### Confirm the admin's VDOM

Each administrator account, other than the `super_admin` account, is tied to one specific VDOM. That administrator is not able to access any other VDOM. It may be possible they are trying to access the wrong VDOM.

### Confirm the VDOM's interfaces

An administrator can only access their VDOM through interfaces that are assigned to that VDOM. If interfaces on that VDOM are disabled or unavailable there will be no method of accessing that VDOM by its local administrator. The `super_admin` will be required to either bring up the interfaces, fix the interfaces, or move another interface to that VDOM to restore access.

### Confirm the VDOMs admin access

As with all FortiGate units, administration access on the VDOM's interfaces must be enabled for that VDOM's administrators to gain access. For example if SSH is not enabled, that is not available to administrators.

To enable admin access, the `super_admin` will go to the *Global > Network > Interfaces* page, and for the interface in question enable the admin access.

## FortiGate unit running very slowly

You may experience a number of problems resulting from your FortiGate unit being overloaded. These problems may appear as:

- CPU and memory threshold limits exceeded on a continual basis
- AV failopen happening on a regular basis
- dropped traffic or sessions due to lack of resources

These problems are caused by a lack of system resources. There are a number of possible reasons for this.



## Too many VDOMs

If you have configured many VDOMs on your system, past the default ten VDOMs, this could easily be your problem.

Each VDOM you create on your FortiGate unit requires system resources to function - CPU cycles, memory, and disk space. When there are too many VDOMs configured there are not enough resources for operation. This may be a lack of memory in the session table, or no CPU cycles for processing incoming IPS traffic, or even a full disk drive.

Go to *System > VDOM* and see the number of configured VDOMs on your system. If you are running 500 or more VDOMs, you must have a FortiGate 5000 chassis. Otherwise you need to reduce the number of VDOMs on your system to fix the problem. Even if you have the proper hardware, you may encounter noticeably slow throughput if you are using advanced features such as security profiles or deep content inspection with many configured VDOMs.

## One or more VDOMs are consuming all the resources

If you have sufficient hardware to support the number of VDOMs you are running, check the global resources on your FortiGate unit. At a glance it will tell you if you are running out of a particular resource such as sessions, or users. If this is the case, you can then check your VDOMs to see if one particular VDOM is using more than its share of resources. If that is the case you can change the resource settings to allow that VDOM (or those VDOMs) fewer resources and in turn allow the other VDOMs access to those resources.

## Too many Security Features in use

It is likely that reducing the Security Features in use regardless of number of VDOMs will greatly improve overall system performance and should be considered as an option.

Finally it is possible that your FortiGate unit configuration is incorrect in some other area, which is using up all your resources. For example, forgetting that you are running a network sniffer on an interface will create significant amounts of traffic that may prevent normal operation.

## General VDOM tips and troubleshooting

Besides ping and traceroute, there are additional tools for troubleshooting your VDOM configurations. These include packet sniffing and debugging the packet flow.

### Perform a sniffer trace

When troubleshooting networks, it helps to look inside the headers of packets to determine if they are traveling along the route you expect that they are. Packet sniffing can also be called a network tap, packet capture, or logic analyzing.



If your FortiGate unit has NP interfaces that are offloading traffic, this will change the sniffer trace. Before performing a trace on any NP interfaces, you should disable offloading on those interfaces.

---

### What sniffing packets can tell you

If you are running a constant traffic application such as ping, packet sniffing can tell you if the traffic is reaching the destination, what the port of entry is on the FortiGate unit, if the ARP resolution is correct, and if the traffic is being sent back to the source as expected.

Sniffing packets can also tell you if the Fortigate unit is silently dropping packets for reasons such as RPF (Reverse Path Forwarding), also called Anti Spoofing, which prevents an IP packet from being forwarded if its Source IP does not either belong to a locally attached subnet (local interface), or be part of the routing between the FortiGate and another source (static route, RIP, OSPF, BGP). Note that RPF can be disabled by turning on asymmetric routing in the CLI (`config system setting, set asymmetric enable`), however this will disable stateful inspection on the FortiGate unit and cause many features to be turned off.

**Note** If you configure virtual IP addresses on your Fortigate unit, it will use those addresses in preference to the physical IP addresses. You will notice this when you are sniffing packets because all the traffic will be using the virtual IP addresses. This is due to the ARP update that is sent out when the VIP address is configured.

## How to sniff packets

When you are using VDOMs, you must be in a VDOM to access the `diag sniffer` command. At the global level, the command is not available. This is limit the packets only to the ones on your VDOM, and protects the privacy of other VDOM clients.

The general form of the internal FortiOS packet sniffer command is:

```
diag sniffer packet <interface_name> <'filter'> <verbose> <count>
```

To stop the sniffer, type `CTRL+C`.

<b>&lt;interface_name&gt;</b>	The name of the interface to sniff, such as “port1” or “internal”. This can also be “any” to sniff all interfaces.
<b>&lt;'filter'&gt;</b>	What to look for in the information the sniffer reads. “none” indicates no filtering, and all packets will be displayed as the other arguments indicate.  The filter must be inside single quotes (').
<b>&lt;verbose&gt;</b>	The level of verbosity as one of:  1 - print header of packets 2 - print header and data from IP of packets 3 - print header and data from Ethernet of packets
<b>&lt;count&gt;</b>	The number of packets the sniffer reads before stopping. If you don't put a number here, the sniffer will run forever until you stop it with <code>&lt;CTRL C&gt;</code> .

For a simple sniffing example, enter the CLI command `diag sniffer packet port1 none 1 3`. This will display the next 3 packets on the port1 interface using no filtering, and using verbose level 1. At this verbosity level you can see the source IP and port, the destination IP and port, action (such as ack), and sequence numbers.

In the output below, port 443 indicates these are HTTPS packets, and 172.20.120.17 is both sending and receiving traffic.

```
Head_Office_620b # diag sniffer packet port1 none 1 3
interfaces=[port1]
filters=[none]
0.545306 172.20.120.17.52989 -> 172.20.120.141.443: psh 3177924955
 ack 1854307757

0.545963 172.20.120.141.443 -> 172.20.120.17.52989: psh 1854307757
 ack 3177925808

0.562409 172.20.120.17.52988 -> 172.20.120.141.443: psh 4225311614
 ack 3314279933
```

For a more advanced example of packet sniffing, the following commands will report packets on any interface travelling between a computer with the host name of PC1 and the computer with the host name of PC2. With verbosity 4 and above, the sniffer trace will display the interface names where traffic enters or leaves the FortiGate unit. Remember to stop the sniffer, type CTRL+C. Note that PC1 and PC2 may be VDOMs.

```
FGT# diagnose sniffer packet any "host <PC1> or host <PC2>" 4
```

or

```
FGT# diagnose sniffer packet any "(host <PC1> or host <PC2>) and
 icmp" 4
```

The following sniffer CLI command includes the ARP protocol in the filter which may be useful to troubleshoot a failure in the ARP resolution (for instance PC2 may be down and not responding to the FortiGate ARP requests).

```
FGT# diagnose sniffer packet any "host <PC1> or host <PC2> or arp" 4
```

## Debugging the packet flow

Traffic should come in and leave the VDOM. If you have determined that network traffic is not entering and leaving the VDOM as expected, debug the packet flow.

Debugging can only be performed using CLI commands. Debugging the packet flow requires a number of debug commands to be entered as each one configures part of the debug action, with the final command starting the debug.



If your FortiGate unit has NP interfaces that are offloading traffic, this will change the packet flow. Before performing the debug on any NP interfaces, you should disable offloading on those interfaces.

---

The following configuration assumes that PC1 is connected to the internal interface of the FortiGate unit and has an IP address of 10.11.101.200. PC1 is the host name of the computer.

To debug the packet flow in the CLI, enter the following commands:

```
FGT# diag debug enable
FGT# diag debug flow filter add <PC1>
FGT# diag debug flow show console enable
FGT# diag debug flow trace start 100
FGT# diag debug enable
```

The `start 100` argument in the above list of commands will limit the output to 100 packets from the flow. This is useful for looking at the flow without flooding your log or your display with too much information.

To stop all other debug activities, enter the command:

```
FGT# diag debug flow trace stop
```

The following is an example of debug flow output for traffic that has no matching Firewall Policy, and is in turn blocked by the FortiGate unit. The denied message indicates the traffic was blocked. Note that even with VDOMs not enabled, `vd-root` is still shown.

```
id=20085 trace_id=319 func=resolve_ip_tuple_fast line=2825
 msg="vd-root received a packet(proto=6,
 192.168.129.136:2854->192.168.96.153:1863) from port3."

id=20085 trace_id=319 func=resolve_ip_tuple line=2924 msg="allocate
 a new session-013004ac"

id=20085 trace_id=319 func=vf_ip4_route_input line=1597 msg="find a
 route: gw-192.168.150.129 via port1"

id=20085 trace_id=319 func=fw_forward_handler line=248 msg=" Denied
 by forward policy check"
```

# Chapter 20 Virtual FortiGate Units for FortiOS 5.0

This document describes how to deploy a FortiGate virtual appliance in several virtualization server environments. This includes how to configure the virtual hardware settings of the virtual appliance.

This document assumes:

- you have already successfully installed the virtualization server on the physical machine,
- you have installed appropriate VM management software on either the physical server or a computer to be used for VM management.

This document does not cover configuration and operation of the virtual appliance after it has been successfully installed and started. For these issues, see the [FortiGate 5.0 Handbook](#).

This document includes the following sections:

- [FortiGate VM Overview](#)
- [Deployment example: VMware](#)
- [Deployment example: MS Hyper-V](#)
- [Deployment example: KVM](#)
- [Deployment example: OpenXen](#)
- [Deployment example: Citrix XenServer](#)
- [FortiGate-AWS Cloud-based Appliance](#)

# FortiGate VM Overview

The following topics are included in this section:

- [FortiGate VM models and licensing](#)
- [Registering FortiGate VM with Customer Service & Support](#)
- [Downloading the FortiGate VM deployment package](#)
- [Deployment package contents](#)
- [Deploying the FortiGate VM appliance](#)

## FortiGate VM models and licensing

Fortinet offers the FortiGate VM in five virtual appliance models determined by license. When configuring your FortiGate VM, be sure to configure hardware settings within the ranges outlined in [Table 111](#). Contact your Fortinet Authorized Reseller for more information.

**Table 111:**FortiGate VM model information

Technical Specification	FG-VM00	FG-VM01	FG-VM02	FG-VM04	FG-VM08
Virtual CPUs (min/max)	1/1	1/1	1/2	1/4	1/8
Virtual Network Interfaces (min/max)	2 / 10				
Virtual Memory (min/max)	1 GB / 1 GB	1 GB / 2 GB	1 GB / 4 GB	1 GB / 6 GB	1 GB / 12 GB
Virtual Storage (min/max)	30 GB / 2 TB				
Managed Wireless Access Points (tunnel mode / global)	32 / 32	32 / 64	256 / 512	256 / 512	1024 / 4096
Virtual Domains (default / max)	1 / 1	10 / 10	10 / 25	10 / 50	10 / 250

After placing an order for FortiGate VM, a license registration code is sent to the email address used on the order form. Use the registration number provided to register the FortiGate VM with Customer Service & Support and then download the license file. Once the license file is uploaded to the FortiGate VM and validated, your FortiGate VM appliance is fully functional.

## FortiGate VM evaluation license

FortiGate VM includes a limited embedded 15-day trial license that supports:

- 1 CPU maximum
- 1024 MB memory maximum
- low encryption only (no HTTPS administrative access)
- all features except FortiGuard updates

You cannot upgrade the firmware, doing so will lock the Web-based Manager until a license is uploaded. Technical support is not included. The trial period begins the first time you start

FortiGate VM. After the trial license expires, functionality is disabled until you upload a license file.

## Registering FortiGate VM with Customer Service & Support

To obtain the FortiGate VM license file you must first register your FortiGate VM with [Customer Service & Support](#).

### To register your FortiGate VM:

1. Log in to the Customer Service & Support portal using an existing support account or select *Sign Up* to create a new account.
2. In the main page, under *Asset*, select *Register/Renew*.  
The *Registration* page opens.
3. Enter the registration code that was emailed to you and select *Register*. A registration form will display.
4. After completing the form, a registration acknowledgement page will appear.
5. Select the *License File Download* link.
6. You will be prompted to save the license file (.lic) to your local computer. See [“Upload the FortiGate VM license file” on page 2484](#) for instructions on uploading the license file to your FortiGate VM via the Web-based Manager.

## Downloading the FortiGate VM deployment package

FortiGate VM deployment packages are included with FortiGate firmware images on the [Customer Service & Support](#) site. First, see [Table 112](#) to determine the appropriate VM deployment package for your VM platform.

**Table 112:** Selecting the correct FortiGate VM deployment package for your VM platform

VM Platform	FortiGate VM Deployment File
Citrix XenServer v5.6sp2, 6.0 and later	FGT_VM64-v500-buildnnnn-FORTINET.out.CitrixXen.zip
OpenXen v3.4.3, 4.1	FGT_VM64-v500-buildnnnn-FORTINET.out.OpenXen.zip
Microsoft Hyper-V Server 2008R2 and 2012	FGT_VM64-v500-buildnnnn-FORTINET.out.hyperv.zip
KVM (qemu 0.12.1)	FGT_VM64-v500-buildnnnn-FORTINET.out.kvm.zip
VMware ESX 4.0, 4.1 ESXi 4.0/4.1/5.0/5.1/5.5	FGT_VM32-v500-buildnnnn-FORTINET.out.ovf.zip (32-bit) FGT_VM64-v500-buildnnnn-FORTINET.out.ovf.zip

For more information see the FortiGate product datasheet available on the Fortinet web site, <http://www.fortinet.com/products/fortigate/virtualappliances.html>.

The firmware images FTP directory is organized by firmware version, major release, and patch release. The firmware images in the directories follow a specific naming convention and each firmware image is specific to the device model. For example, the FGT\_VM32-v500-build0151-FORTINET.out.ovf.zip image found in the v5.0 Patch Release 2 directory is specific to the FortiGate VM 32-bit environment.



You can also download the [FortiOS Release Notes](#), FORTINET-FORTIGATE MIB file, FSSO images, and SSL VPN client in this directory. The Fortinet Core MIB file is located in the main FortiGate v5.00 directory.

---

**To download the FortiGate VM deployment package:**

1. In the main page of the Customer Service & Support site, select *Download > Firmware Images*.  
The *Firmware Images* page opens.
2. In the *Firmware Images* page, select *FortiGate*.
3. Browse to the appropriate directory on the FTP site for the version that you would like to download.
4. Download the appropriate .zip file for your VM server platform.  
You can also download the [FortiGate Release Notes](#).
5. Extract the contents of the deployment package to a new file folder.

## Deployment package contents

### Citrix XenServer

The FORTINET.out.CitrixXen.zip file contains:

- fortios.vhd: the FortiGate VM system hard disk in VHD format
- fortios.xva: binary file containing virtual hardware configuration settings
- in the ovf folder:
  - FortiGate-VM64.ovf: Open Virtualization Format (OVF) template file, containing virtual hardware settings for Xen
  - fortios.vmdk: the FortiGate VM system hard disk in VMDK format
  - datadrive.vmdk: the FortiGate VM log disk in VMDK format

The ovf folder and its contents is an alternative method of installation to the .xva and VHD disk image.

### OpenXEN

The FORTINET.out.OpenXen.zip file contains only fortios.qcow2, the FortiGate VM system hard disk in qcow2 format. You will need to manually:

- create a 30GB log disk
- specify the virtual hardware settings



## Microsoft Hyper-V

The FORTINET.out.hyperv.zip file contains:

- in the Virtual Hard Disks folder:
  - fortios.vhd: the FortiGate VM system hard disk in VHD format
  - DATADRIVE.vhd: the FortiGate VM log disk in VHD format
- In the Virtual Machines folder:
  - fortios.xml: XML file containing virtual hardware configuration settings for Hyper-V. This is compatible with Windows Server 2012.
- Snapshots folder: optionally, Hyper-V stores snapshots of the FortiGate VM state here

## KVM

The FORTINET.out.kvm.zip contains only fortios.qcow2, the FortiGate VM system hard disk in qcow2 format. You will need to manually:

- create a 30GB log disk
- specify the virtual hardware settings

## VMware ESX/ESXi

The FORTINET.out.ovf.zip file contains:

- fortios.vmdk: the FortiGate VM system hard disk in VMDK format
- datadrive.vmdk: the FortiGate VM log disk in VMDK format
- Open Virtualization Format (OVF) template files:
  - FortiGate-VM64.ovf: OVF template based on Intel e1000 NIC driver
  - FortiGate-VM64.hw04.ovf: OVF template file for older (v3.5) VMware ESX server
  - FortiGate-VMxx.hw07\_vmxnet2.ovf: OVF template file for VMware vmxnet2 driver
  - FortiGate-VMxx.hw07\_vmxnet3.ovf: OVF template file for VMware vmxnet3 driver

## Deploying the FortiGate VM appliance

Prior to deploying the FortiGate VM appliance, the VM platform must be installed and configured so that it is ready to create virtual machines. The installation instructions for FortiGate VM assume that

- You are familiar with the management software and terminology of your VM platform.
- An Internet connection is available for FortiGate VM to contact FortiGuard to validate its license or, for closed environments, a FortiManager can be contacted to validate the FortiGate VM license. See [“Validate the FortiGate VM license with FortiManager” on page 2485](#).

For assistance in deploying FortiGate VM, refer to the deployment chapter in this guide that corresponds to your VM environment. You might also need to refer to the documentation provided with your VM server. The deployment chapters are presented as examples because for any particular VM server there are multiple ways to create a virtual machine. There are command line tools, APIs, and even alternative graphical user interface tools.

Before you start your FortiGate VM appliance for the first time, you might need to adjust virtual disk sizes and networking settings. The first time you start FortiGate VM, you will have access only through the console window of your VM server environment. After you configure one

FortiGate network interface with an IP address and administrative access, you can access the FortiGate VM web-based manager.

After deployment and license validation, you can upgrade your FortiGate VM appliance's firmware by downloading either FGT\_VM32-v500-buildnnnn-FORTINET.out (32-bit) or FGT\_VM64-v500-buildnnnn-FORTINET.out (64-bit) firmware. Firmware upgrading on a VM is very similar to upgrading firmware on a hardware FortiGate unit.

# Deployment example: VMware

Once you have downloaded the FGT\_VMxx-v500-build0xxx-FORTINET.out.ovf.zip file and extracted the package contents to a folder on your local computer, you can use the vSphere client to create the virtual machine from the deployment package OVF template.

The following topics are included in this section:

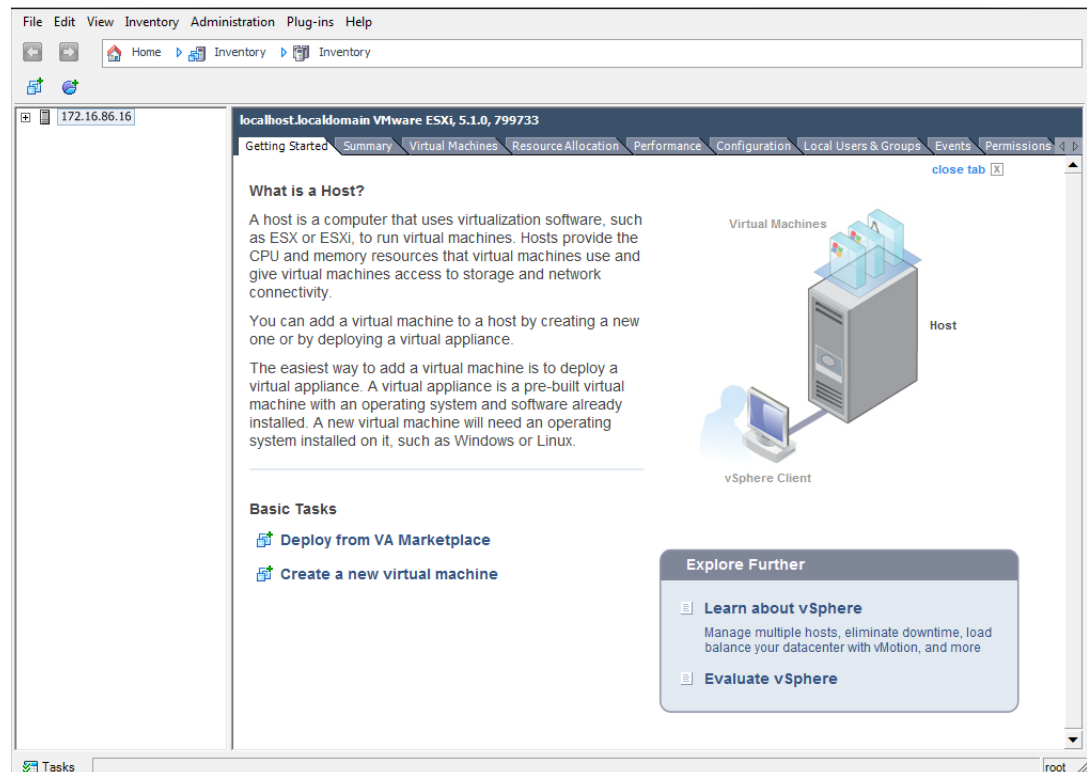
- [Open the FortiGate VM OVF file with the vSphere client](#)
- [Configure FortiGate VM hardware settings](#)
- [Power on your FortiGate VM](#)

## Open the FortiGate VM OVF file with the vSphere client

**To deploy the FortiGate VM OVF template:**

1. Launch the VMware vSphere client, enter the IP address or host name of your server, enter your user name and password and select *Login*.

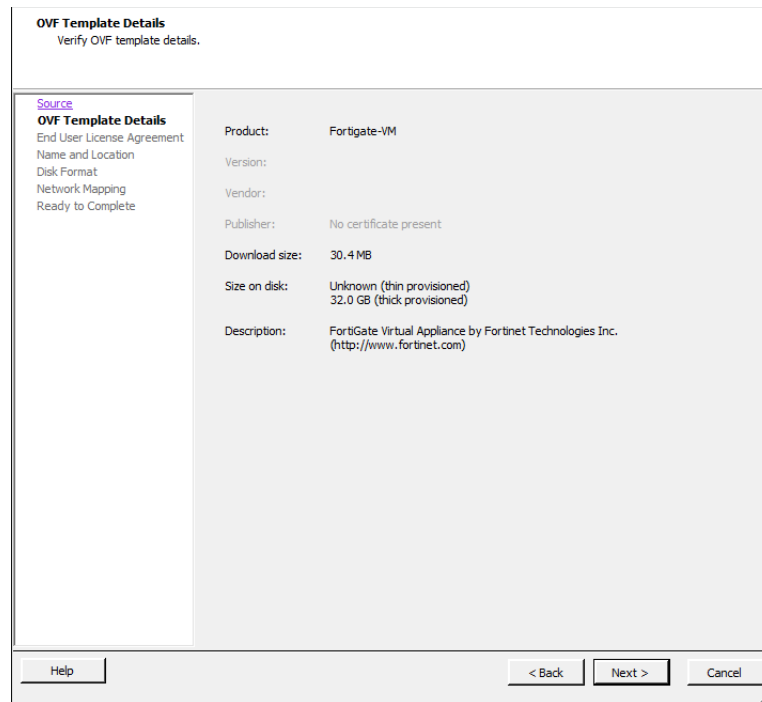
The vSphere client home page opens.



2. Select *File > Deploy OVF Template* to launch the OVF Template wizard.  
The OVF Template *Source* page opens.

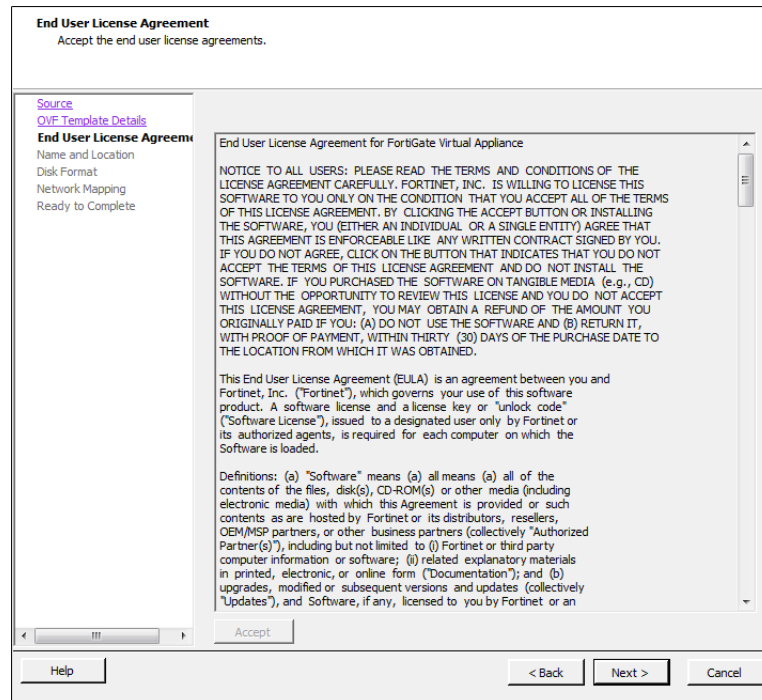
3. Select the source location of the OVF file. Select *Browse* and locate the OVF file on your computer. Select *Next* to continue.

The OVF Template *Details* page opens.



4. Verify the OVF template details. This page details the product name, download size, size on disk, and description. Select *Next* to continue.

The OVF Template *End User License Agreement* page opens.



5. Read the end user license agreement for FortiGate VM. Select *Accept* and then select *Next* to continue.

The OVF Template *Name and Location* page opens.

The screenshot shows the 'Name and Location' page of the OVF template wizard. The title is 'Name and Location' with the subtitle 'Specify a name and location for the deployed template'. On the left, there is a navigation pane with links: 'Source', 'OVF Template Details', 'End User License Agreement', 'Name and Location' (which is highlighted), 'Disk Format', 'Network Mapping', and 'Ready to Complete'. The main area contains a 'Name:' text box with the value 'Fortigate-VM-01'. Below the text box, it says 'The name can contain up to 80 characters and it must be unique within the inventory folder.' At the bottom, there are three buttons: 'Help', '< Back', and 'Next >', and a 'Cancel' button.

6. Enter a name for this OVF template. The name can contain up to 80 characters and it must be unique within the inventory folder. Select *Next* to continue.

The OVF Template *Disk Format* page opens.

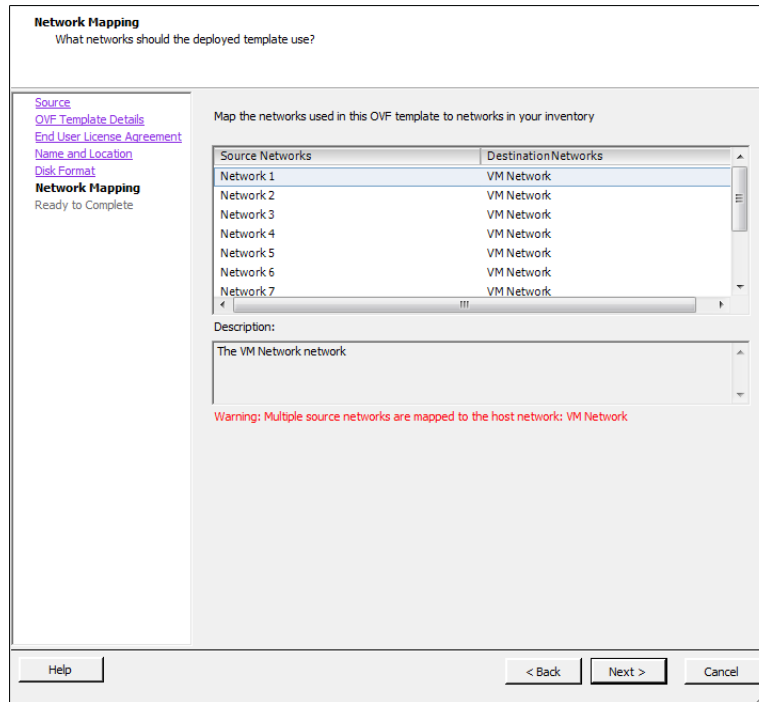
The screenshot shows the 'Disk Format' page of the OVF template wizard. The title is 'Disk Format' with the subtitle 'In which format do you want to store the virtual disks?'. On the left, there is a navigation pane with links: 'Source', 'OVF Template Details', 'End User License Agreement', 'Name and Location', 'Disk Format' (which is highlighted), 'Network Mapping', and 'Ready to Complete'. The main area contains a 'Datastore:' text box with the value 'datastore1' and an 'Available space (GB):' text box with the value '394.6'. Below these, there are three radio button options: 'Thick Provision Lazy Zeroed' (which is selected), 'Thick Provision Eager Zeroed', and 'Thin Provision'. At the bottom, there are three buttons: 'Help', '< Back', and 'Next >', and a 'Cancel' button.

7. Select one of the following:
  - *Thick Provision Lazy Zeroed*: Allocates the disk space statically (no other volumes can take the space), but does not write zeros to the blocks until the first write takes place to that block during runtime (which includes a full disk format).

- *Thick Provision Eager Zeroed*: Allocates the disk space statically (no other volumes can take the space), and writes zeros to all the blocks.
- *Thin Provision*: Allocates the disk space only when a write occurs to a block, but the total volume size is reported by VMFS to the OS. Other volumes can take the remaining space. This allows you to float space between your servers, and expand your storage when your size monitoring indicates there is a problem. Note that once a Thin Provisioned block is allocated, it remains on the volume regardless if you have deleted data, etc.

8. Select *Next* to continue.

The OVF Template *Network Mapping* page opens.



9. Map the networks used in this OVF template to networks in your inventory. Network 1 maps to port1 of the FortiGate VM. You must set the destination network for this entry to access the device console. Select *Next* to continue.


The OVF Template *Ready to Complete* page opens.

10. Review the template configuration. Make sure that *Power on after deployment* is not enabled. You might need to configure the FortiGate VM hardware settings prior to powering on the FortiGate VM.

11. Select *Finish* to deploy the OVF template. You will receive a *Deployment Completed Successfully* dialog box once the FortiGate VM OVF template wizard has finished.

## Configure FortiGate VM hardware settings

Before powering on your FortiGate VM you must configure the virtual memory, virtual CPU, and virtual disk configuration to match your FortiGate VM license. See [Table 111 on page 2446](#) for FortiGate VM model information.



## Transparent mode configuration

If you want to use your FortiGate-VM in transparent mode, your VMware server's virtual switches must operate in promiscuous mode. This permits these interfaces to receive traffic that will pass through the FortiGate unit but was not addressed to the FortiGate unit.

In VMware, promiscuous mode must be explicitly enabled:

1. In the vSphere client, select your VMware server in the left pane and then select the *Configuration* tab in the right pane.
2. In *Hardware*, select *Networking*.
3. Select *Properties* of vSwitch0.
4. In the *Properties* window left pane, select *vSwitch* and then select *Edit*.
5. Select the *Security* tab, set *Promiscuous Mode* to *Accept*, then select *OK*.
6. Select *Close*.
7. Repeat steps 3 through 6 for other vSwitches that your transparent mode FortiGate-VM uses.

## Power on your FortiGate VM

You can now proceed to power on your FortiGate VM. There are several ways to do this:

- Select the name of the FortiGate VM you deployed in the inventory list and select *Power on the virtual machine* in the *Getting Started* tab.
- In the inventory list, right-click the name of the FortiGate VM you deployed, and select *Power > Power On*.
- Select the name of the FortiGate VM you deployed in the inventory list. Click the *Power On* button on the toolbar.

Select the Console tab to view the console. To enter text, you must click in the console pane. The mouse is then captured and cannot leave the console screen. As the FortiGate console is text-only, no mouse pointer is visible. To release the mouse, press Ctrl-Alt.



# Deployment example: MS Hyper-V

Once you have downloaded the `.hyperv.zip` file and extracted the package contents to a folder on your Microsoft server, you can deploy the VHD package to your Microsoft Hyper-V environment.

The following topics are included in this section:

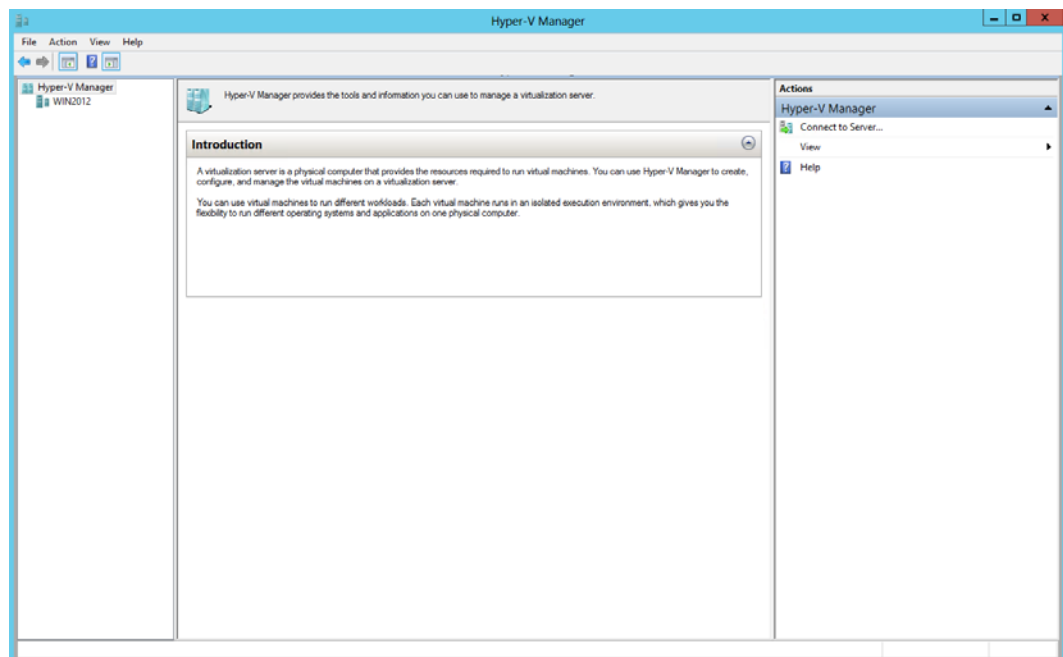
- [Create the FortiGate VM virtual machine](#)
- [Configure FortiGate VM hardware settings](#)
- [Start the FortiGate VM](#)

## Create the FortiGate VM virtual machine

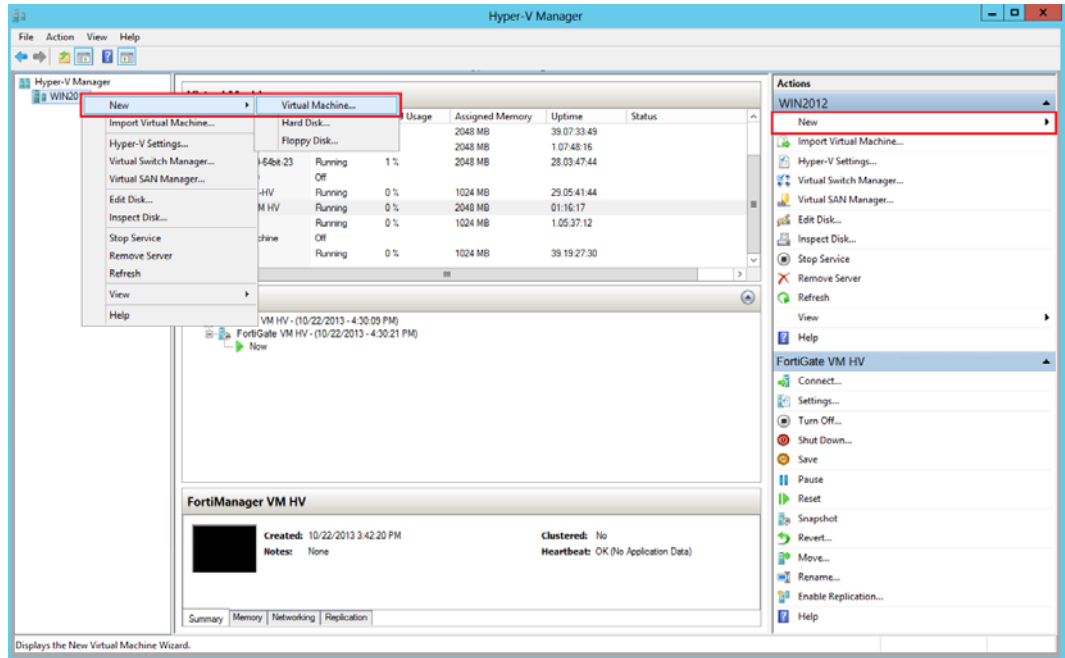
**To create the FortiGate VM virtual machine:**

1. Launch the Hyper-V Manager in your Microsoft server.

The *Hyper-V Manager* home page opens.

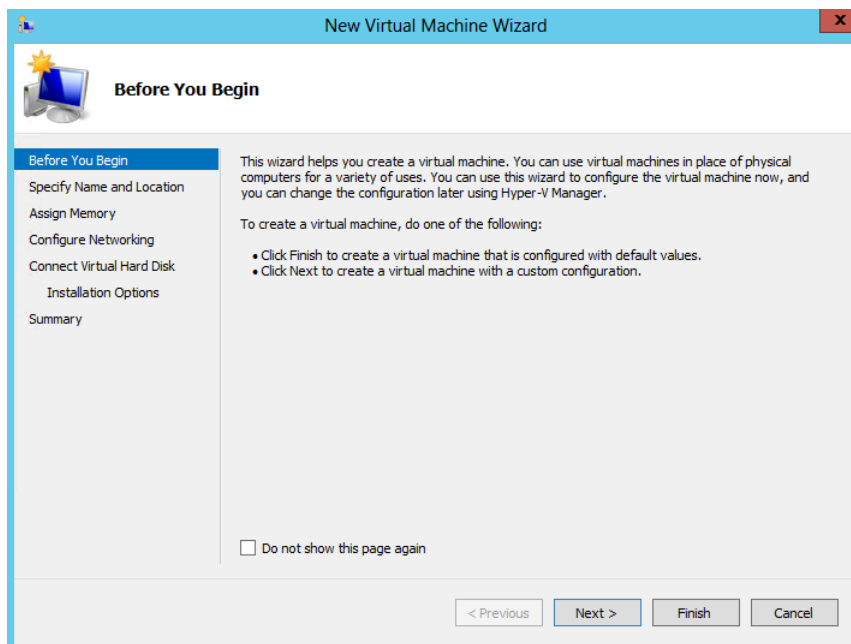


2. Select the server in the right-tree menu. The server details page is displayed.

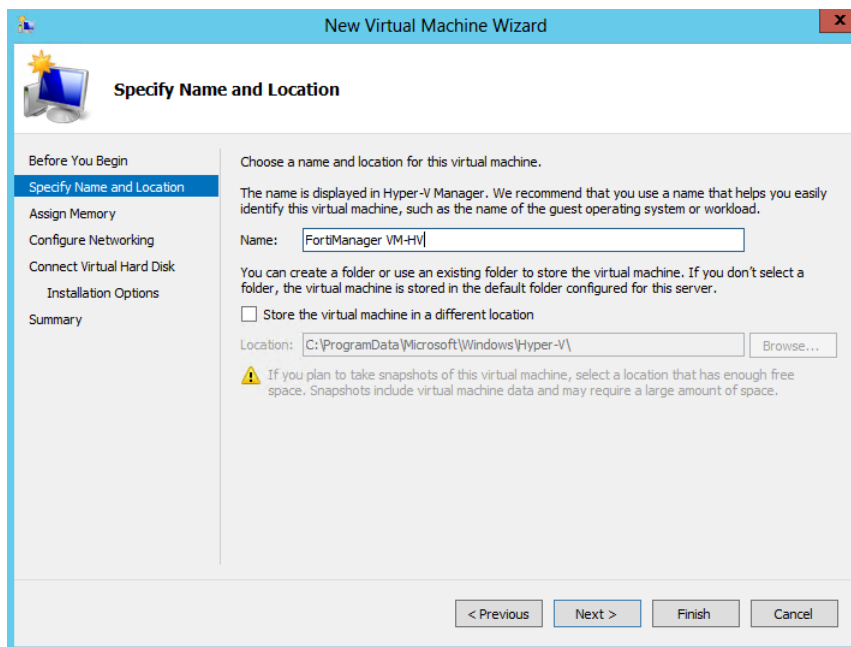


3. Right-click the server and select *New* and select *Virtual Machine* from the menu. Optionally, in the *Actions* menu, select *New* and select *Virtual Machine* from the menu.

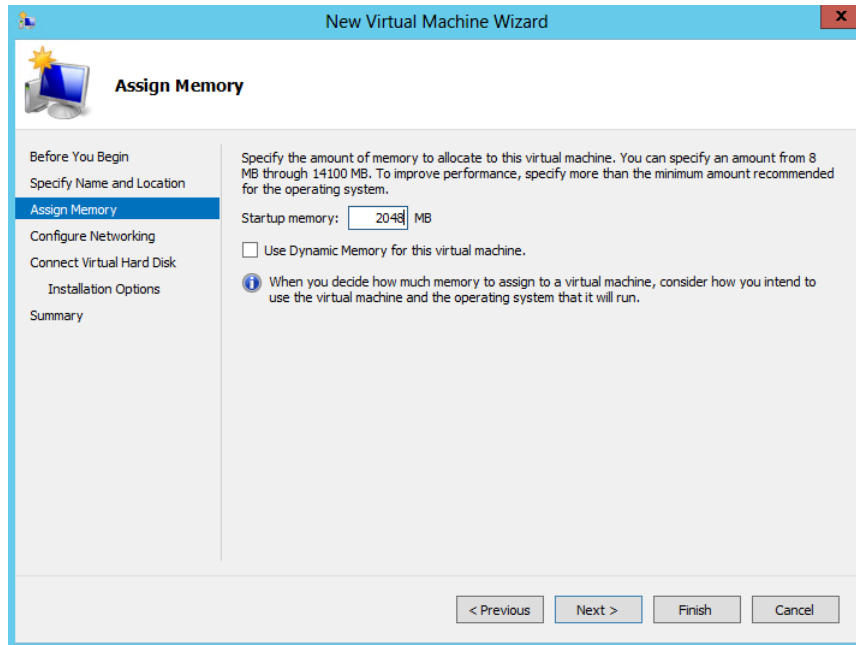
The *New Virtual Machine Wizard* opens.



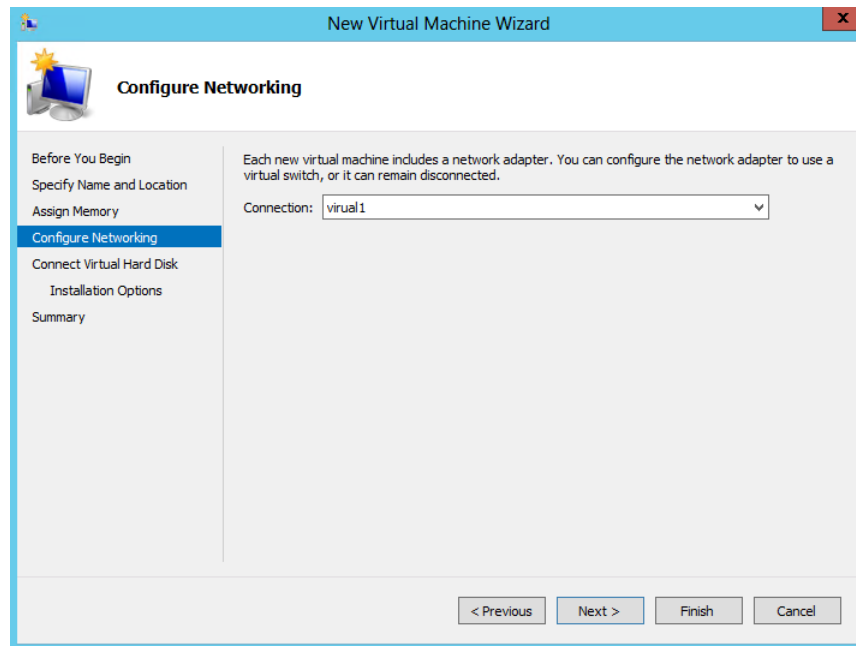
4. Select *Next* to create a virtual machine with a custom configuration. The *Specify Name and Location* page is displayed.



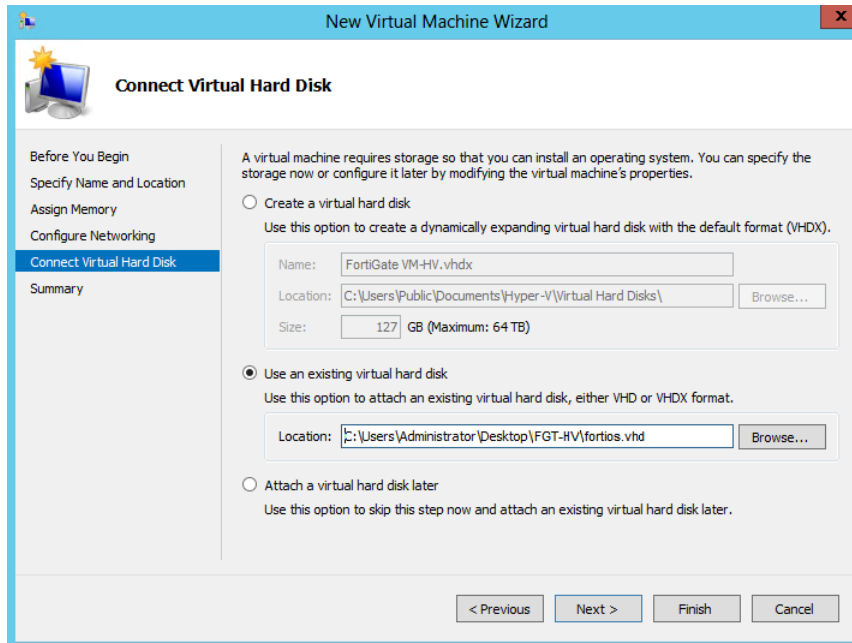
5. Enter a name for this virtual machine. The name is displayed in the Hyper-V Manager. Select *Next* to continue. The *Assign Memory* page is displayed.



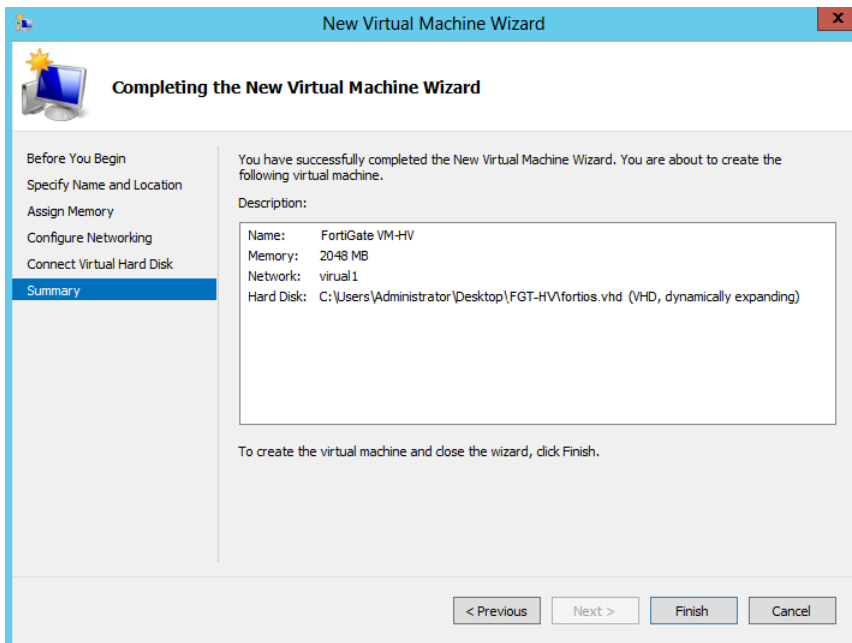
6. Specify the amount of memory to allocate to this virtual machine. The default memory for FortiGate VM is 1GB (1024MB). Select *Next* to continue. The *Configure Networking* page is displayed.



- Each new virtual machine includes a network adapter. You can configure the network adapter to use a virtual switch, or it can remain disconnected. FortiGate VM requires four network adapters. You must configure network adapters in the *Settings* page. Select *Next* to continue. The *Connect Virtual Hard Disk* page is displayed.



- Select to use an existing virtual hard disk and browse for the `fmg.vhd` file that you downloaded from the [Fortinet Customer Service & Support portal](#). Select *Next* to continue. The *Summary* page is displayed.



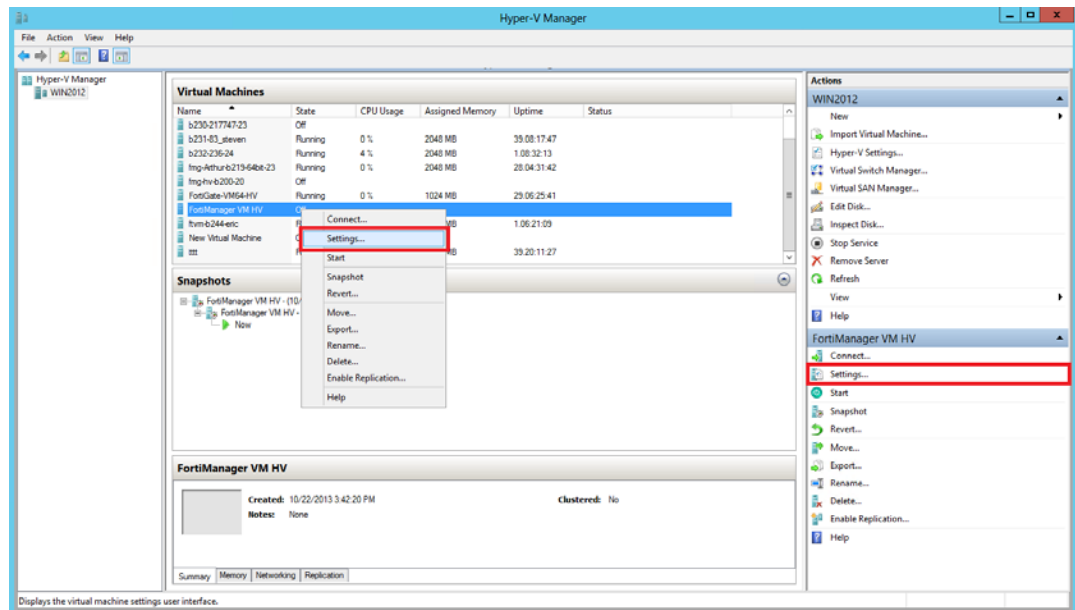
- To create the virtual machine and close the wizard, select *Finish*.

## Configure FortiGate VM hardware settings


Before powering on your FortiGate VM you must configure the virtual memory, virtual CPU, and virtual disk configuration to match your FortiGate VM license. See [Table 111 on page 2446](#) for FortiGate VM model information.

### To configure settings for FortiGate VM on the server:

1. In the Hyper-V Manager, locate the name of the virtual machine, right-click the entry, and select *Settings* from the menu. Optionally, you can select the virtual machine and select *Settings* in the *Actions* menu.



The *Settings* page is displayed.



2. Configure virtual processors, network adapters, and virtual hard drive settings.
3. Select *Apply* to save the settings and then select *OK* to close the settings page.

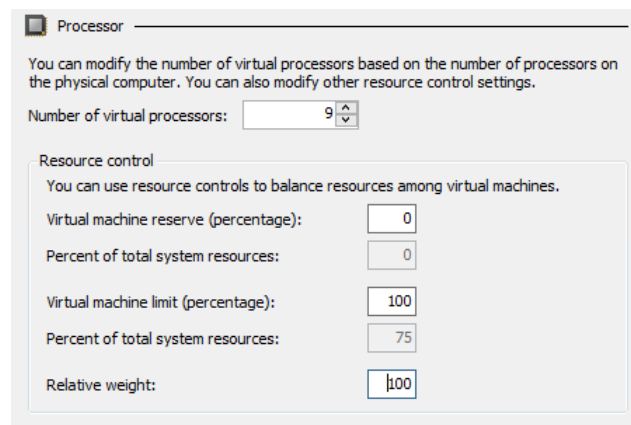
## FortiGate VM virtual processors

You must configure FortiGate VM virtual processors in the server settings page. The number of processors is dependent on your server environment.

### Configure FortiGate VM virtual processors:

1. In the *Settings* page, select *Processor* from the *Hardware* menu.

The *Processor* page is displayed.



The screenshot shows the 'Processor' configuration page. At the top, it says 'Processor' and provides instructions: 'You can modify the number of virtual processors based on the number of processors on the physical computer. You can also modify other resource control settings.' Below this, there is a 'Number of virtual processors' field with a value of 9 and a dropdown arrow. Underneath is a 'Resource control' section with the instruction: 'You can use resource controls to balance resources among virtual machines.' This section contains five fields: 'Virtual machine reserve (percentage)' with a value of 0, 'Percent of total system resources' with a value of 0, 'Virtual machine limit (percentage)' with a value of 100, 'Percent of total system resources' with a value of 75, and 'Relative weight' with a value of 100.

2. Configure the number of virtual processors for the FortiGate VM virtual machine. Optionally, you can use resource controls to balance resources among virtual machines.
3. Select *Apply* to save the settings.

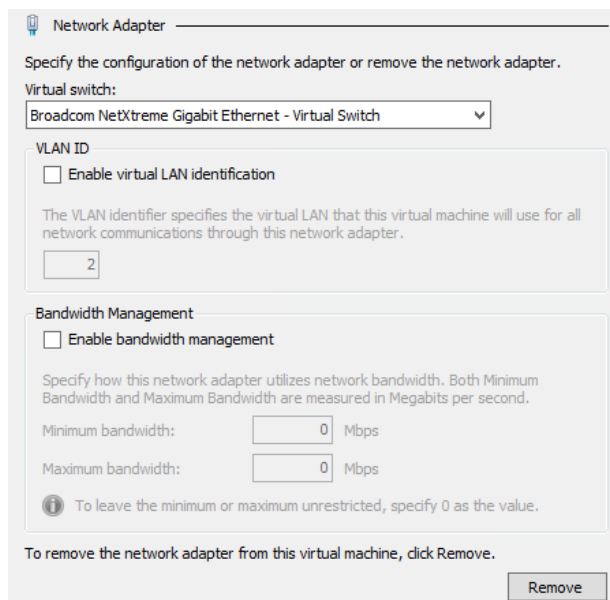
## FortiGate VM network adapters

You must configure FortiGate VM network adapters in the server settings page. FortiGate VM supports four network adapters.

### Configure FortiGate VM network adapters:

1. In the *Settings* page, select *Add Hardware* from the *Hardware* menu, select *Network Adapter* in the device list, and select the *Add* button.

The *Network Adapter* page is displayed.



The screenshot shows the 'Network Adapter' configuration page. It starts with the title 'Network Adapter' and the instruction: 'Specify the configuration of the network adapter or remove the network adapter.' Below this is a 'Virtual switch:' dropdown menu with 'Broadcom NetXtreme Gigabit Ethernet - Virtual Switch' selected. Underneath is a 'VLAN ID' section with a checkbox for 'Enable virtual LAN identification' which is unchecked. Below the checkbox is a text box containing the value '2'. The next section is 'Bandwidth Management' with a checkbox for 'Enable bandwidth management' which is unchecked. Below this is a text box for 'Minimum bandwidth:' with a value of '0' and 'Mbps' next to it. Below that is a text box for 'Maximum bandwidth:' with a value of '0' and 'Mbps' next to it. At the bottom of this section is an information icon and the text: 'To leave the minimum or maximum unrestricted, specify 0 as the value.' At the very bottom of the page, there is a 'Remove' button and the text: 'To remove the network adapter from this virtual machine, click Remove.'

2. You must manually configure four network adapters for FortiGate VM in the settings page. For each network adapter, select the virtual switch from the drop-down list.
3. Select *Apply* to save the settings.

## FortiGate VM virtual hard disk

You must configure the FortiGate VM virtual hard disk in the server settings page.

If you know your environment will expand in the future, it is recommended to increase the hard disk size beyond 30GB. The VM license limit is 2TB.

### Configure a FortiGate VM virtual hard drive:

1. In the *Settings* page, select *IDE Controller 0 > Hard Drive* from the *Hardware* menu.

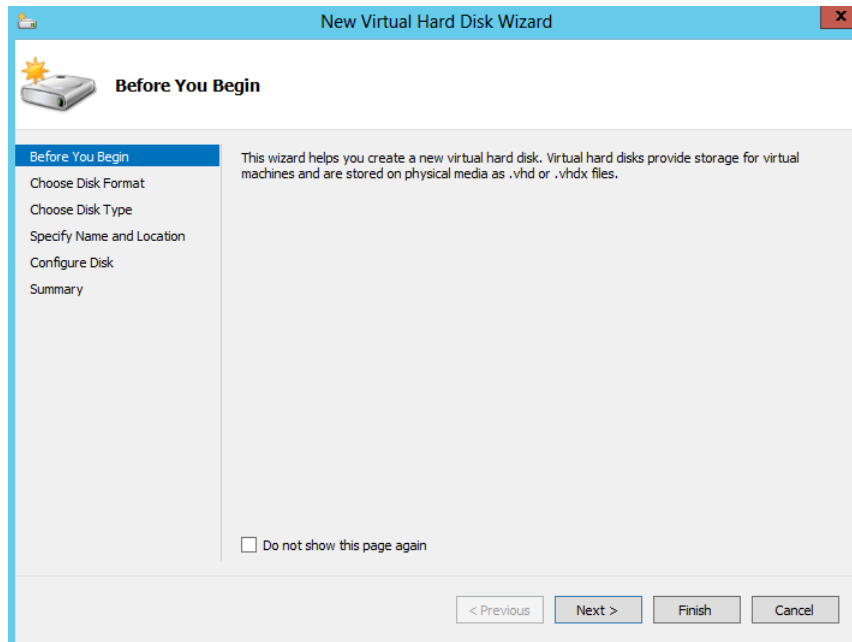
The *Hard Drive* page is displayed.

The screenshot shows the 'Hard Drive' configuration window. At the top, it says 'Hard Drive' and provides a warning: 'You can change how this virtual hard disk is attached to the virtual machine. If an operating system is installed on this disk, changing the attachment might prevent the virtual machine from starting.' Below this, there are two dropdown menus: 'Controller:' set to 'IDE Controller 0' and 'Location:' set to '0 (in use)'. The 'Media' section has a sub-header 'Media' and a note: 'You can compact or convert a virtual hard disk by editing the associated file. Specify the full path to the file.' There are two radio button options: 'Virtual hard disk:' (selected) and 'Physical hard disk:'. The 'Virtual hard disk:' option has a text input field containing 'C:\Users\Administrator\Desktop\FMG-HV\fmv.vhd' and four buttons: 'New', 'Edit', 'Inspect', and 'Browse...'. The 'Physical hard disk:' option has a dropdown menu. At the bottom, there is an information icon and a note: 'If the physical hard disk you want to use is not listed, make sure that the disk is offline. Use Disk Management on the physical computer to manage physical hard disks.' Finally, at the very bottom, there is a 'Remove' button and a note: 'To remove the virtual hard disk, click Remove. This disconnects the disk but does not delete the associated file.'



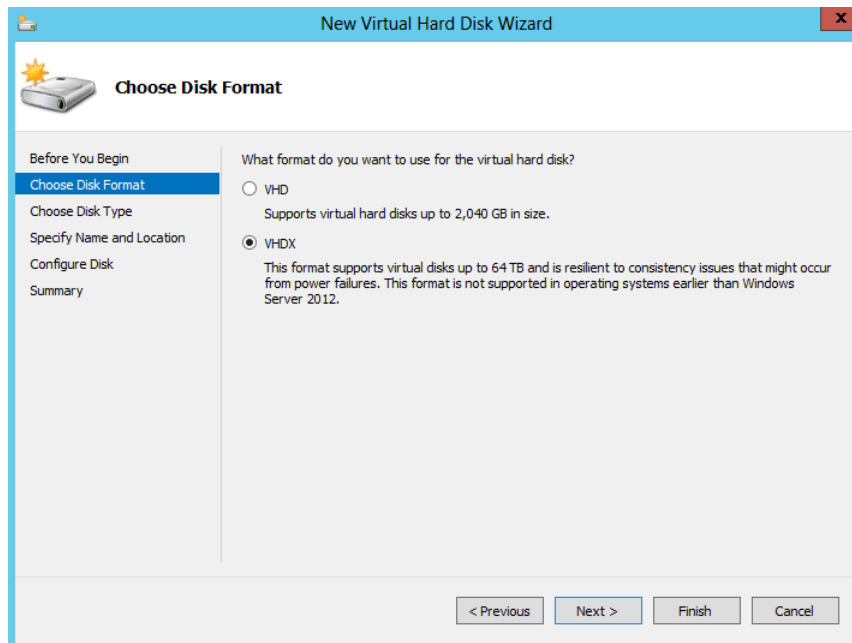
2. Select *New* to create a new virtual hard disk.

The *New Virtual Hard Disk Wizard* opens.



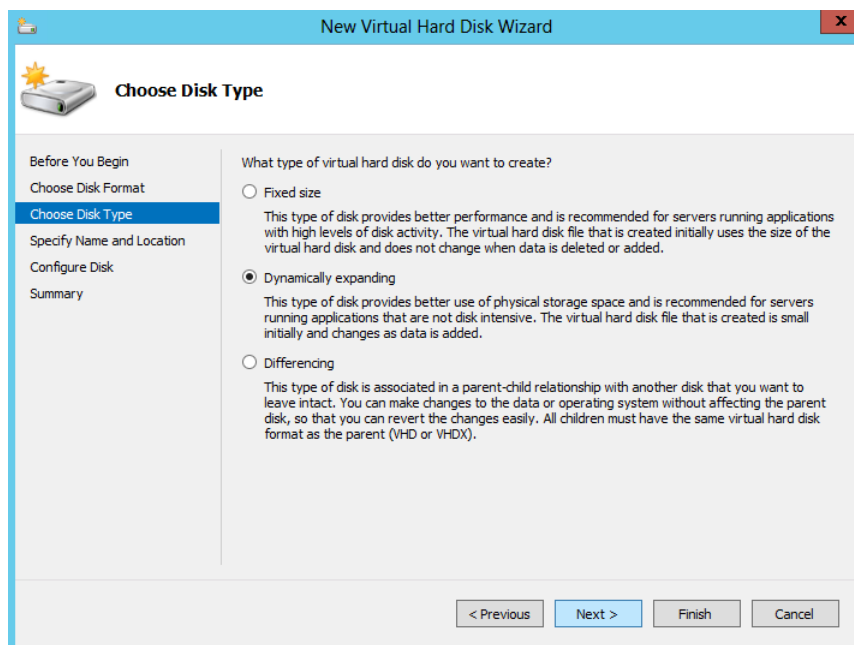
3. This wizard helps you to create a new virtual hard disk.

Select *Next* to continue. The *Choose Disk Format* page opens.



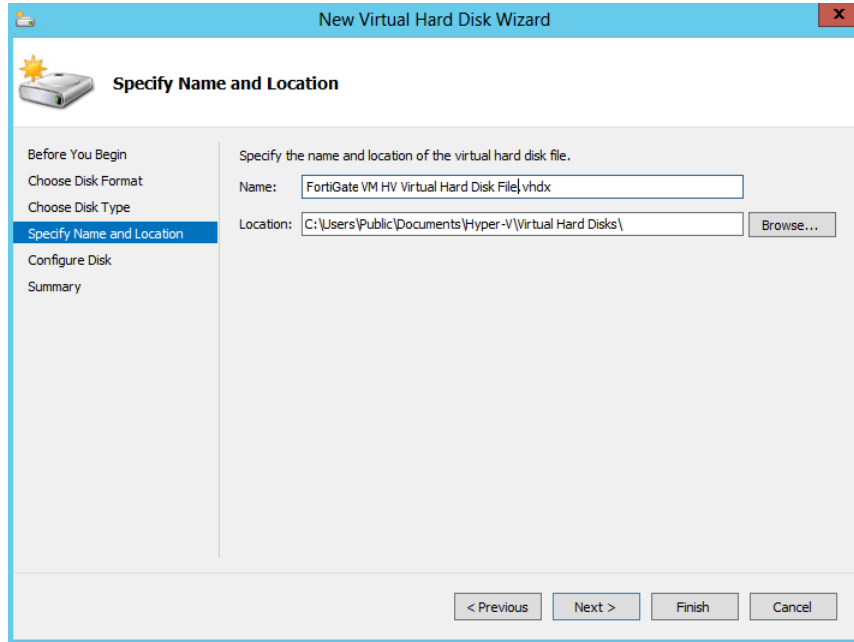
4. Select to use VHDX format virtual hard disks. This format supports virtual disks up to 64TB and is resilient to consistency issues that might occur from power failures. This format is not supported in operating systems earlier than Windows Server 2012. Note that FortiGate-VM does not support hard disks larger than 2TB.

Select *Next* to continue. The *Choose Disk Type* page opens.



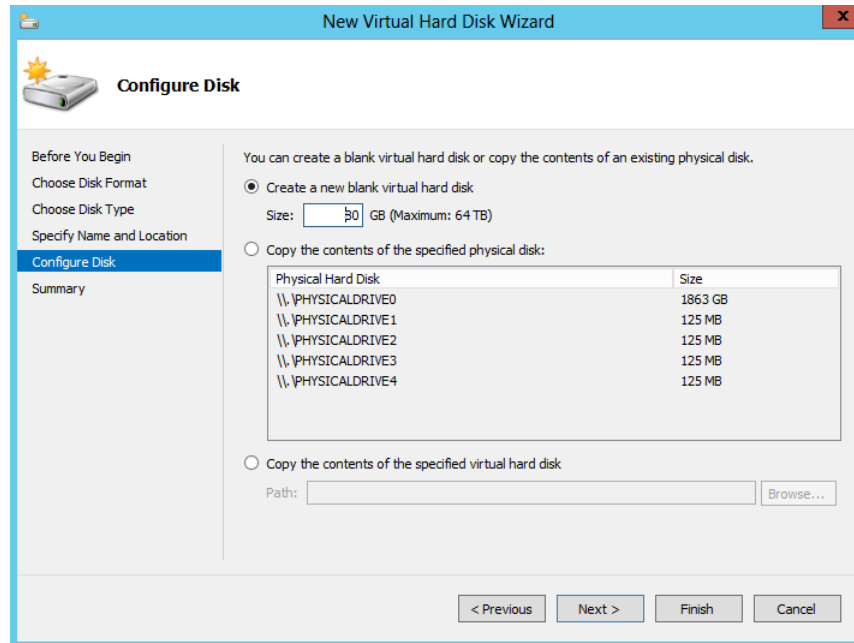
5. Select the type of virtual disk you want to use. Select one of the following disk types:
  - Fixed size: This type of disk provides better performance and is recommended for servers running applications with high levels of disk activity. The virtual hard disk file that is created initially uses the size of the virtual hard disk and does not change when data is deleted or added.
  - Dynamic expanding: This type of disk provides better use of physical storage space and is recommended for servers running applications that are not disk intensive. The virtual disk file that is created is small initially and changes as data is added.
  - Differencing: This type of disk is associated in a parent-child relationship with another disk that you want to leave intact. You can make changes to the data or operating system without affecting the parent disk, so that you can revert the changes easily. All children must have the same virtual hard disk format as the parent (VHD or VHDX).

Select *Next* to continue. The *Specify Name and Location* page opens.

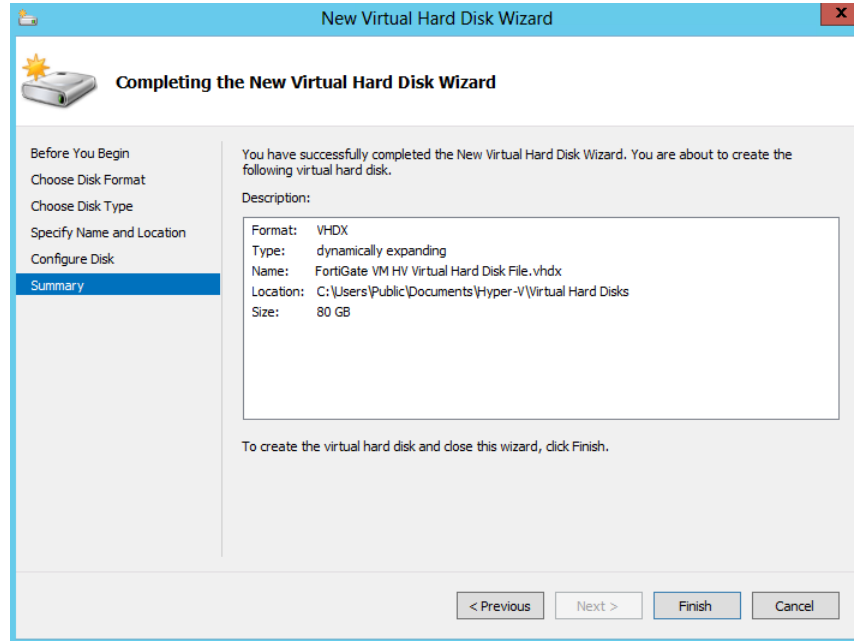


6. Specify the name and location of the virtual hard disk file. Use the *Browse* button to select a specific file folder on your server.

Select *Next* to continue. The *Configure Disk* page opens.



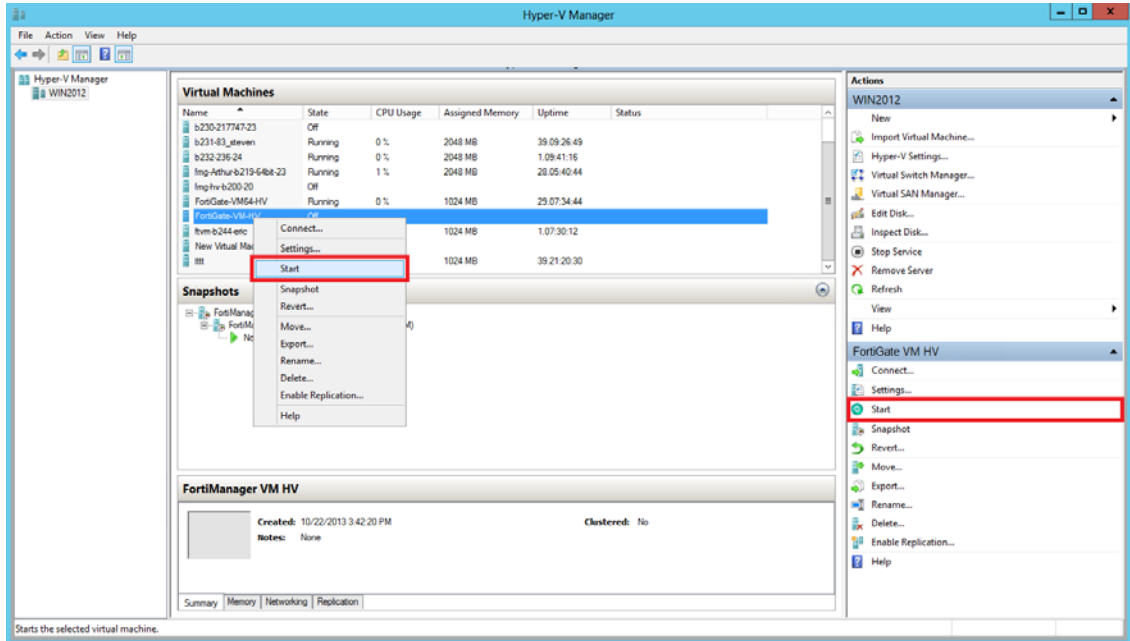
7. Select to *Create a new blank virtual hard disk* and enter the size of the disk in GB. The maximum size is dependent on your server environment. Select *Next* to continue. The *Summary* page opens.



8. The summary page provides details of the virtual hard disk. Select *Finish* to create the virtual hard disk.
9. Select *Apply* to save the settings and select *OK* to exit the settings page.

## Start the FortiGate VM

You can now proceed to power on your FortiGate VM. Select the name of the FortiGate VM in the list of virtual machines, right-click, and select *Start* in the menu. Optionally, you can select the name of the FortiGate VM in the list of virtual machines and select *Start* in the *Actions* menu.



# Deployment example: KVM

Once you have downloaded the FORTINET.out.kvm.zip file and extracted virtual hard drive image file fortios.qcow2, you can create the virtual machine in your KVM environment.

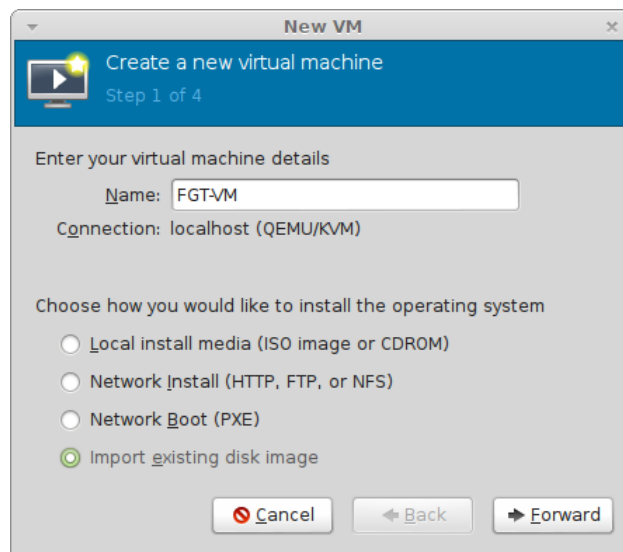
The following topics are included in this section:

- [Create the FortiGate VM virtual machine](#)
- [Configure FortiGate VM hardware settings](#)
- [Start the FortiGate VM](#)

## Create the FortiGate VM virtual machine

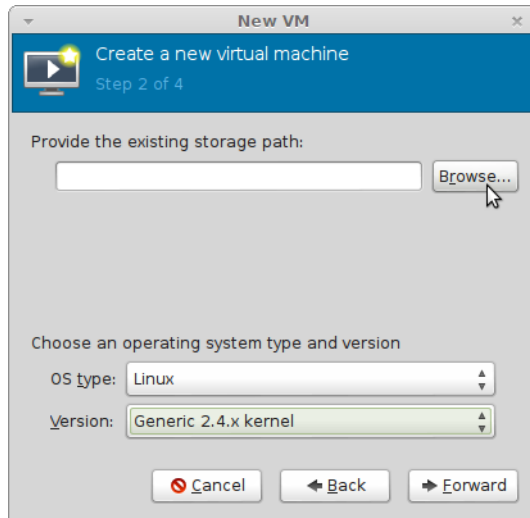
**To create the FortiGate VM virtual machine:**

1. Launch Virtual Machine Manager (virt-manager) on your KVM host server.  
The *Virtual Machine Manager* home page opens.
2. In the toolbar, select *Create a new virtual machine*.



3. Enter a *Name* for the VM, FGT-VM for example.
4. Ensure that *Connection* is localhost. (This is the default.)
5. Select *Import existing disk image*.


6. Select *Forward*.



7. In *OS Type* select *Linux*.

8. In *Version*, select *Generic 2.4.x.kernel*.

9. Select *Browse*.



10. If you copied the `fortios.qcow2` file to `/var/lib/libvirt/images`, it will be visible on the right. If you saved it somewhere else on your server, select *Browse Local* and find it.

11. Choose *Choose Volume*.

12. Select *Forward*.

13. Specify the amount of memory and number of CPUs to allocate to this virtual machine. The amounts must not exceed your license limits. See [“FortiGate VM models and licensing” on page 2446](#).

14. Select *Forward*.

15. Expand *Advanced options*. A new virtual machine includes one network adapter by default. Select a network adapter on the host computer. Optionally, set a specific MAC address for the virtual network interface. Set *Virt Type* to *virtio* and *Architecture* to *qcow2*.

16. Select *Finish*.

## Configure FortiGate VM hardware settings

Before powering on your FortiGate VM you must add the log disk and configure the virtual hardware of your FortiGate VM.

### To configure settings for FortiGate VM on the server:

1. In the Virtual Machine Manager, locate the name of the virtual machine and then select *Open* from the toolbar.
2. Select *Add Hardware*. In the *Add Hardware* window select *Storage*.
3. Select *Create a disk image on the computer's harddrive* and set the size to 30GB.



If you know your environment will expand in the future, it is recommended to increase the hard disk size beyond 30GB. The VM license limit is 2TB.

- 
4. Enter:

<b>Device type</b>	Virtio disk
<b>Cache mode</b>	Default
<b>Storage format</b>	raw

5. Select *Network* to configure add more the network interfaces. The *Device type* must be *Virtio*.

A new virtual machine includes one network adapter by default. You can add more through the *Add Hardware* window. FortiGate VM requires four network adapters. You can configure network adapters to connect to a virtual switch or to network adapters on the host computer.

6. Select *Finish*.

## Start the FortiGate VM

You can now proceed to power on your FortiGate VM. Select the name of the FortiGate VM in the list of virtual machines. In the toolbar, select *Console* and then select *Start*.



# Deployment example: OpenXen

Once you have downloaded the FORTINET.out.OpenXen.zip file and extracted virtual hard drive image file fortios.qcow2, you can create the virtual machine in your OpenXen environment.

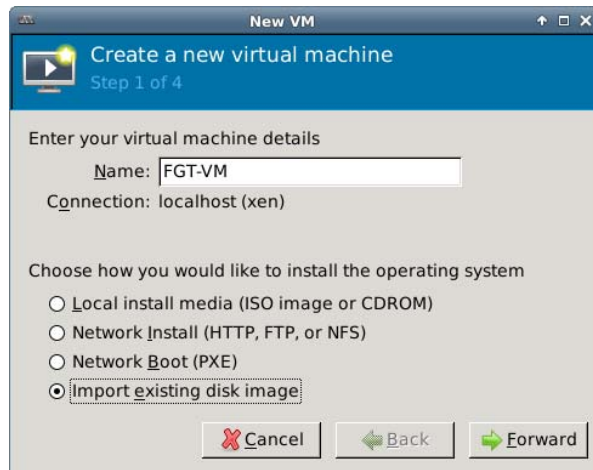
The following topics are included in this section:

- [Create the FortiGate VM virtual machine \(VMM\)](#)

## Create the FortiGate VM virtual machine (VMM)

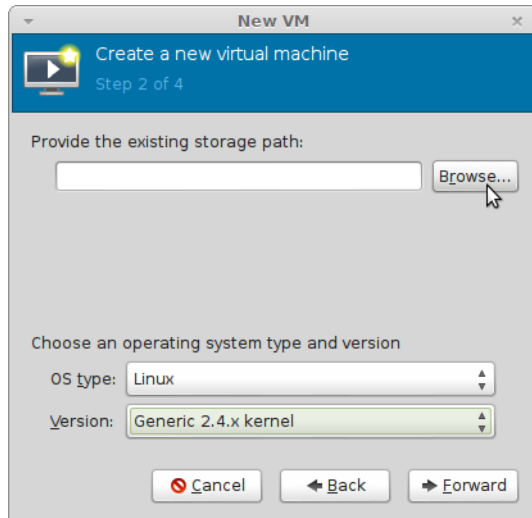
**To create the FortiGate VM virtual machine:**

1. Launch Virtual Machine Manager (virt-manager) on your OpenXen host server.  
The *Virtual Machine Manager* home page opens.
2. In the toolbar, select *Create a new virtual machine*.



3. Enter a *Name* for the VM, FGT-VM for example.
4. Ensure that *Connection* is localhost. (This is the default.)
5. Select *Import existing disk image*.

6. Select *Forward*.



7. In *OS Type* select *Linux*.

8. In *Version*, select *Generic 2.4.x.kernel*.

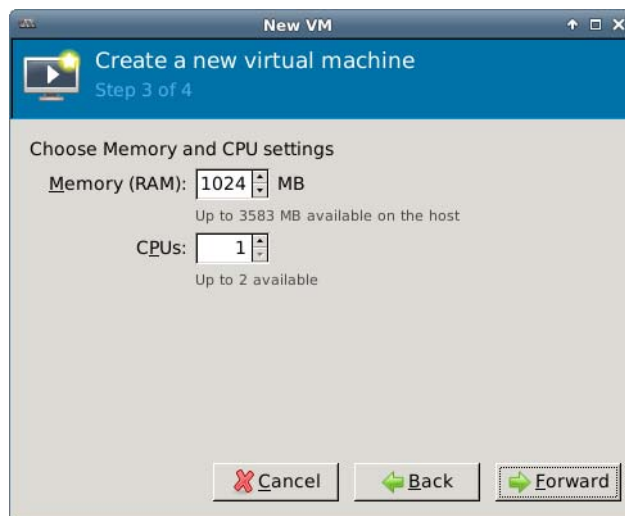
9. Select *Browse*.

The *Locate or create storage volume* window opens.

10. Select *Browse Local*, find the *fortios.qcow2* disk image file.

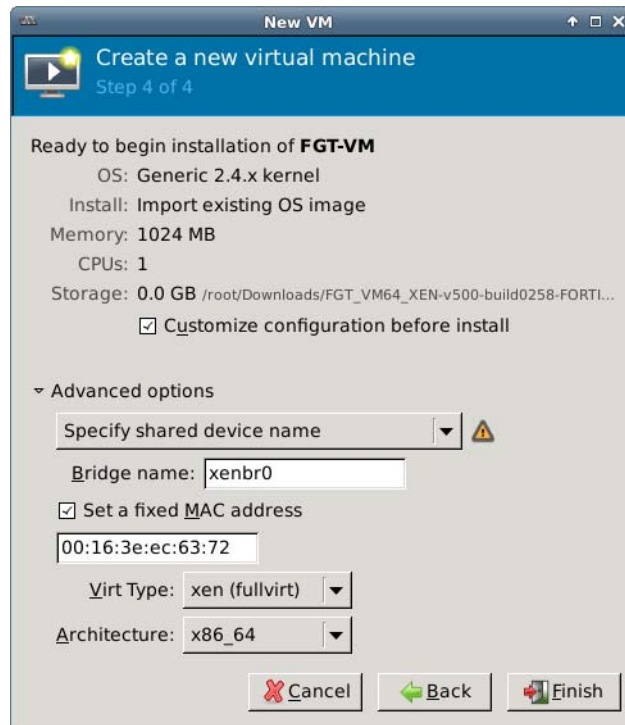
11. Select *fortios.qcow2* and select *Choose Volume*.

12. Select *Forward*.



13. Specify the amount of memory and number of CPUs to allocate to this virtual machine. The amounts must not exceed your license limits. See [“FortiGate VM models and licensing”](#) on page 2446.

14. Select *Forward*.



15. Select *Customize configuration before install*. This enables you to make some hardware configuration changes before VM creation is started.

16. Expand *Advanced options*. A new virtual machine includes one network adapter by default. Select *Specify shared device name* and enter the name of the bridge interface on the OpenXen host. Optionally, set a specific MAC address for the virtual network interface. *Virt Type* and *Architecture* are set by default and should be correct.

17. Select *Finish*.

The virtual machine hardware configuration window opens.



You can use this window to add hardware such as network interfaces and disk drives.

18. Select *Add Hardware*. In the *Add Hardware* window select *Storage*.

19. Select *Create a disk image on the computer's harddrive* and set the size to 30GB.



If you know your environment will expand in the future, it is recommended to increase the hard disk size beyond 30GB. The VM license limit is 2TB.

20. Enter:

<b>Device type</b>	Virtio disk
<b>Cache mode</b>	Default
<b>Storage format</b>	raw

21. Select *Network* to configure add more the network interfaces. The *Device type* must be *Virtio*.

A new virtual machine includes one network adapter by default. You can add more through the Add Hardware window. FortiGate VM requires four network adapters. You can configure network adapters to connect to a virtual switch or to network adapters on the host computer.

22. Select *Finish*.

23. Select *Begin Installation*. After the installation completes successfully, the VM starts and the console window opens.

# Deployment example: Citrix XenServer

Once you have downloaded the FORTINET.out.CitrixXen.zip file and extracted the files, you can create the virtual machine in your Citrix Xen environment.

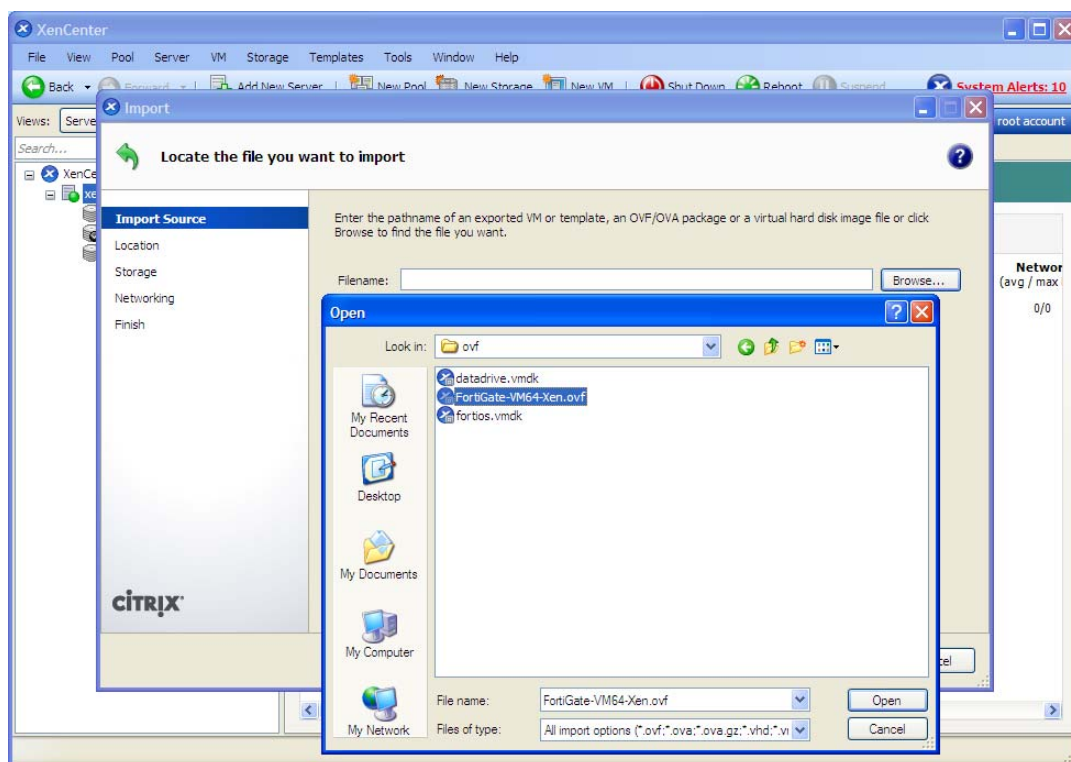
The following topics are included in this section:

- [Create the FortiGate VM virtual machine \(XenCenter\)](#)

## Create the FortiGate VM virtual machine (XenCenter)

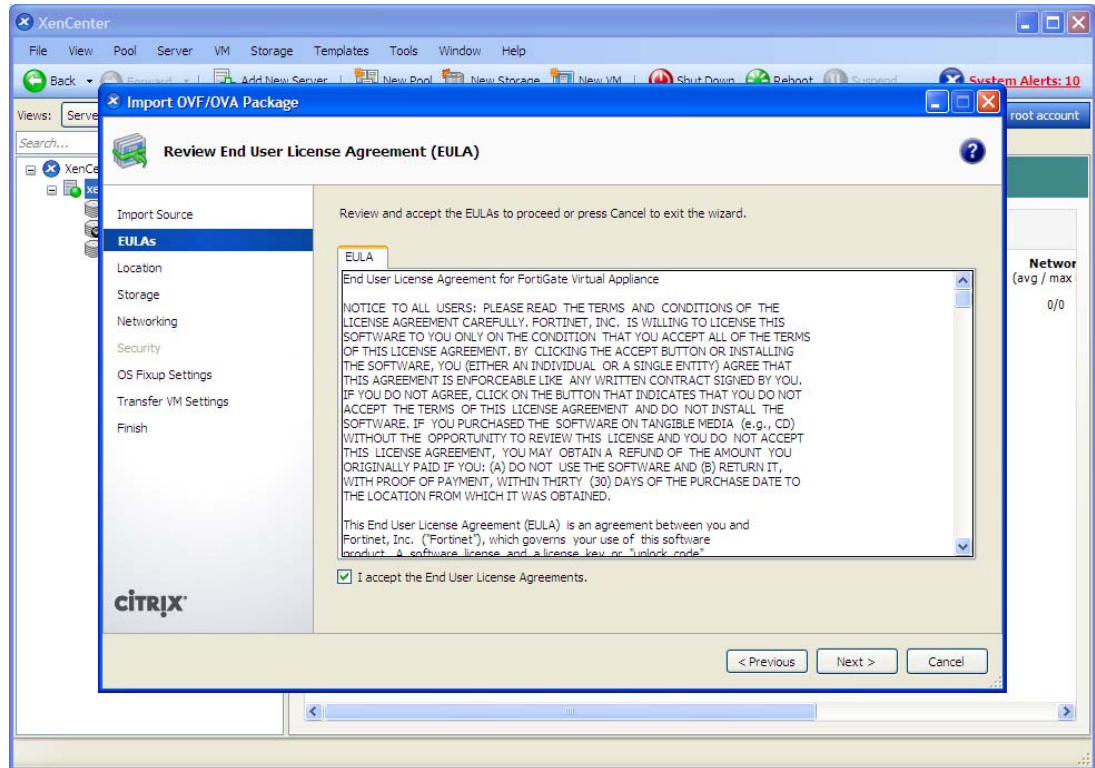
### To create the FortiGate VM virtual machine from the OVF file

1. Launch XenCenter on your management computer.  
The management computer can be any computer that can run Citrix XenCenter, a Windows application.
2. If you have not already done so, select *ADD a server*. Enter your Citrix XenServer IP address and the root logon credentials required to manage that server.  
Your Citrix XenServer is added to the list in the left pane.  
The *Virtual Machine Manager* home page opens.
3. Go to *File > Import*. An import dialog will appear.
4. Click the *Browse* button, find the FortiGate-VM64-Xen.ovf template file, then click *Open*.

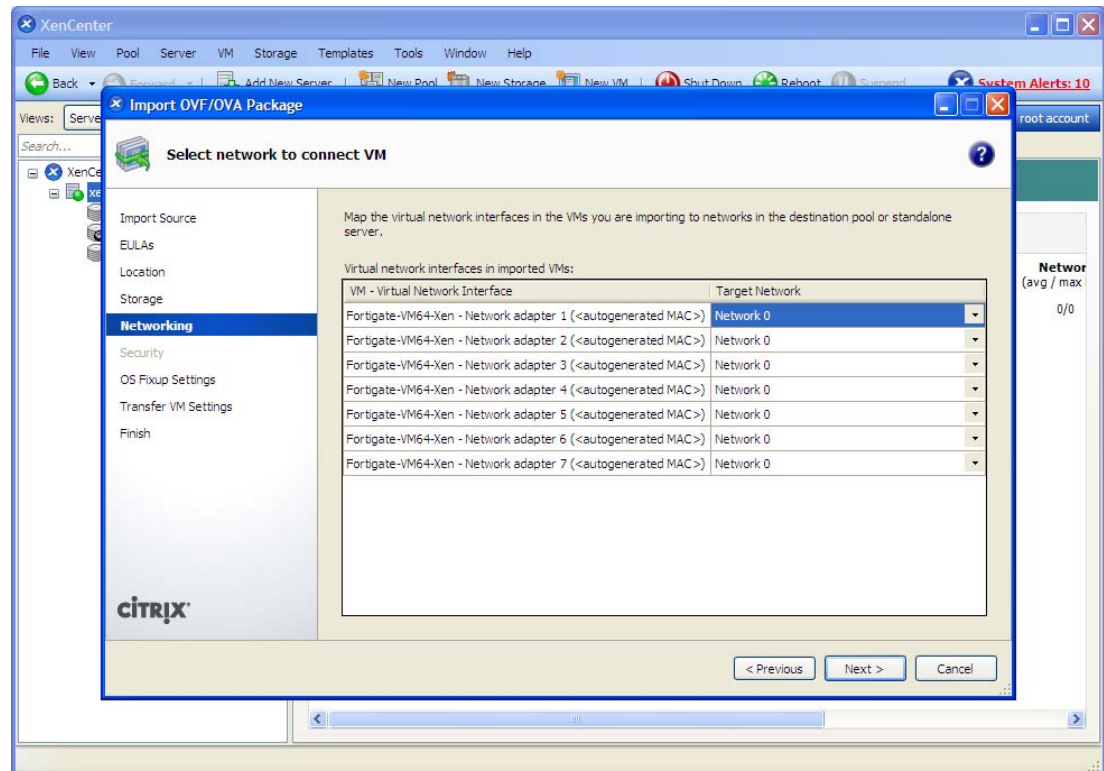


5. Select *Next*.

- Accept the FortiGate Virtual Appliance EULA, then select *Next*.



- Choose the pool or standalone server that will host the VM, then select *Next*.
- Select the storage location for FortiGate VM disk drives or accept the default. Select *Next*.
- Configure how each vNIC (virtual network adapter) in FortiGate VM will be mapped to each vNetwork on the Citrix XenServer, then click *Next*.



10. Click *Next* to skip OS fixup.
11. Select *Next* to use the default network settings for transferring the VM to the host.
12. select *Finish*.

The Citrix XenServer imports the FortiGate VM files and configures the VM as specified in the OVF template. Depending on your computer's hardware speed and resource load, and also on the file size and speed of the network connection, this might take several minutes to complete.



When VM import is complete, the XenCenter left pane includes the FortiGate VM in the list of deployed VMs for your Citrix XenServer.

## Configure virtual hardware

Before you start your FortiGate-VM for the first time, you need to adjust your virtual machine's virtual hardware settings to meet your network requirements.

### Configuring number of CPUs and memory size

Your FortiGate-VM license limits the number CPUs and amount of memory that you can use. The amounts you allocate must not exceed your license limits. See [“FortiGate VM models and licensing” on page 2446](#).

#### To access virtual machine settings

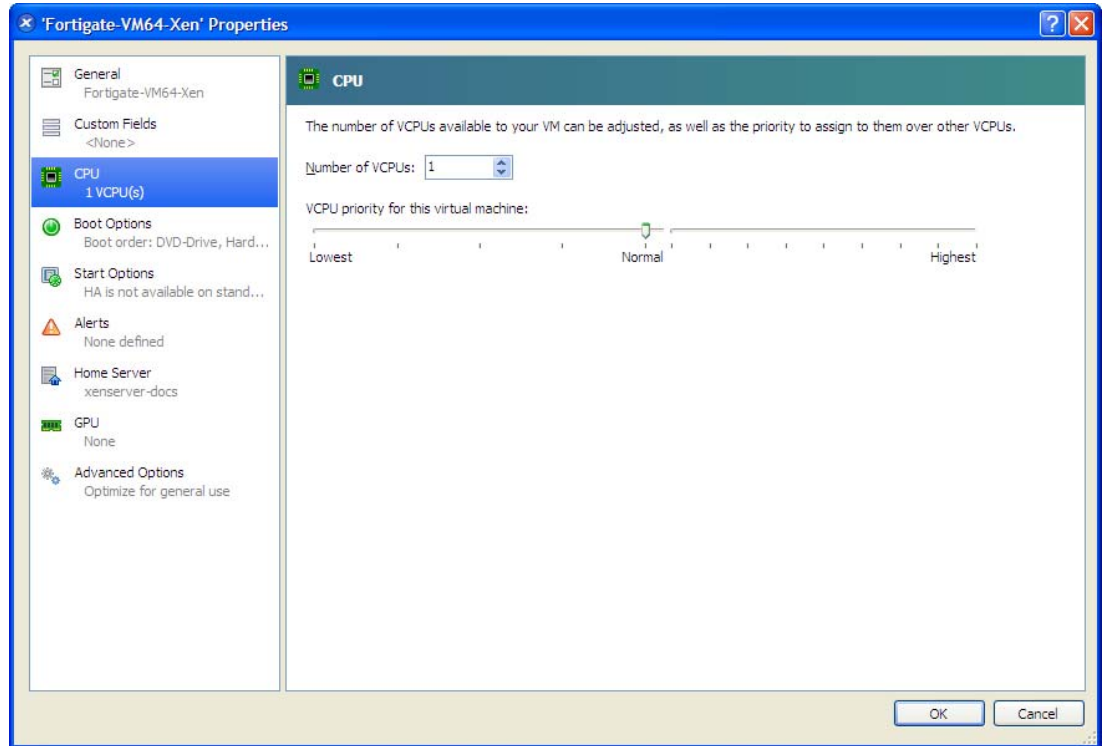
1. Open XenCenter.
2. Select your FortiGate VM in the left pane.

The tabs in the right pane provide access to the virtual hardware configuration. The Console tab provides access to the FortiGate console.

#### To set the number of CPUs

1. In the XenCenter left pane, right-click the FortiGate VM and select Properties.  
The Properties window opens.

2. In the left pane, select CPU.

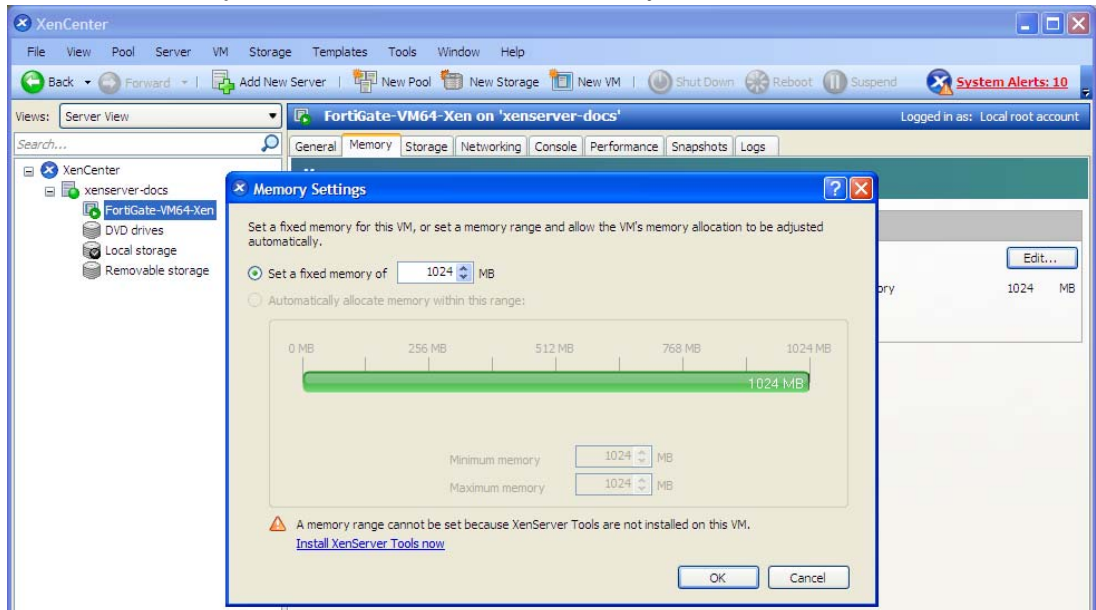


3. Adjust *Number of CPUs* and then select *OK*.

XenCenter will warn if you select more CPUs than the Xen host computer contains. Such a configuration might reduce performance.

### To set memory size

1. In the XenCenter left pane, select the FortiGate VM.
2. In the right pane, select the *Memory* tab.
3. Select *Edit*, modify the value in the *Set a fixed memory of* field and select *OK*.



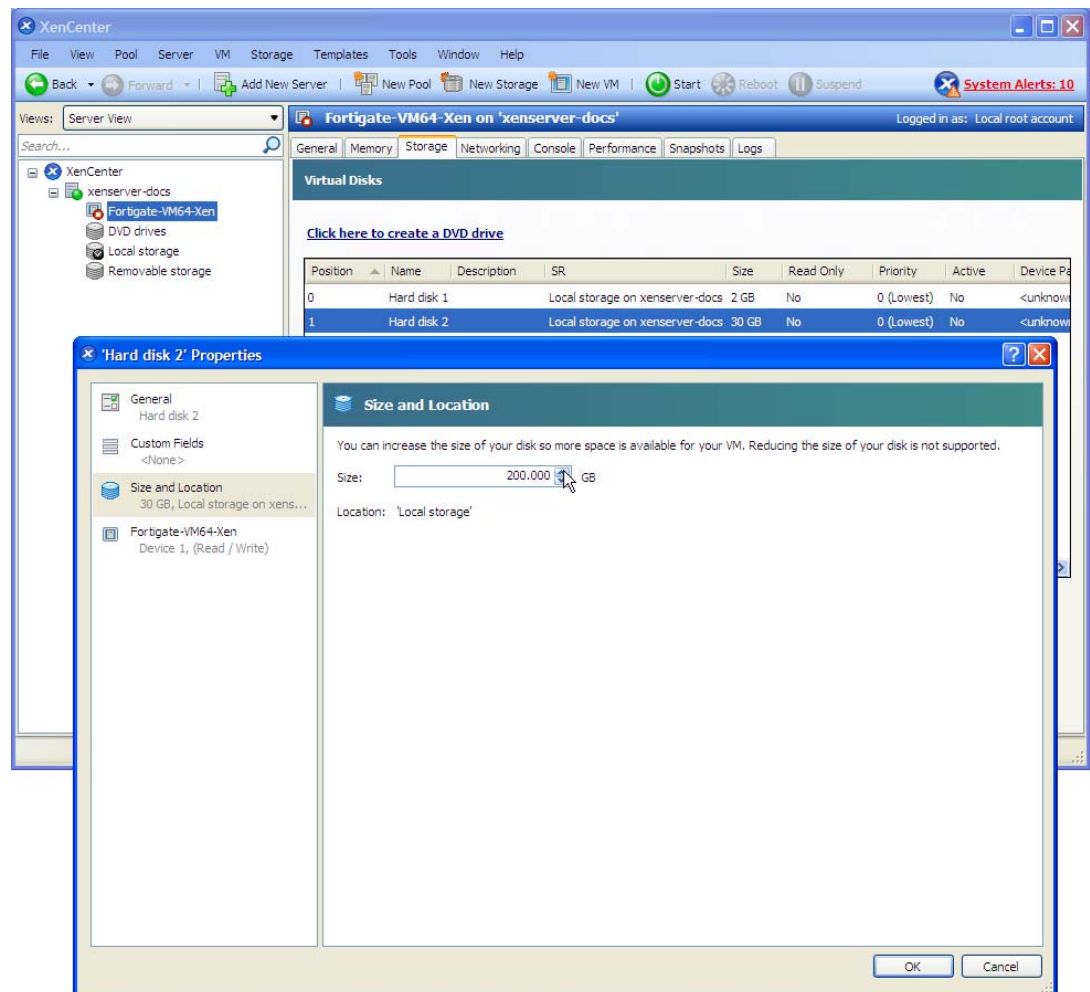


## Configuring disk storage

By default the FortiGate VM data disk 30GB. You will probably want to increase this. Disk resizing must be done before you start the VM for the first time.

### To resize the FortiGate data disk

1. In the XenCenter left pane, select the FortiGate VM.
2. Select the *Storage* tab. Select *Hard disk 2* (the 30GB drive), then select *Properties*.  
The 'Hard disk 2' Properties window opens.
3. Select *Size and Location*. Adjust *Size* and select *OK*.



# FortiGate VM Initial Configuration

Before you can connect to the FortiGate VM web-based manager you must configure a network interface in the FortiGate VM console. Once an interface with administrative access is configured, you can connect to the FortiGate VM web-based Manager and upload the FortiGate VM license file that you downloaded from the [Customer Service & Support](#) website.

The following topics are included in this section:

- [Set FortiGate VM port1 IP address](#)
- [Connect to the FortiGate VM Web-based Manager](#)
- [Upload the FortiGate VM license file](#)
- [Configure your FortiGate VM](#)

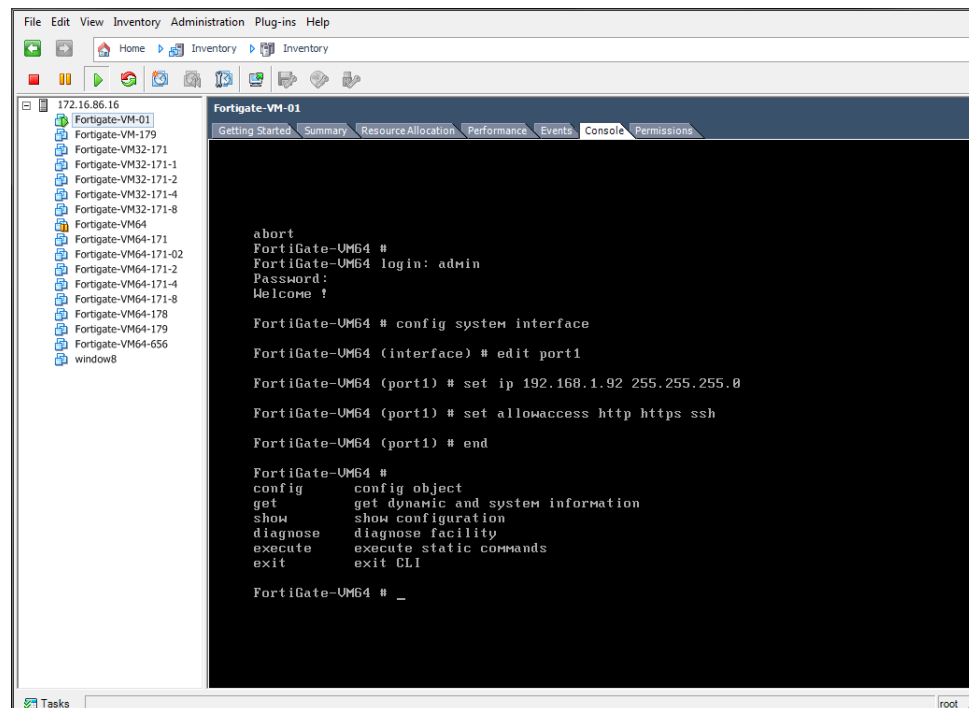
## Set FortiGate VM port1 IP address

Hypervisor management environments include a guest console window. On the FortiGate VM, this provides access to the FortiGate console, equivalent to the console port on a hardware FortiGate unit. Before you can access the Web-based manager, you must configure FortiGate VM port1 with an IP address and administrative access.

**To configure the port1 IP address:**

1. In your hypervisor manager, start the FortiGate VM and access the console window. You might need to press Return to see a login prompt.

**Figure 344:**Example of FortiGate VM console access



The screenshot shows a hypervisor management interface with a list of VMs on the left and a console window for 'Fortigate-VM-01' on the right. The console displays the following text:

```
172.16.86.16
Fortigate-VM-01
Fortigate-VM-179
Fortigate-VM32-171
Fortigate-VM32-171-1
Fortigate-VM32-171-2
Fortigate-VM32-171-4
Fortigate-VM32-171-8
Fortigate-VM64
Fortigate-VM64-171
Fortigate-VM64-171-02
Fortigate-VM64-171-2
Fortigate-VM64-171-4
Fortigate-VM64-171-8
Fortigate-VM64-178
Fortigate-VM64-179
Fortigate-VM64-656
window8

Fortigate-VM-01
Getting Started Summary Resource Allocation Performance Events Console Permissions

abort
FortiGate-UM64 #
FortiGate-UM64 login: admin
Password:
Welcome !

FortiGate-UM64 # config system interface
FortiGate-UM64 (interface) # edit port1
FortiGate-UM64 (port1) # set ip 192.168.1.92 255.255.255.0
FortiGate-UM64 (port1) # set allowaccess http https ssh
FortiGate-UM64 (port1) # end

FortiGate-UM64 #
config config object
get get dynamic and system information
show show configuration
diagnose diagnose facility
execute execute static commands
exit exit CLI

FortiGate-UM64 # _
```

2. At the FortiGate VM login prompt enter the username `admin`. By default there is no password. Just press Return.
3. Using CLI commands, configure the `port1` IP address and netmask. Also, HTTP access must be enabled because until it is licensed the FortiGate VM supports only low-strength encryption. HTTPS access will not work.

For example:

```
config system interface
 edit port1
 set ip 192.168.0.100 255.255.255.0
 append allowaccess http
 end
```



You can also use the `append allowaccess` CLI command to enable other access protocols, such as `auto-ipsec`, `http`, `probe-response`, `radius-acct`, `snmp`, and `telnet`. The `ping`, `https`, `ssh`, and `fgfm` protocols are enabled on the `port1` interface by default.

---

4. To configure the default gateway, enter the following CLI commands:

```
config router static
 edit 1
 set device port1
 set gateway <class_ip>
 end
```



You must configure the default gateway with an IPv4 address. FortiGate VM needs to access the Internet to contact the FortiGuard Distribution Network (FDN) to validate its license.

---

5. To configure your DNS servers, enter the following CLI commands:

```
config system dns
 set primary <Primary DNS server>
 set secondary <Secondary DNS server>
end
```



The default DNS servers are `208.91.112.53` and `208.91.112.52`.

---

6. To upload the FortiGate VM license from an FTP or TFTP server, use the following CLI command:

```
execute restore vmlicense {ftp | tftp} <VM license file name>
 <Server IP or FQDN>[:server port]
```



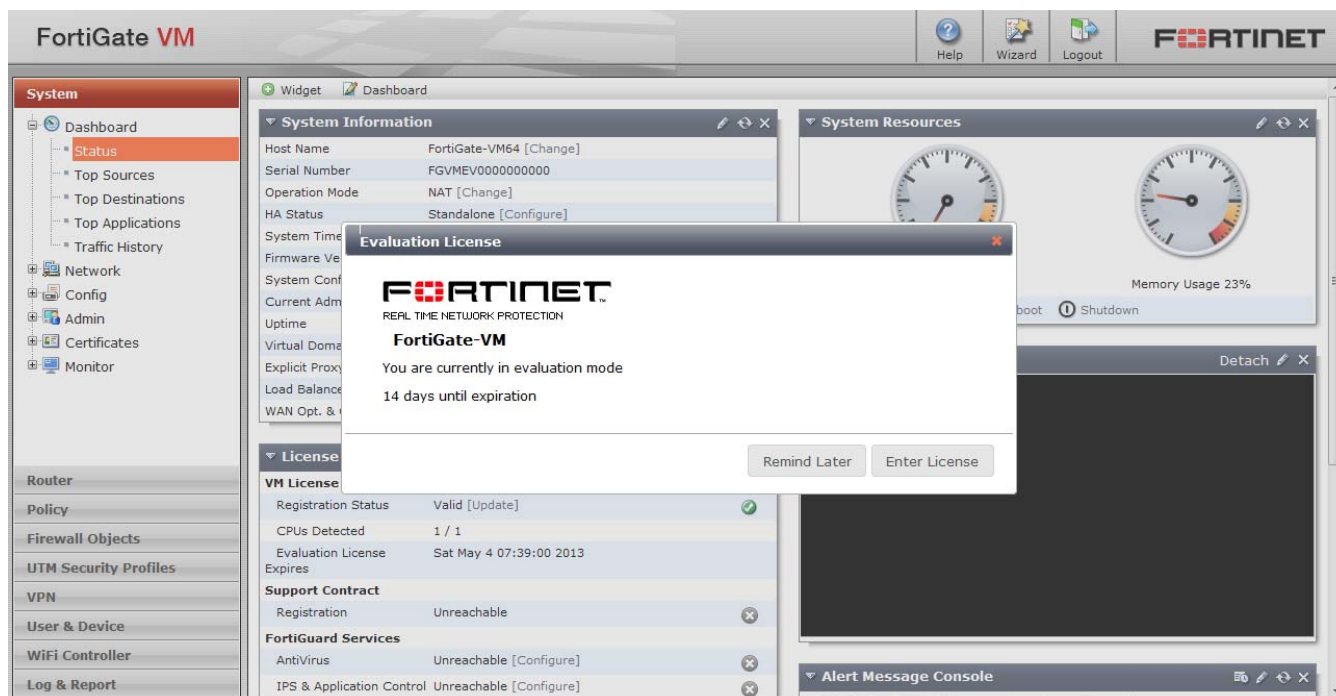
You can also upload the license in the FortiGate VM Web-based Manager. See [“Upload the FortiGate VM license file”](#) on page 2484.

---

## Connect to the FortiGate VM Web-based Manager

When you have configured the port1 IP address and netmask, launch a web browser and enter the IP address that you configured for port1. At the login page, enter the username `admin` and password field and select *Login*. The default password is no password. The Web-based Manager will appear with an *Evaluation License* dialog box, see [Figure 345](#).

**Figure 345:** Web-based Manager and Evaluation License dialog box



## Upload the FortiGate VM license file

Every Fortinet VM includes a 15-day trial license. During this time the FortiGate VM operates in evaluation mode. Before using the FortiGate VM you must enter the license file that you downloaded from the [Customer Service & Support](#) website upon registration.

### To upload the FortiGate VM licence file:

1. In the *Evaluation License* dialog box, select *Enter License*.




You can also upload the license file via the CLI using the following CLI command:

```
execute restore vmlicense [ftp | tftp] <filename string> <ftp server>[:ftp port]
```

The license upload page opens.

**Figure 346:**License upload page



2. Select *Browse* and locate the license file (.lic) on your computer. Select *OK* to upload the license file.
3. Refresh the browser to login.
4. Enter `admin` in the Name field and select *Login*. The VM registration status appears as valid in the License Information widget once the license has been validated by the FortiGuard Distribution Network (FDN) or FortiManager for closed networks.

## Validate the FortiGate VM license with FortiManager

You can validate your FortiGate VM license with some models of FortiManager. To determine whether your FortiManager unit has the VM Activation feature, see Features section of the [FortiManager Product Data sheet](#).

### To validate your FortiGate VM with your FortiManager:

1. To configure your FortiManager as a closed network, enter the following CLI command on your FortiManager:

```
config fmupdate publicnetwork
 set status disable
end
```
2. To configure FortiGate VM to use FortiManager as its override server, enter the following CLI commands on your FortiGate VM:

```
config system central-management
 set mode normal
 set type fortimanager
 set fmg <IPv4 address of the FortiManager device>
```

```

set fmg-source-ip <Source IPv4 address when connecting to the
FortiManager device>
set fortimanager-fds-override enable
set vdom <Enter the name of the VDOM to use when communicating
with the FortiManager device>
end

```

3. Load the FortiGate VM license file in the Web-based Manager. Go to *System > Dashboard > Status*. In the *License Information* widget, in the *Registration Status* field, select *Update*. Browse for the *.lic* license file and select *OK*.
4. To activate the FortiGate VM license, enter the following CLI command on your FortiGate VM:

```
execute update-now
```

5. To check the FortiGate VM license status, enter the following CLI commands on your FortiGate VM:

```
get system status
```

The following output is displayed:

```

Version: Fortigate-VM v5.0,build0099,120910 (Interim)
Virus-DB: 15.00361(2011-08-24 17:17)
Extended DB: 15.00000(2011-08-24 17:09)
Extreme DB: 14.00000(2011-08-24 17:10)
IPS-DB: 3.00224(2011-10-28 16:39)
FortiClient application signature package: 1.456(2012-01-17 18:27)
Serial-Number: FGVM02Q105060000
License Status: Valid
BIOS version: 04000002
Log hard disk: Available
Hostname: Fortigate-VM
Operation Mode: NAT
Current virtual domain: root
Max number of virtual domains: 10
Virtual domains status: 1 in NAT mode, 0 in TP mode
Virtual domain configuration: disable
FIPS-CC mode: disable
Current HA mode: standalone
Distribution: International
Branch point: 511
Release Version Information: MR3 Patch 4
System time: Wed Jan 18 11:24:34 2012

```

```
diagnose hardware sysinfo vm full
```

The following output is displayed:

```

UUID: 564db33a29519f6b1025bf8539a41e92
valid: 1
status: 1
code: 200 (If the license is a duplicate, code 401 will be
displayed)
warn: 0
copy: 0
received: 45438

```

```
warning: 0
recv: 201201201918
dup:
```

## Configure your FortiGate VM

Once the FortiGate VM license has been validated you can begin to configure your device. You can use the *Wizard* located in the top toolbar for basic configuration including enabling central management, setting the admin password, setting the time zone, and port configuration.

For more information on configuring your FortiGate VM see the FortiOS Handbook at <http://docs.fortinet.com>.

# Chapter 21 VoIP Solutions: SIP for FortiOS 5.0

This FortiOS Handbook chapter contains the following sections:

[FortiGate VoIP solutions: SIP](#) describes FortiGate SIP support.



# FortiGate VoIP solutions: SIP

This chapter includes the following sections:

- [SIP overview](#)
- [Common SIP VoIP configurations](#)
- [SIP messages and media protocols](#)
- [The SIP session helper](#)
- [The SIP ALG](#)
- [How the SIP ALG performs NAT](#)
- [Enhancing SIP pinhole security](#)
- [Hosted NAT traversal](#)
- [SIP over IPv6](#)
- [Deep SIP message inspection](#)
- [Blocking SIP request messages](#)
- [SIP rate limiting](#)
- [SIP logging and DLP archiving](#)
- [Inspecting SIP over SSL/TLS \(secure SIP\)](#)
- [SIP and HA: session failover and geographic redundancy](#)
- [SIP and IPS](#)
- [SIP debugging](#)

## SIP overview

The Session Initiation Protocol (SIP) is an IETF application layer signaling protocol used for establishing, conducting, and terminating multiuser multimedia sessions over TCP/IP networks using any media. SIP is often used for Voice over IP (VoIP) calls but can be used for establishing streaming communication between end points.

SIP employs a request and response transaction model similar to HTTP for communicating between endpoints. SIP sessions begin with a SIP client sending a SIP request message to another client to initiate a multimedia session. The other client responds with a SIP response message. Using these request and response messages, the clients engage in a SIP dialog to negotiate how to communicate and then start, maintain, and end the communication session.

SIP commonly uses TCP or UDP port 5060 and/or 5061. Port 5060 is used for non-encrypted SIP signaling sessions and port 5061 is typically used for SIP sessions encrypted with SSL or TLS.

Devices involved in SIP communications are called SIP User Agents (UAs) (also sometimes called a User Element (UE)). UAs include User Agent Clients (UACs) that communicate with each other and User Agent Servers (UASs) that facilitate communication between UACs. For a VoIP application, an example of a UAC would be a SIP phone and an example of a UAS would be a SIP proxy server.

A SIP message contains headers that include client and server names and addresses required for the communication sessions. The body of a SIP message contains Session Description Protocol (SDP) statements that establish the media communication (port numbers, protocols and codecs) that the SIP UAs use. SIP VoIP most commonly uses the Real Time Protocol (RTP) and the Real Time Control Protocol (RTCP) for voice communication. Once the SIP dialog establishes the SIP call the VoIP stream can run independently, although SIP messages can affect the VoIP stream by changing port numbers or addresses and by ending it.

Once SIP communication and media settings are established, the UAs communicate with each other using the established media settings. When the communication session is completed, one of the UAs ends the session by sending a final SIP request message and the other UA sends a SIP response message and both UAs end the SIP call and stop the media stream.

FortiGate units provide security for SIP communications using the SIP session helper and the SIP ALG:

- The SIP session-helper provides basic high-performance support for SIP calls passing through the FortiGate unit by opening SIP and RTP pinholes and performing source and destination IP address and port translation for SIP and RTP packets and for the IP addresses and port numbers in the SIP headers and the SDP body of the SIP messages. For more about the SIP session helper, see [“The SIP session helper” on page 2505](#).
- The SIP Application Layer Gateway (ALG) provides the same features as the session helper plus additional advanced features such as deep SIP message inspection, SIP logging, SIP IPv6 support, SIP message checking, HA failover of SIP sessions, and SIP rate limiting. For more about the SIP ALG, see [“The SIP ALG” on page 2512](#).

There are a large number of SIP-related Internet Engineering Task Force (IETF) documents (Request for Comments) that define behavior of SIP and related applications. FortiGate units provide complete support of [RFC 3261](#) for SIP, [RFC 4566](#) for SDP and [RFC 3262](#) for Provisional Response Acknowledgement (PRACK). FortiGate units also provide support for other SIP and SIP-related RFCs and performs [“Deep SIP message inspection” on page 2558](#) for SIP statements defined in other SIP RFCs.


## Common SIP VoIP configurations

This section describes some common SIP VoIP configurations and simplified SIP dialogs for these configurations. This section also shows some examples of how adding a FortiGate unit affects SIP processing.

### Peer to peer configuration

In the peer to peer configuration shown in [Figure 347](#), two SIP phones (in the example, FortiPhones) communicate directly with each other. The phones send SIP request and response messages back and forth between each other to establish the SIP session.

**Figure 347:** SIP peer to peer configuration

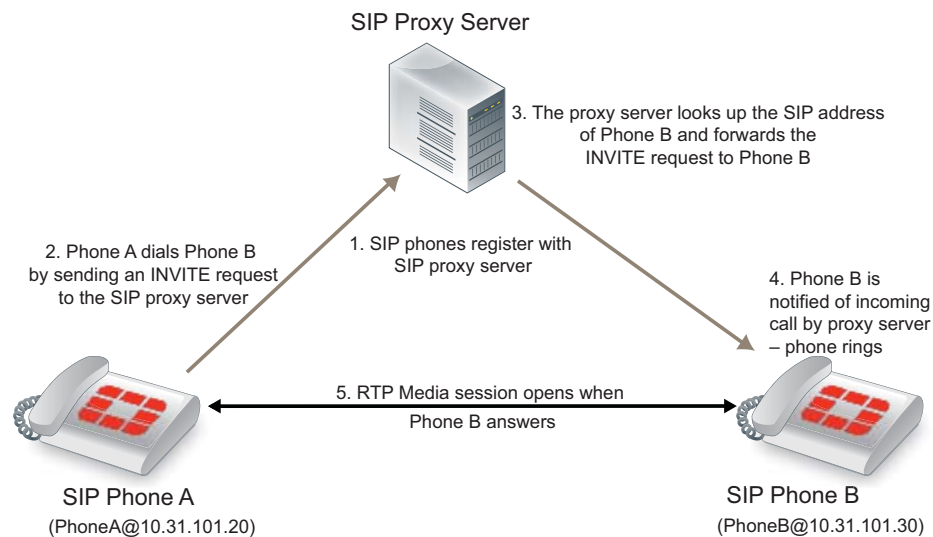


Peer to peer configurations are not very common because they require the SIP phones to keep track of the names and addresses of all of the other SIP phones that they can communicate with. In most cases a SIP proxy or re-direct server maintains addresses of a large number of SIP phones and a SIP phone starts a call by contacting the SIP proxy server.

## SIP proxy server configuration

A SIP proxy server act as intermediary between SIP phones and between SIP phones (for example, two FortiPhones) and other SIP servers. As shown in [Figure 348](#), SIP phones send request and response messages the SIP proxy server. The proxy server forwards the messages to other clients or to other SIP proxy servers. Proxy servers can hide SIP phones by proxying the signaling messages. To the other users on the VoIP network, the signaling invitations look as if they come from the SIP proxy server.

**Figure 348:**SIP in proxy mode

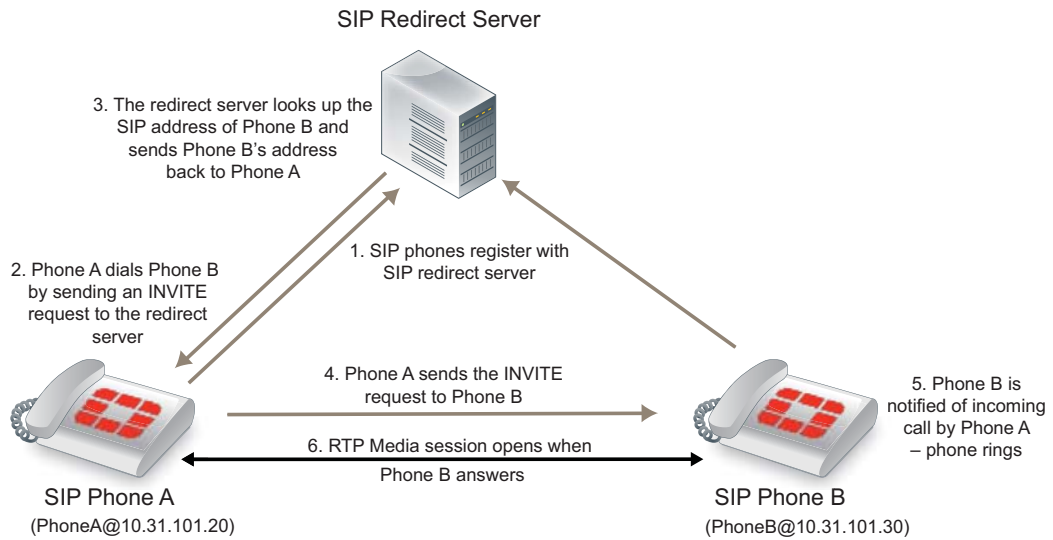


A common SIP configuration would include multiple networks of SIP phones. Each of the networks would have its own SIP server. Each SIP server would proxy the communication between phones on its own network and between phones in different networks.

## SIP redirect server configuration

A SIP redirect server accepts SIP requests, maps the addresses in the request into zero or more new addresses and returns those addresses to the client. The redirect server does not initiate SIP requests or accept calls. As shown in [Figure 349](#), SIP clients send INVITE requests to the redirect server, which then looks up the destination address. The redirect server returns the destination address to the client. The client uses this address to send the INVITE request directly to the destination SIP client.


**Figure 349:**SIP in redirect model



### SIP registrar configuration

A SIP registrar accepts SIP REGISTER requests from SIP phones for the purpose of updating a location database with this contact information. This database can then become a SIP location service that can be used by SIP proxy servers and redirect servers to locate SIP clients. As shown in [Figure 350](#), SIP clients send REGISTER requests to the SIP registrar.

**Figure 350:**SIP registrar and proxy servers

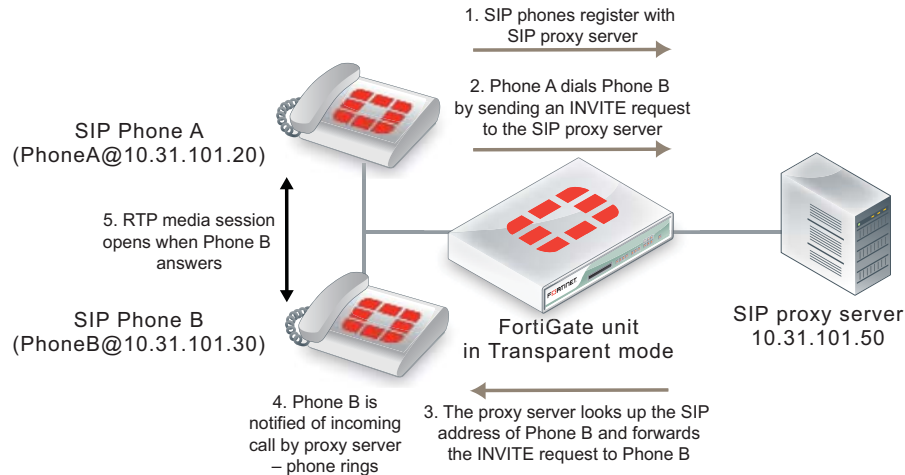


## SIP with a FortiGate unit

Depending on your security requirements and network configuration FortiGate units may be in many different places in a SIP configuration. This section shows a few examples.

Figure 351 shows a FortiGate unit installed between a SIP proxy server and SIP phones on the same network. The FortiGate unit is operating in Transparent mode so both the proxy server and the phones are on the same subnet. In this configuration, called SIP inspection without address translation, the FortiGate unit could be protecting the SIP proxy server on the private network by implementing SIP security features for SIP sessions between the SIP phones and the SIP proxy server.

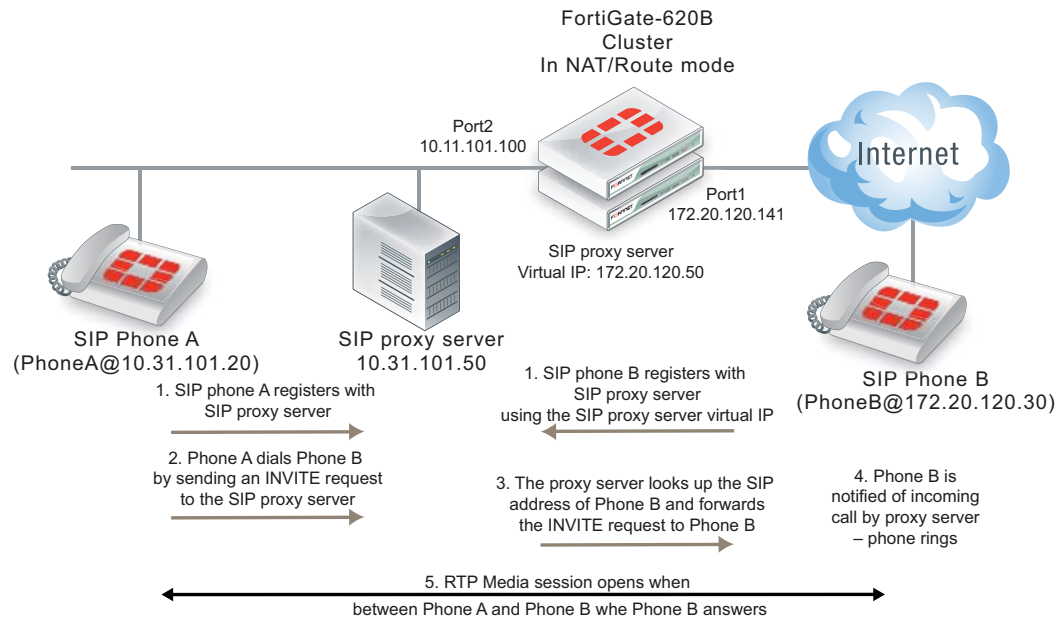
**Figure 351:** SIP network with FortiGate unit in Transparent mode



The phones and server use the same SIP dialogs as they would if the FortiGate unit was not present. However, the FortiGate unit can be configured to control which devices on the network can connect to the SIP proxy server and can also protect the SIP proxy server from SIP vulnerabilities.

Figure 352 shows a FortiGate unit operating in NAT/Route mode and installed between a private network and the Internet. Some SIP phones and the SIP proxy server are connected to the private network and some SIP phones are connected to the Internet. The SIP phones on the Internet can connect to the SIP proxy server through the FortiGate unit and communication between SIP phones on the private network and SIP phones on the Internet must pass through the FortiGate unit.

**Figure 352:**SIP network with FortiGate unit in NAT/Route mode



The phones and server use the same SIP dialog as they would if the FortiGate unit was not present. However, the FortiGate unit can be configured to control which devices on the network can connect to the SIP proxy server and can also protect the SIP proxy server from SIP vulnerabilities. In addition, the FortiGate unit has a firewall virtual IP that forwards packets sent to the SIP proxy server Internet IP address (172.20.120.50) to the SIP proxy server internal network IP address (10.31.101.30).

Since the FortiGate unit is operating in NAT/Route mode it must translate packet source and destination IP addresses (and optionally ports) as the sessions pass through the FortiGate unit. Also, the FortiGate unit must translate the addresses contained in the SIP headers and SDP body of the SIP messages. As well the FortiGate unit must open SIP and RTP pinholes through the FortiGate unit. SIP pinholes allow SIP signalling sessions to pass through the FortiGate between phones and between phones and SIP servers. RTP pinholes allow direct RTP communication between the SIP phones once the SIP dialog has established the SIP call. Pinholes are opened automatically by the FortiGate unit. Administrators do not add security policies for pinholes or for RTP sessions. All that is required is a security policy that accepts SIP traffic.

Opening an RTP pinhole means opening a port on a FortiGate interface to allow RTP traffic to use that port to pass through the FortiGate unit between the SIP phones on the Internet and SIP phones on the internal network. A pinhole only accepts packets from one RTP session. Since a SIP call involves at least two media streams (one from Phone A to Phone B and one from Phone B to Phone A) the FortiGate unit opens two RTP pinholes. Phone A sends RTP packets through a pinhole in port2 and Phone B sends RTP packets through a pinhole in port1. The FortiGate unit opens the pinholes when required by the SIP dialog and closes the pinholes when the SIP call is completed. The FortiGate unit opens new pinholes for each SIP call.

Each RTP pinhole actually includes two port numbers. The RTP port number as defined in the SIP message and an RTCP port number, which is the RTP port number plus 1. For example, if the SIP call used RTP port 3346 the FortiGate unit would create a pinhole for ports 3346 and 3347.

## SIP messages and media protocols

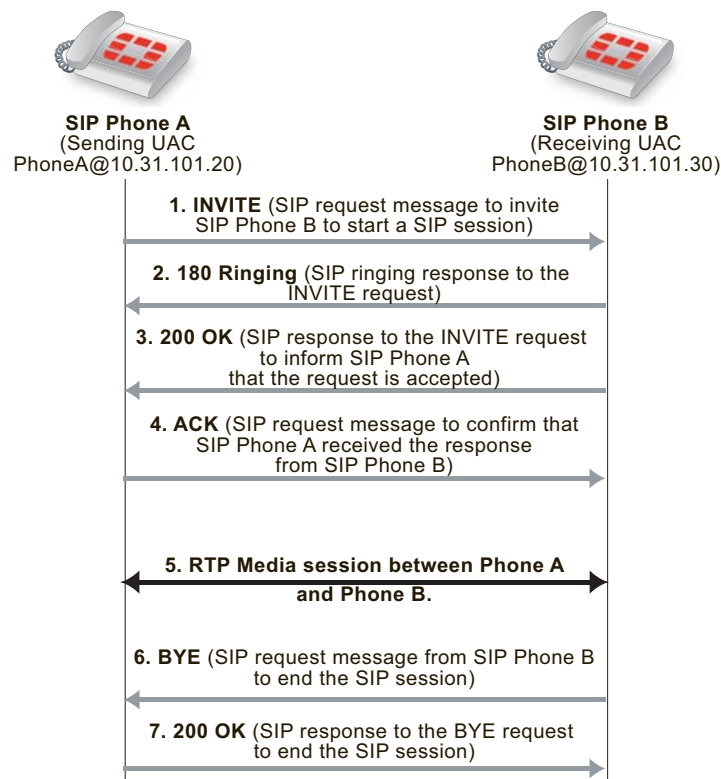
This section provides an overview of SIP messages and how they communicate information about SIP sessions and how SDP, RTP, and RTCP fits in with SIP communications.

SIP uses clear text messages to start, maintain, and end media sessions between SIP user agent clients (UACs) and user agent servers (UASs). These messages form a SIP dialog. A typical SIP dialog begins with an INVITE request message sent from a UAC to another UAC or to a UAS. The first INVITE request message attempts to start a SIP call and includes information about the sending UAC and the receiving UAC as well as information about the communication session.

If only two UACs are involved as shown in [Figure 353](#), the receiving UAC (Phone B) responds with a 180 Ringing and then a 200 OK SIP response message that informs Phone A that Phone B received and accepted the request. Phone A then sends an ACK message to notify Phone B that the SIP response was received. Phone A and Phone B can then participate in the RTP media session set up by the SIP messages.

When the phone call is complete, one of the UACs (in the example Phone B) hangs up sending a BYE request message to Phone A. Phone A then sends a 200 OK response to Phone B acknowledging that the session has ended.

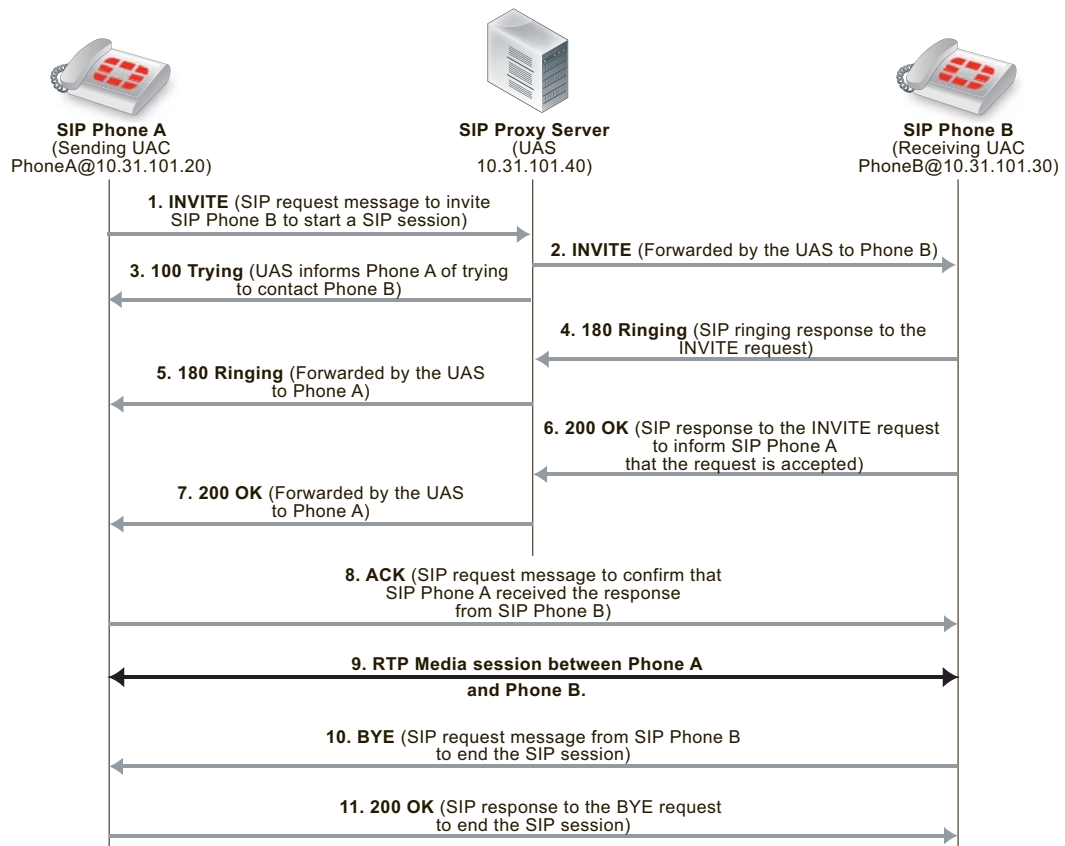
**Figure 353:**Basic SIP dialog between two UACs



If a UAS in the form of a SIP proxy server is involved, similar messages are sent and received, but the proxy server participates as an intermediary in the initial call setup. In the example in [Figure 354](#) the SIP proxy server receives the INVITE request from Phone A and forwards it to Phone B. The proxy server then sends a 100 Trying response to Phone A. Phone B receives the INVITE request and responds with a 180 Ringing and then a 200 OK SIP response message. These messages are received by the proxy server and forwarded to Phone A to notify Phone A that Phone B received and accepted the request. Phone A then sends an ACK message to notify Phone B that the SIP response was received. This response is received by the proxy server and forwarded to Phone B. Phone A and Phone B can then participate in the media session independently of the proxy server.

When the phone call is complete Phone B hangs up sending a BYE request message to Phone A. Phone A then sends a 200 OK response to Phone B acknowledging that the session has ended.

**Figure 354:**Basic SIP dialog between UACs with a SIP proxy server UAS



The SIP messages include SIP headers that contain names and addresses of Phone A, Phone B and the proxy server. This addressing information is used by the UACs and the proxy server during the call set up.

The SIP message body includes Session Description Protocol (SDP) statements that Phone A and Phone B use to establish the media session. The SDP statements specify the type of media stream to use for the session (for example, audio for SIP phone calls) and the protocol to use for the media stream (usually the Real Time Protocol (RTP) media streaming protocol).

Phone A includes the media session settings that it would like to use for the session in the INVITE message. Phone B includes its response to these media settings in the 200 OK response. Phone A's ACK response confirms the settings that Phone A and Phone B then use for the media session.



## Hardware accelerated RTP processing

FortiGate units can offload RTP packet processing to network processor (NP) interfaces. This acceleration greatly enhance the overall throughput and resulting in near speed RTP performance.

## SIP request messages

SIP sessions always start with a SIP request message (also just called a SIP request). SIP request messages also establish, maintain, and terminate SIP communication sessions. [Table 113](#) lists some common SIP request message types.

**Table 113:**Common SIP request message types

Message Type	Description
<b>INVITE</b>	A client sends an INVITE request to invite another client to participate in a multimedia session. The INVITE request body usually contains the description of the session.
<b>ACK</b>	The originator of an INVITE message sends an ACK request to confirm that the final response to an INVITE request was received. If the INVITE request did not contain the session description, it must be included in the ACK request.
<b>PRACK</b>	In some cases, SIP uses provisional response messages to report on the progress of the response to a SIP request message. The provisional response messages are sent before the final SIP response message. Similar to an ACK request message, a PRACK request message is sent to acknowledge that a provisional response message has been received.
<b>OPTIONS</b>	The UA uses OPTIONS messages to get information about the capabilities of a SIP proxy. The SIP proxy server replies with a description of the SIP methods, session description protocols, and message encoding that are supported.
<b>BYE</b>	A client sends a BYE request to end a session. A BYE request from either end of the SIP session terminates the session.
<b>CANCEL</b>	A client sends a CANCEL request to cancel a previous INVITE request. A CANCEL request has no effect if the SIP server processing the INVITE sends a final response to the INVITE before receiving the CANCEL.
<b>REGISTER</b>	A client sends a REGISTER request to a SIP registrar server with information about the current location (IP address and so on) of the client. A SIP registrar server saves the information it receives in REGISTER requests and makes this information available to any SIP client or server attempting to locate the client.
<b>Info</b>	For distributing mid-session signaling information along the signaling path for a SIP call. I
<b>Subscribe</b>	For requesting the current state and state updates of a remote node.
<b>Notify</b>	Informs clients and servers of changes in state in the SIP network.
<b>Refer</b>	Refers the recipient (identified by the Request-URI) to a third party according to the contact information in the request.

**Table 113:**Common SIP request message types (continued)

Message Type	Description
<b>Update</b>	Opens a pinhole for new or updated SDP information.
<b>Response codes (1xx, 202, 2xx, 3xx, 4xx, 5xx, 6xx)</b>	Indicates the status of a transaction. For example: 200 OK, 202 Accepted, or 400 Bad Request.

## SIP response messages

SIP response messages (often just called SIP responses) provide status information in response to SIP request messages. All SIP response messages include a response code and a reason phrase. There are five SIP response message classes. They are described below.

There are also two types of SIP response messages, provisional and final. Final response messages convey the result of the request processing, and are sent reliably. Provisional responses provide information on the progress of the request processing, but may not be sent reliably. Provisional response messages start with 1xx and are also called informational response messages.

### Informational (or provisional)

Informational or provisional responses indicate that a request message was received and imply that the endpoint is going to process the request. Information messages may not be sent reliably and may not require an acknowledgement.

If the SIP implementation uses Provisional Response Acknowledgement (PRACK) ([RFC 3262](#)) then informational or provisional messages are sent reliably and require a PRACK message to acknowledge that they have been received.

Informational responses can contain the following reason codes and reason phrases:

```
100 Trying
180 Ringing
181 Call is being forwarded
182 Queued
183 Session progress
```

### Success

Success responses indicate that a request message was received, understood, and accepted. Success responses can contain the following reason codes and reason phrases:

```
200 OK
202 Accepted
```

## Redirection

Redirection responses indicate that more information is required for the endpoint to respond to a request message. Redirection responses can contain the following reason codes and reason phrases:

- 300 Multiple choices
- 301 Moved permanently
- 302 Moved temporarily
- 305 Use proxy
- 380 Alternative service

## Client error

Client error responses indicate that a request message was received by a server that contains syntax that the server cannot understand (i.e. contains a syntax error) or cannot comply with. Client error responses include the following reason codes and reason phrases:

- 400 Bad request
- 401 Unauthorized
- 402 Payment required
- 403 Forbidden
- 404 Not found
- 405 Method not allowed
- 406 Not acceptable
- 407 Proxy authentication required
- 408 Request time-out
- 409 Conflict
- 410 Gone
- 411 Length required
- 413 Request entity too large
- 414 Request-URL too large
- 415 Unsupported media type
- 420 Bad extension
- 480 Temporarily not available
- 481 Call leg/transaction does not exist
- 482 Loop detected
- 483 Too many hops
- 484 Address incomplete
- 485 Ambiguous
- 486 Busy here
- 487 Request canceled
- 488 Not acceptable here

## Server error

Server error responses indicate that a server was unable to respond to a valid request message. Server error responses include the following reason codes and reason phrases:

- 500 Server internal error
- 501 Not implemented
- 502 Bad gateway
- 502 Service unavailable
- 504 Gateway time-out
- 505 SIP version not supported

## Global failure

Global failure responses indicate that there are no servers available that can respond to a request message. Global failure responses include the following reason codes and reason phrases:

- 600 Busy everywhere
- 603 Decline
- 604 Does not exist anywhere
- 606 Not acceptable

## SIP message start line

The first line in a SIP message is called the start line. The start line in a request message is called the request-line and the start line in a response message is called the status-line.

---

**Request-line** The first line of a SIP request message. The request-line includes the SIP message type, the SIP protocol version, and a Request URI that indicates the user or service to which this request is being addressed. The following example request-line specifies the INVITE message type, the address of the sender of the message (inviter@example.com), and the SIP version:

```
INVITE sip:inviter@example.com SIP/2.0
```

---

**Status-line** The first line of a SIP response message. The status-line includes the SIP protocol version, the response code, and the reason phrase. The example status-line includes the SIP version, the response code (200) and the reason phrase (OK).

```
SIP/2.0 200 OK
```

---

## SIP headers

Following the start line, SIP messages contain SIP headers (also called SIP fields) that convey message attributes and to modify message meaning. SIP headers are similar to HTTP header fields and always have the following format:

```
<header_name>:<value>
```

SIP messages can include the SIP headers listed in [Table 114](#):

**Table 114:**SIP headers

SIP Header	Description
<b>Allow</b>	Lists the set of SIP methods supported by the UA generating the message. All methods, including ACK and CANCEL, understood by the UA MUST be included in the list of methods in the Allow header field, when present. For example: <pre>Allow: INVITE, ACK, OPTIONS, CANCEL, BYE</pre>
<b>Call-ID</b>	A globally unique identifier for the call, generated by the combination of a random string and the sender's host name or IP address. The combination of the To, From, and Call-ID headers completely defines a peer-to-peer SIP relationship between the sender and the receiver. This relationship is called a SIP dialog. <pre>Call-ID: ddeg45e793@10.31.101.30</pre>
<b>Contact</b>	Included in SIP request messages, the Contact header contains the SIP URI of the sender of the SIP request message. The receiver uses this URI to contact the sender. For example: <pre>Contact: Sender &lt;sip:sender@10.31.100.20&gt;t</pre>
<b>Content-Length</b>	The number of bytes in the message body (in bytes). <pre>Content-Length: 126</pre>

**Table 114:**SIP headers (continued)

SIP Header	Description
<b>Content-Type</b>	<p>In addition to SIP headers, SIP messages include a message body that contains information about the content or communication being managed by the SIP session. The Content-Type header specifies what the content of the SIP message is. For example, if you are using SIP with SDP, the content of the SIP message is SDP code.</p> <p style="text-align: center;">Content-Type: application/sdp</p>
<b>CSeq</b>	<p>The command sequence header contains a sequence integer that is increased for each new SIP request message (but is not incremented in the response message). This header also includes the request name found in the request message request-line. For example:</p> <p style="text-align: center;">CSeq: 1 INVITE</p>
<b>Expires</b>	<p>Gives the relative time after which the message (or content) expires. The actual time and how the header is used depends on the SIP method. For example:</p> <p style="text-align: center;">Expires: 5</p>
<b>From</b>	<p>Identifies the sender of the message. Responses to a message are sent to the address of the sender. The following example includes the sender's name (Sender) and the sender's SIP address (sender@10.31.101.20.):</p> <p style="text-align: center;">From: Sender &lt;sip:sender@10.31.101.20&gt;</p>
<b>Max-forwards</b>	<p>An integer in the range 0-255 that limits the number of proxies or gateways that can forward the request message to the next downstream server. Also called the number of hops, this value is decreased every time the message is forwarded. This can also be useful when the client is attempting to trace a request chain that appears to be failing or looping in mid-chain. For example:</p> <p style="text-align: center;">Max-Forwards: 30</p>
<b>P-Asserted-Identity</b>	<p>The P-Asserted-Identity header is used among trusted SIP entities to carry the identity of the user sending a SIP message as it was verified by authentication. See <a href="#">RFC 3325</a>. The header contains a SIP URI and an optional display-name, for example:</p> <p style="text-align: center;">P-Asserted-Identity: "Example Person" &lt;sip:10.31.101.50&gt;</p>
<b>RAck</b>	<p>Sent in a PRACK request to support reliability of information or provisional response messages. It contains two numbers and a method tag. For example:</p> <p style="text-align: center;">RAck: 776656 1 INVITE</p>
<b>Record-Route</b>	<p>Inserted into request messages by a SIP proxy to force future requests to be routed through the proxy. In the following example, the host at IP address 10.31.101.50 is a SIP proxy. The <code>lr</code> parameter indicates the URI of a SIP proxy in Record-Route headers.</p> <p style="text-align: center;">Record-Route: &lt;sip:10.31.101.50;lr&gt;</p>

**Table 114:**SIP headers (continued)

SIP Header	Description
<b>Route</b>	Forces routing for a request message through one or more SIP proxies. The following example includes two SIP proxies:  Route: <sip:172.20.120.10;lr>, <sip:10.31.101.50;lr>
<b>RSeq</b>	The RSeq header is used in information or provisional response messages to support reliability of informational response messages. The header contains a single numeric value. For example:  RSeq: 33456
<b>To</b>	Identifies the receiver of the message. The address in this field is used to send the message to the receiver. The following example includes the receiver's name (Receiver) and the receiver's SIP address (receiver@10.31.101.30.):  To: Receiver <sip:receiver@10.31.101.30>
<b>Via</b>	Indicates the SIP version and protocol to be used for the SIP session and the address to which to send the response to the message that contains the Via field. The following example Via field indicates to use SIP version 2, UDP for media communications, and to send the response to 10.31.101.20 using port 5060.  Via: SIP/2.0/UDP 10.31.101.20:5060

## The SIP message body and SDP session profiles

The SIP message body describes the session to be initiated. For example, in a SIP phone call the body usually includes audio codec types, sampling rates, server IP addresses and so on. For other types of SIP session the body could contain text or binary data of any type which relates in some way to the session. The message body is included in request and response messages.

Two possible SIP message body types:

- Session Description Protocol (SDP), most commonly used for SIP VoIP.
- Multipurpose Internet Mail Extensions (MIME)

SDP is most often used for VoIP and FortiGate units support SDP content in SIP message bodies. SDP is a text-based protocol used by SIP to control media sessions. SDP does not deliver media but provides a session profile that contains media details, transport addresses, parameter negotiation, and other session description metadata for the participants in a media session. The participants use the information in the session profile to negotiate how to communicate and to manage the media session. SDP is described by [RFC 4566](#).

An SDP session profile always contains session information and may contain media information. Session information appears at the start of the session profile and media information (using the `m=` attribute) follows.

SDP session profiles can include the attributes listed in [Table 115](#).

**Table 115:**SDP session profile attributes

Attribute	Description
<b>a=</b>	Attributes to extend SDP in the form <code>a=&lt;attribute&gt;</code> or <code>a=&lt;attribute&gt;:&lt;value&gt;</code> .
<b>b=</b>	Contains information about the bandwidth required for the session or media in the form <code>b=&lt;bandwidth_type&gt;:&lt;bandwidth&gt;</code> .
<b>c=</b>	Connection data about the session including the network type (usually IN for Internet), address type (IPv4 or IPv6), the connection source address, and other optional information. For example:  <code>c=IN IPv4 10.31.101.20</code>
<b>i=</b>	A text string that contains information about the session. For example:  <code>i=A audio presentation about SIP</code>
<b>k=</b>	Can be used to convey encryption keys over a secure and trusted channel. For example:  <code>k=clear:444gdduudjffdee</code>
<b>m=</b>	Media information, consisting of one or more lines all starting with <code>m=</code> and containing details about the media including the media type, the destination port or ports used by the media, the protocol used by the media, and a media format description.  <code>m=audio 49170 RTP 0 3</code> <code>m-video 3345/2 udp 34</code> <code>m-video 2910/2 RTP/AVP 3 56</code>  Multiple media lines are needed if SIP is managing multiple types of media in one session (for example, separate audio and video streams).  Multiple ports for a media stream are indicated using a slash. <code>3345/2 udp</code> means UDP ports 3345 and 3346. Usually RTP uses even-numbered ports for data with the corresponding one-higher odd ports used for the RTCP session belonging to the RTP session. So <code>2910/2 RTP/AVP</code> means ports 2910 and 2912 are used for RTP and 2911 and 2913 are used for RTCP.  Media types include <code>udp</code> for an unspecified protocol that uses UDP, <code>RTP</code> or <code>RTP/AVP</code> for standard RTP and <code>RTP/SAVP</code> for secure RTP.
<b>o=</b>	The sender's username, a session identifier, a session version number, the network type (usually IN for Internet), the address type (for example, IPv4 or IPv6), and the sending device's IP address. The <code>o=</code> field becomes a universal identifier for this version of this session description. For example:  <code>o=PhoneA 5462346 332134 IN IP4 10.31.101.20</code>
<b>r=</b>	Repeat times for a session. Used if a session will be repeated at one or more timed intervals. Not normally used for VoIP calls. The times can be in different formats. For example.  <code>r=7d 1h 0 25h</code> <code>r=604800 3600 0 90000</code>

**Table 115:**SDP session profile attributes (continued)

Attribute	Description
<b>s=</b>	Any text that describes the session or s= followed by a space. For example: s=Call from inviter
<b>t=</b>	The start and stop time of the session. Sessions with no time restrictions (most VoIP calls) have a start and stop time of 0. t=0 0
<b>v=</b>	SDP protocol version. The current SDP version is 0 so the v= field is always: v=0
<b>z=</b>	Time zone adjustments. Used for scheduling repeated sessions that span the time between changing from standard to daylight savings time. z=2882844526 -1h 2898848070 0

### Example SIP messages

The following example SIP INVITE request message was sent by PhoneA to PhoneB. The first nine lines are the SIP headers. The SDP profile starts with v=0 and the media part of the session profile is the last line, starting with m=.

```

INVITE sip:PhoneB@172.20.120.30 SIP/2.0
Via: SIP/2.0/UDP 10.31.101.50:5060
From: PhoneA <sip:PhoneA@10.31.101.20>
To: PhoneB <sip:PhoneB@172.20.120.30>
Call-ID: 314159@10.31.101.20
CSeq: 1 INVITE
Contact: sip:PhoneA@10.31.101.20
Content-Type: application/sdp
Content-Length: 124
v=0
o=PhoneA 5462346 332134 IN IP4 10.31.101.20
s=Let's Talk
t=0 0
c=IN IP4 10.31.101.20
m=audio 49170 RTP 0 3

```



The following example shows a possible 200 OK SIP response message in response to the previous INVITE request message. The response includes 200 OK which indicates success, followed by an echo of the original SIP INVITE request followed by PhoneB's SDP profile.

```
SIP/2.0 200 OK
Via: SIP/2.0/UDP 10.31.101.50:5060
From: PhoneA <sip:PhoneA@10.31.101.20>
To: PhoneB <sip:PhoneB@172.20.120.30>
Call-ID: 314159@10.31.101.20
CSeq: 1 INVITE
Contact: sip:PhoneB@10.31.101.30
Content-Type: application/sdp
Content-Length: 107
v=0
o=PhoneB 124333 67895 IN IP4 172.20.120.30
s=Hello!
t=0 0
c=IN IP4 172.20.120.30
m=audio 3456 RTP 0
```

SIP can support multiple media streams for a single SIP session. Each media stream will have its own c= and m= lines in the body of the message. For example, the following message includes three media streams:

```
INVITE sip:PhoneB@172.20.120.30 SIP/2.0
Via: SIP/2.0/UDP 10.31.101.20:5060
From: PhoneA <sip:PhoneA@10.31.101.20>
To: PhoneB <sip:PhoneB@172.20.120.30>
Call-ID: 314159@10.31.101.20
CSeq: 1 INVITE
Contact: sip:PhoneA@10.31.101.20
Content-Type: application/sdp
Content-Length: 124
v=0
o=PhoneA 5462346 332134 IN IP4 10.31.101.20
s=Let's Talk
t=0 0
c=IN IP4 10.31.101.20
m=audio 49170 RTP 0 3
c=IN IP4 10.31.101.20
m=audio 49172 RTP 0 3
c=IN IP4 10.31.101.20
m=audio 49174 RTP 0 3
```

## The SIP session helper

The SIP session-helper is a high-performance solution that provides basic support for SIP calls passing through the FortiGate unit by opening SIP and RTP pinholes and by performing NAT of the addresses in SIP messages.

The SIP session helper:

- Understands SIP dialog messages.
- Keeps the states of the SIP transactions between SIP UAs and SIP servers.
- Translates SIP header and SDP information to account for NAT operations performed by the FortiGate unit.
- Opens up and closes dynamic SIP pinholes for SIP signalling traffic.
- Opens up and closes dynamic RTP and RTSP pinholes for RTP and RTSP media traffic.
- Provides basic SIP security as an access control device.
- Uses the intrusion protection (IPS) engine to perform basic SIP protocol checks.

## SIP session helper configuration overview

The SIP session helper is enabled by default and set to listen for SIP traffic on TCP or UDP port 5060. SIP sessions using port 5060 accepted by a security policy that does not include a VoIP profile are processed by the SIP session helper.

You can enable and disable the SIP session helper, change the TCP or UDP port that the session helper listens on for SIP traffic, and enable or disable SIP NAT tracing. If the FortiGate unit is operating with multiple VDOMs, each VDOM can have a different SIP session helper configuration.

To have the SIP session helper process SIP sessions you need to add a security policy that accepts SIP sessions on the configured SIP UDP or TCP ports. The security policies can have service set to ANY, or to the SIP pre-defined firewall service, or a custom firewall service. The SIP pre-defined firewall service restricts the security policy to only accepting sessions on UDP port 5060.

If NAT is enabled for security policies that accept SIP traffic, the SIP session helper translates addresses in SIP headers and in the RDP profile and opens up pinholes as required for the SIP traffic. This includes security policies that perform source NAT and security policies that contain virtual IPs that perform destination NAT and port forwarding. No special SIP configuration is required for this address translation to occur, it is all handled automatically by the SIP session helper according to the NAT configuration of the security policy that accepts the SIP session.

To use the SIP session helper you must not add a VoIP profile to the security policy. If you add a VoIP profile, SIP traffic bypasses the SIP session helper and is processed by the SIP ALG.



In most cases you would want to use the SIP ALG since the SIP session helper provides limited functionality. However, the SIP session helper is available and can be useful for high-performance solutions where a high level of SIP security is not a requirement.

---

## Disabling and enable the SIP session helper

You can use the following steps to disable the SIP session helper. You might want to disable the SIP session helper if you don't want the FortiGate unit to apply NAT or other SIP session help features to SIP traffic. With the SIP session helper disabled, the FortiGate unit can still accept SIP sessions if they are allowed by a security policy, but the FortiGate unit will not be able to open pinholes or NAT the addresses in the SIP messages.

### To disable the sip session helper

1. Enter the following command to find the sip session helper entry in the session-helper list:

```
show system session-helper
.
.
.
edit 13
 set name sip
 set port 5060
 set protocol 17
next
.
.
.
```

This command output shows that the sip session helper listens in UDP port 5060 for SIP sessions.

2. Enter the following command to delete session-helper list entry number 13 to disable the sip session helper:

```
config system session-helper
 delete 13
end
```

If you want to use the SIP session helper you can verify whether it is enabled or disabled using the `show system session-helper` command.



You do not have to disable the SIP session helper to use the SIP ALG.

---

If the SIP session helper has been disabled by being removed from the session-helper list you can use the following command to enable the SIP session helper by adding it back to the session helper list:

```
config system session-helper
 edit 0
 set name sip
 set port 5060
 set protocol 17
 end
```

## Changing the port numbers that the SIP session helper listens on

You can use the following command to change the port number that the SIP session helper listens on for SIP traffic to 5064. The SIP session helper listens on the same port number for UDP and TCP SIP sessions. In this example, the SIP session helper is session helper 13:

```
config system session-helper
 edit 13
 set port 5064
 end
```



The `config system settings options sip-tcp-port`, `sip-udp-port`, and `sip-ssl-port` control the ports that the SIP ALG listens on for SIP sessions. See [“Changing the port numbers that the SIP ALG listens on” on page 2516](#).

Your FortiGate unit may use a different session helper number for SIP. Enter the following command to view the session helpers:


```
show system session-helper
.
.
.
edit 13
 set name sip
 set port 5060
 set protocol 17
end
.
.
.
```

## Configuration example: SIP session helper in Transparent Mode

Figure 355 shows an example SIP network consisting of a FortiGate unit operating in Transparent mode between two SIP phones. Since the FortiGate unit is operating in Transparent mode both phones are on the same network and the FortiGate unit and the SIP session helper does not perform NAT. Even though the SIP session helper is not performing NAT you can use this configuration to apply SIP session helper security features to the SIP traffic.

The FortiGate unit requires two security policies that accept SIP packets. One to allow SIP Phone A to start a session with SIP Phone B and one to allow SIP Phone B to start a session with SIP Phone A.

**Figure 355:** SIP network with FortiGate unit in Transparent mode



## General configuration steps

The following general configuration steps are required for this SIP configuration that uses the SIP session helper. This example includes security policies that specifically allow SIP sessions using UDP port 5060 from Phone A to Phone B and from Phone B to Phone A. In most cases you would have more than two phones so would use more general security policies. Also, you can set the firewall service to ANY to allow traffic other than SIP on UDP port 5060.

1. Add firewall addresses for Phone A and Phone B.
2. Add a security policy that accepts SIP sessions initiated by Phone A.
3. Add a security policy that accepts SIP sessions initiated by Phone B.

## Configuration steps - web-based manager

### To add firewall addresses for the SIP phones

1. Go to *Firewall Objects > Address > Addresses*.
2. Select *Create New* to add the following addresses for Phone A and Phone B:

<b>Address Name</b>	Phone_A
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	10.31.101.20/255.255.255.255
<b>Interface</b>	port1
<b>Address Name</b>	Phone_B
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	10.31.101.30/255.255.255.255
<b>Interface</b>	port2

### To add security policies to accept SIP sessions

1. Go to *Policy > Policy > Policy*.
2. Select *Create New* to add a security policy.
3. Add a security policy to allow Phone A to send SIP request messages to Phone B:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	port1
<b>Source Address</b>	Phone_A
<b>Outgoing Interface</b>	port2
<b>Destination Address</b>	Phone_B
<b>Schedule</b>	always
<b>Service</b>	SIP
<b>Action</b>	ACCEPT

4. Select *OK*.
5. Add a security policy to allow Phone B to send SIP request messages to Phone A:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	port2
<b>Source Address</b>	Phone_B
<b>Outgoing Interface</b>	port1
<b>Destination Address</b>	Phone_A
<b>Schedule</b>	always
<b>Service</b>	SIP
<b>Action</b>	ACCEPT

6. Select *OK*.

### Configuration steps - CLI

#### To add firewall addresses for Phone A and Phone B and security policies to accept SIP sessions

1. Enter the following command to add firewall addresses for Phone A and Phone B.

```
config firewall address
 edit Phone_A
 set associated interface port1
 set type ipmask
 set subnet 10.31.101.20 255.255.255.255
 next
 edit Phone_B
 set associated interface port2
 set type ipmask
 set subnet 10.31.101.30 255.255.255.255
 end
```

2. Enter the following command to add security policies to allow Phone A to send SIP request messages to Phone B and Phone B to send SIP request messages to Phone A.

```
config firewall policy
 edit 0
 set srcintf port1
 set dstintf port2
 set srcaddr Phone_A
 set dstaddr Phone_B
 set action accept
 set schedule always
 set service SIP
 next
 edit 0
 set srcintf port2
 set dstintf port1
 set srcaddr Phone_B
 set dstaddr Phone_A
 set action accept
 set schedule always
 set service SIP
 set utm-status enable
 end
```

## SIP session helper diagnose commands

You can use the `diagnose sys sip` commands to display diagnostic information for the SIP session helper.

Use the following command to set the debug level for the SIP session helper. Different debug masks display different levels of detail about SIP session helper activity.

```
diagnose sys sip debug-mask <debug_mask_int>
```

Use the following command to display the current list of SIP dialogs being processed by the SIP session helper. You can also use the `clear` option to delete all active SIP dialogs being processed by the SIP session helper.

```
diagnose sys sip dialog {clear | list}
```

Use the following command to display the current list of SIP NAT address mapping tables being used by the SIP session helper.

```
diagnose sys sip mapping list
```

Use the following command to display the current SIP session helper activity including information about the SIP dialogs, mappings, and other SIP session helper counts. This command can be useful to get an overview of what the SIP session helper is currently doing.

```
diagnose sys sip status
```

## The SIP ALG

In most cases you should use the SIP Application Layer Gateway (ALG) for processing SIP sessions. The SIP ALG provides the same basic SIP support as the SIP session helper. Additionally, the SIP ALG provides a wide range of features that protect your network from SIP attacks, can apply rate limiting to SIP sessions, can check the syntax of SIP and SDP content of SIP messages, and provide detailed logging and reporting of SIP activity.

You apply the SIP ALG to SIP traffic by adding a VoIP profile with SIP enabled to a security policy that accepts SIP traffic. The SIP session helper is automatically bypassed by traffic accepted by a security policy that includes a VoIP profile.

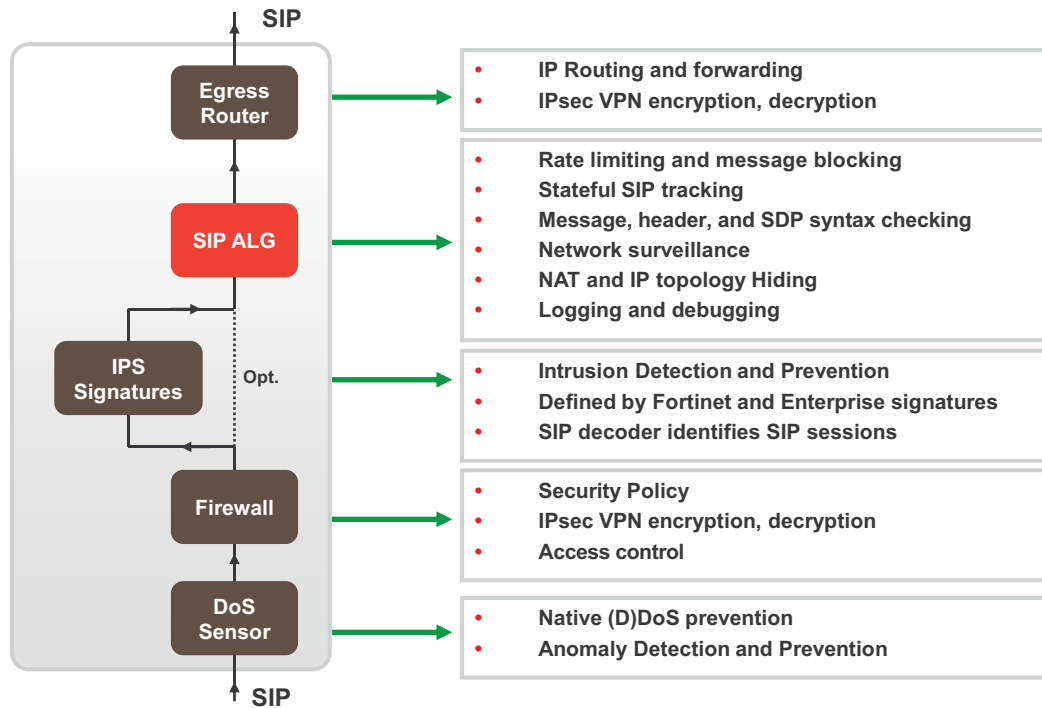
As shown in [Figure 356](#), the FortiGate SIP ALG intercepts SIP packets after they have been routed by the routing module, accepted by a security policy and passed through DoS and IPS Sensors (if DoS and IPS are enabled). The ALG raises SIP packets to the application layer, analyzes the SIP and SDP addressing information in the SIP messages, makes adjustments (for example, NAT) to this addressing if required, and then sends the packets out the egress interface to their destination.

The SIP ALG provides:

- All the same features as the SIP session help including NAT and SIP and RTP Pinholes.
- In addition for the ALG you can enable or disable RTP pinholing, SIP register pinholing and SIP contact pinholing. In a signalling only environment where the RTP stream bypasses the FortiGate unit, you can disable RTP pinholing to improve performance.
- SIP TCP and UDP support
- SIP Message order checking
- Configurable Header line length maximums



**Figure 356:**The SIP ALG works at the application level after ingress packets are accepted by a security policy



- Message fragment assembly (TCP)
- If SIP messages are fragmented across multiple packets, the FortiGate unit assembles the fragments, does inspection and pass the message in its entirety to the SIP server as one packet. This offloads the server from doing all the TCP processing of fragments.
- L4 Protocol Translation
- Message Flood Protection
- Protects a SIP server from intentional or unintentional DoS of flooding INVITE, REGISTER, and other SIP methods by allowing control of the rate that these messages pass through the FortiGate unit.
- SIP message type filtering
- The FortiGate unit can prevent specified SIP message types from passing through the FortiGate unit to a SIP server. For example In a voice only SIP implementation, there may be no need to permit a SUBSCRIBE message to ever make it's way to the SIP call processor. Also, if a SIP server cannot process some SIP message types you can use SIP message type filtering to block them. For example, a SIP server could have a bug that prevents it from

processing certain SIP messages. In this case you can temporarily block these message types until problem with the SIP server has been fixed.

- SIP statistics and logging
- SIP over IPv6
- SIP over SSL/TLS
- Deep SIP message syntax checking (also called deep SIP header inspection or SIP fuzzing protection). Prevents attacks that use malformed SIP messages. Can check many SIP headers and SDP statements. Configurable bypass and modification options.
- Hosted NAT traversal, Resolves IP address issue in SIP and SDP lines due to NAT-PT in far end firewall. Important feature for VoIP access networks.
- SIP High Availability (HA), including active-passive clustering and session pickup (session failover) for SIP sessions.
- Geographical Redundancy. In an HA configuration, if the active SIP server fails (missing SIP heartbeat messages or SIP traffic) SIP sessions can be redirected to a secondary SIP server in another location.
- SIP per request method message rate limitation with configurable threshold for SIP message rates per request method. Protects SIP servers from SIP overload and DoS attacks.
- RTP Bypass, Supports configurations with and without RTP pinholing. May inspect and protect SIP signaling only.
- SIP NAT with IP address conservation. Performs SIP and RTP aware IP Network Address translation. Preserves the lost IP address information in the SDP profile `ip` line for later processing/debugging in the SIP server. See [“NAT with IP address conservation” on page 2544](#).
- IP topology hiding
- The IP topology of a network can be hidden through NAT and NATPT manipulation of IP and SIP level addressing. For example, see [“SIP NAT configuration example: destination address translation \(destination NAT\)” on page 2539](#).
- SIP inspection without address translation
- The SIP ALG inspects SIP messages but addresses in the messages are not translated. This feature can be applied to a FortiGate unit operating in Transparent mode or in NAT/Route mode. In Transparent mode you add normal Transparent mode security policies that enable the SIP ALG and include a VoIP profile that causes the SIP ALG to inspect SIP traffic as required. For an example configuration, see [“Configuration example: SIP in Transparent Mode” on page 2522](#).
- For a FortiGate unit operating in NAT/Route mode, if SIP traffic can pass between different networks without requiring NAT because is supported by the routing configuration, you can add security policies that accept SIP traffic without enabling NAT. In the VoIP profile you can configure the SIP ALG to inspect SIP traffic as required.

## SIP ALG configuration overview

To apply the SIP ALG, you add a SIP VoIP profile to a security policy that accepts SIP sessions. All SIP sessions accepted by the security policy will be processed by the SIP ALG using the settings in the VoIP profile. The VoIP profile contains settings that are applied to SIP, Session Initiation Protocol for Instant Messaging and Presence Leveraging Extensions (SIMPLE) and Skinny Call Control Protocol (SCCP) sessions. You configure SIP and SCCP settings separately. SIP settings also apply to SIMPLE sessions.

## Enabling VoIP support on the web-based manager

Before you begin adding VoIP profiles to security policies you may have to enable VoIP support on the web-based manager. To do this, on the web-based manager go to *System > Admin > Settings* and make sure that the *VoIP* checkbox is selected.

From the CLI you can also enter the following command enable VoIP support on the GUI:

```
config system global
 set gui-voip-profile enable
end
```

## VoIP profiles

To add a new VoIP profile from the web-based manager go to *UTM Security Profiles > VoIP > Profile* and select *Create New*.

For SIP, from the web-based manager you can configure the VoIP profile to limit the number of SIP REGISTER and INVITE requests and enable logging of SIP sessions and SIP violations. Many additional options for configuring how the ALG processes SIP sessions are available from the CLI.

For SCCP you can limit the call setup time. Additional SCCP options are available from the CLI.

Use the following command to add a VoIP profile named *VoIP\_Pro\_1* from the CLI:

```
config voip profile
 edit VoIP_Pro_1
end
```

FortiGate units include two pre-defined VoIP profiles. On the web-based manager these profiles look identical. However, the CLI-only settings result in the following functionality.

<b>default</b>	<p>The most commonly used VoIP profile. This profile enables both SIP and SCCP and places the minimum restrictions on what calls will be allowed to negotiate. This profile allows normal SCCP, SIP and RTP sessions and enables the following security settings:</p> <ul style="list-style-type: none"><li>• <code>block-long-lines</code> to block SIP messages with lines that exceed maximum line lengths.</li><li>• <code>block-unknown</code> to block unrecognized SIP request messages.</li><li>• <code>log-violations</code> to write log messages that record SIP violations.</li><li>• <code>log-call-summary</code> to write log messages that record SIP call progress (similar to DLP archiving).</li><li>• <code>nat-trace</code> (see <a href="#">“NAT with IP address conservation” on page 2544</a>).</li><li>• <code>contact-fixup</code> perform NAT on the IP addresses and port numbers in SIP headers in SIP CONTACT messages even if they don't match the session's IP address and port numbers.</li></ul>
<b>strict</b>	<p>This profile is available for users who want to validate SIP messages and to only allow SIP sessions that are compliant with RFC 3261. In addition to the settings in the default VoIP profile, the strict profile sets all SIP deep message inspection header checking to block and drop SIP messages that contain malformed SIP or SDP lines that can be detected by the ALG. For more information about SIP deep header inspection, see <a href="#">“Deep SIP message inspection” on page 2558</a>.</p>

Neither of the default profiles applies SIP rate limiting or message blocking. To apply more ALG features to SIP sessions you can clone (copy) the pre-defined VoIP profiles and make your own modifications to them. For example, to clone the default profile and configure the limit for SIP NOTIFY request messages to 1000 messages per second per security policy and block SIP INFO request messages.

```
config voip profile
 clone default to my_voip_pro
 edit my_voip_pro
 config sip
 set notify-rate 1000
 set block-info enable
 end
 end
end
```

### Changing the port numbers that the SIP ALG listens on

Most SIP configurations use TCP or UDP port 5060 for SIP sessions and port 5061 for SIP SSL sessions. If your SIP network uses different ports for SIP sessions you can use the following command to configure the SIP ALG to listen on a different TCP, UDP, or SSL ports. For example, to change the TCP port to 5064, the UDP port to 5065, and the SSL port to 5066.

```
config system settings
 set sip-tcp-port 5064
 set sip-udp-port 5065
 set sip-ssl-port 5066
end
```

You also configure the SIP ALG to listen in two different TCP ports and two different UDP ports for SIP sessions. For example, if you receive SIP TCP traffic on port 5060 and 5064 and UDP traffic on ports 5061 and 5065 you can enter the following command to receive the SIP traffic on all of these ports:

```
config system settings
 set sip-tcp-port 5060 5064
 set sip-udp-port 5061 5065
end
```

### Disabling the SIP ALG in a VoIP profile

SIP is enabled by default in a VoIP profile. Usually you would want SIP to be enabled in a VoIP profile. But in some cases if you are just using the VoIP profile for SCCP you can use the following command to disable SIP in a VoIP profile.

```
config voip profile
 edit VoIP_Pro_2
 config sip
 set status disable
 end
 end
end
```

### SIP ALG get and diagnose commands

You can use the following commands to display diagnostic information for the SIP ALG.

Use the following command to list all active SIP calls being processed by the SIP ALG. You can also use the `clear` option to delete all active SIP calls being processed by the SIP ALG.

```
diagnose sys sip-proxy calls {clear | list}
```

Use the following commands to use filters to display specific information about the SIP ALG and the session that it is processing.

```
diagnose sys sip-proxy filter <filter_options>
diagnose sys sip-proxy log-filter <filter_options>
```

Use the following command to display the active SIP rate limiting meters and their current settings.

```
diagnose sys sip-proxy meters list
```

Use the following command to display status information about the SIP sessions being processed by the SIP ALG. You can also clear all SIP ALG statistics.

```
diagnose sys sip-proxy stats {clear | list}
```

## Conflicts between the SIP ALG and the session helper

Even if the SIP session helper is enabled, if a security policy with a VoIP profile that has SIP enabled accepts a SIP session on the TCP or UDP port that the SIP ALG listens on the ALG is used. You don't need to turn off the session helper to use the ALG.

You may find that the session helper is being used for some SIP sessions even when you only want to use the ALG. This happens if a policy that does not include a VoIP profile is accepting SIP sessions. The VoIP profile could have been left out of the policy by mistake or the wrong policy could be accepting SIP sessions.

Consider a configuration with a SIP server on a private network that is contacted by SIP phones on the Internet and on the private network (similar to the configuration in [Figure 352 on page 2494](#)). The FortiGate unit that provides NAT between the private network and the Internet requires a security policy with a firewall virtual IP that allows the SIP phones on the Internet to contact the SIP server. The FortiGate unit also requires outgoing security policies to allow the SIP phones and the SIP server to contact the SIP phones on the Internet.

If a VoIP profile is not added to one of the outgoing security policies the SIP sessions accepted by that policy will be processed by the SIP session helper instead of the SIP ALG. Also, it's possible that some of the SIP sessions could be accepted by a general outgoing policy instead of the policy intended for SIP traffic. You can fix the first problem by adding a VoIP profile to the policy. You can fix the second problem by reviewing the security policy order and source and destination addresses in the security policies and determining if there is a conflict between these and the IP addresses of the SIP server or SIP phones on the Internal network.

You can use `diagnose sys sip` commands to determine if the SIP session helper is processing SIP sessions. For example, the following command displays the overall status of the SIP sessions being processed by the SIP session helper:



The `diagnose sys sip` commands only display current status information. To see activity the SIP session helper has to actually be processing SIP sessions when you enter the command. For example, if the SIP session helper had been used for processing calls that ended 5 minutes ago, the command output would show no SIP session helper activity.

---

```
diagnose sys sip status
dialogs: max=32768, used=0
mappings: used=0
dialog hash by ID: size=2048, used=0, depth=0
dialog hash by RTP: size=2048, used=0, depth=0
mapping hash: size=2048, used=0, depth=0
count0: 0
count1: 0
count2: 0
count3: 0
count4: 0
```

This command output shows that the session helper is not processing SIP sessions because all of the used and count fields are 0. If any of these fields contains non-zero values then the SIP session helper may be processing SIP sessions.

Also, you can check to see if some ALG-only features are not being applied to all SIP sessions. For example, the VoIP usage widget on the FortiGate dashboard displays statistics for SIP and SCCP calls processed by the ALG but not for calls processed by the session helper. So if you see fewer calls than expected the session helper may be processing some of them.

Other logging and monitoring features such as log messages and DLP archiving are only supported by the ALG.

Finally, you can check the policy usage and session information dashboard widgets to see if SIP sessions are being accepted by the wrong security policies.

## Stateful SIP tracking, call termination, and session inactivity timeout

The SIP ALG tracks SIP dialogs over their lifespan between the first INVITE message and the Final 200 OK and ACK messages. For every SIP dialog, stateful SIP tracking reviews every SIP message and makes adjustment to SIP tracking tables as required. These adjustments include source and destination IP addresses, address translation, dialog expiration information, and media stream port changes. Such changes can also result in dynamically opening and closing pinholes. You can use the `diagnose sys sip-proxy stats list` and the `diagnose sys sip-proxy filter` command to view the SIP call data being tracked by the SIP ALG.

The SIP ALG uses the SIP Expires header line to time out a SIP dialog if the dialog is idle and a Re-INVITE or UPDATE message is not received. The SIP ALG gets the Session-Expires value, if present, from the 200 OK response to the INVITE message. If the SIP ALG receives an INVITE before the session times out, all timeout values are reset to the settings in the new INVITE message or to default values. As a precautionary measure, the SIP ALG uses hard timeout values to set the maximum amount of time a call can exist. This ensures that the FortiGate unit is protected if a call ends prematurely.

When a SIP dialog ends normally, the SIP ALG deletes the SIP call information and closes open pinholes. A SIP call can also end abnormally due to an unexpected signaling or transport event that cuts off the call. When a call ends abnormally the SIP messages to end the call may not be sent or received. A call can end abnormally for the following reasons:

- Phones or servers crash during a call and a BYE message is not received.
- To attack a SIP system, a malicious user never send a BYE message.
- Poor implementations of SIP fail to process Record-Route messages and never send a BYE message.
- Network failures prevent a BYE message from being received.

Any phone or server in a SIP call can cancel the call by sending a CANCEL message. When a CANCEL message is received by the FortiGate unit, the SIP ALG closes open pinholes. Before terminating the call, the ALG waits for the final 200 OK message.

The SIP ALG can be configured to terminate SIP calls if the SIP dialog message flow or the call RTP (media) stream is interrupted and does not recover. You can use the following commands to configure terminating inactive SIP sessions and to set timers or counters to control when the call is terminated by the SIP ALG.

### Adding a media stream timeout for SIP calls

Use the following command in a VoIP profile to terminate SIP calls accepted by a security policy containing the VoIP profile when the RTP media stream is idle for 100 seconds.

```
config voip profile
 edit VoIP_Pro_Name
 config sip
 set call-keepalive 100
 end
 end
```

You can adjust this setting between 1 and 10,080 seconds. The default call keepalive setting of 0 disables terminating a call if the media stream is interrupted. Set call keepalive higher if your network has latency problems that could temporarily interrupt media streams. If you have configured call keepalive and the FortiGate unit terminates calls unexpectedly you can increase the call keepalive time to resolve the problem.



Call keep alive should be used with caution because enabling this feature results in extra FortiGate CPU overhead and can cause delay/jitter for the VoIP call. Also, the FortiGate unit terminates the call without sending SIP messages to end the call. And if the SIP endpoints send SIP messages to terminate the call they will be blocked by the FortiGate unit if they are sent after the FortiGate unit terminates the call.

---

### Adding an idle dialog setting for SIP calls

Use the following command in a VoIP profile to terminate SIP calls when for a single security policy, when the configured number of SIP calls (or dialogs) has stopped receiving SIP messages or has not received legitimate SIP messages. Using this command you can configure how many dialogs that have been accepted by a security policy that the VoIP profile is added to become idle before the SIP ALG deletes the oldest ones. The following command sets the maximum number of idle dialogs to 200:

```
config voip profile
 edit VoIP_Pro_Name
 config sip
 set max-idle-dialogs 200
 end
 end
```

Idle dialogs would usually be dialogs that have been interrupted because of errors or problems or as the result of a SIP attack that opens a large number of SIP dialogs without closing them. This command provides a way to remove these dialogs from the dialog table and recover memory and resources being used by these open and idle dialogs.

You can adjust this setting between 1 and a very high number. The default maximum idle dialogs setting of 0 disables this feature. Set maximum dialogs higher if your network has latency problems that could temporarily interrupt SIP messaging. If you have configured max idle dialogs and the FortiGate unit terminates calls unexpectedly you can increase the max idle dialogs number to resolve the problem.

## Changing how long to wait for call setup to complete

In some cases and some configurations your SIP system may experience delays during call setup. If this happens, some SIP ALG timers may expire before call setup is complete and drop the call. In some cases you may also want to reduce the amount of time the SIP ALG allows for call setup to complete.

You can use the `provisional-invite-expiry-time` SIP VoIP profile option to control how long the SIP ALG waits for provisional INVITE messages before assuming that the call setup has been interrupted and the SIP call should be dropped. The default value for this timer is 210 seconds. You can change it to between 10 and 3600 seconds.

Use the following command to change the expiry time to 100 seconds.

```
config voip profile
 edit Profile_name
 config sip
 set provisional-invite-expiry-time 100
 end
 end
```

## SIP and RTP/RTCP

FortiGate units support the Real Time Protocol (RTP) application layer protocol for the VoIP call audio stream. RTP uses dynamically assigned port numbers that can change during a call. SIP control messages that start a call and that are sent during the call inform callers of the port number to use and of port number changes during the call.

During a call, each RTP session will usually have a corresponding Real Time Control Protocol (RTCP) session. By default, the RTCP session port number is one higher than the RTP port number.

The RTP port number is included in the `m=` part of the SDP profile. In the example above, the SIP INVITE message includes RTP port number is 49170 so the RTCP port number would be 49171. In the SIP response message the RTP port number is 3456 so the RTCP port number would be 3457.

## How the SIP ALG creates RTP pinholes

The SIP ALG requires the following information to create a pinhole. The SIP ALG finds this information in SIP messages and some is provided by the SIP ALG:

<b>Protocol</b>	UDP (Extracted from SIP messages by the SIP ALG.)
<b>Source IP</b>	Any
<b>Source port</b>	Any



<b>Destination IP</b>	The SIP ALG extracts the destination IP address from the c= line in the SDP profile. The c= line can appear in either the session or media part of the SDP profile. The SIP ALG uses the IP address in the c= line of the media part of the SDP profile first. If the media part does not contain a c= line, the SIP ALG checks the c= line in the session part of the SDP profile. If the session part of the profile doesn't contain a c= line the packet is dropped. Pinholes for RTP and RTCP sessions share the same destination IP address.
<b>Destination port</b>	The SIP ALG extracts the destination port number for RTP from the m= field and adds 1 to this number to get the RTCP port number.
<b>Lifetime</b>	The length of time during which the pinhole will be open. When the lifetime ends, the SIP ALG removes the pinhole.

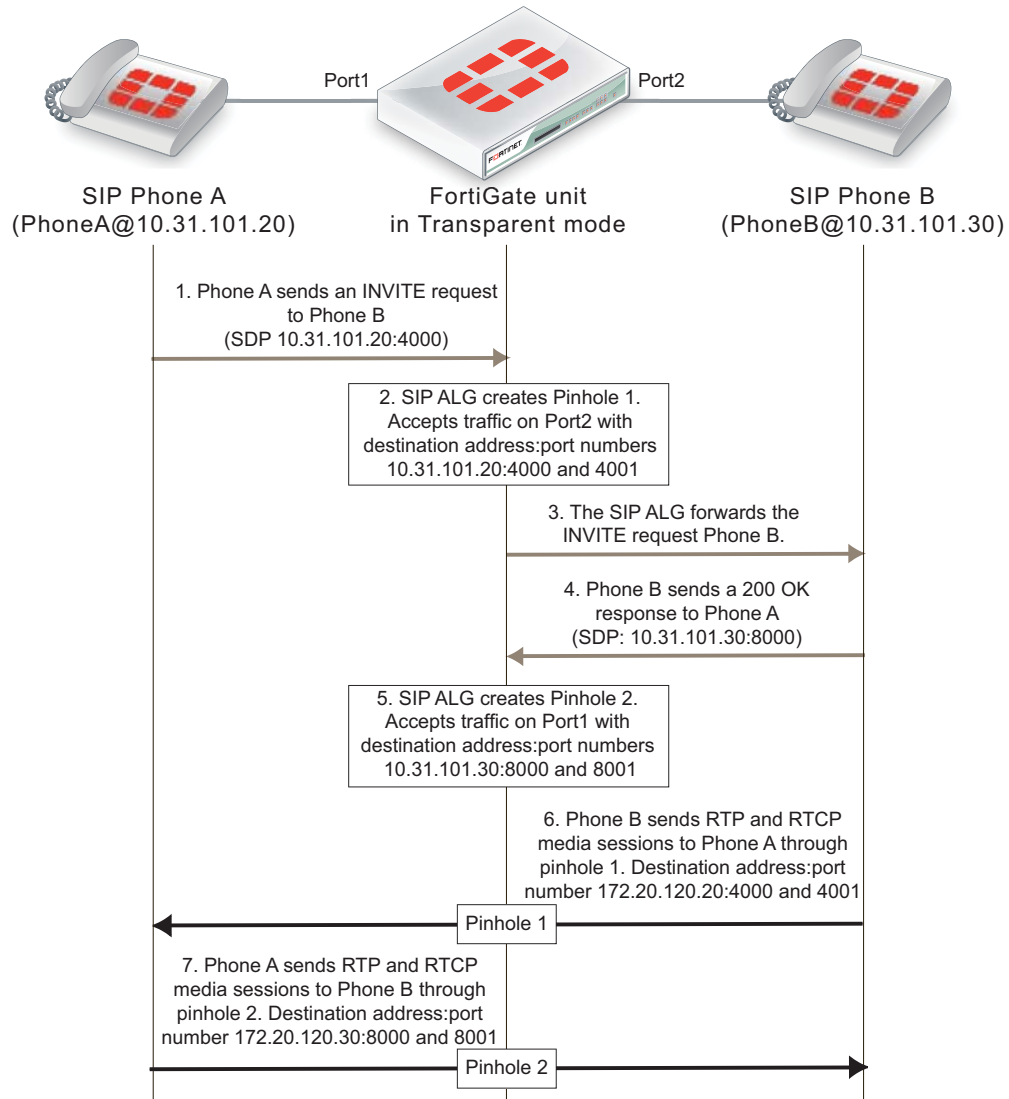
The SIP ALG keeps RTP pinholes open as long as the SIP session is alive. When the associated SIP session is terminated by the SIP ALG or the SIP phones or servers participating in the call, the RTP pinhole is closed.

[Figure 357](#) shows a simplified call setup sequence that shows how the SIP ALG opens pinholes. Phone A and Phone B are installed on either side of a FortiGate unit operating in Transparent mode. Phone A and Phone B are on the same subnet. The FortiGate unit includes a security policy that accepts SIP sessions from port1 to port2 and from port2 to port1. The FortiGate unit does not require an RTP security policy, just the SIP policy.

You can see from this diagram that the SDP profile in the INVITE request from Phone A indicates that Phone A is expecting to receive a media stream sent to its IP address using port 4000 for RTP and port 4001 for RTCP. The SIP ALG creates pinhole 1 to allow this media traffic to pass through the FortiGate unit. Pinhole 1 is opened on the Port2 interface and will accept media traffic sent from Phone B to Phone A.

When Phone B receives the INVITE request from Phone A, Phone B will know to send media streams to Phone A using destination IP address 10.31.101.20 and ports 4000 and 4001. The 200 OK response sent from Phone B indicates that Phone B is expecting to receive a media stream sent to its IP address using ports 8000 and 8001. The SIP ALG creates pinhole 2 to allow this media traffic to pass through the FortiGate unit. Pinhole 2 is opened on the Port1 interface and will accept media traffic sent from Phone A to Phone B.

**Figure 357:**SIP call setup with a FortiGate unit in Transparent mode




### Configuration example: SIP in Transparent Mode

Figure 358 shows an example SIP network consisting of a FortiGate unit operating in Transparent mode between two SIP phones. Since the FortiGate unit is operating in Transparent mode both phones are on the same network and the FortiGate unit and the SIP ALG does not perform NAT. Even though the SIP ALG is not performing NAT you can use this configuration to apply SIP security features to the SIP traffic.

The FortiGate unit requires two security policies that accept SIP packets. One to allow SIP Phone A to start a session with SIP Phone B and one to allow SIP Phone B to start a session with SIP Phone A.

**Figure 358:**SIP network with FortiGate unit in Transparent mode



## General configuration steps

The following general configuration steps are required for this SIP configuration. This example uses the default VoIP profile. The example also includes security policies that specifically allow SIP sessions using UDP port 5060 from Phone A to Phone B and from Phone B to Phone A. In most cases you would have more than two phones so would use more general security policies. Also, you can set the security service to ANY to allow traffic other than SIP on UDP port 5060.

1. Add firewall addresses for Phone A and Phone B.
2. Add a security policy that accepts SIP sessions initiated by Phone A and includes the default VoIP profile.
3. Add a security policy that accepts SIP sessions initiated by Phone B and includes the default VoIP profile.

## Configuration steps - web-based manager



Before you begin this procedure you may have to enable VoIP support on the web-based manager by going to *System > Admin > Settings* and selecting the *VoIP* checkbox.

### To add firewall addresses for the SIP phones

1. Go to *Firewall Objects > Address > Addresses*.
2. Add the following addresses for Phone A and Phone B:

<b>Address Name</b>	Phone_A
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	10.31.101.20/255.255.255.255
<b>Interface</b>	port1
<b>Address Name</b>	Phone_B
<b>Type</b>	Subnet / IP Range
<b>Subnet / IP Range</b>	10.31.101.30/255.255.255.255
<b>Interface</b>	port2

### To add security policies to apply the SIP ALG to SIP sessions

1. Go to *Policy > Policy > Policy*.
2. Select Create New to add a security policy.

3. Add a security policy to allow Phone A to send SIP request messages to Phone B:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	port1
<b>Source Address</b>	Phone_A
<b>Outgoing Interface</b>	port2
<b>Destination Address</b>	Phone_B
<b>Schedule</b>	always
<b>Service</b>	SIP
<b>Action</b>	ACCEPT

4. Under *UTM Security Profiles*, select *Use Standard UTM Profiles*.
5. Turn on *VoIP* and select the *default* VoIP profile.
6. Select *OK*.
7. Add a security policy to allow Phone B to send SIP request messages to Phone A:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	port2
<b>Source Address</b>	Phone_B
<b>Outgoing Interface</b>	port1
<b>Destination Address</b>	Phone_A
<b>Schedule</b>	always
<b>Service</b>	SIP
<b>Action</b>	ACCEPT

8. Under *UTM Security Profiles*, select *Use Standard UTM Profiles*.
9. Turn on *VoIP* and select the *default* VoIP profile.
10. Select *OK*.

## Configuration steps - CLI

### To add firewall addresses for Phone A and Phone B and security policies to apply the SIP ALG to SIP sessions

1. Enter the following command to add firewall addresses for Phone A and Phone B.

```
config firewall address
 edit Phone_A
 set associated interface port1
 set type ipmask
 set subnet 10.31.101.20 255.255.255.255
 next
 edit Phone_B
 set associated interface port2
 set type ipmask
 set subnet 10.31.101.30 255.255.255.255
end
```

2. Enter the following command to add security policies to allow Phone A to send SIP request messages to Phone B and Phone B to send SIP request messages to Phone A.

```
config firewall policy
 edit 0
 set srcintf port1
 set dstintf port2
 set srcaddr Phone_A
 set dstaddr Phone_B
 set action accept
 set schedule always
 set service SIP
 set utm-status enable
 set voip-profile default
 next
 edit 0
 set srcintf port2
 set dstintf port1
 set srcaddr Phone_B
 set dstaddr Phone_A
 set action accept
 set schedule always
 set service SIP
 set utm-status enable
 set voip-profile default
end
```

## RTP enable/disable (RTP bypass)

You can configure the SIP ALG to stop from opening RTP pinholes. Called RTP bypass, this configuration can be used when you want to apply SIP ALG features to SIP signalling messages but do not want the RTP media streams to pass through the FortiGate unit. The FortiGate unit only acts as a signalling firewall and RTP media session bypass the FortiGate unit and no pinholes need to be created.

Enter the following command to enable RTP bypass in a VoIP profile by disabling opening RTP pinholes:

```
config voip profile
 edit VoIP_Pro_1
 config sip
 set rtp disable
 end
 end
```

## Opening and closing SIP register, contact, via and record-route pinholes

You can use the `open-register-pinhole`, `open-contact-pinhole`, `open-via-port`, and `open-record-route-pinhole` VoIP profile CLI options to control whether the FortiGate unit opens various pinholes.

If `open-register-pinhole` is enabled (the default setting) the FortiGate unit opens pinholes for SIP Register request messages. You can disable `open-register-pinhole` so that the FortiGate unit does not open pinholes for SIP Register request messages.

If `open-contact-pinhole` is enabled (the default setting) the FortiGate unit opens pinholes for non-Register SIP request messages. You can disable `open-contact-pinhole` so that the FortiGate unit does not open pinholes for non-register requests. Non-register pinholes are usually opened for SIP INVITE requests.

If `open-via-pinhole` is disabled (the default setting) the FortiGate unit does not open pinholes for Via messages. You can enable `open-via-pinhole` so that the FortiGate unit opens pinholes for Via messages.

If `open-record-route-pinhole` is enabled (the default setting) the FortiGate unit opens pinholes for Record-Route messages. You can disable `open-record-route-pinhole` so that the FortiGate unit does not open pinholes for Record-Route messages.

Usually you would want to open these pinholes. Keeping them closed may prevent SIP from functioning properly through the FortiGate unit. They can be disabled, however, for interconnect scenarios (where all SIP traffic is between proxies and traveling over a single session). In some cases these settings can also be disabled in access scenarios if it is known that all users will be registering regularly so that their contact information can be learned from the register request.

You might want to prevent pinholes from being opened to avoid creating a pinhole for every register or non-register request. Each pinhole uses additional system memory, which can affect system performance if there are hundreds or thousands of users, and requires refreshing which can take a relatively long amount of time if there are thousands of active calls.

To configure a VoIP profile to prevent opening register and non-register pinholes:

```
config voip profile
 edit VoIP_Pro_1
 config sip
 set open-register-pinhole disable
 set open-contact-pinhole disable
 end
 end
```

In some cases you may not want to open pinholes for the port numbers specified in SIP Contact headers. For example, in an interconnect scenario when a FortiGate unit is installed between two SIP servers and the only SIP traffic through the FortiGate unit is between these SIP servers pinholes may not need to be opened for the port numbers specified in the Contact header lines.

If you disable `open-register-pinhole` then pinholes are not opened for ports in Contact header lines in SIP Register messages. If you disable `open-contact-pinhole` then pinholes are not opened for ports in Contact header lines in all SIP messages except SIP Register messages.

## Accepting SIP register responses

You can enable the VoIP profile `open-via-pinhole` options to accept a SIP Register response message from a SIP server even if the source port of the Register response message is different from the destination port.

Most SIP servers use 5060 as the source port in the SIP register response. Some SIP servers, however, may use a different source port. If your SIP server uses a different source port, you can enable `open-via-pinhole` and the SIP ALG will create a temporary pinhole when the Register request from a SIP client includes a different source port. The FortiGate unit will accept a SIP Register response with any source port number from the SIP server.

Enter the following command to enable accepting any source port from a SIP server:

```
config voip profile
 edit VoIP_Pro_1
 config sip
 set open-via-pinhole enable
 end
 end
```

## How the SIP ALG performs NAT

In most Network Address Translation (NAT) configurations, multiple hosts in a private network share a single public IP address to access the Internet. For sessions originating on the private network for the Internet, NAT replaces the private IP address of the PC in the private subnet with the public IP address of the NAT device. The NAT device converts the public IP address for responses from the Internet back into the private address before sending the response over the private network to the originator of the session.

Using NAT with SIP is more complex because of the IP addresses and media stream port numbers used in SIP message headers and bodies. When a caller on the private network sends a SIP message to a phone or SIP server on the Internet, the SIP ALG must translate the private network addresses in the SIP message to IP addresses and port numbers that are valid on the Internet. When the response message is sent back to the caller, the SIP ALG must translate these addresses back to valid private network addresses.

In addition, the media streams generated by the SIP session are independent of the SIP message sessions and use varying port numbers that can also change during the media session. The SIP ALG opens pinholes to accept these media sessions, using the information in the SIP messages to determine the pinholes to open. The ALG may also perform port translation on the media sessions.

When an INVITE message is received by the SIP ALG, the FortiGate unit extracts addressing and port number information from the message header and stores it in a SIP dialog table. Similar to an IP session table the data in the dialog table is used to translate addresses in subsequent SIP messages that are part of the same SIP call.

When the SIP ALG receives a response to the INVITE message arrives, (for example, an ACK or 200 OK), the SIP ALG compares the addresses in the message fields against the entries in the SIP dialog table to identify the call context of the message. The SIP ALG then translates addresses in the SIP message before forwarding them to their destination.

The addressing and port number information in SDP fields is used by the ALG to reserve ports for the media session and create a NAT mapping between them and the ports in the SDP fields. Because SDP uses sequential ports for the RTP and RTCP channels, the ALG provides consecutive even-odd ports.

## Source address translation

When a SIP call is started by a phone on a private network destined for a phone on the Internet, only source address translation is required. The phone on the private network attempts to contact the actual IP address of the phone on the Internet. However, the source address of the phone on the private network is not routable on the Internet so the SIP ALG must translate all private IP addresses in the SIP message into public IP addresses.

To configure the FortiGate for source address translation you add security policy that accepts sessions from the internal network destined for the Internet. You must enable NAT for the security policy and add a VoIP profile.

When a SIP request is received from the internal to the external network, the SIP ALG replaces the private network IP addresses and port numbers in the SIP message with the IP address of the FortiGate interface connected to the Internet. Depending on the content of the message, the ALG translates addresses in the Via:, Contact:, Route:, and Record-Route: SIP header fields. The message is then forwarded to the destination (either a VoIP phone or a SIP server on the Internet).

The VoIP phone or server in the Internet sends responses to these SIP messages to the external interface of the FortiGate unit. The addresses in the response messages are translated back into private network addresses and the response is forwarded to the originator of the request.

For the RTP communication between the SIP phones, the SIP ALG opens pinholes to allow media through the FortiGate unit on the dynamically assigned ports negotiated based on information in the SDP and the Via:, Contact:, and Record-Route: header fields. The pinholes also allow incoming packets to reach the Contact:, Via:, and Record-Route: IP addresses and ports. When processing return traffic, the SIP ALG inserts the original Contact:, Via:, Route:, and Record-Route: SIP fields back into the packets.

## Destination address translation

Incoming calls are directed from a SIP phone on the Internet to the interface of the FortiGate unit connected to the Internet. To receive these calls you must add a security policy to accept SIP sessions from the Internet. The security policy requires a firewall virtual IP. SIP INVITE messages from the Internet connect to the external IP address of the virtual IP. The SIP ALG uses the destination address translation defined in the virtual IP to translated the addresses in the SIP message to addresses on the private network.

When a 200 OK response message arrives from the private network, the SIP ALG translates the addresses in the message to Internet addresses and opens pinholes for media sessions from the private network to the Internet.

When the ACK message is received for the 200 OK, it is also intercepted by the SIP ALG. If the ACK message contains SDP information, the SIP ALG checks to determine if the IP addresses and port numbers are not changed from the previous INVITE. If they are, the SIP ALG deletes pinholes and creates new ones as required. The ALG also monitors the Via:, Contact:, and Record-Route: SIP fields and opens new pinholes as required.



## Call Re-invite messages

SIP Re-INVITE messages can dynamically add and remove media sessions during a call. When new media sessions are added to a call the SIP ALG opens new pinholes and update SIP dialog data. When media sessions are ended, the SIP ALG closes pinholes that are no longer needed and removes SIP dialog data.

## How the SIP ALG translates IP addresses in SIP headers

The SIP ALG applies NAT to SIP sessions by translating the IP addresses contained in SIP headers. For example, the following SIP message contains most of the SIP fields that contain addresses that need to be translated:

```
INVITE PhoneB@172.20.120.30 SIP/2.0
Via: SIP/2.0/UDP 172.20.120.50:5434
From: PhoneA@10.31.101.20
To: PhoneB@172.20.120.30
Call-ID: a12abcde@172.20.120.50
Contact: PhoneA@10.31.101.20:5434
Route: <sip:example@172.20.120.50:5060>
Record-Route: <sip:example@172.20.120.50:5060>
```

How IP address translation is performed depends on whether source NAT or destination NAT is applied to the session containing the message:

### Source NAT translation of IP addresses in SIP messages

Source NAT translation occurs for SIP messages sent from a phone or server on a private network to a phone or server on the Internet. The source addresses in the SIP header fields of the message are typically set to IP addresses on the private network. The SIP ALG translates these addresses to the address the FortiGate unit interface connected to the Internet.

**Table 116:**Source NAT translation of IP addresses in SIP request messages

SIP header	NAT action
<b>To:</b>	None
<b>From:</b>	Replace private network address with IP address of FortiGate unit interface connected to the Internet.
<b>Call-ID:</b>	Replace private network address with IP address of FortiGate unit interface connected to the Internet.
<b>Via:</b>	Replace private network address with IP address of FortiGate unit interface connected to the Internet.
<b>Request-URI:</b>	None
<b>Contact:</b>	Replace private network address with IP address of FortiGate unit interface connected to the Internet.
<b>Record-Route:</b>	Replace private network address with IP address of FortiGate unit interface connected to the Internet.
<b>Route:</b>	Replace private network address with IP address of FortiGate unit interface connected to the Internet.

Response messages from phones or servers on the Internet are sent to the FortiGate unit interface connected to the Internet where the destination addresses are translated back to addresses on the private network before forwarding the SIP response message to the private network.

**Table 117:**Source NAT translation of IP addresses in SIP response messages

SIP header	NAT action
<b>To:</b>	None
<b>From:</b>	Replace IP address of FortiGate unit interface connected to the Internet with private network address.
<b>Call-ID:</b>	Replace IP address of FortiGate unit interface connected to the Internet with private network address.
<b>Via:</b>	Replace IP address of FortiGate unit interface connected to the Internet with private network address.
<b>Request-URI:</b>	N/A
<b>Contact:</b>	None
<b>Record-Route:</b>	Replace IP address of FortiGate unit interface connected to the Internet with private network address.
<b>Route:</b>	Replace IP address of FortiGate unit interface connected to the Internet with private network address.

### Destination NAT translation of IP addresses in SIP messages

Destination NAT translation occurs for SIP messages sent from a phone or server on the Internet to a firewall virtual IP address. The destination addresses in the SIP header fields of the message are typically set to the virtual IP address. The SIP ALG translates these addresses to the address of a SIP server or phone on the private network on the other side of the FortiGate unit.

**Table 118:**Destination NAT translation of IP addresses in SIP request messages

SIP header	NAT action
<b>To:</b>	Replace VIP address with address on the private network as defined in the firewall virtual IP.
<b>From:</b>	None
<b>Call-ID:</b>	None
<b>Via:</b>	None
<b>Request-URI:</b>	Replace VIP address with address on the private network as defined in the firewall virtual IP.
<b>Contact:</b>	None
<b>Record-Route:</b>	None
<b>Route:</b>	None

SIP response messages sent in response to the destination NAT translated messages are sent from a server or a phone on the private network back to the originator of the request messages on the Internet. These reply messages are accepted by the same security policy that accepted the initial request messages. The firewall VIP in the original security policy contains the information that the SIP ALG uses to translate the private network source addresses in the SIP headers into the firewall virtual IP address.

**Table 119:** Destination NAT translation of IP addresses in SIP response messages

SIP header	NAT action
To:	None
From:	Replace private network address with firewall VIP address.
Call-ID:	None
Via:	None
Request-URI:	N/A
Contact:	Replace private network address with firewall VIP address.
Record-Route:	Replace private network address with firewall VIP address.
Route:	None

## How the SIP ALG translates IP addresses in the SIP body

The SDP session profile attributes in the SIP body include IP addresses and port numbers that the SIP ALG uses to create pinholes for the media stream.

The SIP ALG translates IP addresses and port numbers in the o=, c=, and m= SDP lines. For example, in the following lines the ALG could translate the IP addresses in the o= and c= lines and the port number (49170) in the m= line.

```
o=PhoneA 5462346 332134 IN IP4 10.31.101.20
c=IN IP4 10.31.101.20
m=audio 49170 RTP 0 3
```

If the SDP session profile includes multiple RTP media streams, the SIP ALG opens pinholes and performs the required address translation for each one.

The two most important SDP attributes for the SIP ALG are c= and m=. The c= attribute is the connection information attribute. This field can appear at the session or media level. The syntax of the connection attribute is:

```
c=IN {IPV4 | IPV6} <destination_ip_address>
```

Where

- IN is the network type. FortiGate units support the IN or Internet network type.
- {IPV4 | IPV6} is the address type. FortiGate units support IPv4 or IPv6 addresses in SDP statements. However, FortiGate units do not support all types of IPv6 address translation. See [“SIP over IPv6” on page 2558](#).
- <destination\_IP\_address> is the unicast numeric destination IP address or domain name of the connection in either IPv4 or IPv6 format.

The syntax of the media attribute is:

```
m=audio <port_number> RTP <format_list>
```

Where

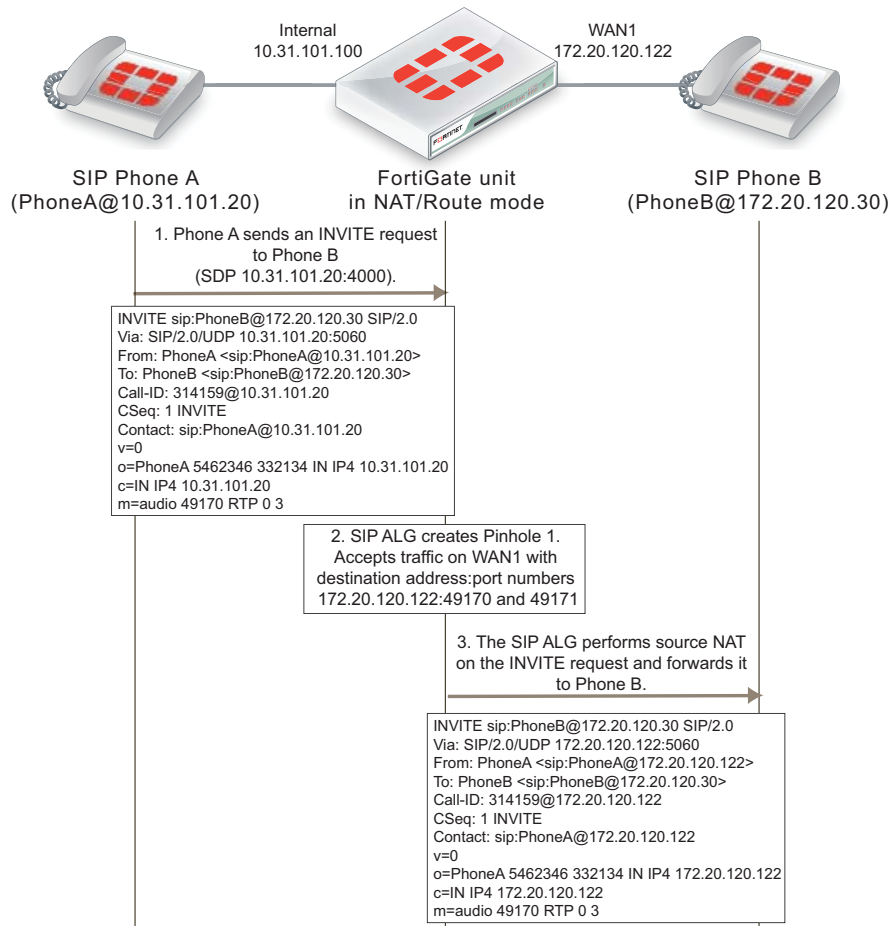
- `audio` is the media type. FortiGate units support the `audio` media type.
- `<port_number>` is the destination port number used by the media stream.
- RTP is the application layer transport protocol used for the media stream. FortiGate units support the Real Time Protocol (RTP) transport protocol.
- `<format_list>` is the format list that provides information about the application layer protocol that the media uses.

## SIP NAT scenario: source address translation (source NAT)

Figure 359 and Figure 360 show a source address translation scenario involving two SIP phones on different networks, separated by a FortiGate unit. In the scenario, SIP Phone A sends an INVITE request to SIP Phone B and SIP Phone B replies with a 200 OK response and then the two phones start media streams with each other.

To simplify the diagrams, some SIP messages are not included (for example, the Ringing and ACK response messages) and some SIP header lines and SDP profile lines have been removed from the SIP messages.

**Figure 359:** SIP source NAT scenario part 1: INVITE request sent from Phone A to Phone B

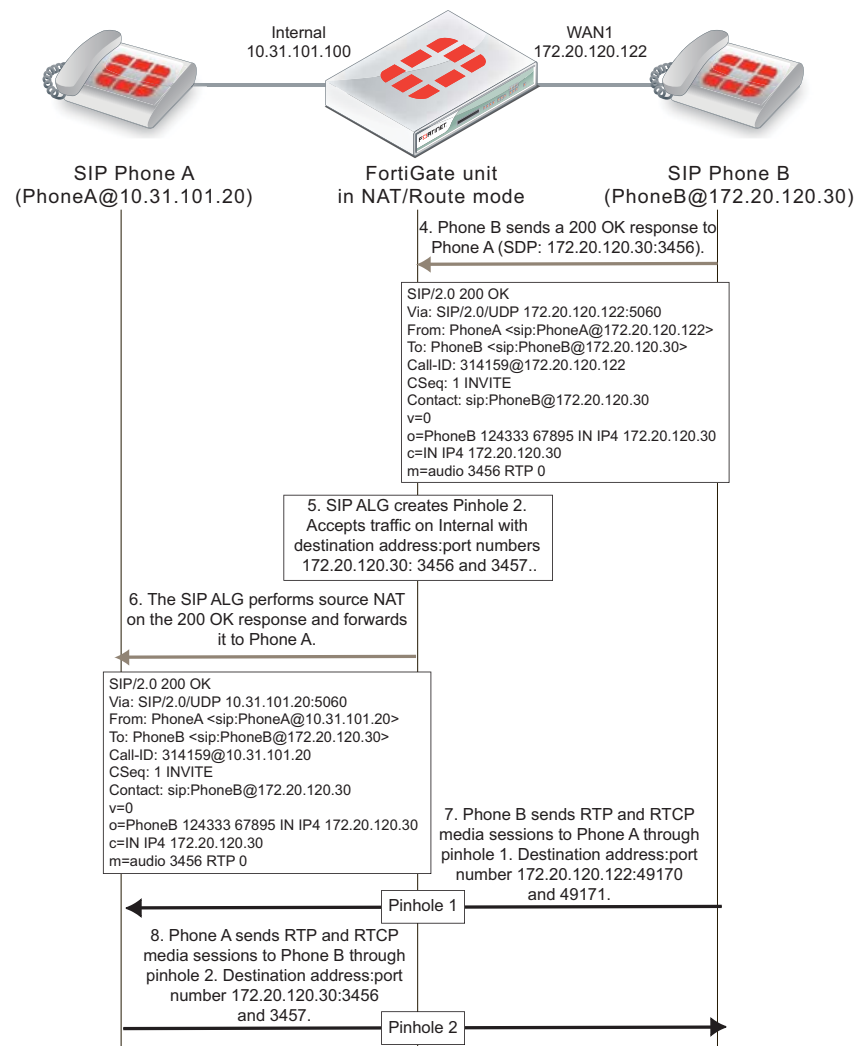


For the replies to SIP packets sent by Phone A to be routable on Phone B's network, the FortiGate unit uses source NAT to change their source address to the address of the WAN1 interface. The SIP ALG makes similar changes to the source addresses in the SIP headers and SDP profile. For example, the original INVITE request from Phone A includes the address of Phone A (10.31.101.20) in the from header line. After the INVITE request passes through the FortiGate unit, the address of Phone A in the From SIP header line is translated to 172.20.120.122, the address of the FortiGate unit WAN1 interface. As a result, Phone B will reply to SIP messages from Phone A using the WAN1 interface IP address.

The FortiGate unit also opens a pinhole so that it can accept media sessions sent to the WAN1 IP address using the port number in the m= line of the INVITE request and forward them to Phone A after translating the destination address to the IP address of Phone A.

Phone B sends the 200 OK response to the INVITE message to the WAN1 interface. The SDP profile includes the port number that Phone B wants to use for its media stream. The FortiGate unit forwards the 200 OK response to Phone A after translating the addresses in the SIP and SDP lines back to the IP address of Phone A. The SIP ALG also opens a pinhole on the Internal interface that accepts media stream sessions from Phone A with destination address set to the IP address of Phone B and using the port that Phone B added to the SDP m= line.

**Figure 360:**SIP source NAT scenario part 2: 200 OK returned and media streams established



## SIP NAT scenario: destination address translation (destination NAT)

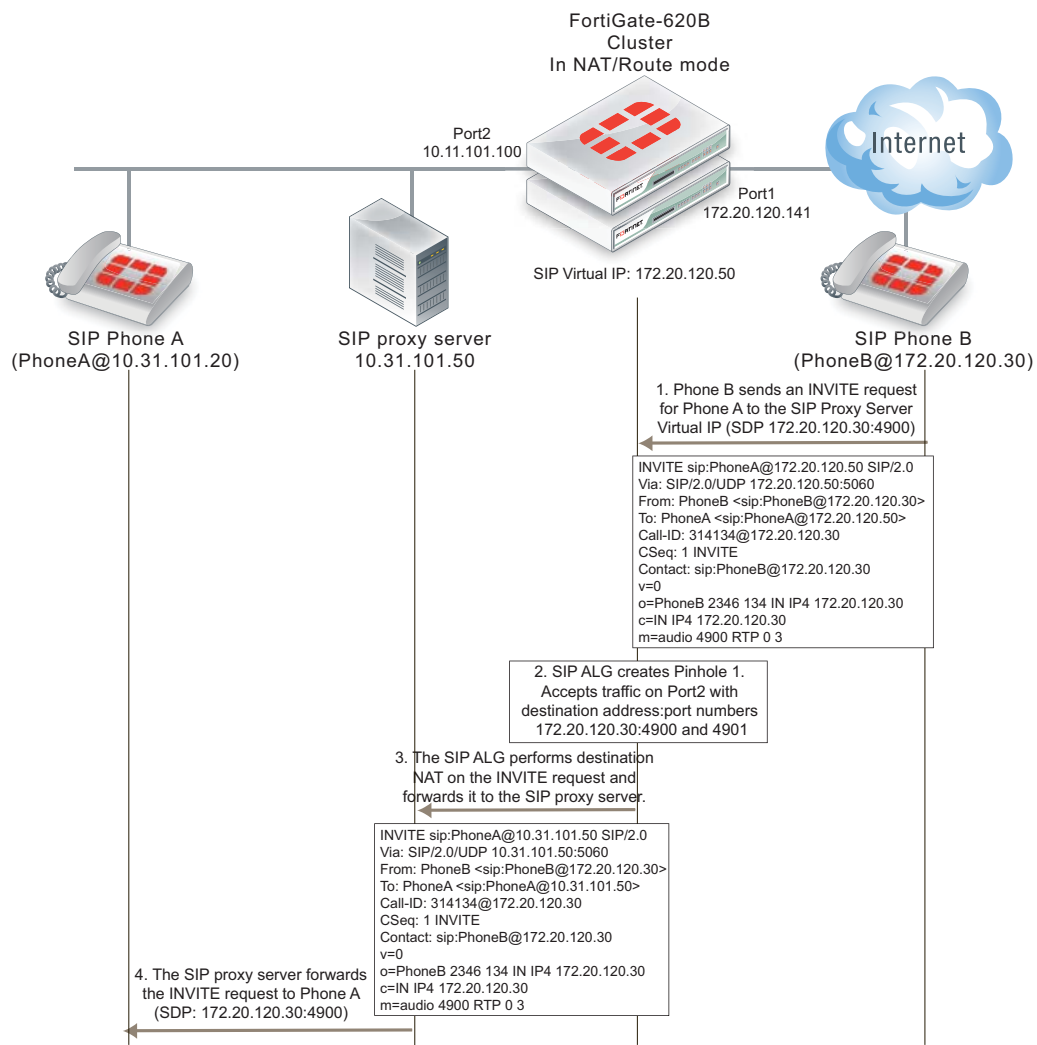
Figure 361 and Figure 362 show how the SIP ALG translates addresses in a SIP INVITE message sent from SIP Phone B on the Internet to SIP Phone A on a private network using the SIP proxy server. Because the addresses on the private network are not visible from the Internet, the security policy on the FortiGate unit that accepts SIP sessions includes a virtual IP. Phone A sends SIP the INVITE message to the virtual IP address. The FortiGate unit accepts the INVITE message packets and using the virtual IP, translates the destination address of the packet to the IP address of the SIP proxy server and forwards the SIP message to it.

To simplify the diagrams, some SIP messages are not included (for example, the Ringing and ACK response messages) and some SIP header lines and SDP profile lines have been removed from the SIP messages.

The SIP ALG also translates the destination addresses in the SIP message from the virtual IP address (172.20.120.50) to the SIP proxy server address (10.31.101.50). For this configuration to work, the SIP proxy server must be able to change the destination addresses for Phone A in the SIP message from the address of the SIP proxy server to the actual address of Phone A.

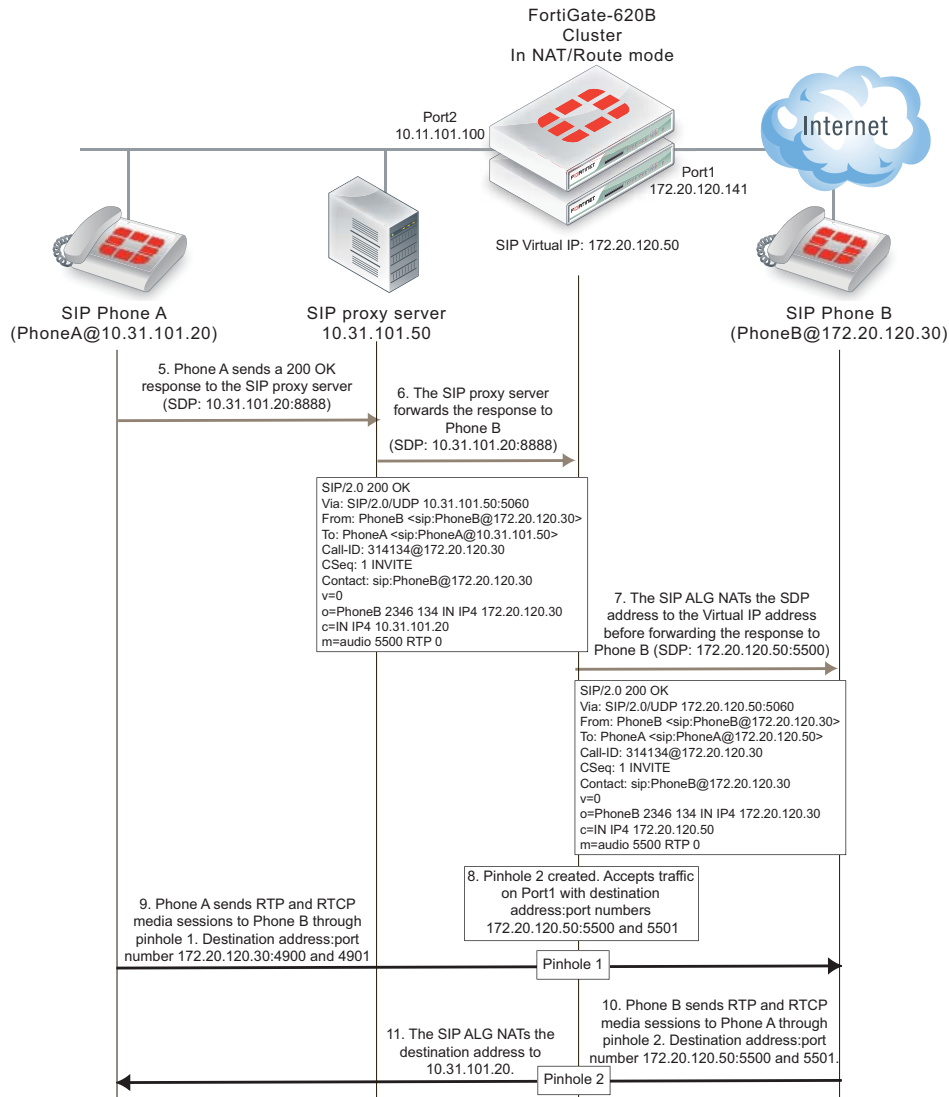
The SIP ALG also opens a pinhole on the Port2 interface that accepts media sessions from the private network to SIP Phone B using ports 4900 and 4901.

**Figure 361:**SIP destination NAT scenario part 1: INVITE request sent from Phone B to Phone A



Phone A sends a 200 OK response back to the SIP proxy server. The SIP proxy server forwards the response to Phone B. The FortiGate unit accepts the 100 OK response. The SIP ALG translates the Phone A addresses back to the SIP proxy server virtual IP address before forwarding the response back to Phone B. The SIP ALG also opens a pinhole using the SIP proxy server virtual IP which is the address in the o= line of the SDP profile and the port number in the m= line of the SDP code.

**Figure 362:**SIP destination NAT scenario part 2: 200 OK returned to Phone B and media streams established




The media stream from Phone A is accepted by pinhole one and forwarded to Phone B. The source address of this media stream is changed to the SIP proxy server virtual IP address. The media stream from Phone B is accepted by pinhole 2 and forwarded to Phone B. The destination address of this media stream is changed to the IP address of Phone A.

## SIP NAT configuration example: source address translation (source NAT)

This configuration example shows how to configure the FortiGate unit to support the source address translation scenario shown in [Figure 363](#). The FortiGate unit requires two security policies that accept SIP packets. One to allow SIP Phone A to start a session with SIP Phone B and one to allow SIP Phone B to start a session with SIP Phone A. Both of these policies must include source NAT. In this example the networks are not hidden from each other so destination NAT is not required.

**Figure 363:**SIP source NAT configuration



### General configuration steps

The following general configuration steps are required for this SIP configuration. This example uses the default VoIP profile. The example also includes security policies that specifically allow SIP sessions using UDP port 5060 from Phone A to Phone B and from Phone B to Phone A. In most cases you would have more than two phones so would use more general security policies. Also, you can set the firewall service to ANY to allow traffic other than SIP on UDP port 5060.

1. Add firewall addresses for Phone A and Phone B.
2. Add a security policy that accepts SIP sessions initiated by Phone A and includes the default VoIP profile.
3. Add a security policy that accepts SIP sessions initiated by Phone B and includes the default VoIP profile.

### Configuration steps - web-based manager

#### To add firewall addresses for the SIP phones

1. Go to *Firewall Objects > Address*.
2. Add the following addresses for Phone A and Phone B:

<b>Address Name</b>	Phone_A
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	10.31.101.20/255.255.255.255
<b>Interface</b>	Internal
<b>Address Name</b>	Phone_B
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	172.20.120.30/255.255.255.255
<b>Interface</b>	wan1



### To add security policies to apply the SIP ALG to SIP sessions

1. Go to *Policy > Policy > Policy*.
2. Select Create New to add a security policy.
3. Add a security policy to allow Phone A to send SIP request messages to Phone B:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	internal
<b>Source Address</b>	Phone_A
<b>Outgoing Interface</b>	wan1
<b>Destination Address</b>	Phone_B
<b>Schedule</b>	always
<b>Service</b>	SIP
<b>Action</b>	ACCEPT

4. Select *Enable NAT* and select *Use Destination Interface Address*.
5. Under *UTM Security Profiles*, select *Use Standard UTM Profiles*.
6. Turn on *VoIP* and select the *default* VoIP profile.
7. Select *OK*.
8. Add a security policy to allow Phone B to send SIP request messages to Phone A:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	wan1
<b>Source Address</b>	Phone_B
<b>Outgoing Interface</b>	internal
<b>Destination Address</b>	Phone_A
<b>Schedule</b>	always
<b>Service</b>	SIP
<b>Action</b>	ACCEPT

9. Select *Enable NAT* and select *Use Destination Interface Address*.
10. Under *UTM Security Profiles*, select *Use Standard UTM Profiles*.
11. Turn on *VoIP* and select the *default* VoIP profile.
12. Select *OK*.

## Configuration steps - CLI

### To add firewall addresses for Phone A and Phone B and security policies to apply the SIP ALG to SIP sessions

- 1 Enter the following command to add firewall addresses for Phone A and Phone B.

```
config firewall address
 edit Phone_A
 set associated interface internal
 set type ipmask
 set subnet 10.31.101.20 255.255.255.255
 next
 edit Phone_B
 set associated interface wan1
 set type ipmask
 set subnet 172.20.120.30 255.255.255.255
end
```

- 2 Enter the following command to add security policies to allow Phone A to send SIP request messages to Phone B and Phone B to send SIP request messages to Phone A.

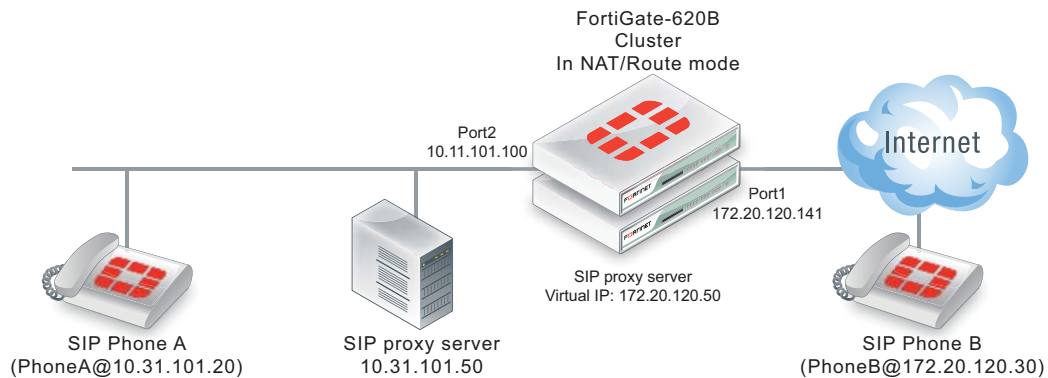
```
config firewall policy
 edit 0
 set srcintf internal
 set dstintf wan1
 set srcaddr Phone_A
 set dstaddr Phone_B
 set action accept
 set schedule always
 set service SIP
 set nat enable
 set utm-status enable
 set voip-profile default
 next
 edit 0
 set srcintf wan1
 set dstintf internal
 set srcaddr Phone_B
 set dstaddr Phone_A
 set action accept
 set schedule always
 set service SIP
 set nat enable
 set utm-status enable
 set voip-profile default
end
```

## SIP NAT configuration example: destination address translation (destination NAT)

This configuration example shows how to configure the FortiGate unit to support the destination address translation scenario shown in Figure 364. The FortiGate unit requires two SIP security policies:

- A destination NAT security policy that allows SIP messages to be sent from the Internet to the private network. This policy must include destination NAT because the addresses on the private network are not routable on the Internet.
- A source NAT security policy that allows SIP messages to be sent from the private network to the Internet.

**Figure 364:**SIP destination NAT scenario part two: 200 OK returned to Phone B and media streams established



### General configuration steps

The following general configuration steps are required for this destination NAT SIP configuration. This example uses the default VoIP profile.

- 1 Add the SIP proxy server firewall virtual IP.
- 2 Add a firewall address for the SIP proxy server on the private network.
- 3 Add a destination NAT security policy that accepts SIP sessions from the Internet destined for the SIP proxy server virtual IP and translates the destination address to the IP address of the SIP proxy server on the private network.
- 4 Add a security policy that accepts SIP sessions initiated by the SIP proxy server and destined for the Internet.

### Configuration steps - web-based manager

#### To add the SIP proxy server firewall virtual IP

1. Go to *Firewall Objects > Virtual IP > Virtual IP* and select *Create New*.
2. Add the SIP proxy server virtual IP.

<b>Name</b>	SIP_Proxy_VIP
<b>External Interface</b>	port1
<b>Type</b>	Static NAT
<b>External IP Address/Range</b>	172.20.120.50
<b>Mapped IP Address/Range</b>	10.31.101.50

### To add a firewall address for the SIP proxy server

1. Go to *Firewall Objects > Address > Addresses*.
2. Add the following for the SIP proxy server:

<b>Address Name</b>	SIP_Proxy_Server
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	10.31.101.50/255.255.255.255
<b>Interface</b>	port2

### To add the security policies

1. Go to *Policy > Policy > Policy*.
2. Add a destination NAT security policy that includes the SIP proxy server virtual IP that allows Phone B (and other SIP phones on the Internet) to send SIP request messages to the SIP proxy server.

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	port1
<b>Source Address</b>	all
<b>Outgoing Interface</b>	port2
<b>Destination Address</b>	SIP_Proxy_VIP
<b>Schedule</b>	always
<b>Service</b>	SIP
<b>Action</b>	ACCEPT

3. Select *Enable NAT* and select *Use Destination Interface Address*.
4. Under *UTM Security Profiles*, select *Use Standard UTM Profiles*.
5. Turn on *VoIP* and select the *default* VoIP profile.
6. Select *OK*.
7. Add a source NAT security policy to allow the SIP proxy server to send SIP request messages to Phone B and the Internet:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	port2
<b>Source Address</b>	SIP_Proxy_Server
<b>Outgoing Interface</b>	port1
<b>Destination Address</b>	all

<b>Schedule</b>	always
<b>Service</b>	SIP
<b>Action</b>	ACCEPT

8. Select *Enable NAT* and select *Use Destination Interface Address*.
9. Under *UTM Security Profiles*, select *Use Standard UTM Profiles*.
10. Turn on *VoIP* and select the *default* VoIP profile.
11. Select *OK*.

## Configuration steps - CLI

### To add the SIP proxy server firewall virtual IP and firewall address

1. Enter the following command to add the SIP proxy server firewall virtual IP.

```
config firewall vip
 edit SIP_Proxy_VIP
 set type static-nat
 set extip 172.20.120.50
 set mappedip 10.31.101.50
 set extintf port1
 end
```

2. Enter the following command to add the SIP proxy server firewall address.

```
config firewall address
 edit SIP_Proxy_Server
 set associated interface port2
 set type ipmask
 set subnet 10.31.101.50 255.255.255.255
 end
```

### To add security policies

1. Enter the following command to add a destination NAT security policy that includes the SIP proxy server virtual IP that allows Phone B (and other SIP phones on the Internet) to send SIP request messages to the SIP proxy server.

```
config firewall policy
 edit 0
 set srcintf port1
 set dstintf port2
 set srcaddr all
 set dstaddr SIP_Proxy_VIP
 set action accept
 set schedule always
 set service SIP
 set nat enable
 set utm-status enable
 set voip-profile default
 end
```

2. Enter the following command to add a source NAT security policy to allow the SIP proxy server to send SIP request messages to Phone B and the Internet:

```
config firewall policy
 edit 0
 set srcintf port2
 set dstintf port1
 set srcaddr SIP_Proxy_Server
 set dstaddr all
 set action accept
 set schedule always
 set service SIP
 set nat enable
 set utm-status enable
 set voip-profile default
 end
```

## Additional SIP NAT scenarios

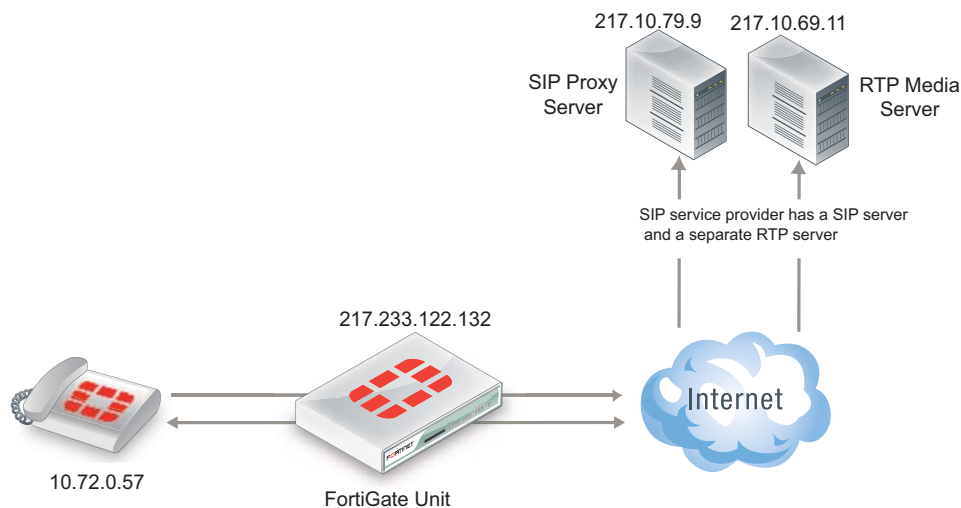
This section lists some additional SIP NAT scenarios.

### Source NAT (SIP and RTP)

In the source NAT scenario shown in [Figure 365](#), a SIP phone connects to the Internet through a FortiGate unit with an IP address configured using PPPoE. The SIP ALG translates all private IPs in the SIP contact header into public IPs.

You need to configure an internal to external SIP security policy with NAT selected, and include a VoIP profile with SIP enabled.


**Figure 365:**SIP source NAT



### Destination NAT (SIP and RTP)

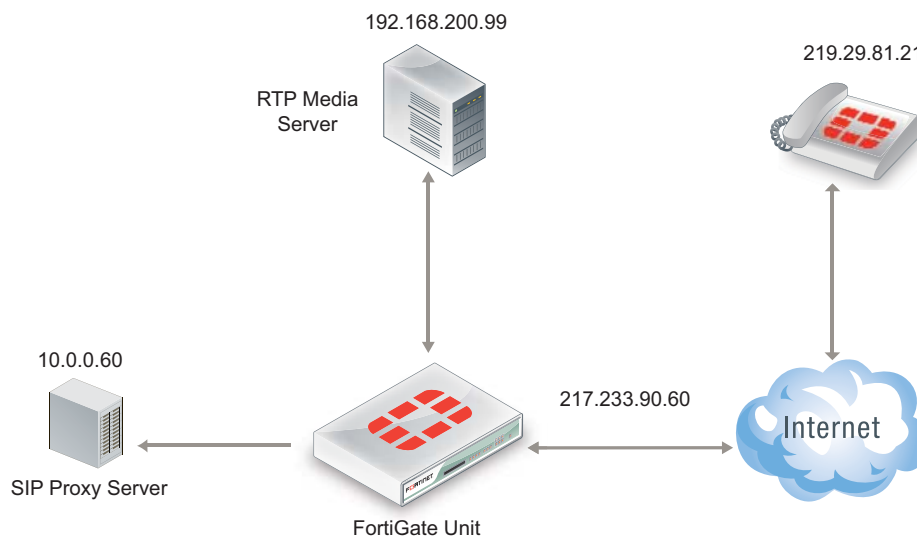
In the following destination NAT scenario, a SIP phone can connect through the FortiGate unit to private IP address using a firewall virtual IP (VIP). The SIP ALG translates the SIP contact header to the IP of the real SIP proxy server located on the Internet.

**Figure 366:**SIP destination NAT



In the scenario, shown in [Figure 366](#), the SIP phone connects to a VIP (10.72.0.60). The SIP ALG translates the SIP contact header to 217.10.79.9, opens RTP pinholes, and manages NAT. The FortiGate unit also supports a variation of this scenario where the RTP media server's IP address is hidden on a private network or DMZ.

**Figure 367:**SIP destination NAT-RTP media server hidden



In the scenario shown in [Figure 367](#), a SIP phone connects to the Internet. The VoIP service provider only publishes a single public IP. The FortiGate unit is configured with a firewall VIP. The SIP phone connects to the FortiGate unit (217.233.90.60) and using the VIP the FortiGate unit translates the SIP contact header to the SIP proxy server IP address (10.0.0.60). The SIP proxy server changes the SIP/SDP connection information (which tells the SIP phone which RTP media server IP it should contact) also to 217.233.90.60.

### Source NAT with an IP pool

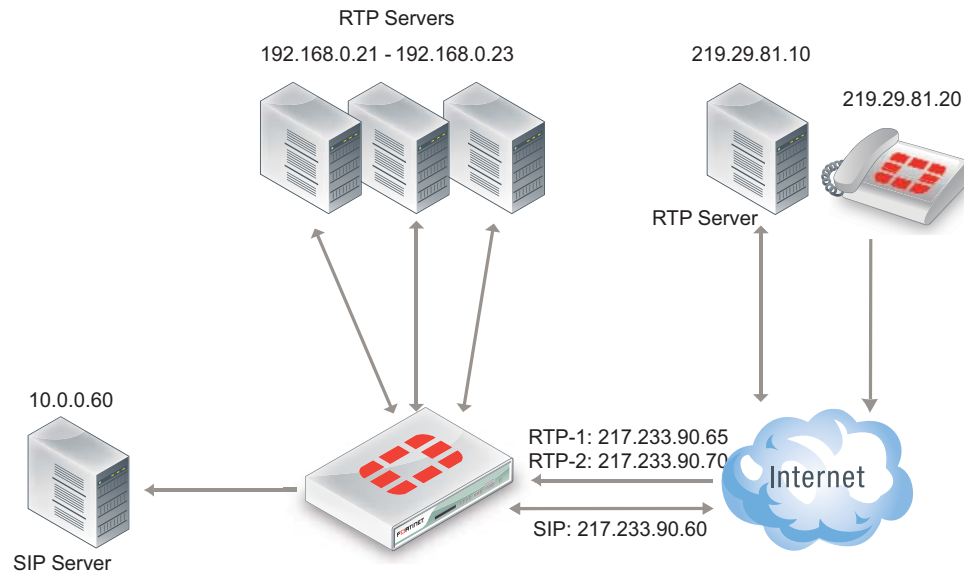
You can choose *NAT* with the *Dynamic IP Pool* option when configuring a security policy if the source IP of the SIP packets is different from the interface IP. The FortiGate ALG interprets this configuration and translates the SIP header accordingly.

This configuration also applies to destination NAT.

## Different source and destination NAT for SIP and RTP

This is a more complex scenario that a SIP service provider may use. It can also be deployed in large-scale SIP environments where RTP has to be processed by the FortiGate unit and the RTP server IP has to be translated differently than the SIP server IP.

**Figure 368:** Different source and destination NAT for SIP and RTP



In this scenario, shown in [Figure 368](#), assume there is a SIP server and a separate media gateway. The SIP server is configured so that the SIP phone (219.29.81.20) will connect to 217.233.90.60. The media gateway (RTP server: 219.29.81.10) will connect to 217.233.90.65.

What happens is as follows:

1. The SIP phone connects to the SIP VIP. The FortiGate ALG translates the SIP contact header to the SIP server: 219.29.81.20 > 217.233.90.60 (> 10.0.0.60).
2. The SIP server carries out RTP to 217.233.90.65.
3. The FortiGate ALG opens pinholes, assuming that it knows the ports to be opened.
4. RTP is sent to the RTP-VIP (217.233.90.65.) The FortiGate ALG translates the SIP contact header to 192.168.0.21.

## NAT with IP address conservation

In a source or destination NAT security policy that accepts SIP sessions, you can configure the SIP ALG or the SIP session helper to preserve the original source IP address of the SIP message in the `i=` line of the SDP profile. NAT with IP address conservation (also called SIP NAT tracing) changes the contents of SIP messages by adding the source IP address of the originator of the message into the SDP `i=` line of the SIP message. The SDP `i=` line is used for free-form text. However, if your SIP server can retrieve information from the SDP `i=` line, it can be useful for keeping a record of the source IP address of the originator of a SIP message when operating in a NAT environment. You can use this feature for billing purposes by extracting the IP address of the originator of the message.



## Configuring SIP IP address conservation for the SIP ALG

You can use the following command to enable or disable SIP IP address conservation in a VoIP profile for the SIP ALG. SIP IP address conservation is enabled by default in a VoIP profile.

```
config voip profile
 edit VoIP_Pro_1
 config sip
 set nat-trace disable
 end
 end
```

If the SIP message does not include an `i=` line and if the original source IP address of the traffic (before NAT) was 10.31.101.20 then the FortiGate unit would add the following `i=` line.

```
i=(o=IN IP4 10.31.101.20)
```

You can also use the `preserve-override` option to configure the SIP ALG to either add the original `o=` line to the end of the `i=` line or replace the `i=` line in the original message with a new `i=` line in the same form as above for adding a new `i=` line.

By default, `preserve-override` is disabled and the SIP ALG adds the original `o=` line to the end of the original `i=` line. Use the following command to configure the SIP ALG to replace the original `i=` line:

```
config voip profile
 edit VoIP_Pro_1
 config sip
 set preserve-override enable
 end
 end
```

## Configuring SIP IP address conservation for the SIP session helper

You can use the following command to enable or disable SIP IP address conservation for the SIP session helper. IP address conservation is enabled by default for the SIP session helper.

```
config system settings
 set sip-nat-trace disable
end
```

If the SIP message does not include an `i=` line and if the original source IP address of the traffic (before NAT) was 10.31.101.20 then the FortiGate unit would add the following `i=` line.

```
i=(o=IN IP4 10.31.101.20)
```

## Controlling how the SIP ALG NATs SIP contact header line addresses

You can enable `contact-fixup` so that the SIP ALG performs normal SIP NAT translation to SIP contact headers as SIP messages pass through the FortiGate unit.

Disable `contact-fixup` if you do not want the SIP ALG to perform normal NAT translation of the SIP contact header if a Record-Route header is also available. If `contact-fixup` is disabled, the FortiGate ALG does the following with contact headers:

- For Contact in Requests, if a Record-Route header is present and the request comes from the external network, the SIP Contact header is not translated.
- For Contact in Responses, if a Record-Route header is present and the response comes from the external network, the SIP Contact header is not translated.

If `contact-fixup` is disabled, the SIP ALG must be able to identify the external network. To identify the external network, you must use the `config system interface` command to set the `external` keyword to `enable` for the interface that is connected to the external network.

Enter the following command to perform normal NAT translation of the SIP contact header:

```
config voip profile
 edit VoIP_Pro_1
 config sip
 set contact-fixup enable
 end
 end
```

## Controlling NAT for addresses in SDP lines

You can use the `no-sdp-fixup` option to control whether the FortiGate unit performs NAT on addresses in SDP lines in the SIP message body.

The `no-sdp-fixup` option is disabled by default and the FortiGate unit performs NAT on addresses in SDP lines. Enable this option if you don't want the FortiGate unit to perform NAT on the addresses in SDP lines.

```
config voip profile
 edit VoIP_Pro_1
 config sip
 set no-sdp-fixup enable
 end
 end
```

## Translating SIP session destination ports

Using port forwarding virtual IPs you can change the destination port of SIP sessions as they pass through the FortiGate unit.


This section describes:

- [Translating SIP sessions to a different destination port](#)
- [Translating SIP sessions to multiple destination ports](#)

### Translating SIP sessions to a different destination port

To configure translating SIP sessions to a different destination port you must add a static NAT virtual IP that translates the SIP destination port to another port destination. In the example the destination port is translated from 5060 to 50601. This configuration can be used if SIP sessions use different destination ports on different networks.

**Figure 369:**Example translating SIP sessions to a different destination port



### To translate SIP sessions to a different destination port

#### 1. Add the static NAT virtual IP.

This virtual IP forwards traffic received at the port1 interface for IP address 172.20.120.20 and destination port 5060 to the SIP server at IP address 192.168.10.20 with destination port 5061.

```
config firewall vip
 edit "sip_port_trans_vip"
 set type static-nat
 set portforward enable
 set protocol tcp
 set extip 172.20.120.20
 set extport 5060
 set extintf "port1"
 set mappedip 192.168.10.20
 set mappedport 50601
 set comment "Translate SIP destination port"
 end
```


#### 2 Add a security policy that includes the virtual IP and the default VoIP profile.

```
config firewall policy
 edit 1
 set srcintf "port1"
 set dstintf "port2"
 set srcaddr "all"
 set dstaddr "sip_port_trans_vip"
 set action accept
 set schedule "always"
 set service "ANY"
 set utm-status enable
 set profile-protocol-options default
 set comments "Translate SIP destination port"
 end
```

## Translating SIP sessions to multiple destination ports

You can use a load balance virtual IP to translate SIP session destination ports to a range of destination ports. In this example the destination port is translated from 5060 to the range 50601 to 50603. This configuration can be used if your SIP server is configured to receive SIP traffic on multiple ports.

**Figure 370:**Example translating SIP traffic to multiple destination ports



### To translated SIP sessions to multiple destination ports

#### 1. Add the load balance virtual IP.

This virtual IP forwards traffic received at the port1 interface for IP address 172.20.120.20 and destination port 5060 to the SIP server at IP address 192.168.10.20 with destination port 5061.

```
config firewall vip
 edit "sip_port_ldbl_vip"
 set type load-balance
 set portforward enable
 set protocol tcp
 set extip 172.20.120.20
 set extport 5060
 set extintf "port1"
 set mappedip 192.168.10.20
 set mappedport 50601-50603
 set comment "Translate SIP destination port range"
 end
```

- 2 Add a security policy that includes the virtual IP and VoIP profile.

```
config firewall policy
 edit 1
 set srcintf "port1"
 set dstintf "port2"
 set srcaddr "all"
 set dstaddr "sip_port_ldbl_vip"
 set action accept
 set schedule "always"
 set service "ANY"
 set utm-status enable
 set voip-profile default
 set comments "Translate SIP destination port"
 end
```

## Adding the original IP address and port to the SIP message header after NAT

In some cases your SIP configuration may require that the original IP address and port from the SIP contact request is kept after NAT. For example, the original SIP contact request could include the following:

```
Contact: <sip:0150302438@172.20.120.110:5060>;
```

After the packet goes through the FortiGate unit and NAT is performed, the contact request could normally look like the following (the IP address translated to a different IP address and the port to a different port):

```
Contact: <sip:0150302438@10.10.10.21:33608>;
```

You can enable `register-contact-trace` in a VoIP profile to have the SIP ALG add the original IP address and port in the following format:

```
Contact: <sip:0150302438@<nated-ip>:<nated-port>;o=<original-ip>:
<original-port>>;
```

So the contact line after NAT could look like the following:

```
Contact: <sip:0150302438@10.10.10.21:33608;o=172.20.120.110:5060>;
```

Enter the following command to enable keeping the original IP address and port:

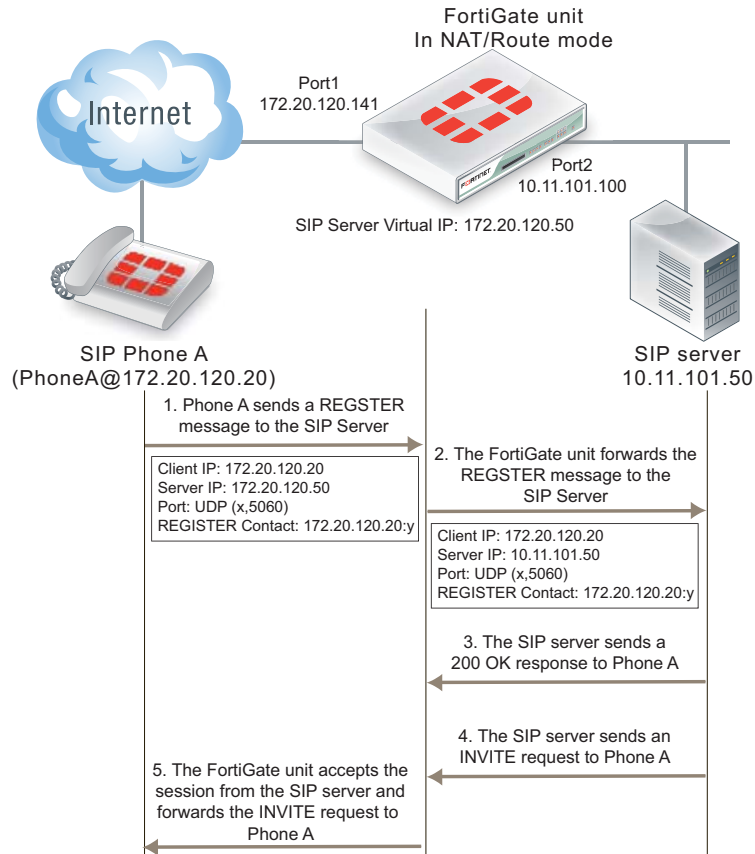
```
config voip profile
 edit Profile_name
 config sip
 set register-contract-trace enable
 end
```

## Enhancing SIP pinhole security

You can use the `strict-register` option in a SIP VoIP profile to open smaller pinholes.

As shown in [Figure 371](#) when FortiGate unit is protecting a SIP server on a private network, the FortiGate unit does not have to open a pinhole for the SIP server to send INVITE requests to a SIP Phone on the Internet after the SIP Phone has registered with the server.

**Figure 371:**FortiGate unit protecting a SIP server on a private network



In the example, a client (SIP Phone A) sends a REGISTER request to the SIP server with the following information:

```
Client IP: 10.31.101.20
Server IP: 10.21.101.50
Port: UDP (x,5060)
REGISTER Contact: 10.31.101.20:y
```

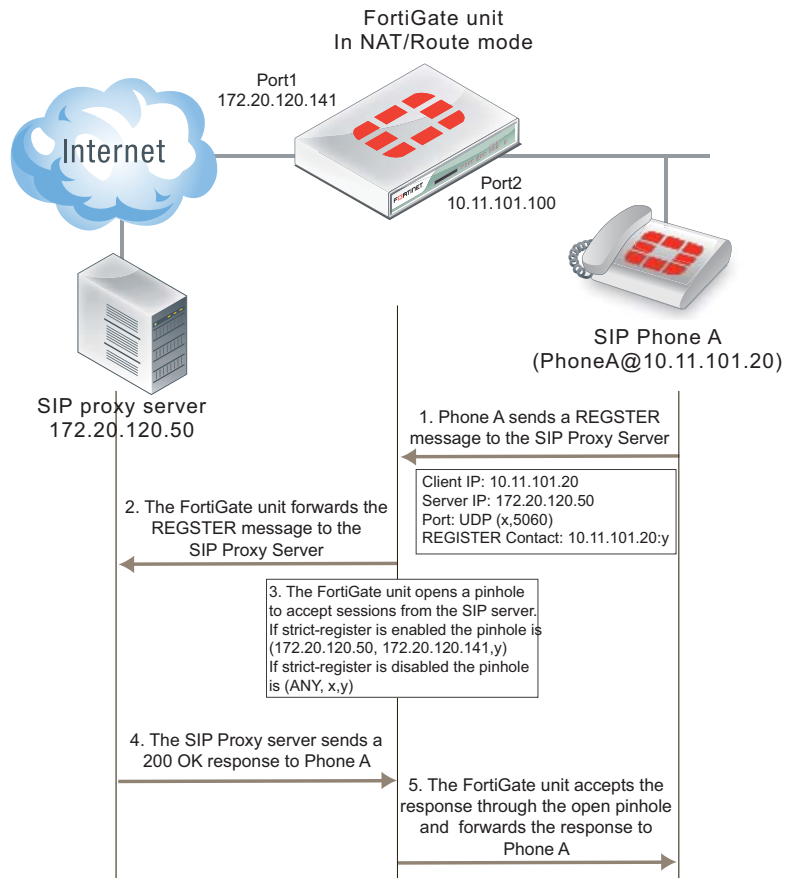
Where  $x$  and  $y$  are ports chosen by Phone A.

As soon as the server sends the 200 OK reply it can forward INVITE requests from other SIP phones to SIP Phone A. If the SIP proxy server uses the information in the REGISTER message received from SIP Phone A the INVITE messages sent to Phone A will only get through the FortiGate unit if a policy has been added to allow the server to send traffic from the private network to the Internet. Or the SIP ALG must open a pinhole to allow traffic from the server to the Internet. In most cases the FortiGate unit is protecting the SIP server so there is no reason not to add a security policy to all the SIP server to send outbound traffic to the Internet.

In a typical SOHO scenario shown in [Figure 372](#), SIP Phone A is being protected from the Internet by a FortiGate unit. In most cases the FortiGate unit would not allow incoming traffic from the Internet to reach the private network. So the only way that an INVITE request from the SIP server can reach SIP Phone A is if the SIP ALG creates an incoming pinhole. All pinholes have three attributes:

(source address, destination address, destination port)

**Figure 372:**SOHO configuration, FortiGate unit protecting a network with SIP phones



The more specific a pinhole is the more secure it is because it will accept less traffic. In this situation, the pinhole would be more secure if it only accepted traffic from the SIP server. This is what happens if `strict-register` is enabled in the VoIP profile that accepts the REGISTER request from Phone A.

(SIP server IP address, client IP address, destination port)

If `strict-register` is disabled (the default configuration) the pinhole is set up with the following attributes

(ANY IP address, client IP address, destination port)

This pinhole allows connections through the FortiGate unit from ANY source address which is a much bigger and less secure pinhole. In most similar network configurations you should enable `strict-register` to improve pinhole security.

Enabling `strict-register` can cause problems when the SIP registrar and SIP proxy server are separate entities with separate IP addresses.

Enter the following command to enable `strict-register` in a VoIP profile.

```
config voip profile
 edit Profile_name
 config sip
 set strict-register enable
 end
```

## Hosted NAT traversal

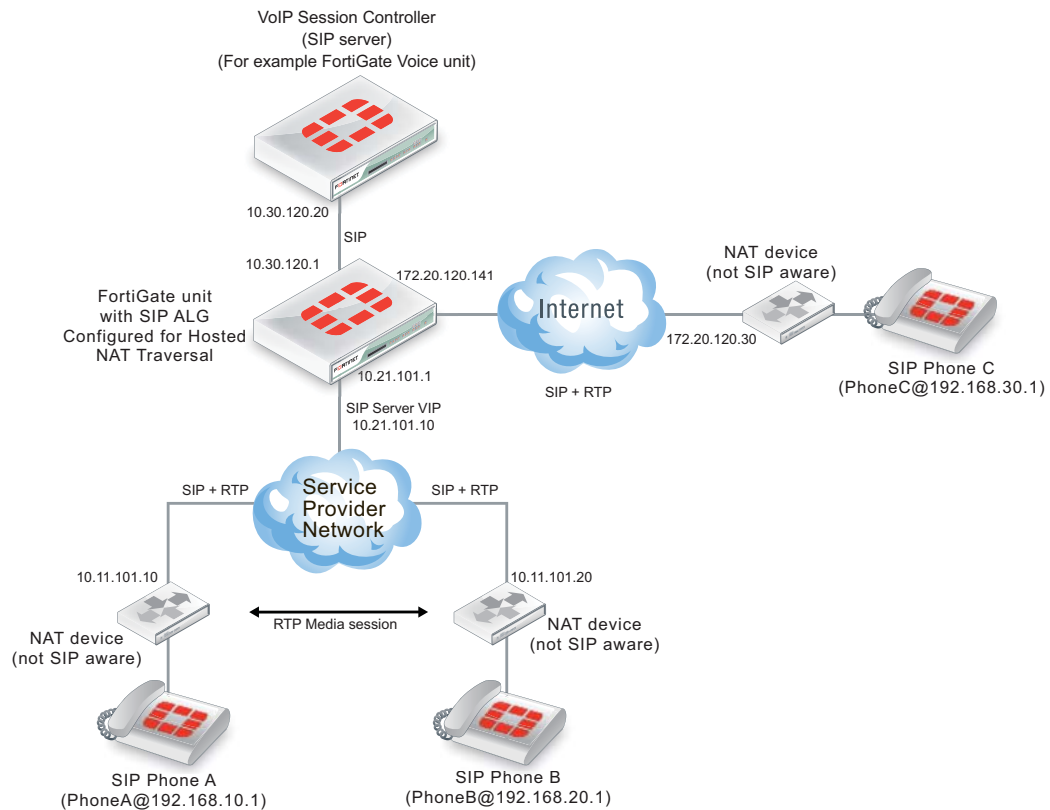
With the increase in the use of VoIP and other media traffic over the Internet, service provider network administrators must defend their networks from threats while allowing voice and multimedia traffic to flow transparently between users and servers and among users. A common scenario could involve providing SIP VoIP services for customers with SIP phones installed behind NAT devices that are not SIP aware. NAT devices that are not SIP aware cannot translate IP addresses in SIP headers and SDP lines in SIP packets but can and do perform source NAT on the source or addresses of the packets. In this scenario the user's SIP phones would communicate with a SIP proxy server to set up calls between SIP phones. Once the calls are set up RTP packets would be communicated directly between the phones through each user's NAT device.

The problem with this configuration is that the SIP headers and SDP lines in the SIP packets sent from the phones and received by the SIP proxy server would contain the private network addresses of the VoIP phones that would not be routable on the service provider network or on the Internet. One solution could be to for each customer to install and configure SIP aware NAT devices. If this is not possible, another solution requires implement hosted NAT traversal.

In a hosted NAT traversal (HNT) configuration (for example, see [Figure 373](#)), a FortiGate unit is installed between the NAT device and the SIP proxy server and configured with a VoIP profile that enables SIP hosted NAT traversal. Security policies that include the VoIP profile also support destination NAT using a firewall virtual IP. When the SIP phones connect to the SIP server IP address the security policy accepts the SIP packets, the virtual IP translates the destination addresses of the packets to the SIP server IP address, and the SIP ALG NAT traversal configuration translates the source IP addresses on the SIP headers and SDP lines to the source address of the SIP packets (which would be the external IP address of the NAT devices). The SIP server then sees the SIP phone IP address as the external IP address of the NAT device. As a result SIP and RTP media sessions are established using the external IP addresses of the NAT devices instead of the actual IP addresses of the SIP phones.

**Figure 373:**FortiGate SIP Hosted NAT Traversal configuration





## Configuration example: Hosted NAT traversal for calls between SIP Phone A and SIP Phone B

The following address translation takes place to allow a SIP call from SIP Phone A to SIP Phone B in [Figure 373](#).

1. SIP Phone A sends a SIP Invite message to the SIP server. Packet source IP address: 192.168.10.1, destination IP address: 10.21.101.10.
2. The SIP packets are received by the NAT device which translates the source address of the SIP packets from 192.168.10.1 to 10.11.101.20.
3. The SIP packets are received by the FortiGate unit which translates the packet destination IP address to 10.30.120.20. The SIP ALG also translates the IP address of the SIP phone in the SIP header and SDP lines from 192.168.10.1 to 10.11.101.20.
4. The SIP server accepts the Invite message and forwards it to SIP Phone B at IP address 10.11.101.20. The SIP server has this address for SIP Phone B because SIP packets from SIP Phone B have also been translated using the hosted NAT traversal configuration of the SIP ALG.
5. When the SIP call is established, the RTP session is between 10.11.101.10 and 10.11.101.20 and does not pass through the FortiGate unit. The NAT devices translated the destination address of the RTP packets to the private IP addresses of the SIP phones.

### General configuration steps

The following general configuration steps are required for this destination NAT SIP configuration. This example uses the default VoIP profile.

1. Add a VoIP profile that enables hosted NAT translation.
2. Add a SIP proxy server firewall virtual IP.

3. Add a firewall address for the SIP proxy server on the private network.
4. Add a destination NAT security policy that accepts SIP sessions from the Internet destined for the SIP proxy server virtual IP and translates the destination address to the IP address of the SIP proxy server on the private network.
5. Add a security policy that accepts SIP sessions initiated by the SIP proxy server and destined for the Internet.

### Configuration steps - web-based manager

#### To add the SIP proxy server firewall virtual IP

1. Go to *Firewall Objects > Virtual IP > Virtual IP*.
2. Add the SIP proxy server virtual IP.

<b>Name</b>	SIP_Proxy_VIP
<b>External Interface</b>	port1
<b>Type</b>	Static NAT
<b>External IP Address/Range</b>	172.20.120.50
<b>Mapped IP Address/Range</b>	10.31.101.50

#### To add a firewall address for the SIP proxy server

1. Go to *Firewall Objects > Address > Addresses*.
2. Add the following for the SIP proxy server:

<b>Address Name</b>	SIP_Proxy_Server
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	10.31.101.50/255.255.255.255
<b>Interface</b>	port2

#### To add the security policies

1. Go to *Policy > Policy > Policy*.
2. Add a destination NAT security policy that includes the SIP proxy server virtual IP that allows Phone B (and other SIP phones on the Internet) to send SIP request messages to the SIP proxy server.

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	port1
<b>Source Address</b>	all
<b>Outgoing Interface</b>	port2
<b>Destination Address</b>	SIP_Proxy_VIP
<b>Schedule</b>	always

<b>Service</b>	SIP
<b>Action</b>	ACCEPT

3. Select *Enable NAT* and select *Use Destination Interface Address*.
4. Under *UTM Security Profiles*, select *Use Standard UTM Profiles*.
5. Turn on *VoIP* and select the *default* VoIP profile.
6. Select *OK*.
7. Add a source NAT security policy to allow the SIP proxy server to send SIP request messages to Phone B and the Internet:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	port2
<b>Source Address</b>	SIP_Proxy_Server
<b>Outgoing Interface</b>	port1
<b>Destination Address</b>	all
<b>Schedule</b>	always
<b>Service</b>	SIP
<b>Action</b>	ACCEPT

8. Select *Enable NAT* and select *Use Destination Interface Address*.
9. Under *UTM Security Profiles*, select *Use Standard UTM Profiles*.
10. Turn on *VoIP* and select the *default* VoIP profile.
11. Select *OK*.

## Configuration steps - CLI

### To add a VoIP profile that enables hosted NAT translation

- 1 Enter the following command to add a VoIP profile named HNT that enables hosted NAT traversal. This command shows how to clone the default VoIP profile and enable hosted NAT traversal.

```
config voip profile
 clone default to HNT
 edit HNT
 config sip
 set hosted-nat-traversal enable
 end
 end
end
```

### To add the SIP proxy server firewall virtual IP and firewall address

1. Enter the following command to add the SIP proxy server firewall virtual IP.

```
config firewall vip
 edit SIP_Proxy_VIP
 set type static-nat
 set extip 10.21.101.10
 set mappedip 10.30.120.20
 set extintf port1
 end
```

2. Enter the following command to add the SIP proxy server firewall address.

```
config firewall address
 edit SIP_Proxy_Server
 set associated interface port2
 set type ipmask
 set subnet 10.30.120.20 255.255.255.255
 end
```

### To add security policies

1. Enter the following command to add a destination NAT security policy that includes the SIP proxy server virtual IP that allows Phone A to send SIP request messages to the SIP proxy server.

```
config firewall policy
 edit 0
 set srcintf port1
 set dstintf port2
 set srcaddr all
 set dstaddr SIP_Proxy_VIP
 set action accept
 set schedule always
 set service SIP
 set nat enable
 set utm-status enable
 set profile-protocol-options default
 set voip-profile HNT
 end
```

2. Enter the following command to add a source NAT security policy to allow the SIP proxy server to send SIP request messages to Phone B:

```
config firewall policy
 edit 0
 set srcintf port2
 set dstintf port1
 set srcaddr SIP_Proxy_Server
 set dstaddr all
 set action accept
 set schedule always
 set service SIP
 set nat enable
 set utm-status enable
 set profile-protocol-options default
 set voip-profile default
 end
```

## Hosted NAT traversal for calls between SIP Phone A and SIP Phone C

The following address translation takes place to allow a SIP call from SIP Phone A to SIP Phone C in [Figure 373 on page 2552](#).

1. SIP Phone A sends a SIP Invite message to the SIP server. Packet source IP address: 192.168.10.1 and destination IP address: 10.21.101.10.
2. The SIP packets are received by the NAT device which translates the source address of the SIP packets from 192.168.10.1 to 10.11.101.20.
3. The SIP packets are received by the FortiGate unit which translates the packet destination IP address to 10.30.120.20. The SIP ALG also translates the IP address of the SIP phone in the SIP header and SDP lines from 192.168.10.1 to 10.11.101.20.
4. The SIP server accepts the Invite message and forwards it to SIP Phone C at IP address 172.20.120.30. The SIP server has this address for SIP Phone C because SIP packets from SIP Phone C have also been translated using the hosted NAT traversal configuration of the SIP ALG.
5. When the SIP call is established, the RTP session is between 10.11.101.10 and 172.20.120.30. The packets pass through the FortiGate unit which performs NAT as required.

## Restricting the RTP source IP

Use the following command in a VoIP profile to restrict the RTP source IP to be the same as the SIP source IP when hosted NAT traversal is enabled.

```
config voip profile
 edit VoIP_HNT
 config sip
 set hosted-nat-traversal enable
 set hnt-restrict-source-ip enable
 end
 end
```


## SIP over IPv6

FortiGate units operating in NAT/Route and in Transparent mode support SIP over IPv6. The SIP ALG can process SIP messages that use IPv6 addresses in the headers, bodies, and in the transport stack. The SIP ALG cannot modify the IPv6 addresses in the SIP headers so FortiGate units cannot perform SIP or RTP NAT over IPv6 and also cannot translate between IPv6 and IPv4 addresses.

In the scenario shown in [Figure 374](#), a SIP phone connects to the Internet through a FortiGate unit operating. The phone and the SIP and RTP servers all have IPv6 addresses.

The FortiGate unit has IPv6 security policies that accept SIP sessions. The SIP ALG understands IPv6 addresses and can forward IPv6 sessions to their destinations. Using SIP application control features the SIP ALG can also apply rate limiting and other settings to SIP sessions.

**Figure 374:**SIP support for IPv6

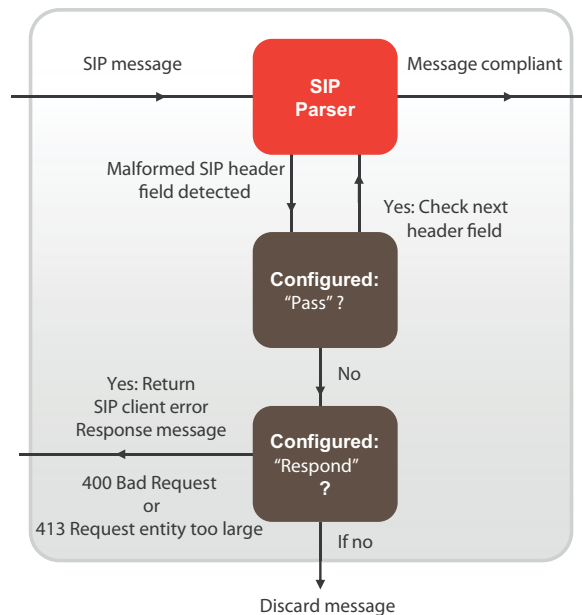


To enable SIP support for IPv6 add an IPv6 security policy that accepts SIP packets and includes a VoIP profile.

## Deep SIP message inspection

Deep SIP message syntax inspection (also called Deep SIP header inspection or SIP fuzzing protection) provides protection against malicious SIP messages by applying SIP header and SDP profile syntax checking. SIP Fuzzing attacks can be used by attackers to discover and exploit vulnerabilities of a SIP entity (for example a SIP proxy server). Most often these attacks could crash or compromise the SIP entity.

**Figure 375:**Deep SIP message inspection



- Checks the SIP request message Request-line
- Checks the following SIP header fields:
  - Allow, Call-id, Contact, Content-length, Content-type, CSeq, Expires, From, Max-Forwards, P-asserted-identity, Rack, Record-Route, Route, Rseq, To, Via
- Checks all SDP profile lines
- Configurable header and body length checks
- Optional logging of message violations

Deep SIP message inspection checks the syntax of each SIP header and SDP profile line to make sure they conform to the syntax defined in the relevant RFC and IETF standard. You can also configure the SIP ALG to inspect for:

- Unknown SIP message types (message types not defined in a SIP RFC) this option is enabled by default and can be disabled. When enabled unknown message types are discarded. Configured using the `block-unknown` option.
- Unknown line types (message line types that are not defined in any SIP or SDP RFC). Configured using the `unknown-header` option.
- Messages that are longer than a configured maximum size. Configured using the `max-body-length` option.
- Messages that contain one or more lines that are longer than a set maximum line length (default 998 characters). Configured using the `max-line-length` option.

## Actions taken when a malformed message line is found

When a malformed message line or other error is found the SIP ALG can be configured to discard the message containing the error, pass the message without any other actions, or responding to the message with a 400 Bad Request or 413 Request entity too large client error SIP response message and then discard the message. (For information about client error SIP response messages, see [“Client error” on page 2499](#).)

If a message line is longer than the configured maximum, the SIP ALG sends the following message:

```
SIP/2.0 413 Request Entity Too Large, <optional_info>
```

If a message line is incorrect or in an unknown message line is found, the SIP ALG sends the following message:

```
SIP/2.0 400 Bad Request, <optional_info>
```

The `<optional_info>` provides more information about why the message was rejected. For example, if the SIP ALG finds a malformed Via header line, the response message may be:

```
SIP/2.0 400 Bad Request, malformed Via header
```

If the SIP ALG finds a malformed message line, and the action for this message line type is discard, the message is discarded with no further checking or responses. If the action is pass, the SIP ALG continues parsing the SIP message for more malformed message lines. If the action is respond, the SIP ALG sends the SIP response message and discards the message containing the malformed line with no further checking or response. If only malformed message line types with action set to pass are found, the SIP ALG extracts as much information as possible from the message (for example for NAT and opening pinholes, and forwards the message to its destination).

If a SIP message containing a malformed line is discarded the SIP ALG will not use the information in the message for call processing. This could result in the call being terminated. If a malformed line in a SIP message includes information required for the SIP call that the SIP ALG cannot interpret (for example, if an IP address required for SIP NAT is corrupted) the SIP ALG may not be able to continue processing the call and it could be terminated. Discarded messages are counted by SIP ALG static message counters.

## Logging and statistics

To record a log message each time the SIP ALG finds a malformed header, enable logging SIP violations in a VoIP profile. In all cases, when the SIP ALG finds an error the FortiGate unit records a malformed header log message that contains information about the error. This happens even if the action is set to pass.

If, because of recording log messages for deep message inspection, the CPU performance is affected by a certain amount, the FortiGate unit records a critical log message about this event and stops writing log messages for deep SIP message inspection.

The following information is recorded in malformed header messages:

- The type of message line in which the error was found.
- The content of the message line in which the error was found (it will be truncated if it makes the log message too long)
- The column or character number in which the error was found (to make it easier to determine what caused the error)

## Deep SIP message inspection best practices

Because of the risks imposed by SIP header attacks or incorrect data being allowed and because selecting drop or respond does not require more CPU overhead than pass you would want to set all tests to drop or respond. However, in some cases malformed lines may be less of a threat or risk. For example, the SDP `i=` does not usually contain information that is parsed by any SIP device so a malformed `i=` line may not pose a threat.

You can also use the pre-defined VoIP profiles to apply different levels of deep message inspection. The default VoIP profile sets all deep message inspection options to pass and the strict VoIP profile sets all deep message inspection options to discard. From the CLI you can use the `clone` command to copy these pre-defined VoIP profiles and then customize them for your requirements.

## Configuring deep SIP message inspection

You configure deep SIP message inspection in a VoIP profile. All deep SIP message inspection options are available only from the CLI.



Enter the following command to configure deep SIP message inspection to discard messages with malformed Request-lines (the first line in a SIP request message):

```
config voip profile
 edit VoIP_Pro_Name
 config sip
 set malformed-request-line respond
 end
 end
```



You cannot configure message inspection for the Status-line, which is the first line in a SIP response message.

[Table 120](#) lists the SIP header lines that the SIP ALG can inspect and the CLI command for configuring the action for each line type. The table also lists the RFC that the header line is defined in.

**Table 120:**SIP header lines that the SIP ALG can inspect for syntax errors

SIP Header line	VoIP profile option	RFC
<b>Allow</b>	malformed-header-allow	RFC 3261
<b>Call-ID</b>	malformed-header-call-id	RFC 3261
<b>Contact</b>	malformed-header-contact	RFC 3261
<b>Content-Length</b>	malformed-header-content-length	RFC 3261
<b>Content-Type</b>	malformed-header-content-type	RFC 3261
<b>CSeq</b>	malformed-header-cseq	RFC 3261
<b>Expires</b>	malformed-header-expires	RFC 3261
<b>From</b>	malformed-header-from	RFC 3261
<b>Max-forwards</b>	malformed-header-max-forwards	RFC 3261
<b>P-Asserted-Identity</b>	malformed-header-p-asserted-identity	RFC 3325
<b>RAck</b>	malformed-header-rack	RFC 3262
<b>Record-Route</b>	malformed-header-record-route	RFC 3261
<b>Route</b>	malformed-header-route	RFC 3261
<b>RSeq</b>	malformed-header-rseq	RFC 3262
<b>To</b>	malformed-header-to	RFC 3261
<b>Via</b>	malformed-header-via	RFC 3261

Table 121 lists the SDP profile lines that the SIP ALG inspects and the CLI command for configuring the action for each line type. SDP profile lines are defined by RFC 4566 and RFC 2327.

**Table 121:**SDP profile lines that the SIP ALG can inspect for syntax errors

Attribute	VoIP profile option
<b>a=</b>	malformed-header-sdb-a
<b>b=</b>	malformed-header-sdp-b
<b>c=</b>	malformed-header-sdp-c
<b>i=</b>	malformed-header-sdp-i
<b>k=</b>	malformed-header-sdp-k
<b>m=</b>	malformed-header-sdp-m
<b>o=</b>	malformed-header-sdp-o
<b>r=</b>	malformed-header-sdp-r
<b>s=</b>	malformed-header-sdp-s
<b>t=</b>	malformed-header-sdp-t
<b>v=</b>	malformed-header-sdp-v
<b>z=</b>	malformed-header-sdp-z

### Discarding SIP messages with some malformed header and body lines

Enter the following command to configure deep SIP message inspection to discard SIP messages with a malformed Via line, a malformed route line or a malformed m= line but to pass messages with a malformed i= line or a malformed Max-Forwards line

```

config voip profile
 edit VoIP_Pro_Name
 config sip
 set malformed-header-via discard
 set malformed-header-route discard
 set malformed-header-sdp-m discard
 set malformed-header-sdp-i pass
 set malformed-header-max-forwards pass
 end
 end
end

```

## Discarding SIP messages with an unknown SIP message type

Enter the following command to discard SIP messages with an unknown SIP message line type as defined in all current SIP RFCs:

```
config voip profile
 edit VoIP_Pro_Name
 config sip
 set unknown-header discard
 end
 end
end
```

## Discarding SIP messages that exceed a message size

Enter the following command to set the maximum size of a SIP message to 200 bytes. Messages longer than 200 bytes are discarded.

```
config voip profile
 edit VoIP_Pro_Name
 config sip
 set max-body-length 200
 end
 end
end
```

The `max-body-length` option checks the value in the SIP Content-Length header line to determine body length. The Content-Length can be larger than the actual size of a SIP message if the SIP message content is split over more than one packet. SIP message sizes vary widely. The size of a SIP message can also change with the addition of Via and Record-Route headers as the message is transmitted between users and SIP servers.

## Discarding SIP messages with lines longer than 500 characters

Enter the following command to set the length of a SIP message line to 500 characters and to block messages that include lines with 500 or more characters:

```
config voip profile
 edit VoIP_Pro_Name
 config sip
 set max-line-length 500
 set block-long-lines enable
 end
 end
end
```

## Blocking SIP request messages

You may want to block different types of SIP requests:

- to prevent SIP attacks using these messages.
- If your SIP server cannot process some SIP messages because of a temporary issue (for example a bug that crashes or compromises the server when it receives a message of a certain type).
- Your SIP implementation does not use certain message types.

When you enable message blocking for a message type in a VoIP profile, whenever a security policy containing the VoIP profile accepts a SIP message of this type, the SIP ALG silently discards the message and records a log message about the action.

Use the following command to configure a VoIP profile to block SIP CANCEL and Update request messages:

```
config voip profile
 edit VoIP_Pro_Name
 config sip
 set block-cancel enable
 set block-update enable
 end
 end
```

SIP uses a variety of text-based messages or requests to communicate information about SIP clients and servers to the various components of the SIP network. Since SIP requests are simple text messages and since the requests or their replies can contain information about network components on either side of the FortiGate unit, it may be a security risk to allow these messages to pass through.

[Table 122](#) lists all of the VoIP profile SIP request message blocking options. All of these options are disabled by default.



Blocking SIP OPTIONS messages may prevent a redundant configuration from operating correctly. See [“Supporting geographic redundancy when blocking OPTIONS messages” on page 2571](#) for information about resolving this problem.

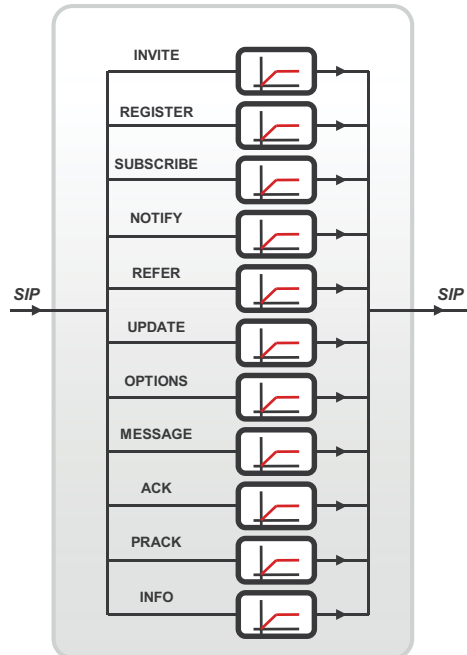
**Table 122:**Options for blocking SIP request messages

SIP request message	SIP message blocking CLI Option
<b>ACK</b>	block-ack
<b>BYE</b>	block-bye
<b>Cancel</b>	block-cancel
<b>INFO</b>	block-info
<b>INVITE</b>	block-invite
<b>Message</b>	block-message
<b>Notify</b>	block-notify
<b>Options</b>	block-options
<b>PRACK</b>	block-prack
<b>Publish</b>	block-publish
<b>Refer</b>	block-refer
<b>Register</b>	block-register
<b>Subscribe</b>	block-subscribe
<b>Update</b>	block-update

## SIP rate limiting

Configurable threshold for SIP message rates per request method. Protects SIP servers from SIP overload and DoS attacks.

**Figure 376:**SIP rate limiting



- **SIP message rate limitation**
- **Individually configurable per SIP method**
- **When threshold is hit additional messages with this method will be discarded**
- **Prevents SIP server from getting overloaded by flash crowds or Denial-of-Service attacks.**
- **May block some methods at all (with extra "block" option)**
- **Can be disabled (unlimited rate)**

FortiGate units support rate limiting for the following types of VoIP traffic:

- Session Initiation Protocol (SIP)
- Skinny Call Control Protocol (SCCP) (most versions)
- Session Initiation Protocol for Instant Messaging and Presence Leveraging Extensions (SIMPLE).

You can use rate limiting of these VoIP protocols to protect the FortiGate unit and your network from SIP and SCCP Denial of Service (DoS) attacks. Rate limiting protects against SIP DoS attacks by limiting the number of SIP REGISTER and INVITE requests that the FortiGate unit receives per second. Rate limiting protects against SCCP DoS attacks by limiting the number of SCCP call setup messages that the FortiGate unit receives per minute.

You configure rate limiting for a message type by specifying a limit for the number of messages that can be received per second. The rate is limited per security policy. When VoIP rate limiting is enabled for a message type, if a single security policy accepts more messages per second than the configured rate, the extra messages are dropped and log messages are written when the messages are dropped.

Use the following command to configure a VoIP profile to limit the number of INVITE messages accepted by each security policy that the VoIP profile is added to 100 INVITE messages a second:

```
config voip profile
 edit VoIP_Pro_Name
 config sip
 set invite-rate 100
 end
 end
end
```

If you are experiencing denial of service attacks from traffic using these VoIP protocols, you can enable VoIP rate limiting and limit the rates for your network. Limit the rates depending on the amount of SIP and SCCP traffic that you expect the FortiGate unit to be handling. You can adjust the settings if some calls are lost or if the amount of SIP or SCCP traffic is affecting FortiGate unit performance.

[Table 123](#) lists all of the VoIP profile SIP rate limiting options. All of these options are set to 0 so are disabled by default.



Blocking SIP OPTIONS messages may prevent a redundant configuration from operating correctly. See [“Supporting geographic redundancy when blocking OPTIONS messages” on page 2571](#) for information about resolving this problem.

**Table 123:**Options for SIP rate limiting

SIP request message	Rate Limiting CLI Option
ACK	ack-rate
BYE	bye-rate
Cancel	cancel-rate
INFO	info-rate
INVITE	invite-rate
Message	message-rate
Notify	notify-rate
Options	options-rate
PRACK	prack-rate
Publish	publish-rate
Refer	refer-rate
Register	register-rate
Subscribe	subscribe-rate
Update	update-rate

### Limiting the number of SIP dialogs accepted by a security policy

In addition to limiting the rates for receiving SIP messages, you can use the following command to limit the number of SIP dialogs (or SIP calls) that the FortiGate unit accepts.

```
config voip profile
 edit VoIP_Pro_Name
 config sip
 set max-dialogs 2000
 end
 end
```

This command sets the maximum number of SIP dialogs that can be open for SIP sessions accepted by any security policy that you add the VoIP profile to. The default setting of 0 does not limit the number of dialogs. You can add a limit to control the number of open dialogs and raise and lower it as required. You might want to limit the number of open dialogs for protection against SIP-based attackers opening large numbers of SIP dialogs. Every dialog takes memory and FortiGate CPU resources to process. Limiting the number of dialogs may improve the overall performance of the FortiGate unit. Limiting the number of dialogs will not drop calls in progress but may prevent new calls from connecting.

## SIP logging and DLP archiving

You can enable SIP logging and logging of SIP violations, and SIP DLP archiving a VoIP profile. To record SIP log messages you must also enable VoIP event logging in the FortiGate unit event logging configuration.

To view SIP log messages go to *Log&Report > Log Access > Event*.

To view SIP DLP archive messages go to *Log&Report > Archive Access > VoIP*.

Use the following command enable SIP logging, SIP archiving, and logging of SIP violations in a VoIP profile:

```
config voip profile
 edit VoIP_Pro_Name
 config sip
 set log-call-summary enable
 set log-violations enable
 end
 end
end
```

## Inspecting SIP over SSL/TLS (secure SIP)

Some SIP phones and SIP servers can communicate using SSL or TLS to encrypt the SIP signalling traffic. To allow SIP over SSL/TLS calls to pass through the FortiGate unit, the encrypted signalling traffic has to be unencrypted and inspected. To do this, the FortiGate SIP ALG intercepts and unencrypts and inspects the SIP packets. The packets are then re-encrypted and forwarded to their destination.

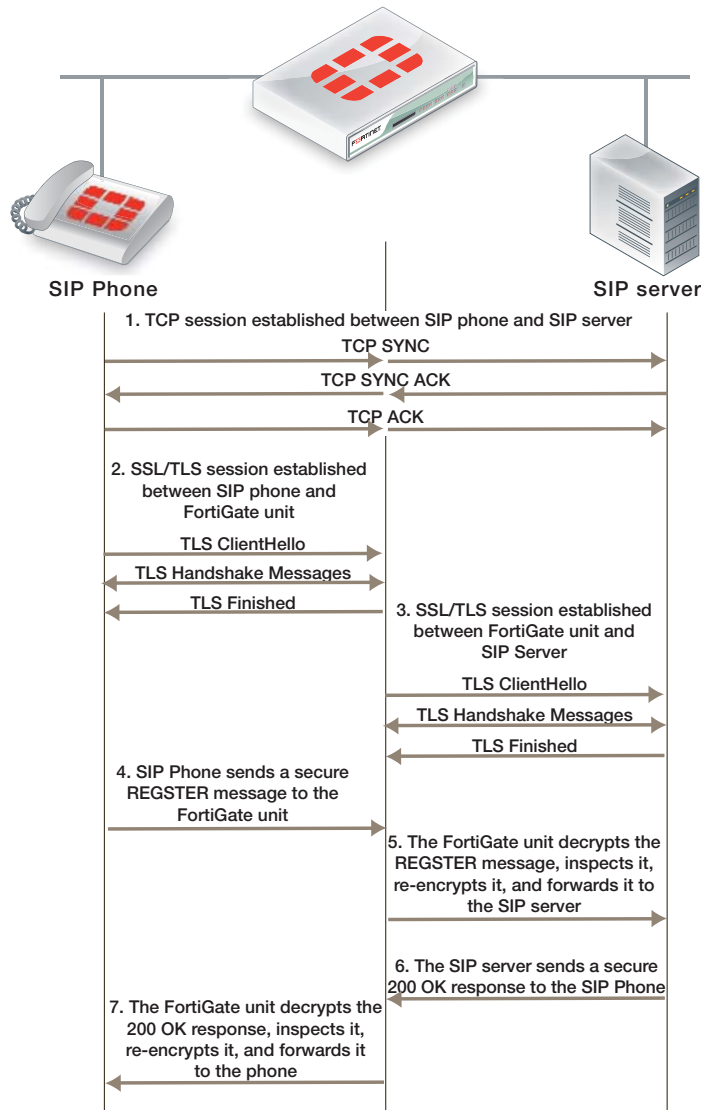
Normally SIP over SSL/TLS uses port 5061. You can use the following command to change the port that the FortiGate listens on for SIP over SSL/TLS sessions to port 5066:

```
config system settings
 set sip-ssl-port 5066
end
```

The SIP ALG supports full mode SSL/TLS only. Traffic between SIP phones and the FortiGate unit and between the FortiGate unit and the SIP server is always encrypted.

You enable SSL/TLS SIP communication by enabling SSL mode in a VoIP profile. You also need to install the SIP server and client certificates on your FortiGate unit and add them to the SSL configuration in the VoIP profile.

**Figure 377:**SIP over SSL/TLS between a SIP phone and a SIP server



Other than enabling SSL mode and making sure the security policies accept the encrypted traffic, the FortiGate configuration for SSL/TLS SIP is the same as any SIP configuration. SIP over SSL/TLS is supported for all supported SIP configurations.

### Adding the SIP server and client certificates

A VoIP profile that supports SSL/TLS SIP requires one certification for the SIP server and one certificate that is used by all of the clients. Use the following steps to add these certificates to the FortiGate unit. Before you start, make sure the client and server certificate files and their key files are accessible from the management computer.

1. Go to *System > Certificates > Local Certificates* and select *Import*.
2. Set *Type* to *Certificate*.
3. Browse to the *Certificate file* and the *Key file* and select *OK*.
4. Enter a password for the certificate and select *OK*.

The certificate and key are uploaded to the FortiGate unit and added to the *Local Certificates* List.

5. Repeat to upload the other certificate.



The certificates are added to the list of Local Certificates as the filenames you uploaded. You can add comments to make it clear where its from and how it is intended to be used.

## Adding SIP over SSL/TLS support to a VoIP profile

Use the following commands to add SIP over SSL/TLS support to the default VoIP profile. The following command enables SSL mode and adds the client and server certificates and passwords, the same ones you entered when you imported the certificates:

```
config voip profile
 edit default
 config sip
 set ssl-mode full
 set ssl-client-certificate "Client_cert"
 set ssl-server-certificate "Server_cert"
 set ssl-auth-client "check-server"
 set ssl-auth-server "check-server-group"
 end
 end
```

Other SSL mode options are also available:

<code>ssl-send-empty-frags</code> {disable   enable}	Enable to send empty fragments to avoid CBC IV attacks. Compatible with SSL 3.0 and TLS 1.0 only. Default is enable.
<code>ssl-client-renegotiation</code> {allow   deny   secure}	Control how the ALG responds when a client attempts to renegotiate the SSL session. You can allow renegotiation or block sessions when the client attempts to renegotiate. You can also select <code>secure</code> to reject an SSL connection that does not support <a href="#">RFC 5746</a> secure renegotiation indication. Default is <code>allow</code> .
<code>ssl-algorithm</code> {high   low   medium}	Select the relative strength of the algorithms that can be selected. You can select <code>high</code> , the default, to allow only AES or 3DES, <code>medium</code> , to allow AES, 3DES, or RC4 or <code>low</code> , to allow AES, 3DES, RC4, or DES.
<code>ssl-pfs</code> {allow   deny   require}	Select whether to allow, deny, or require perfect forward secrecy (PFS). Default is <code>allow</code> .
<code>ssl-min-version</code> {ssl-3.0   tls-1.0   tls-1.1}	Select the minimum level of SSL support to allow. The default is <code>ssl-3.0</code> .
<code>ssl-max-version</code> {ssl-3.0   tls-1.0   tls-1.1}	Select the maximum level of SSL support to allow. The default is <code>tls-1.1</code> .

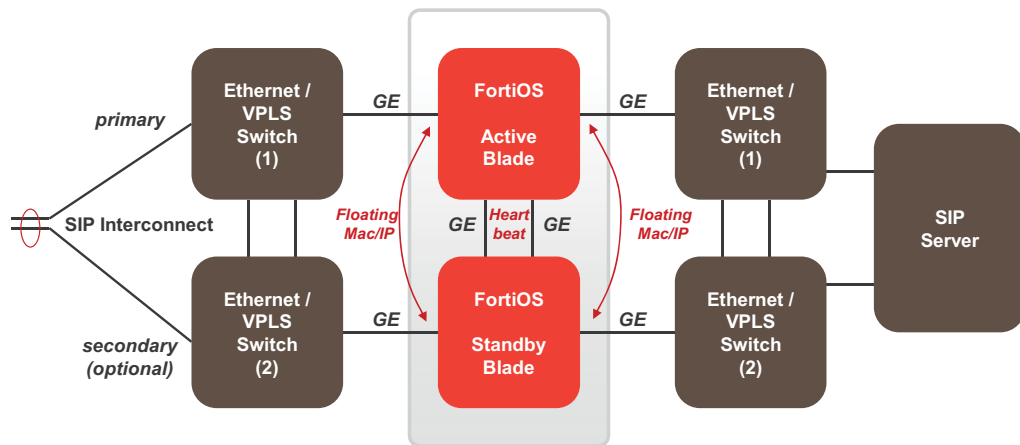
## SIP and HA: session failover and geographic redundancy

FortiGate high availability supports SIP session failover (also called stateful failover) for active-passive HA. To support SIP session failover, create a standard HA configuration and select the Enable Session Pick-up option.

SIP session failover replicates SIP states to all cluster units. If an HA failover occurs, all in progress SIP calls (setup complete) and their RTP flows are maintained and the calls will continue after the failover with minimal or no interruption.

SIP calls being set up at the time of a failover may lose signaling messages. In most cases the SIP clients and servers should use message retransmission to complete the call setup after the failover has completed. As a result, SIP users may experience a delay if their calls are being set up when an HA a failover occurs. But in most cases the call setup should be able to continue after the failover.

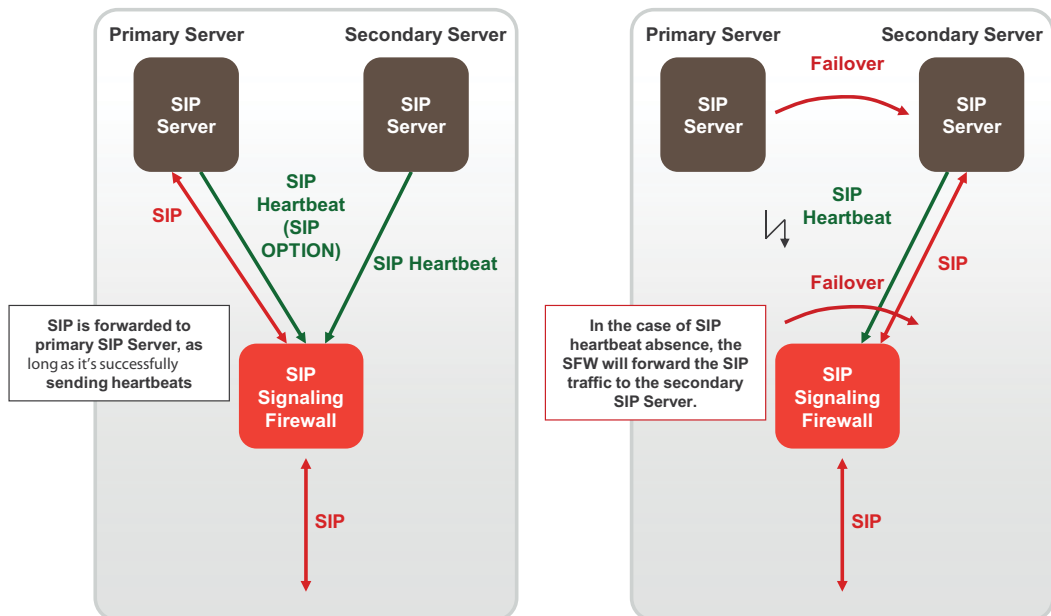
**Figure 378:**SIP HA session failover



## SIP geographic redundancy

Maintains a active-standby SIP server configuration, which even supports geographical distribution. If the active SIP server fails (missing SIP heartbeat messages or SIP traffic) FortiOS will redirect the SIP traffic to a secondary SIP server. SIP geographic redundancy

**Figure 379:**Geographic redundancy



## Supporting geographic redundancy when blocking OPTIONS messages

For some geographic redundant SIP configurations, the SIP servers may use SIP OPTIONS messages as heartbeats to notify the FortiGate unit that they are still operating (or alive). This is a kind of passive SIP monitoring mechanism where the FortiGate unit isn't actively monitoring the SIP servers and instead the FortiGate unit passively receives and analyzes OPTIONS messages from the SIP servers.

If FortiGate units block SIP OPTIONS messages because `block-options` is enabled, the configuration may fail to operate correctly because the OPTIONS messages are blocked by one or more FortiGate units.

However, you can work around this problem by enabling the `block-geo-red-options` application control list option. This option causes the FortiGate unit to refresh the local SIP server status when it receives an OPTIONS message before dropping the message. The end result is the heartbeat signals between geographically redundant SIP servers are maintained but OPTIONS messages do not pass through the FortiGate unit.

Use the following command to block OPTIONS messages while still supporting geographic redundancy:

```
config voip profile
 edit VoIP_Pro_Name
 config sip
 set block-options disable
 set block-geo-red-options enable
 end
 end
end
```



The `block-options` option setting overrides the `block-geo-red-options` option. If `block-options` is enabled the FortiGate unit only blocks SIP OPTIONS messages and does not refresh local SIP server status.

---

## Support for RFC 2543-compliant branch parameters

RFC 3261 is the most recent SIP RFC, it obsoletes RFC 2543. However, some SIP implementations may use RFC 2543-compliant SIP calls.

The `rfc2543-branch` VoIP profile option allows the FortiGate unit to support SIP calls that include an RFC 2543-compliant branch parameter in the SIP Via header. This option also allows FortiGate units to support SIP calls that include Via headers that are missing the branch parameter.

```
config voip profile
 edit VoIP_Pro_Name
 config sip
 set rfc2543-branch enable
 end
 end
end
```

## SIP and IPS

You can enable IPS in security policies that also accept SIP sessions to protect the SIP traffic from SIP-based attacks. If you enable IPS in this way then by default the pinholes that the SIP ALG creates to allow RTP and RTCP to flow through the firewall will also have IPS enabled.

This inheritance of the IPS setting can cause performance problems if the RTP traffic volume is high since IPS checking may reduce performance in some cases. Also if you are using network processor (NP) interfaces to accelerate VoIP performance, when IPS is enabled for the pinhole traffic is diverted to the IPS and as a result is not accelerated by the network processors.

You can use the following CLI command to disable IPS for the RTP pinhole traffic.

```
config voip profile
 edit VoIP_Pro_Name
 config sip
 set ips-rtp disable
 end
 end
```

## SIP debugging

### SIP debug log format

Assuming that `diagnose debug console timestamp` is enabled then the following shows the debug that is generated for an INVITE if `diag debug appl sip -1` is enabled:

```
2010-01-04 21:39:59 sip port 26 locate session for 192.168.2.134:5061 ->
172.16.67.192:5060
2010-01-04 21:39:59 sip sess 0x979df38 found for 192.168.2.134:5061 ->
172.16.67.192:5060
2010-01-04 21:39:59 sip port 26 192.168.2.134:5061 -> 172.16.67.192:5060
2010-01-04 21:39:59 sip port 26 read [(0,515)
(494e56495445207369703a73657276696365403139322e3136382e322e3130303a35303630205349502f322e300d0a566961
3a205349502f322e302f554450203132372e302e312e313a353036313b6272616e63683d7a39684734624b2d363832
372d3632302d300d0a46726f6d3a2073697070203c7369703a73697070403132372e302e312e313a353036313e3b74
61673d363832375349507054616730303632300d0a546f3a20737574203c7369703a73657276696365403139322e31
36382e322e3130303a353036303e0d0a43616c6c2d49443a203632302d36383237403132372e302e312e310d0a4353
65713a203120494e564954450d0a436f6e746163743a207369703a73697070403132372e302e312e313a353036310d
0a4d61782d466f7277617264733a2037300d0a5375626a6563743a20506572666f726d616e636520546573740d0a43
6f6e74656e742d547970653a206170706c69636174696f6e2f7364700d0a436f6e74656e742d4c656e6774683a2020
3132390d0a0d0a763d300d0a6f3d7573657231203533363535373635203233353336383736333720494e2049503420
```

```

3132372e302e312e310d0a733d2d0d0a633d494e20495034203132372e302e312e310d0a743d3020300d0a6d3d6175
64696f2036303031205254502f41565020300d0a613d7274706d61703a302050434d552f383030300d0a) (INVITE
sip:service@192.168.2.100:5060 SIP/2.0..Via: SIP/2.0/UDP
127.0.1.1:5061;branch=z9hG4bK-6827-620-0..From: sipp
%lt;sip:sipp@127.0.1.1:5061>;tag=6827SIPpTag00620..To: sut
%lt;sip:service@192.168.2.100:5060>..Call-ID: 620-6827@127.0.1.1..CSeq: 1
INVITE..Contact: sip:sipp@127.0.1.1:5061..Max-Forwards: 70..Subject: Performance
Test..Content-Type: application/sdp..Content-Length: 129...v=0..o=user1 53655765
2353687637 IN IP4 127.0.1.1..s=-..c=IN IP4 127.0.1.1..t=0 0..m=audio 6001 RTP/AVP
0..a=rtpmap:0 PCMU/8000..)]
2010-01-04 21:39:59 sip port 26 len 515
2010-01-04 21:39:59 sip port 26 INVITE '192.168.2.100:5060' addr 192.168.2.100:5060
2010-01-04 21:39:59 sip port 26 CSeq: 1 INVITE
2010-01-04 21:39:59 sip port 26 Via: UDP 127.0.1.1:5061 len 14 received 0 rport 0 0 branch
'z9hG4bK-6827-620-0'
2010-01-04 21:39:59 sip port 26 From: 'sipp ;tag=6827SIPpTag00620' URI 'sip:sipp@127.0.1.1:5061' tag
'6827SIPpTag00620'
2010-01-04 21:39:59 sip port 26 To: 'sut ' URI 'sip:service@192.168.2.100:5060' tag ''
2010-01-04 21:39:59 sip port 26 Call-ID: '620-6827@127.0.1.1'
2010-01-04 21:39:59 sip port 26 Contact: '127.0.1.1:5061' addr 127.0.1.1:5061 expires 0
2010-01-04 21:39:59 sip port 26 Content-Length: 129 len 3
2010-01-04 21:39:59 sip port 26 sdp o=127.0.1.1 len=9
2010-01-04 21:39:59 sip port 26 sdp c=127.0.1.1 len=9
2010-01-04 21:39:59 sip port 26 sdp m=6001 len=4
2010-01-04 21:39:59 sip port 26 find call 0 '620-6827@127.0.1.1'
2010-01-04 21:39:59 sip port 26 not found
2010-01-04 21:39:59 sip port 26 call 0x97a47c0 open (collision (nil))
2010-01-04 21:39:59 sip port 26 call 0x97a47c0 open txn 0x979f7f8 INVITE dir 0
2010-01-04 21:39:59 sip port 26 sdp i: 127.0.1.1:6001
2010-01-04 21:39:59 sip port 26 policy id 1 is_client_vs_policy 1 policy_dir_rev 0
2010-01-04 21:39:59 sip port 26 policy 1 not RTP policy
2010-01-04 21:39:59 sip port 26 learn sdp from stream address
2010-01-04 21:39:59 sip port 26 call 0x97a47c0 sdp 172.16.67.198:43722
2010-01-04 21:39:59 sip port 26 call 0x97a47c0 txn 0x979f7f8 127.0.1.1:5061 find new address and port
2010-01-04 21:39:59 sip port 26 call 0x97a47c0 txn 0x979f7f8 127.0.1.1:5061 find new address and port
2010-01-04 21:39:59 sip port 26 call 0x97a47c0 txn 0x979f7f8 127.0.1.1:5061 find new address and port
2010-01-04 21:39:59 sip port 30 write 192.168.2.134:5061 -> 172.16.67.192:5060 (13,539)
2010-01-04 21:39:59 sip port 30 write [(13,539)
(494e56495445207369703a73657276696365403137322e31362e36372e3139323a35303630205349502f322e300d0a566961
3a205349502f322e302f554450203137322e31362e36372e3139383a35323036353b6272616e63683d7a3968473462
4b2d363832372d3632302d300d0a46726f6d3a207369707070203c7369703a73697070403137322e31362e36372e3139
383a34333732343e3b7461673d363832375349507054616730303632300d0a546f3a20737574203c7369703a736572
76696365403137322e31362e36372e3139323a353036303e0d0a43616c6c2d49443a203632302d3638323740313237
2e302e312e310d0a435365713a203120494e564954450d0a436f6e746163743a207369703a73697070403137322e31
362e36372e3139383a34333732350d0a4d61782d466f7277617264733a2037300d0a5375626a6563743a2050657266
6f726d616e636520546573740d0a436f6e74656e742d547970653a206170706c69636174696f6e2f7364700d0a436f
6e74656e742d4c656e6774683a20203133380d0a0d0a763d300d0a6f3d757365723120353336353537363520323335
3336383736333720494e20495034203137322e31362e36372e3139380d0a733d2d0d0a633d494e2049503420313732
2e31362e36372e3139380d0a743d3020300d0a6d3d617564696f203433373232205254502f41565020300d0a613d72
74706d61703a302050434d552f383030300d0a) (INVITE sip:service@172.16.67.192:5060 SIP/2.0..Via:
SIP/2.0/UDP 172.16.67.198:52065;branch=z9hG4bK-6827-620-0..From: sipp
;tag=6827SIPpTag00620..To: sut ..Call-ID: 620-6827@127.0.1.1..CSeq: 1 INVITE..Contact:
sip:sipp@172.16.67.198:43725..Max-Forwards: 70..Subject: Performance Test..Content-Type:
application/sdp..Content-Length: 138...v=0..o=user1 53655765 2353687637 IN IP4
172.16.67.198..s=-..c=IN IP4 172.16.67.198..t=0 0..m=audio 43722 RTP/AVP 0..a=rtpmap:0
PCMU/8000..)]

```

## SIP-proxy filter per VDOM

You can use the diagnose sys sip-proxy xxx command in a VDOM to get info about how SIP is operating in each VDOM.

## SIP-proxy filter command

Use the `diagnose system sip-proxy filter` to filter diagnose information for the SIP ALG. The following filters are available:

```
diag sys sip-proxy filter vd
diag sys sip-proxy filter dst-addr4
diag sys sip-proxy filter dst-addr6
diag sys sip-proxy filter dst-port
diag sys sip-proxy filter identity-policy
diag sys sip-proxy filter negate
diag sys sip-proxy filter policy
diag sys sip-proxy filter policy-type
diag sys sip-proxy filter profile-group
diag sys sip-proxy filter src-addr4
diag sys sip-proxy filter src-addr6
diag sys sip-proxy filter src-port
diag sys sip-proxy filter vd
diag sys sip-proxy filter voip-profile
```

You can clear, view and negate/invert the sense of a filter using these commands:

```
diag sys sip-proxy filter clear
diag sys sip-proxy filter list
diag sys sip-proxy filter negate
```

## SIP debug log filtering

You can filter by VDOM/IP/PORT and by policy and VoIP profile. The filtering can be controlled by:

```
diagnose system sip-proxy log-filter
```

The list of filters is:

```
diag sys sip-proxy log-filter vd
diag sys sip-proxy log-filter dst-addr4
diag sys sip-proxy log-filter dst-addr6
diag sys sip-proxy log-filter dst-port
diag sys sip-proxy log-filter identity-policy
diag sys sip-proxy log-filter policy
diag sys sip-proxy log-filter policy-type
diag sys sip-proxy log-filter profile-group
diag sys sip-proxy log-filter src-addr4
diag sys sip-proxy log-filter src-addr6
diag sys sip-proxy log-filter src-port
diag sys sip-proxy log-filter vd
diag sys sip-proxy log-filter voip-profile
```

You can clear, view and negate/invert the sense of a filter using these commands:

```
diag sys sip-proxy log-filter clear
diag sys sip-proxy log-filter list
diag sys sip-proxy log-filter negate
```

## SIP debug setting

Control of the SIP debug output is governed by the following command

```
diagnose debug application sip <debug_level_int>
```

Where the <debug\_level\_int> is a bitmask and the individual values determine whether the listed items are logged or not. The <debug\_level\_int> can be

- 1 - configuration changes. Mainly addition/deletion/modification of virtual domains.
- 2 - (TCP) connection accepts or connects, redirect creation
- 4 - create or delete a session
- 16 - any IO read or write
- 32 - an ASCII dump of all data read or written
- 64 - Include HEX dump in the above output
- 128 - any activity related to the use of the FortiCarrier dynamic profile feature to determine the correct profile-group to use
- 256 - log summary of interesting fields in a SIP call
- 1024 - any activity related to SIP geo-redundancy.
- 2048 - any activity related to HA syncing of SIP calls.

## Display SIP rate-limit data

You can use the `diagnose sys sip-proxy meters` command to display SIP rate limiting data.

For the following command output `rate 1` shows that the current (over last second) measured rate for INVITE/ACK and BYE was 1 per second, the `peak 1` shows that the peak rate recorded is 1 per second, the `max 0` shows that there is no maximum limit set, the `count 18` indicates that 18 messages were received and `drop 0` indicates that none were dropped due to being over the limit.

```
diag sys sip-proxy meters
sip
sip vd: 0
sip policy: 1
sip identity-policy: 0
sip policy-type: IPv4
sip profile-group:
sip dialogs: 18
sip dialog-limit: 0
sip UNKNOWN: rate 0 peak 0 max 0 count 0 drop 0
sip ACK: rate 1 peak 1 max 0 count 18 drop 0
sip BYE: rate 1 peak 1 max 0 count 18 drop 0
sip CANCEL: rate 0 peak 0 max 0 count 0 drop 0
sip INFO: rate 0 peak 0 max 0 count 0 drop 0
sip INVITE: rate 1 peak 1 max 0 count 18 drop 0
sip MESSAGE: rate 0 peak 0 max 0 count 0 drop 0
sip NOTIFY: rate 0 peak 0 max 0 count 0 drop 0
sip OPTIONS: rate 0 peak 0 max 0 count 0 drop 0
sip PRACK: rate 0 peak 0 max 0 count 0 drop 0
sip PUBLISH: rate 0 peak 0 max 0 count 0 drop 0
sip REFER: rate 0 peak 0 max 0 count 0 drop 0
sip REGISTER: rate 0 peak 0 max 0 count 0 drop 0
sip SUBSCRIBE: rate 0 peak 0 max 0 count 0 drop 0
sip UPDATE: rate 0 peak 0 max 0 count 0 drop 0
sip PING: rate 0 peak 0 max 0 count 0 drop 0
sip YAHOOREF: rate 0 peak 0 max 0 count 0 drop 0
```



# Chapter 22 WAN Optimization, Web

## Cache, Explicit Proxy, and WCCP for FortiOS 5.0

Welcome and thank you for selecting Fortinet products for your network protection.

You can use FortiGate WAN optimization and web caching to improve performance and security of traffic passing between locations on your wide area network (WAN) or from the Internet to your web servers. You can also use the FortiGate unit as an explicit FTP and web proxy server. If your FortiGate unit supports web caching, you can also add web caching to any HTTP sessions including WAN optimization, explicit web proxy and other HTTP sessions.

This document describes how FortiGate WAN optimization, web caching, explicit web proxy, explicit FTP proxy and WCCP work and also describes how to configure these features.

### Before you begin

Before you begin to configure WAN optimization, Web caching, explicit proxies or WCCP, take a moment to note the following:

- To use WAN optimization and web caching your FortiGate unit must support these features and not all do. In general your FortiGate unit must include a hard disk to support these features. See [“FortiGate models that support WAN optimization” on page 2578](#). Most FortiGate units support the explicit web and FTP proxies.
- To be able to configure WAN optimization and web caching from the web-based manager you should begin by going to the System Information dashboard widget and enabling *WAN Opt. & Cache*.
- To be able to configure the Explicit Web and FTP proxies from the web-based manager you should begin by going to the System Information dashboard widget and enabling *Explicit Proxy*.
- If you enable virtual domains (VDOMs) on the FortiGate unit, WAN optimization, web caching, and the explicit web and FTP proxies are available separately for each VDOM.
- This guide is based on the assumption that you are a FortiGate administrator. It is not intended for others who may also use the FortiGate unit, such as FortiClient administrators or end users.
- FortiGate WAN optimization is proprietary to Fortinet. FortiGate WAN optimization is compatible only with FortiClient WAN optimization, and will not work with other vendors' WAN optimization or acceleration features.
- FortiGate web caching, explicit web and FTP proxies, and WCCP support known standards for these features. See the appropriate chapters of this document for details.

At this stage, the following installation and configuration conditions are assumed:

- For WAN optimization you have already successfully installed two or more FortiGate units at various locations across your WAN by following the instructions in the appropriate FortiGate unit QuickStart or Installation Guide. You can download FortiGate installation guides from the FortiGate documentation page: <http://docs.fortinet.com/fgt.html>.
- For web caching, the explicit proxies and WCCP you have already successfully installed one or more FortiGate units on your network by following the instructions in the appropriate FortiGate unit QuickStart Guide. You can download FortiGate installation guides from the FortiGate documentation page: <http://docs.fortinet.com/fgt.html>.
- You have administrative access to the web-based manager and/or CLI.
- The FortiGate units are integrated into your WAN or other networks
- The operation mode has been configured.
- The system time, DNS settings, administrator password, and network interfaces have been configured.
- Firmware, FortiGuard Antivirus and FortiGuard Antispam updates are completed.
- You have added security policies to allow your FortiGate units to process traffic.
- You Fortinet products have been registered. Register your Fortinet products at the Fortinet Technical Support web site, <https://support.fortinet.com>.

## FortiGate models that support WAN optimization

WAN optimization is available on FortiGate models with internal storage that also support SSL acceleration. Internal storage includes high-capacity internal hard disks, AMC hard disk modules, FortiGate Storage Modules (FSMs) or over 4 Gbytes of internal flash storage. All of these storage locations can provide similar web caching and byte caching performance. If you add more than one storage location (for example, by creating multiple partitions on a storage device, by using more than one FSM, or by using an FSM and AMC hard disk in the same FortiGate unit) you can configure different storage locations for web caching and byte caching.

You can configure WAN optimization storage options from the FortiGate CLI. See “[Storage](#)” on [page 2716](#).

## How this chapter is organized

This FortiOS Handbook chapter describes how to implement WAN optimization, web caching and the web proxy on supported FortiGate units.

The FortiOS Handbook chapter contains the following sections:

**Example network topologies:** Provides an overview of FortiGate WAN optimization best practices and technologies and some of the concepts and rules for using them. We recommend that you begin with this chapter before attempting to configure your FortiGate unit to use WAN optimization.

**Storage:** Describes how to configure WAN optimization storage settings to control how data is stored for web caching and byte caching.

**Peers and authentication groups:** Describes how to use WAN optimization peers and authentication groups to control access to WAN optimization tunnels.

**Configuring WAN optimization:** Provides basic configuration for WAN optimization rules, including adding rules, organizing rules in the rule list and using WAN optimization addresses. This chapter also explains how WAN optimization accepts sessions, as well as how and when you can apply security profile to WAN optimization traffic.

**Configuration examples:** Describes basic active-passive and peer-to-peer WAN optimization configuration examples. This chapter is a good place to start learning how to put an actual WAN optimization network together.

**Web caching and SSL offloading:** Describes how web caching works to cache HTTP and HTTPS, how to use SSL offloading to improved performance of HTTPS websites, and includes web caching configuration examples.

**FortiClient WAN optimization:** Describes how FortiGate and FortiClient WAN optimization work together and includes an example configuration.

**The FortiGate explicit web proxy:** Describes how to configure the FortiGate explicit web proxy, how users connect to the explicit web proxy, and how to add web caching to the explicit web proxy.

**The FortiGate explicit FTP proxy:** Describes how to configure the FortiGate explicit FTP proxy and how users connect to the explicit FTP proxy.

**FortiGate WCCP:** Describes FortiGate WCCP and how to configure WCCP and the WCCP client.

**Diagnose commands:** describes get and diagnose commands available for troubleshooting WAN optimization, web cache, and WCCP.

# Example network topologies

FortiGate WAN optimization consists of a number of techniques that you can apply to improve the efficiency of communication across your WAN. These techniques include protocol optimization, byte caching, web caching, SSL offloading, and secure tunnelling. Protocol optimization can improve the efficiency of traffic that uses the CIFS, FTP, HTTP, or MAPI protocol, as well as general TCP traffic. Byte caching caches files and other data on FortiGate units to reduce the amount of data transmitted across the WAN. Web caching stores web pages on FortiGate units to reduce latency and delays between the WAN and web servers. SSL offloading offloads SSL decryption and encryption from web servers onto FortiGate SSL acceleration hardware. Secure tunnelling secures traffic as it crosses the WAN.

You can apply different combinations of these WAN optimization techniques to a single traffic stream depending on the traffic type. For example, you can apply byte caching and secure tunneling to any TCP traffic. For HTTP and HTTPS traffic, you can also apply protocol optimization and web caching.

You can configure a FortiGate unit to be an explicit web proxy server for both IPv4 and IPv6 traffic and an explicit FTP proxy server. Users on your internal network can browse the Internet through the explicit web proxy server or connect to FTP servers through the explicit FTP proxy server. You can also configure these proxies to protect access to web or FTP servers behind the FortiGate unit using a reverse proxy configuration.

Web caching can be applied to any HTTP or HTTPS traffic, this includes normal traffic accepted by a security policy, explicit web proxy traffic, and WAN optimization traffic.

You can also configure a FortiGate unit to operate as a Web Cache Communication Protocol (WCCP) client or server. WCCP provides the ability to offload web caching to one or more redundant web caching servers.

FortiGate units can also apply security profiles to traffic as part of a WAN optimization, explicit web proxy, explicit FTP proxy, web cache and WCCP configuration. Security policies that include any of these options can also include settings to apply all forms of security profile inspection supported by your FortiGate unit.

This chapter describes:

- [WAN optimization topologies](#)
- [Explicit Web proxy topologies](#)
- [Explicit FTP proxy topologies](#)
- [Web caching topologies](#)
- [WCCP topologies](#)

## WAN optimization topologies


This section describes some common WAN optimization topologies:

- [“Basic WAN optimization topologies” on page 2581](#)
- [“Out-of-path topology” on page 2581](#)
- [“Topology for multiple networks” on page 2583](#)
- [“WAN optimization with web caching” on page 2583](#)
- [“WAN optimization and web caching with FortiClient peers” on page 2584](#)

## Basic WAN optimization topologies

The basic FortiGate WAN optimization topology consists of two FortiGate units operating as WAN optimization peers intercepting and optimizing traffic crossing the WAN between the private networks.

**Figure 380:** Security device and WAN optimization topology



As shown in [Figure 380](#), FortiGate units can be deployed as security devices that protect private networks connected to the WAN and also perform WAN optimization. In this configuration, the FortiGate units are configured as typical security devices for the private networks and are also configured for WAN optimization. The WAN optimization configuration intercepts traffic to be optimized as it passes through the FortiGate unit and uses a WAN optimization tunnel with another FortiGate unit to optimize the traffic that crosses the WAN.

As shown in [Figure 381](#), you can also deploy WAN optimization on single-purpose FortiGate units that only perform WAN optimization. In [Figure 381](#), the WAN optimization FortiGate units are located on the WAN outside of the private networks. You can also install the WAN optimization FortiGate units behind the security devices on the private networks.

The WAN optimization configuration is the same for FortiGate units deployed as security devices and for single-purpose WAN optimization FortiGate units. The only differences would result from the different network topologies.

## Out-of-path topology

In an out-of-path topology, one or both of the FortiGate units configured for WAN optimization are not directly in the main data path. Instead, the out-of-path FortiGate unit is connected to a device on the data path, and the device is configured to redirect sessions to be optimized to the out-of-path FortiGate unit.

**Figure 381:**Single-purpose WAN optimization topology





Figure 382 shows out-of-path FortiGate units configured for WAN optimization and connected directly to FortiGate units in the data path. The FortiGate units in the data path use a method such as policy routing to redirect traffic to be optimized to the out-of-path FortiGate units. The out-of-path FortiGate units establish a WAN optimization tunnel between each other and optimize the redirected traffic.

**Figure 382:**Out-of-path WAN optimization



One of the benefits of out-of-path WAN optimization is that out-of-path FortiGate units only perform WAN optimization and do not have to process other traffic. An in-path FortiGate unit configured for WAN optimization also has to process other non-optimized traffic on the data path.


The out-of-path FortiGate units can operate in NAT/Route or Transparent mode.

Other out-of-path topologies are also possible. For example, you can install the out-of-path FortiGate units on the private networks instead of on the WAN. Also, the out-of-path FortiGate units can have one connection to the network instead of two. In a one-arm configuration such as this, security policies and routing have to be configured to send the WAN optimization tunnel out the same interface as the one that received the traffic.

## Topology for multiple networks

As shown in [Figure 383](#), you can create multiple WAN optimization configurations between many private networks. Whenever WAN optimization occurs, it is always between two FortiGate units, but you can configure any FortiGate unit to perform WAN optimization with any of the other FortiGate units that are part of your WAN.

**Figure 383:**WAN optimization among multiple networks




You can also configure WAN optimization between FortiGate units with different roles on the WAN. FortiGate units configured as security devices and for WAN optimization can perform WAN optimization as if they are single-purpose FortiGate units just configured for WAN optimization.

## WAN optimization with web caching

You can add web caching to a WAN optimization topology when users on a private network communicate with web servers located across the WAN on another private network.

**Figure 384:**WAN optimization with web caching topology




The topology in [Figure 384](#) is the same as that of [Figure 380 on page 2581](#) with the addition of web caching to the FortiGate unit in front of the private network that includes the web servers. You can also add web caching to the FortiGate unit that is protecting the private network. In a similar way, you can add web caching to all of the topologies shown in “[WAN optimization topologies](#)” on [page 2580](#).

## WAN optimization and web caching with FortiClient peers

FortiClient WAN optimization works with FortiGate WAN optimization to accelerate remote user access to the private networks behind FortiGate units. The FortiClient application requires a simple WAN optimization configuration to automatically detect if WAN optimization is enabled on the FortiGate unit. Once WAN optimization is enabled, the FortiClient application transparently makes use of the WAN optimization and web caching features available.




**Figure 385:**FortiClient WAN optimization topology



## Explicit Web proxy topologies


You can configure a FortiGate unit to be an explicit web proxy server for Internet web browsing of IPv4 and IPv6 web traffic. To use the explicit web proxy, users must add the IP address of the FortiGate interface configured for the explicit web proxy to their web browser proxy configuration.

**Figure 386:**Explicit web proxy topology



If the FortiGate unit supports web caching, you can also add web caching to the security policy that accepts explicit web proxy sessions. The FortiGate unit then caches Internet web pages on a hard disk to improve web browsing performance.


**Figure 387:**Explicit web proxy with web caching topology



## Explicit FTP proxy topologies


You can configure a FortiGate unit to be an explicit FTP proxy server for FTP users. To use the explicit web proxy, FTP users must connect to and authenticate with the explicit FTP proxy before connecting to an FTP server.

**Figure 388:**Explicit FTP proxy topology



You can also configure reverse explicit FTP proxy ([Figure 389](#)). In this configuration, users on the Internet connect to the explicit web proxy before connecting to an FTP server installed behind a FortiGate unit.

**Figure 389:**Reverse explicit FTP proxy topology




## Web caching topologies

FortiGate web caching can be added to any security policy and any HTTP or HTTPS traffic accepted by that security policy can be cached on the FortiGate unit hard disk. This includes WAN optimization and explicit web proxy traffic. The network topologies for these scenarios are very similar. They involved a FortiGate unit installed between users and web servers with web caching enabled.


A typical web-caching topology includes one FortiGate unit that acts as a web cache server (Figure 390). Web caching is enabled in a security policy and the FortiGate unit intercepts web page requests accepted by the security policy, requests web pages from the web servers, caches the web page contents, and returns the web page contents to the users. When the FortiGate unit intercepts subsequent requests for cached web pages, the FortiGate unit contacts the destination web server just to check for changes.

**Figure 390:** Web caching topology



You can also configure reverse proxy web-caching (Figure 391). In this configuration, users on the Internet browse to a web server installed behind a FortiGate unit. The FortiGate unit intercepts the web traffic (HTTP and HTTPS) and caches pages from the web server. Reverse proxy web caching on the FortiGate unit reduces the number of requests that the web server must handle, leaving it free to process new requests that it has not serviced before.

**Figure 391:** Reverse proxy web caching topology




## WCCP topologies

You can operate a FortiGate unit as a Web Cache Communication Protocol (WCCP) router or cache engine. As a router, the FortiGate unit intercepts web browsing requests from client web

browsers and forwards them to a WCCP cache engine. The cache engine returns the required cached content to the client web browser. If the cache server does not have the required content it accesses the content, caches it and returns the content to the client web browser.

**Figure 392:**WCCP topology



FortiGate units can also operate as WCCP cache servers, communicating with WCCP routers, caching web content and providing it to client web browsers as required.

WCCP is transparent to client web browsers. The web browsers do not have to be configured to use a web proxy.

# Configuring WAN optimization

This chapter describes FortiGate WAN optimization client server architecture and other concepts you need to understand to be able to configure FortiGate WAN optimization.

This chapter describes:

- Client/server architecture
- WAN optimization peers
- Manual (peer-to-peer) and active-passive WAN optimization
- WAN optimization profiles
- Protocol optimization
- Byte caching
- WAN optimization transparent mode
- FortiClient WAN optimization
- Operating modes and VDOMs
- WAN optimization tunnels
- WAN optimization and user and device identity policies, load balancing and traffic shaping
- WAN optimization and HA
- WAN optimization, web caching and memory usage
- Monitoring WAN optimization performance
- WAN optimization configuration summary
- Best practices

## Client/server architecture

Traffic across a WAN typically consists of clients on a client network communicating across a WAN with a remote server network. The clients do this by starting communication sessions from the client network to the server network. To optimize these sessions, you can add **WAN optimization security policies** to the **client-side FortiGate unit** to accept sessions from the client network that are destined for the server network. The client-side FortiGate unit is located between the client network and the WAN (see [Figure 393](#)). WAN optimization security policies include **WAN optimization profiles** that control how the traffic is optimized.


The client-side FortiGate unit must also include the IP address of the **server-side FortiGate unit** in its WAN optimization **peer** configuration. The server-side FortiGate unit is located between the server network and the WAN, The peer configuration allows the client-side FortiGate unit to find the server-side FortiGate unit and attempt to establish a WAN optimization **tunnel** with it.

For the server-side FortiGate unit you must add a security policy with *wanopt* as the *Incoming Interface*. This security policy allows the FortiGate unit to accept WAN optimization sessions from the client-side FortiGate unit. For the server-side FortiGate unit to accept a WAN optimization connection it must have the client-side FortiGate unit in its WAN optimization peer configuration.



WAN optimization profiles are only added to the client-side WAN optimization security policy. The server-side FortiGate unit employs the WAN optimization settings set in the WAN optimization profile on the client-side FortiGate unit.

**Figure 393:**Client/server architecture



When both peers are identified the FortiGate units attempt to establish a WAN optimization **tunnel** between them. WAN optimization tunnels use port 7810. All optimized data flowing across the WAN between the client-side and server-side FortiGate units use this tunnel. WAN optimization tunnels can be encrypted use SSL encryption to keep the data in the tunnel secure.


In addition to basic identification by peer host ID and IP address you can configure WAN optimization **authentication** using certificates and pre-shared keys to improve security. You can also configure FortiGate units involved in WAN optimization to accept connections from any identified peer or restrict connections to specific peers.

The FortiClient application can act in the same manner as a client-side FortiGate unit to optimize traffic between a computer running FortiClient and a FortiGate unit.

## WAN optimization peers

The client-side and server-side FortiGate units are called WAN optimization peers (see [Figure 394](#)) because all of the FortiGate units in a WAN optimization network have the same peer relationship with each other. The client and server roles just relate to how a session is started. Any FortiGate unit configured for WAN optimization can be a client-side and a server-side FortiGate unit at the same time, depending on the direction of the traffic. Client-side FortiGate units initiate WAN optimization sessions and server-side FortiGate units respond to the session requests. Any FortiGate unit can simultaneously be a client-side FortiGate unit for some sessions and a server-side FortiGate unit for others.

**Figure 394:**WAN optimization peer and tunnel architecture



To identify all of the WAN optimization peers that a FortiGate unit can perform WAN optimization with, you add host IDs and IP addresses of all of the peers to the FortiGate unit configuration. The peer IP address is actually the IP address of the peer unit interface that communicates with the FortiGate unit.

## Manual (peer-to-peer) and active-passive WAN optimization

You can create **manual** (peer-to-peer) and **active-passive** WAN optimization configurations.

### Manual (peer to peer) configurations

Manual configurations allow for WAN optimization between one client-side FortiGate unit and one server-side FortiGate unit. To create a manual configuration you add a **manual mode** WAN optimization security policy to the client-side FortiGate unit. The manual mode policy includes the peer ID of a server-side FortiGate unit.



In a manual mode configuration, the client-side peer can only connect to the named server-side peer. When the client-side peer initiates a tunnel with the server-side peer, the packets that initiate the tunnel include extra information so that the server-side peer can determine that it is a peer-to-peer tunnel request. This extra information is required because the server-side peer does not require a WAN optimization policy; you just need to add the client peer host ID and IP address to the server-side FortiGate unit peer list and a security policy with the wanopt interface as the incoming interface. WAN optimization tunnel requests are accepted by the policy and if the client-side peer is in the server side peer's address list the traffic is forwarded to its destination.

### Manual mode client-side policy

Add a manual mode policy to the client-side FortiGate unit from the CLI. The policy enables WAN optimization, sets wanopt-detection to off, and uses the wanopt-peer option to specify the server-side peer. The following example uses the default WAN optimization profile and also enables virus scanning using the default virus scanning profile.

```
config firewall policy
 edit 2
 set srcintf "internal"
 set dstintf "wan1"
 set srcaddr "client-subnet"
 set dstaddr "server-subnet"
 set action accept
 set schedule "always"
 set service "ANY"
 set utm-status enable
 set av-profile default
 set profile-protocol-options default
 set wanopt enable
 set wanopt-detection off
 set wanopt-profile "default"
 set wanopt-peer "server"
 next
end
```

### Server-side tunnel policy

The server-side policy allows WAN optimization tunnel connections by including the wanopt tunnel interface as the *Incoming Interface*. From the CLI the policy could look like the following:

```
configure firewall policy
 edit 3
 set srcintf "wanopt"
 set dstintf "internal"
 set srcaddr "all"
 set dstaddr "server-subnet"
 set action accept
 set schedule "always"
 set service "ANY"
 next
end
```

## Active-passive configurations

Active-passive WAN optimization requires an **active** WAN optimization policy on the client-side FortiGate unit and a **passive** WAN optimization policy on the server-side FortiGate unit. The server-side FortiGate unit also requires a security policy with the *wanopt* tunnel interface as the *Incoming Interface*.

You can use the passive policy to control WAN optimization address translation by specifying **transparent mode** or non-transparent mode. See [“WAN optimization transparent mode” on page 2599](#). You can also use the passive policy to apply security profiles, web caching, and other FortiGate features at the server-side FortiGate unit. For example, if a server-side FortiGate unit is protecting a web server, the passive policy could enable web caching.

A single passive policy can accept tunnel requests from multiple FortiGate units as long as the server-side FortiGate unit includes their peer IDs and all of the client-side FortiGate units include the server-side peer ID.

### Active client-side policy

Add an active policy to the client-side FortiGate unit by selecting *Enable WAN Optimization* and selecting *active*. Then select a WAN optimization *Profile*. From the CLI the policy could look like the following:

```
config firewall policy
 edit 2
 set srcintf "internal"
 set dstintf "wan1"
 set srcaddr "client-subnet"
 set dstaddr "server-subnet"
 set action accept
 set schedule "always"
 set service "ANY"
 set utm-status enable
 set av-profile default
 set profile-protocol-options default
 set wanopt enable
 set wanopt-detection active
 set wanopt-profile "default"
 next
end
```

## Server-side tunnel policy

The server-side requires a policy that allows WAN optimization tunnel connections by including the *wanopt* tunnel interface as the *Incoming Interface*. From the CLI the policy could look like the following:

```
configure firewall policy
 edit 3
 set srcintf "wanopt"
 set dstintf "internal"
 set srcaddr "all"
 set dstaddr "server-subnet"
 set action accept
 set schedule "always"
 set service "ANY"
 next
end
```

## Server-side passive policy

Add a passive policy to the client-side FortiGate unit by selecting *Enable WAN Optimization* and selecting *passive*. Then set the *Passive Option* to transparent. From the CLI the policy could look like the following:

```
config firewall policy
 edit 2
 set srcintf "wan1"
 set dstintf "internal"
 set srcaddr "all"
 set dstaddr "all"
 set action accept
 set schedule "always"
 set service "ANY"
 set utm-status enable
 set av-profile default
 set profile-protocol-options default
 set wanopt enable
 set wanopt-detection passive
 set wanopt-passive-opt transparent
 next
```

## WAN optimization profiles

Use WAN optimization profiles to apply WAN optimization techniques to traffic to be optimized. In a WAN optimization profile you can select the protocols to be optimized and for each protocol you can enable SSL offloading (if supported), secure tunneling, byte caching and set the port or port range the protocol uses. You can also enable transparent mode and optionally select an authentication group. You can edit the default WAN optimization profile or create new ones.

To configure a WAN optimization profile go to *WAN Opt & Cache > WAN Opt. Profile > Profile* and edit a profile or create a new one.

**Figure 395:**Configuring a WAN optimization profile

Protocol	SSL Offloading	Secure Tunneling	Byte Caching	Port
<input checked="" type="checkbox"/> CIFS		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	445
<input checked="" type="checkbox"/> FTP		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	21
<input checked="" type="checkbox"/> HTTP	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	80
<input checked="" type="checkbox"/> MAPI		<input type="checkbox"/>	<input checked="" type="checkbox"/>	135
<input checked="" type="checkbox"/> TCP	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1-65535

From the CLI you can use the following command to configure a WAN optimization profile to optimize HTTP traffic.

```
config wanopt profile
 edit new-profile
 config http
 set status enable
 end
 end
```

<b>Transparent Mode</b>	Servers receiving packets after WAN optimization “see” different source addresses depending on whether or not you select <i>Transparent Mode</i> .  For more information, see <a href="#">“WAN optimization transparent mode” on page 2599</a> .
<b>Authentication Group</b>	Select this option and select an authentication group so that the client and server-side FortiGate units must authenticate with each other before starting the WAN optimization tunnel. You must also select an authentication group if you select <i>Secure Tunneling</i> for any protocol.  You must add identical authentication groups to both of the FortiGate units that will participate in the WAN optimization tunnel. For more information, see <a href="#">“Configuring authentication groups” on page 2612</a> .
<b>Protocol</b>	Select CIFS, FTP, HTTP or MAPI to apply protocol optimization for the selected protocols. See <a href="#">“Protocol optimization” on page 2598</a> .  Select TCP if the WAN optimization tunnel accepts sessions that use more than one protocol or that do not use the CIFS, FTP, HTTP, or MAPI protocol.
<b>SSL Offloading</b>	Select to apply SSL offloading for HTTPS or other SSL traffic. You can use SSL offloading to offload SSL encryption and decryption from one or more HTTP servers to the FortiGate unit. If you enable this option, you must configure the security policy to accept SSL-encrypted traffic.  If you enable SSL offloading, you must also use the CLI command <code>config wanopt ssl-server</code> to add an SSL server for each HTTP server that you want to offload SSL encryption/decryption for. For more information, see <a href="#">“Turning on web caching and SSL offloading for HTTPS traffic” on page 2644</a> .

<b>Secure Tunnelling</b>	The WAN optimization tunnel is encrypted using SSL encryption. You must also add an authentication group to the profile. For more information, see <a href="#">“Secure tunneling” on page 2615</a> .
<b>Byte Caching</b>	Select to apply WAN optimization byte caching to the sessions accepted by this rule. For more information, see <a href="#">“Byte caching” on page 2598</a> .
<b>Port</b>	Enter a single port number or port number range. Only packets whose destination port number matches this port number or port number range will be optimized.

## Processing non-HTTP sessions accepted by a WAN optimization profile with HTTP optimization

From the CLI, you can use the following command to configure how to process non-HTTP sessions when a rule configured to accept and optimize HTTP traffic accepts a non-HTTP session. This can occur if an application sends non-HTTP sessions using an HTTP destination port.

```
config wanopt profile
 edit default
 config http
 set status enable
 set tunnel-non-http {disable | enable}
 end
```

To drop non-HTTP sessions accepted by the rule set `tunnel-non-http` to `disable`, or set it to `enable` to pass non-HTTP sessions through the tunnel without applying protocol optimization, byte-caching, or web caching. In this case, the FortiGate unit applies TCP protocol optimization to non-HTTP sessions.

## Processing unknown HTTP sessions

Unknown HTTP sessions are HTTP sessions that do not comply with HTTP 0.9, 1.0, or 1.1. From the CLI, use the following command to specify how a rule handles such HTTP sessions.

```
config wanopt profile
 edit default
 config http
 set status enable
 set unknown-http-version {best-effort | reject | tunnel}
 end
```

To assume that all HTTP sessions accepted by the rule comply with HTTP 0.9, 1.0, or 1.1, select `best-effort`. If a session uses a different HTTP version, WAN optimization may not parse it correctly. As a result, the FortiGate unit may stop forwarding the session and the connection may be lost. To reject HTTP sessions that do not use HTTP 0.9, 1.0, or 1.1, select `reject`.

To pass HTTP sessions that do not use HTTP 0.9, 1.0, or 1.1, but without applying HTTP protocol optimization, byte-caching, or web caching, you can also select `tunnel`. TCP protocol optimization is applied to these HTTP sessions.

## Protocol optimization

Protocol optimization techniques optimize bandwidth use across the WAN. These techniques can improve the efficiency of communication across the WAN optimization tunnel by reducing the amount of traffic required by communication protocols. You can apply protocol optimization to Common Internet File System (CIFS), FTP, HTTP, MAPI, and general TCP sessions. You can apply general TCP optimization to MAPI sessions.

For example, CIFS provides file access, record locking, read/write privileges, change notification, server name resolution, request batching, and server authentication. CIFS is a fairly “chatty” protocol, requiring many background transactions to successfully transfer a single file. This is usually not a problem across a LAN. However, across a WAN, latency and bandwidth reduction can slow down CIFS performance.

When you select the CIFS protocol in a WAN optimization profile, the FortiGate units at both ends of the WAN optimization tunnel use a number of techniques to reduce the number of background transactions that occur over the WAN for CIFS traffic.

If a policy accepts a range of different types of traffic, you can set *Protocol* to *TCP* to apply general optimization techniques to TCP traffic. However, applying this TCP optimization is not as effective as applying more protocol-specific optimization to specific types of traffic. TCP protocol optimization uses techniques such as TCP SACK support, TCP window scaling and window size adjustment, and TCP connection pooling to remove TCP bottlenecks.

### Protocol optimization and MAPI

By default the MAPI service uses port number 135 for RPC port mapping and may use random ports for MAPI messages. The random ports are negotiated through sessions using port 135. The FortiOS DCE-RPC session helper learns these ports and opens pinholes for the messages. WAN optimization is also aware of these ports and attempts to apply protocol optimization to MAPI messages that use them. However, to configure protocol optimization for MAPI you should set the WAN optimization profile to a single port number (usually port 135). Specifying a range of ports may reduce performance.

## Byte caching

Byte caching breaks large units of application data (for example, a file being downloaded from a web page) into small chunks of data, labelling each chunk of data with a hash of the chunk and storing those chunks and their hashes in a database. The database is stored on a WAN optimization storage device. Then, instead of sending the actual data over the WAN tunnel, the FortiGate unit sends the hashes. The FortiGate unit at the other end of the tunnel receives the hashes and compares them with the hashes in its local byte caching database. If any hashes match, that data does not have to be transmitted over the WAN optimization tunnel. The data for any hashes that does not match is transferred over the tunnel and added to that byte caching database. Then the unit of application data (the file being downloaded) is reassembled and sent to its destination.

The stored byte caches are not application specific. Byte caches from a file in an email can be used to optimize downloading that same file or a similar file from a web page.

The result is less data transmitted over the WAN. Initially, byte caching may reduce performance until a large enough byte caching database is built up.

To enable byte caching, you select *Byte Caching* in a WAN optimization profile.

Byte caching cannot determine whether or not a file is compressed (for example a zip file), and caches compressed and non-compressed versions of the same file separately.

## Dynamic data chunking for byte caching

Dynamic data chunking can improve byte caching by improving detection of data chunks that are already cached in changed files or in data embedded in traffic using an unknown protocol. Dynamic data chunking is available for HTTP, CIFS and FTP.

Use the following command to enable dynamic data chunking for HTTP in the default WAN optimization profile.

```
config wanopt profile
 edit default
 config http
 set prefer-chunking dynamic
 end
```

By default dynamic data chunking is disabled and `prefer-chunking` is set to `fix`.

## WAN optimization transparent mode

WAN optimization is transparent to users. This means that with WAN optimization in place, clients connect to servers in the same way as they would without WAN optimization. However, servers receiving packets after WAN optimization “see” different source addresses depending on whether or not transparent mode is selected for WAN optimization. If transparent mode is selected, WAN optimization keeps the original source address of the packets, so servers appear to receive traffic directly from clients. Routing on the server network should be configured to route traffic with client source IP addresses from the server-side FortiGate unit to the server and back to the server-side FortiGate unit.



Some protocols, for example CIFS, may not function as expected if transparent mode is **not** selected. In most cases, for CIFS WAN optimization you should select transparent mode and make sure the server network can route traffic as described to support transparent mode.

---

If transparent mode is not selected, the source address of the packets received by servers is changed to the address of the server-side FortiGate unit interface that sends the packets to the servers. So servers appear to receive packets from the server-side FortiGate unit. Routing on the server network is simpler in this case because client addresses are not involved. All traffic appears to come from the server-side FortiGate unit and not from individual clients.



Do not confuse WAN optimization transparent mode with FortiGate transparent mode. WAN optimization transparent mode is similar to source NAT. FortiGate Transparent mode is a system setting that controls how the FortiGate unit (or a VDOM) processes traffic.

---

## FortiClient WAN optimization

PCs running the FortiClient application are client-side peers that initiate WAN optimization tunnels with server-side peer FortiGate units. However, you can have an ever-changing number of FortiClient peers with IP addresses that also change regularly. To avoid maintaining a list of such peers, you can instead configure WAN optimization to accept any peer and use authentication to identify FortiClient peers.

Together, the WAN optimization peers apply the WAN optimization features to optimize the traffic flow over the WAN between the clients and servers. WAN optimization reduces bandwidth requirements, increases throughput, reduces latency, offloads SSL encryption/decryption and improves privacy for traffic on the WAN.

## Operating modes and VDOMs

To use WAN optimization, the FortiGate units can operate in either NAT/Route or Transparent mode. The client-side and server-side FortiGate units do not have to be operating in the same mode.

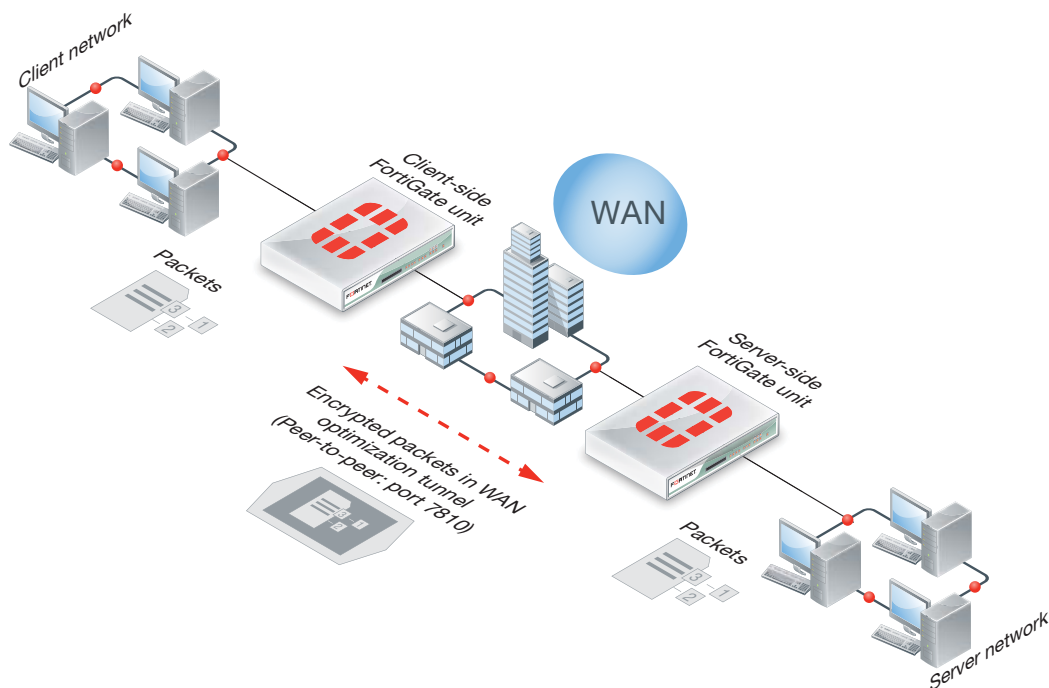
As well, the FortiGate units can be configured for multiple virtual domain (VDOM) operation. You configure WAN optimization for each VDOM and configure one or both of the units to operate with multiple VDOMs enabled.

If a FortiGate unit or VDOM is operating in Transparent mode with WAN optimization enabled, WAN optimization uses the management IP address as the peer IP address of the FortiGate unit instead of the address of an interface.

## WAN optimization tunnels

All optimized traffic passes between the FortiGate units or between a FortiClient peer and a FortiGate unit over a WAN optimization tunnel. Traffic in the tunnel can be sent in plain text or encrypted using AES-128bit-CBC SSL.

**Figure 396:** WAN optimization tunnels



Both plain text and the encrypted tunnels use TCP destination port 7810.

Before a tunnel can be started, the peers must be configured to authenticate with each other. Then, the client-side peer attempts to start a WAN optimization tunnel with the server-side peer. Once the peers authenticate with each other, they bring up the tunnel and WAN optimization communication over the tunnel starts. After a tunnel has been established, multiple WAN optimization sessions can start and stop between peers without restarting the tunnel.



## Tunnel sharing

You can use the `tunnel-sharing` WAN optimization profile CLI keyword to configure tunnel sharing for WAN optimization rules. Tunnel sharing means multiple WAN optimization sessions share the same tunnel. Tunnel sharing can improve performance by reducing the number of WAN optimization tunnels between FortiGate units. Having fewer tunnels means less data to manage. Also, tunnel setup requires more than one exchange of information between the ends of the tunnel. Once the tunnel is set up, each new session that shares the tunnel avoids tunnel setup delays.

Tunnel sharing also uses bandwidth more efficiently by reducing the chances that small packets will be sent down the tunnel. Processing small packets reduces network throughput, so reducing the number of small packets improves performance. A shared tunnel can combine all the data from the sessions being processed by the tunnel and send the data together. For example, suppose a FortiGate unit is processing five WAN optimization sessions and each session has 100 bytes to send. If these sessions use a shared tunnel, WAN optimization combines the packets from all five sessions into one 500-byte packet. If each session uses its own private tunnel, five 100-byte packets will be sent instead. Each packet also requires a TCP ACK reply. The combined packet in the shared tunnel requires one TCP ACK packet. The separate packets in the private tunnels require five.

Use the following command to configure tunnel sharing for HTTP traffic in a WAN optimization profile.

```
config wanopt profile
 edit default
 config http
 set tunnel-sharing {express-shared | private | shared}
 end
```

Tunnel sharing is not always recommended and may not always be the best practice. Aggressive and non-aggressive protocols should not share the same tunnel. An aggressive protocol can be defined as a protocol that is able to get more bandwidth than a non-aggressive protocol. (The aggressive protocols can “starve” the non-aggressive protocols.) HTTP and FTP are considered aggressive protocols. If aggressive and non-aggressive protocols share the same tunnel, the aggressive protocols may take all of the available bandwidth. As a result, the performance of less aggressive protocols could be reduced. To avoid this problem, rules for HTTP and FTP traffic should have their own tunnel. To do this, set `tunnel-sharing` to `private` for WAN optimization rules that accept HTTP or FTP traffic.

It is also useful to set `tunnel-sharing` to `express-shared` for applications, such as Telnet, that are very interactive but not aggressive. Express sharing optimizes tunnel sharing for Telnet and other interactive applications where latency or delays would seriously affect the user’s experience with the protocol.

Set `tunnel-sharing` to `shared` for applications that are not aggressive and are not sensitive to latency or delays. WAN optimization rules set to `shared` and `express-shared` can share the same tunnel.

## WAN optimization and user and device identity policies, load balancing and traffic shaping

Please note the following about WAN optimization and firewall policies:

- WAN optimization is not compatible with firewall load balancing.
- WAN optimization is compatible with source and destination NAT options in firewall policies (including firewall virtual IPs). If a virtual IP is added to a policy the traffic that exits the WAN

optimization tunnel has its destination address changed to the virtual IPs mapped to IP address and port.

- WAN optimization is compatible with user identity-based and device identity security policies. If a session is allowed after authentication or device identification the session can be optimized.

## Traffic shaping

Traffic shaping works for WAN optimization traffic that is not in a WAN optimization tunnel. So traffic accepted by a WAN optimization security policy on a client-side FortiGate unit can be shaped on ingress. However, when the traffic enters the WAN optimization tunnel, traffic shaping is not applied.

In manual mode:

- Traffic shaping works as expected on the client-side FortiGate unit.
- Traffic shaping cannot be applied to traffic on the server-side FortiGate unit.

In active-passive mode:

- Traffic shaping works as expected on the client-side FortiGate unit.
- If transparent mode is enabled in the WAN optimization profile, traffic shaping also works as expected on the server-side FortiGate unit.
- If transparent mode is not enabled, traffic shaping works partially on the server-side FortiGate unit.

## WAN optimization and HA

You can configure WAN optimization on a FortiGate HA cluster. The recommended best practice HA configuration for WAN optimization is active-passive mode. When the cluster is operating, all WAN optimization sessions are processed by the primary unit only. Even if the cluster is operating in active-active mode, HA does not load-balance WAN optimization sessions.

You can also form a WAN optimization tunnel between a cluster and a standalone FortiGate unit or between two clusters.

In a cluster, only the primary unit stores the byte cache database. This database is not synchronized to the subordinate units. So, after a failover, the new primary unit must rebuild its byte cache. Rebuilding the byte cache can happen relatively quickly because the new primary unit gets byte cache data from the other FortiGate unit that it is participating with in WAN optimization tunnels.

## WAN optimization, web caching and memory usage

To accelerate and optimize disk access and to provide better throughput and less latency FortiOS WAN optimization uses provisioned memory to reduce disk I/O and increase disk I/O efficiency. In addition, WAN optimization requires a small amount of additional memory per session for comprehensive flow control logic and efficient traffic forwarding.

When WAN optimization is enabled you will see a reduction in available memory. The reduction increases when more WAN optimization sessions are being processed. If you are thinking of enabling WAN optimization on an operating FortiGate unit, make sure its memory usage is not maxed out during high traffic periods.

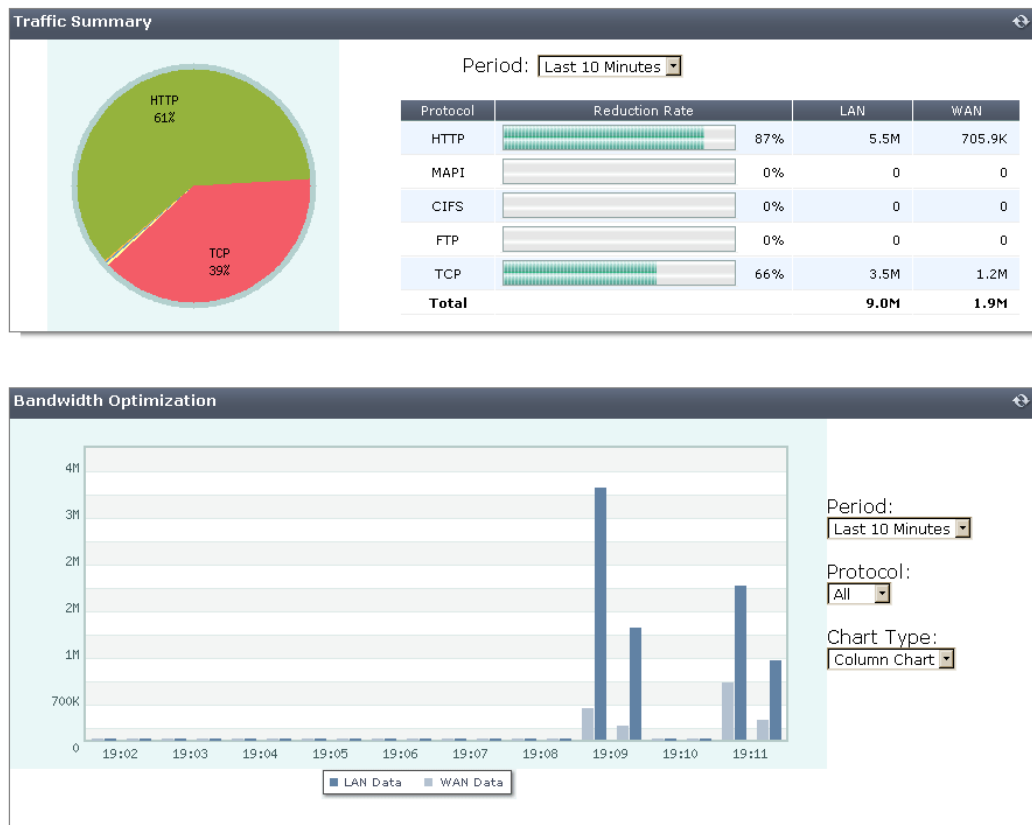
In addition to using the system dashboard to see the current memory usage you can use the `get test wad 1` command to see how much memory is currently being used by WAN optimization. See “[get test {wa\\_cs | wa\\_dbd | wad | wad\\_diskd | wccpd} <test\\_level>](#)” on [page 2719](#) for more information.

## Monitoring WAN optimization performance

Using WAN optimization monitoring, you can confirm that a FortiGate unit is optimizing traffic and view estimates of the amount of bandwidth saved. The WAN optimization monitor presents collected log information in a graphical format to show network traffic summary and bandwidth optimization information.

To view the WAN optimization monitor, go to *WAN Opt. & Cache > Monitor > WAN Opt Monitor*.

**Figure 397:**WAN optimization monitor



### Traffic Summary

The traffic summary shows how WAN optimization is reducing the amount of traffic on the WAN for each WAN optimization protocol by showing the traffic reduction rate as a percentage of the total traffic. The traffic summary also shows the amount of WAN and LAN traffic. If WAN optimization is being effective the amount of WAN traffic should be lower than the amount of LAN traffic.

You can use the refresh icon to update the traffic summary display at any time. You can also set the amount of time for which the traffic summary shows data. The time period can vary from the last 10 minutes to the last month.

## Bandwidth Optimization

This section shows network bandwidth optimization per time period. A line or column chart compares an application's pre-optimized size (LAN data) with its optimized size (WAN data). You can select the chart type, the monitoring time period, and the protocol for which to display data. If WAN optimization is being effective the WAN bandwidth should be lower than the LAN bandwidth.

## WAN optimization configuration summary

This section describes:

- [client-side configuration summary](#)
- [server-side configuration summary](#)

## client-side configuration summary

### WAN optimization profile

Enter the following command to view WAN optimization profile CLI options:

```
tree wanopt profile
-- [profile] --*name (36)
 |- transparent
 |- comments
 |- auth-group (36)
 |- <http> -- status
 |- secure-tunnel
 |- byte-caching
 |- prefer-chunking
 |- tunnel-sharing
 |- log-traffic
 |- port
 |- ssl
 |- ssl-port
 |- unknown-http-version
 +- tunnel-non-http
 |- <cifs> -- status
 |- secure-tunnel
 |- byte-caching
 |- prefer-chunking
 |- tunnel-sharing
 |- log-traffic
 +- port
 |- <mapi> -- status
 |- secure-tunnel
 |- byte-caching
 |- tunnel-sharing
 |- log-traffic
 +- port
 |- <ftp> -- status
 |- secure-tunnel
 |- byte-caching
 |- prefer-chunking
 |- tunnel-sharing
 |- log-traffic
 +- port
 +- <tcp> -- status
 |- secure-tunnel
 |- byte-caching
 |- byte-caching-opt
 |- tunnel-sharing
 |- log-traffic
 |- port
 |- ssl
```

+ - ssl-port

## Local host ID and peer settings

```
config wanopt settings
 set host-id "client"
end
config wanopt peer
 edit "server"
 set ip 10.10.2.82
end
```

## Security policies

Two client-side WAN optimization security policy configurations are possible. One for active-passive WAN optimization and one for manual WAN optimization.

### Active/passive mode on the client-side

```
config firewall policy
 edit 2
 set srcintf "internal"
 set dstintf "wan1"
 set srcaddr "all"
 set dstaddr "all"
 set action accept
 set schedule "always"
 set service "ANY"
 set utm-status enable <<< enable security profiles
 set av-profile default <<< select an antivirus profile
 set profile-protocol-options default
 set wanopt enable <<< enable WAN optimization
 set wanopt-detection active <<< set the mode to active/passive
 set wanopt-profile "default" <<< select the wanopt profile
 next
end
```

## Manual mode on the client-side

```
config firewall policy
 edit 2
 set srcintf "internal"
 set dstintf "wan1"
 set srcaddr "all"
 set dstaddr "all"
 set action accept
 set schedule "always"
 set service "ANY"
 set utm-status enable <<< enable security profiles
 set av-profile default <<< select an antivirus profile
 set profile-protocol-options default
 set wanopt enable <<< enable WAN optimization
 set wanopt-detection off <<< sets the mode to manual
 set wanopt-profile "default" <<< select the wanopt profile
 set wanopt-peer "server" <<< set the only peer to do wanopt with
 (required for manual mode)
 next
end
```

## server-side configuration summary

### Local host ID and peer settings

```
config wanopt settings
 set host-id "server"
end
config wanopt peer
 edit "client"
 set ip 10.10.2.81
 end
```

### Security policies

Two server-side WAN optimization security policy configurations are possible. One for active-passive WAN optimization and one for manual WAN optimization.

### Active/passive mode on server-side

```
config firewall policy
 edit 2 <<< the passive mode policy
 set srcintf "wan1"
 set dstintf "internal"
 set srcaddr "all"
 set dstaddr "all"
 set action accept
 set schedule "always"
 set service "ANY"
 set utm-status enable <<< enable security profiles
 set av-profile default <<< select an antivirus profile
 set profile-protocol-options default
 set wanopt enable
 set wanopt-detection passive
 set wanopt-passive-opt transparent
 next
 edit 3 <<< policy that accepts wanopt tunnel connections from the
 server
 set srcintf "wanopt" <<< wanopt tunnel interface
 set dstintf "internal"
 set srcaddr "all"
 set dstaddr "all"
 set action accept
 set schedule "always"
 set service "ANY"
 next
end
```

### Manual mode on server-side

```
configure firewall policy
 edit 3 <<< wanopt tunnel policy
 set srcintf "wanopt" <<< wanopt tunnel interface
 set dstintf "internal"
 set srcaddr "all"
 set dstaddr "all"
 set action accept
 set schedule "always"
 set service "ANY"
 set utm-status enable <<< enable security profiles
 set av-profile default <<< select an antivirus profile
 set profile-protocol-options default
 next
end
```



## Best practices

This is a short list of WAN optimization and explicit proxy best practices.

- WAN optimization tunnel sharing is recommended for similar types of WAN optimization traffic. However, tunnel sharing for different types of traffic is not recommended. For example, aggressive and non-aggressive protocols should not share the same tunnel. See [“Tunnel sharing” on page 2601](#).
- Active-passive HA is the recommended HA configuration for WAN optimization. See [“WAN optimization and HA” on page 2602](#).
- Configure WAN optimization authentication with specific peers. Accepting any peer is not recommended as this can be less secure. See [“Accepting any peers” on page 2610](#).
- Set the explicit HTTP proxy *Default Policy Action* to *Deny*. This means that a security policy is required to use the explicit web proxy. See [“Explicit web proxy configuration overview” on page 2667](#).
- Set the explicit FTP proxy *Default Policy Action* to *Deny*. This means that a security policy is required to use the explicit FTP proxy. See [“Explicit FTP proxy configuration overview” on page 2692](#).

# Peers and authentication groups

All communication between WAN optimization peers begins with one WAN optimization peer (or client-side FortiGate unit) sending a WAN optimization tunnel request to another peer (or server-side FortiGate unit). During this process, the WAN optimization peers identify and optionally authenticate each other.

This chapter describes:

- [Basic WAN optimization peer requirements](#)
- [How FortiGate units process tunnel requests for peer authentication](#)
- [Configuring peers](#)
- [Configuring authentication groups](#)
- [Secure tunneling](#)
- [Monitoring WAN optimization peer performance](#)

## Basic WAN optimization peer requirements

WAN optimization requires the following configuration on each peer. For information about configuring local and peer host IDs, see [“Configuring peers” on page 2611](#).

- The peer must have a unique host ID.
- Unless authentication groups are used, peers authenticate each other using host ID values. Do not leave the local host ID at its default value.
- The peer must know the host IDs and IP addresses of all of the other peers that it can start WAN optimization tunnels with. This does not apply if you use authentication groups that accept all peers.
- All peers must have the same local certificate installed on their FortiGate units if the units authenticate by local certificate (see [“Certificate-based authentication” on page 532](#)). Similarly, if the units authenticate by pre-shared key (password), administrators must know the password. The type of authentication is selected in the authentication group. This applies only if you use authentication groups.

## Accepting any peers

Strictly speaking, you do not need to add peers. Instead you can configure authentication groups that accept any peer. However, for this to work, both peers must have the same authentication group (with the same name) and both peers must have the same certificate or pre-shared key.

Accepting any peer is useful if you have many peers or if peer IP addresses change. For example, you could have many travelling FortiClient peers with IP addresses that are always changing as the users travel to different customer sites. This configuration is also useful if you have FortiGate units with dynamic external IP addresses (using DHCP or PPPoE). For most other situations, this method is not recommended and is not a best practice as it is less secure than accepting defined peers or a single peer. For more information, see [“Configuring authentication groups” on page 2612](#).

## How FortiGate units process tunnel requests for peer authentication

When a client-side FortiGate unit attempts to start a WAN optimization tunnel with a peer server-side FortiGate unit, the tunnel request includes the following information:

- the client-side local host ID
- the name of an authentication group, if included in the rule that initiates the tunnel
- if an authentication group is used, the authentication method it specifies: pre-shared key or certificate
- the type of tunnel (secure or not).

For information about configuring the local host ID, peers and authentication groups, see [“Configuring peers” on page 2611](#) and [“Configuring authentication groups” on page 2612](#).

The authentication group is optional unless the tunnel is a secure tunnel. For more information, see [“Secure tunneling” on page 2615](#).

If the tunnel request includes an authentication group, the authentication will be based on the settings of this group as follows:

- The server-side FortiGate unit searches its own configuration for the name of the authentication group in the tunnel request. If no match is found, the authentication fails.
- If a match is found, the server-side FortiGate unit compares the authentication method in the client and server authentication groups. If the methods do not match, the authentication fails.
- If the authentication methods match, the server-side FortiGate unit tests the peer acceptance settings in its copy of the authentication group.
- If the setting is *Accept Any Peer*, the authentication is successful.
- If the setting is *Specify Peer*, the server-side FortiGate unit compares the client-side local host ID in the tunnel request with the peer name in the server-side authentication group. If the names match, authentication is successful. If a match is not found, authentication fails.
- If the setting is *Accept Defined Peers*, the server-side FortiGate unit compares the client-side local host ID in the tunnel request with the server-side peer list. If a match is found, authentication is successful. If a match is not found, authentication fails.

If the tunnel request does not include an authentication group, authentication will be based on the client-side local host ID in the tunnel request. The server-side FortiGate unit searches its peer list to match the client-side local host ID in the tunnel request. If a match is found, authentication is successful. If a match is not found, authentication fails.

If the server-side FortiGate unit successfully authenticates the tunnel request, the server-side FortiGate unit sends back a tunnel setup response message. This message includes the server-side local host ID and the authentication group that matches the one in the tunnel request.

The client-side FortiGate unit then performs the same authentication procedure as the server-side FortiGate unit did. If both sides succeed, tunnel setup continues.

## Configuring peers

When you configure peers, you first need to add the local host ID that identifies the FortiGate unit for WAN optimization and then add the peer host ID and IP address of each FortiGate unit with which a FortiGate unit can create WAN optimization tunnels.

### To configure WAN optimization peers - web-based manager

1. Go to *Wan Opt. & Cache > WAN Opt. Peer > Peer*.

- 2 For *Local Host ID*, enter the local host ID of **this** FortiGate unit and select *Apply*. If you add this FortiGate unit as a peer to another FortiGate unit, use this ID as its **peer** host ID.  
The local or host ID can contain up to 25 characters and can include spaces.
- 3 Select *Create New* to add a new peer.
- 4 For *Peer Host ID*, enter the peer host ID of the peer FortiGate unit. This is the local host ID added to the peer FortiGate unit.
- 5 For *IP Address*, add the IP address of the peer FortiGate unit. This is the source IP address of tunnel requests sent by the peer, usually the IP address of the FortiGate interface connected to the WAN.
- 6 Select *OK*.

### To configure WAN optimization peers - CLI

In this example, the local host ID is named `HQ_Peer` and has an IP address of `172.20.120.100`. Three peers are added, but you can add any number of peers that are on the WAN.

1. Enter the following command to set the local host ID to `HQ_Peer`.

```
config wanopt settings
 set host-id HQ_peer
end
```

- 2 Enter the following commands to add three peers.

```
config wanopt peer
 edit Wan_opt_peer_1
 set ip 172.20.120.100
 next
 edit Wan_opt_peer_2
 set ip 172.30.120.100
 next
 edit Wan_opt_peer_3
 set ip 172.40.120.100
end
```

## Configuring authentication groups

You need to add authentication groups to support authentication and secure tunneling between WAN optimization peers.

To perform authentication, WAN optimization peers use a certificate or a pre-shared key added to an authentication group so they can identify each other before forming a WAN optimization tunnel. Both peers must have an authentication group with the same name and settings. You add the authentication group to a peer-to-peer or active rule on the client-side FortiGate unit. When the server-side FortiGate unit receives a tunnel start request from the client-side FortiGate unit that includes an authentication group, the server-side FortiGate unit finds an authentication group in its configuration with the same name. If both authentication groups have the same certificate or pre-shared key, the peers can authenticate and set up the tunnel.

Authentication groups are also required for secure tunneling. See [“Secure tunneling” on page 2615](#).

To add authentication groups, go to *WAN Opt. & Cache > WAN Opt. Peer > Authentication Group*.

## To add an authentication group - web-based manager

Use the following steps to add any kind of authentication group. It is assumed that if you are using a local certificate to authenticate, it is already added to the FortiGate unit. For more information about FortiGate units and certificates, see the [FortiGate Certificate Management Guide](#).

1. Go to *Wan Opt. & Cache > WAN Opt. Peer > Authentication Group*.

2. Select *Create New*.

3. Add a *Name* for the authentication group.

You will select this name when you add the authentication group to a WAN optimization rule.

4. Select the *Authentication Method*.

Select *Certificate* if you want to use a certificate to authenticate and encrypt WAN optimization tunnels. You must select a local certificate that has been added to this FortiGate unit. (To add a local certificate, go to *System > Certificates > Local Certificates*.) Other FortiGate units that participate in WAN optimization tunnels with this FortiGate unit must have an authentication group with the same name and certificate.

Select *Pre-shared key* if you want to use a pre-shared key or password to authenticate and encrypt WAN optimization tunnels. You must add the *Password* (or pre-shared key) used by the authentication group. Other FortiGate units that participate in WAN optimization tunnels with this FortiGate unit must have an authentication group with the same name and password. The password must contain at least 6 printable characters and should be known only by network administrators. For optimum protection against currently known attacks, the key should consist of a minimum of 16 randomly chosen alphanumeric characters.

5. Configure *Peer Acceptance* for the authentication group.

Select *Accept Any Peer* if you do not know the peer host IDs or IP addresses of the peers that will use this authentication group. This setting is most often used for WAN optimization with the FortiClient application or with FortiGate units that do not have static IP addresses, for example units that use DHCP.

Select *Accept Defined Peers* if you want to authenticate with peers added to the peer list only.

Select *Specify Peer* and select one of the peers added to the peer list to authenticate with the selected peer only.

For more information, see [“Configuring peers” on page 2611](#).

6. Select *OK*.

7. Add the authentication group to a WAN optimization rule to apply the authentication settings in the authentication group to the rule.

## To add an authentication group that uses a certificate- CLI

Enter the following command to add an authentication group that uses a certificate and can authenticate all peers added to the FortiGate unit configuration.

In this example, the authentication group is named `auth_grp_1` and uses a certificate named `Example_Cert`.

```
config wanopt auth-group
 edit auth_grp_1
 set auth-method cert
 set cert Example_Cert
 set peer-accept defined
 end
```

### To add an authentication group that uses a pre-shared key - CLI

Enter the following command to add an authentication group that uses a pre-shared key and can authenticate only the peer added to the authentication group.

In this example, the authentication group is named `auth_peer`, the peer that the group can authenticate is named `Server_net`, and the authentication group uses `123456` as the pre-shared key. In practice you should use a more secure pre-shared key.

```
config wanopt auth-group
 edit auth_peer
 set auth-method psk
 set psk 123456
 set peer-accept one
 set peer Server_net
 end
```

### To add an authentication group that accepts WAN optimization connections from any peer - web-based manager

Add an authentication group that accepts any peer for situations where you do not have the *Peer Host IDs or IP Addresses* of the peers that you want to perform WAN optimization with. This setting is most often used for WAN optimization with the FortiClient application or with FortiGate units that do not have static IP addresses, for example units that use DHCP. An authentication group that accepts any peer is less secure than an authentication group that accepts defined peers or a single peer.

The example below sets the authentication method to *Pre-shared key*. You must add the same password to all FortiGate units using this authentication group.

1. Go to *Wan Opt. & Cache > WAN Opt. Peer > Authentication Group*.
2. Select *Create New* to add a new authentication group.
3. Configure the authentication group:

<b>Name</b>	Specify any name.
<b>Authentication Method</b>	Pre-shared key
<b>Password</b>	Enter a pre-shared key.
<b>Peer Acceptance</b>	Accept Any Peer

### To add an authentication group that accepts WAN optimization connections from any peer - CLI

In this example, the authentication group is named `auth_grp_1`. It uses a certificate named `WAN_Cert` and accepts any peer.

```
config wanopt auth-group
 edit auth_grp_1
 set auth-method cert
 set cert WAN_Cert
 set peer-accept any
 end
```

## Secure tunneling

You can configure WAN optimization rules to use AES-128bit-CBC SSL to encrypt the traffic in the WAN optimization tunnel. WAN optimization uses FortiASIC acceleration to accelerate SSL decryption and encryption of the secure tunnel. Peer-to-peer secure tunnels use the same TCP port as non-secure peer-to-peer tunnels (TCP port 7810).

To use secure tunneling, you must select *Enable Secure Tunnel* in a WAN optimization rule and add an authentication group. The authentication group specifies the certificate or pre-shared key used to set up the secure tunnel. The *Peer Acceptance* setting of the authentication group does not affect secure tunneling.

The FortiGate units at each end of the secure tunnel must have the same authentication group with the same name and the same configuration, including the same pre-shared key or certificate. To use certificates you must install the same certificate on both FortiGate units.

For active-passive WAN optimization you can select *Enable Secure Tunnel* only in the active rule. In peer-to-peer WAN optimization you select *Enable Secure Tunnel* in the WAN optimization rule on both FortiGate units. For information about active-passive and peer-to-peer WAN optimization, see [“Configuring WAN optimization” on page 2590](#).

For a secure tunneling configuration example, see [“Example: Adding secure tunneling to an active-passive WAN optimization configuration” on page 2634](#).

## Monitoring WAN optimization peer performance

The WAN optimization peer monitor lists all of the WAN optimization peers that a FortiGate unit can perform WAN optimization with. These include peers manually added to the configuration as well as discovered peers.

The monitor lists each peer’s name, IP address, and peer type. The peer type indicates whether the peer was manually added or discovered. To show WAN optimization performance, for each peer the monitor lists the percent of traffic reduced by the peer in client-side WAN optimization configurations and in server-side configurations (also called gateway configurations).

To view the peer monitor, go to *WAN Opt. & Cache > Monitor > Peer Monitor*.

# Configuration examples

This chapter provides the basic examples to illustrate WAN optimization configurations introduced in the previous chapters. This chapter contains the following sections:

- [Example: Basic manual \(peer-to-peer\) WAN optimization configuration](#)
- [Example: Active-passive WAN optimization](#)
- [Example: Adding secure tunneling to an active-passive WAN optimization configuration](#)

## Example: Basic manual (peer-to-peer) WAN optimization configuration

In a manual (peer to peer) configuration the WAN optimization tunnel can be set up between one client-side FortiGate unit and one server-side FortiGate unit. The peer ID of the server-side FortiGate unit is added to the client-side WAN optimization policy. When the client-side FortiGate unit initiates a tunnel with the server-side FortiGate unit, the packets that initiate the tunnel include information that allows the server-side FortiGate unit to determine that it is a manual tunnel request. The server-side FortiGate unit does not require a WAN optimization profile; you just need to add the client peer host ID and IP address to the server-side FortiGate unit peer list and a security policy to accept WAN optimization tunnel connections.

In a manual WAN optimization configuration, you create a manual WAN optimization security policy on the client-side FortiGate unit. To do this you must use the CLI to set `wanopt-detection` to `off` and to add the peer host ID of the server-side FortiGate unit to the WAN optimization security policy.


### Network topology and assumptions

This example configuration includes a client-side FortiGate unit called Peer-Fgt-1 with a WAN IP address of 172.20.34.12. This unit is in front of a network with IP address 172.20.120.0. The server-side FortiGate unit is called Peer-Fgt-2 with a WAN IP address of 192.168.30.12. This unit is in front of a web server network with IP address 192.168.10.0.

This example customizes the default WAN optimization profile on the client-side FortiGate unit and adds it to the WAN optimization policy. You can also create a new WAN optimization profile. This example also applies virus scanning and application control to WAN optimization traffic on the server-side FortiGate unit by adding the default Antivirus Profile and default Application Control sensor to the WAN optimization security policy.



**Figure 398:**Example manual (peer-to-peer) topology



## General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Configure the client-side FortiGate unit:
  - Add peers.
  - Configure the default WAN optimization profile to optimize HTTP traffic.
  - Add a manual WAN optimization security policy.
2. Configure the server-side FortiGate unit:
  - Add peers.
  - Add a WAN optimization tunnel policy.

## Configuring basic peer-to-peer WAN optimization - web-based manager

Use the following steps to configure the example configuration from the web-based manager.

### To configure the client-side FortiGate unit

1. Go to *WAN Opt. & Cache > WAN Opt. Peer > Peer* and enter a *Local Host ID* for the client-side FortiGate unit:

<b>Local Host ID</b>	Client-Fgt
----------------------	------------

2. Select *Apply*.

3. Select *Create New* and add the server-side FortiGate unit *Peer Host ID* and *IP Address* for the server-side FortiGate:

<b>Peer Host ID</b>	Server-Fgt
<b>IP Address</b>	192.168.30.12

4. Select *OK*.
5. Go to *Firewall Objects > Address > Address* and select *Create New* to add a firewall address for the client network.

<b>Category</b>	Address
<b>Name</b>	Client-Net
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	172.20.120.0/24
<b>Interface</b>	port1

6. Select *Create New* to add a firewall address for the web server network.

<b>Category</b>	Address
<b>Address Name</b>	Web-Server-Net
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	192.168.10.0/24
<b>Interface</b>	port2

7. Go to *WAN Opt. & Cache > WAN Opt. Profile > Profile* and edit the default profile.
8. Select *Transparent Mode*.
9. Under *Protocol*, select *HTTP* and for *HTTP* select *Byte Caching*. Leave the *HTTP Port* set to 80.
10. Select *Apply* to save your changes.
11. Go to *Policy > Policy > Policy* and add a WAN optimization security policy to the client-side FortiGate unit that accepts traffic to be optimized:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	port1
<b>Source Address</b>	all
<b>Outgoing Interface</b>	port2
<b>Destination Address</b>	all
<b>Schedule</b>	always

<b>Service</b>	ALL
<b>Action</b>	ACCEPT

- Under *Security Profiles* turn on Antivirus and select the *default* antivirus profile.
- Turn on *Application Control* and select the *default* application control sensor.
- Select *Enable WAN Optimization* and configure the following settings:

<b>Enable WAN Optimization</b>	active
<b>Profile</b>	default

- Select *OK*.
- Edit the policy from the CLI to turn off `wanopt-detection`, add the peer ID of the server-side FortiGate unit, and the default WAN optimization profile. The following example assumes the ID of the policy is 5:

```
config firewall policy
 edit 5
 set wanopt-detection off
 set wanopt-peer Server-Fgt
 set wanopt-profile default
 end
```

When you set the detection mode to off the policy becomes a manual mode WAN optimization policy. On the web-based manager the WAN optimization part of the policy changes to the following:

<b>Enable WAN Optimization</b>	Manual (Profile: default, Peer: Peer-Fgt-2)
--------------------------------	---------------------------------------------

#### To configure the server-side FortiGate unit

- Go to *WAN Opt. & Cache > WAN Opt. Peer > Peer* and enter a *Local Host ID* for the server-side FortiGate unit:

<b>Local Host ID</b>	Server-Fgt
----------------------	------------

- Select *Apply*.
- Select *Create New* and add a *Peer Host ID* and the *IP Address* for the client-side FortiGate unit:

<b>Peer Host ID</b>	Client-Fgt
<b>IP Address</b>	172.20.34.12

- Select *OK*.
- Go to *Policy > Policy > Policy* and select *Create New* to add a security policy to accept WAN optimization tunnel connections.

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	wanopt

<b>Source Address</b>	all
<b>Outgoing Interface</b>	port1
<b>Destination Address</b>	all
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT

## Configuring basic peer-to-peer WAN optimization - CLI

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit CLI.

### To configure the client-side FortiGate unit

1. Add the Local Host ID to the client-side FortiGate configuration:

```
config wanopt settings
 set host-id Client-Fgt
end
```

2. Add the server-side Local Host ID to the client-side peer list:

```
config wanopt peer
 edit Server-Fgt
 set ip 192.168.30.12
 end
```

3. Add a firewall address for the client network.

```
config firewall address
 edit Client-Net
 set type ipmask
 set subnet 172.20.120.0 255.255.255.0
 set associated-interface port1
 end
```

4. Add a firewall address for the web server network.

```
config firewall address
 edit Web-Server-Net
 set type ipmask
 set subnet 192.168.10.0 255.255.255.0
 set associated-interface port2
 end
```

5. Edit the default WAN optimization profile, select transparent mode, enable HTTP WAN optimization and enable byte caching for HTTP. Leave the HTTP Port set to 80.

```
config wanopt profile
 edit default
 set transparent enable
 config http
 set status enable
 set byte-caching enable
 end
 end
```

6. Add a WAN optimization security policy to the client-side FortiGate unit to accept the traffic to be optimized:

```
config firewall policy
 edit 0
 set srcintf port1
 set dstintf port2
 set srcaddr all
 set dstaddr all
 set action accept
 set service ALL
 set schedule always
 set utm-status enable
 set av-profile default
 set application-list default
 set profile-protocol-options default
 set wanopt enable
 set wanopt-profile default
 set wanopt-detection off
 set wanopt-peer Server-Fgt
 end
```

#### **To configure the server-side FortiGate unit**

1. Add the Local Host ID to the server-side FortiGate configuration:

```
config wanopt settings
 set host-id Server-Fgt
end
```

2. Add the client-side Local Host ID to the server-side peer list:

```
config wanopt peer
 edit Client-Fgt
 set ip 192.168.30.12
 end
```

### 3. Add a WAN optimization tunnel policy.

```
config firewall policy
 edit 0
 set srcintf wanopt
 set dstintf port1
 set srcaddr all
 set dstaddr all
 set action accept
 set service ALL
 set schedule always
 end
```

## Testing and troubleshooting the configuration

To test the configuration attempt to start a web browsing session between the client network and the web server network. For example, from a PC on the client network browse to the IP address of a web server on the web server network, for example `http://192.168.10.100`. Even though this address is not on the client network you should be able to connect to this web server over the WAN optimization tunnel.

If you can connect, check WAN optimization monitoring (go to *WAN Opt. & Cache > Monitor > Monitor*). If WAN optimization has been forwarding the traffic the WAN optimization monitor should show the protocol that has been optimized (in this case HTTP) and the reduction rate in WAN bandwidth usage.

If you can't connect you can try the following to diagnose the problem:

- Review your configuration and make sure all details such as address ranges, peer names, and IP addresses are correct.
- Confirm that the security policy on the client-side FortiGate unit is accepting traffic for the 192.168.10.0 network. You can do this by checking the policy monitor (*Policy > Monitor > Policy Monitor*). Look for sessions that use the policy ID of this policy.
- Check routing on the FortiGate units and on the client and web server networks to make sure packets can be forwarded as required. The FortiGate units must be able to communicate with each other, routing on the client network must allow packets destined for the web server network to be received by the client-side FortiGate unit, and packets from the server-side FortiGate unit must be able to reach the web servers.

You can use the following `get` and `diagnose` commands to display information about how WAN optimization is operating

Enter the following command on the client-side FortiGate unit to display WAN optimization tunnel protocol statistics. The http tunnel and tcp tunnel parts of the command output below shows that WAN optimization has been processing HTTP and TCP packets.

```
get test wad 11
wad tunnel protocol stats:
 http tunnel
 bytes_in=1751767 bytes_out=325468
 ftp tunnel
 bytes_in=0 bytes_out=0
 cifs tunnel
 bytes_in=0 bytes_out=0
 mapi tunnel
 bytes_in=0 bytes_out=0
 tcp tunnel
 bytes_in=3182253 bytes_out=200702
 maintenance tunnel
 bytes_in=11800 bytes_out=15052
```

Enter the following command to display the current WAN optimization peers. You can use this command to make sure all peers are configured correctly. The command output for the client-side FortiGate unit shows one peer with IP address 192.168.20.1, peer name Web-servers, and with 10 active tunnels.

```
get test wad 26
name: Web-servers, vd: 0, ip: 192.168.20.1 ref: 1 type:manual
 traffic:
client: LAN in:0, LAN out:0, WAN in:0, WAN out:0
gateway: LAN in:0, LAN out:0, WAN in:0, WAN out:0
client 0x40e2b4cc, server 0x40e2b4ec
 version=0 tunnels(active/connecting/failover/passive)=10/0/0/0
 ssl tunnels active/connecting/passive)=0/0/0
 sessions=0 n_retries=0 version_valid=true

total peers: 1, manual peers: 1 auto peers: 0
```

Enter the following command to list all of the running WAN optimization tunnels and display information about each one. The command output for the client-side FortiGate unit shows 10 tunnels all created by peer-to-peer WAN optimization rules (auto-detect set to off).

```
diagnose wad tunnel list
```

```
Tunnel: id=100 type=manual
 vd=0 shared=no uses=0 state=3
 peer name=Web-servers id=100 ip=192.168.30.12
 SSL-secured-tunnel=no auth-grp=
 bytes_in=348 bytes_out=384
```

```
Tunnel: id=99 type=manual
 vd=0 shared=no uses=0 state=3
 peer name=Web-servers id=99 ip=192.168.30.12
 SSL-secured-tunnel=no auth-grp=
 bytes_in=348 bytes_out=384
```

```
Tunnel: id=98 type=manual
 vd=0 shared=no uses=0 state=3
 peer name=Web-servers id=98 ip=192.168.30.12
 SSL-secured-tunnel=no auth-grp=
 bytes_in=348 bytes_out=384
```

```
Tunnel: id=39 type=manual
 vd=0 shared=no uses=0 state=3
 peer name=Web-servers id=39 ip=192.168.30.12
 SSL-secured-tunnel=no auth-grp=
 bytes_in=1068 bytes_out=1104
```

```
Tunnel: id=7 type=manual
 vd=0 shared=no uses=0 state=3
 peer name=Web-servers id=7 ip=192.168.30.12
 SSL-secured-tunnel=no auth-grp=
 bytes_in=1228 bytes_out=1264
```

```
Tunnel: id=8 type=manual
 vd=0 shared=no uses=0 state=3
 peer name=Web-servers id=8 ip=192.168.30.12
 SSL-secured-tunnel=no auth-grp=
 bytes_in=1228 bytes_out=1264
```

```
Tunnel: id=5 type=manual
 vd=0 shared=no uses=0 state=3
 peer name=Web-servers id=5 ip=192.168.30.12
 SSL-secured-tunnel=no auth-grp=
 bytes_in=1228 bytes_out=1264
```

```
Tunnel: id=4 type=manual
 vd=0 shared=no uses=0 state=3
```



```
peer name=Web-servers id=4 ip=192.168.30.12
SSL-secured-tunnel=no auth-grp=
bytes_in=1228 bytes_out=1264
```

```
Tunnel: id=1 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web-servers id=1 ip=192.168.30.12
SSL-secured-tunnel=no auth-grp=
bytes_in=1228 bytes_out=1264
```

```
Tunnel: id=2 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web-servers id=2 ip=192.168.30.12
SSL-secured-tunnel=no auth-grp=
bytes_in=1228 bytes_out=1264
```

```
Tunnels total=10 manual=10 auto=0
```

## Example: Active-passive WAN optimization

In active-passive WAN optimization you add an active WAN optimization policy to the client-side FortiGate unit and you add a WAN optimization tunnel policy and a passive WAN optimization policy to the server-side FortiGate unit.

The active policy accepts the traffic to be optimized and sends it down the WAN optimization tunnel to the server-side FortiGate unit. The active policy can also apply security profiles and other features to traffic before it exits the client-side FortiGate unit.

The tunnel policy on the sever-side FortiGate unit allows the server-side FortiGate unit to form a WAN optimization tunnel with the client-side FortiGate unit. The passive WAN optimization policy is required because of the active policy on the client-side FortiGate unit. You can also use the passive policy to apply WAN optimization transparent mode and features such as security profiles, logging, traffic shaping and web caching to the traffic before it exits the server-side FortiGate unit.


### Network topology and assumptions

On the client-side FortiGate unit this example configuration includes a WAN optimization profile that optimizes CIFS, HTTP, and FTP traffic and an active WAN optimization policy. The active policy also applies virus scanning to the WAN optimization traffic.

On the server-side FortiGate unit, the passive policy applies application control to the WAN optimization traffic.

In this example, WAN optimization transparent mode is selected in the WAN optimization profile and the passive WAN optimization policy accepts this transparent mode setting. This means that the optimized packets maintain their original source and destination addresses. As a result, routing on the client network must be configured to route packets for the server network to the client-side FortiGate unit. Also the routing configuration on the server network must be able to route packets for the client network to the server-side FortiGate unit.

**Figure 399:**Example active-passive WAN optimization topology



## General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Configure the client-side FortiGate unit:
  - Add peers.
  - Add a WAN optimization profile to optimize CIFS, FTP, and HTTP traffic.
  - Add firewall addresses for the client and web server networks.
  - Add an active WAN optimization policy that applies virus scanning.
2. Configure the server-side FortiGate unit by:
  - Add peers.
  - Add firewall addresses for the client and web server networks.
  - Add a WAN optimization tunnel policy.
  - Add a passive WAN optimization policy that applies application control.

## Configuring basic active-passive WAN optimization - web-based manager

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit web-based manager.

### To configure the client-side FortiGate unit

1. Go to *WAN Opt. & Cache > WAN Opt. Peer > Peer* and enter a *Local Host ID* for the client-side FortiGate unit:

Local Host ID	Client-Fgt

2. Select *Apply*.
3. Select *Create New* and add a Peer Host ID and the *IP Address* for the server-side FortiGate unit:

<b>Peer Host ID</b>	Server-Fgt
<b>IP Address</b>	192.168.20.1

4. Select *OK*.
5. Go to *WAN Opt. & Cache > WAN Opt. Profile > Profile* and select *Create New* to add a WAN optimization profile to optimize CIFS, HTTP, and FTP traffic:

<b>Name</b>	Custom-wan-opt-pro
<b>Transparent Mode</b>	Select

6. Select the *CIFS* protocol, select *Byte Caching* and set the *Port* to 445.
7. Select the *FTP* protocol, select *Byte Caching* and set the *Port* to 21.
8. Select the *HTTP* protocol, select *Byte Caching* and set the *Port* to 80.
9. Select *OK*.
10. Go to *Firewall Objects > Address > Address* and select *Create New* to add a firewall address for the client network.

<b>Category</b>	Address
<b>Address Name</b>	Client-Net
<b>Type</b>	IP Range
<b>Subnet / IP Range</b>	172.20.120.[100-200]
<b>Interface</b>	port1

11. Select *Create New* to add a firewall address for the web server network.

<b>Category</b>	Address
<b>Address Name</b>	Web-Server-Net
<b>Type</b>	IP Range
<b>Subnet / IP Range</b>	192.168.10.0/24
<b>Interface</b>	port2

12. Go to *Policy > Policy > Policy* and add an active WAN optimization security policy:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	port1
<b>Source Address</b>	Client-Net

<b>Outgoing Interface</b>	port2
<b>Destination Address</b>	Web-Server-Net
<b>Schedule</b>	always
<b>Service</b>	HTTP FTP SMB
<b>Action</b>	ACCEPT

13. Turn on Antivirus and select the *default* antivirus profile.

14. Select *Enable WAN Optimization* and configure the following settings:

<b>Enable WAN Optimization</b>	active
<b>Profile</b>	Custom-wan-opt-pro

15. Select *OK*.

#### To configure the server-side FortiGate unit

1. Go to *WAN Opt. & Cache > WAN Opt. Peer > Peer* and enter a *Local Host ID* for the server-side FortiGate unit:

<b>Local Host ID</b>	Server-Fgt
----------------------	------------

2. Select *Apply*.

3. Select *Create New* and add a *Peer Host ID* and the *IP Address* for the client-side FortiGate unit:

<b>Peer Host ID</b>	Client-Fgt
<b>IP Address</b>	172.30.120.1

4. Select *OK*.

5. Go to *Firewall Objects > Address > Address* and select *Create New* to add a firewall address for the client network.

<b>Category</b>	Address
<b>Address Name</b>	Client-Net
<b>Type</b>	IP Range
<b>Subnet / IP Range</b>	172.20.120.[100-200]
<b>Interface</b>	port1

6. Select *Create New* to add a firewall address for the web server network.

<b>Category</b>	Address
<b>Address Name</b>	Web-Server-Net

<b>Type</b>	IP Range
<b>Subnet / IP Range</b>	192.168.10.0/24
<b>Interface</b>	port2

- Go to *Policy > Policy > Policy* and select *Create New* to add a WAN optimization tunnel policy.

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	wanopt
<b>Source Address</b>	all
<b>Outgoing Interface</b>	port1
<b>Destination Address</b>	all
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT

- Select *OK*.
- Select *Create New* to add a passive WAN optimization policy that applies application control.

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	port2
<b>Source Address</b>	Client-Net
<b>Outgoing Interface</b>	port1
<b>Destination Address</b>	Web-Server-Net
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT

- Turn on *Application Control* and select the *default* application control sensor.
- Select *Enable WAN Optimization* and configure the following settings:

<b>Enable WAN Optimization</b>	passive
<b>Passive Option</b>	default

- Select *OK*.

## Configuring basic active-passive WAN optimization - CLI

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit CLI.

### To configure the client-side FortiGate unit

1. Add the Local Host ID to the client-side FortiGate configuration:

```
config wanopt settings
 set host-id Client-Fgt
end
```

2. Add the server-side Local Host ID to the client-side peer list:

```
config wanopt peer
 edit Server-Fgt
 set ip 192.168.20.1
 end
```

3. Add a WAN optimization profile to optimize CIFS, HTTP, and FTP traffic.

```
config wanopt profile
 edit Custom-wan-opt-pro
 config cifs
 set status enable
 set byte-caching enable
 set port 445
 end
 config http
 set status enable
 set byte-caching enable
 set port 80
 end
 config ftp
 set status enable
 set byte-caching enable
 set port 21
 end
 end
end
```

4. Add a firewall address for the client network.

```
config firewall address
 edit Client-Net
 set type iprange
 set startip 172.20.120.100
 set endip 172.20.120.200
 set associated-interface port1
 end
```

5. Add a firewall address for the web server network.

```
config firewall address
 edit Web-Server-Net
 set type ipmask
 set subnet 192.168.10.0 255.255.255.0
 set associated-interface port2
 end
```

6. Add add an active WAN optimization security policy that applies virus scanning:

```
config firewall policy
 edit 0
 set srcintf port1
 set dstintf port2
 set srcaddr Client-net
 set dstaddr Web-Server-Net
 set action accept
 set service HTTP FTP SMB
 set schedule always
 set wanopt enable
 set wanopt-detection active
 set wanopt-profile Custom-wan-opt-pro
 set utm-status enable
 set av-profile default
 end
```

#### To configure the server-side FortiGate unit

1. Add the Local Host ID to the server-side FortiGate configuration:

```
config wanopt settings
 set host-id Server-Fgt
end
```

2. Add the client-side Local Host ID to the server-side peer list:

```
config wanopt peer
 edit Client-Fgt
 set ip 172.20.120.1
 end
```

3. Add a firewall address for the client network.

```
config firewall address
 edit Client-Net
 set type iprange
 set startip 172.20.120.100
 set endip 172.20.120.200
 set associated-interface port1
 end
```

4. Add a firewall address for the web server network.

```
config firewall address
 edit Web-Server-Net
 set type ipmask
 set subnet 192.168.10.0 255.255.255.0
 set associated-interface port2
 end
```

5. Add a WAN optimization tunnel policy.

```
config firewall policy
 edit 0
 set srcintf wanopt
 set dstintf port1
 set srcaddr all
 set dstaddr all
 set action accept
 set service ALL
 set schedule always
 end
```

6. Add a passive WAN optimization policy that applies application control.

```
config firewall policy
 edit 0
 set srcintf port1
 set dstintf port2
 set srcaddr Client-Net
 set dstaddr Web-Server-Net
 set action accept
 set service ALL
 set schedule always
 set wanopt enable
 set wanopt-detection passive
 set wanopt-passive-opt default
 set utm-status enable
 set application-list default
 end
```

## Testing and troubleshooting the configuration

To test the configuration attempt to start a web browsing session between the client network and the web server network. For example, from a PC on the client network browse to the IP address of a web server on the web server network, for example `http://192.168.10.100`. Even though this address is not on the client network you should be able to connect to this web server over the WAN optimization tunnel.

If you can connect, check WAN optimization monitoring (go to *WAN Opt. & Cache > Monitor > Monitor*). If WAN optimization has been forwarding the traffic the WAN optimization monitor should show the protocol that has been optimized (in this case HTTP) and the reduction rate in WAN bandwidth usage.

If you can't connect you can try the following to diagnose the problem:

- Review your configuration and make sure all details such as address ranges, peer names, and IP addresses are correct.
- Confirm that the security policy on the Client-Side FortiGate unit is accepting traffic for the 192.168.10.0 network and that this security policy does not include security profiles. You can do this by checking the FortiGate session table from the dashboard. Look for sessions that use the policy ID of this policy
- Check routing on the FortiGate units and on the client and web server networks to make sure packets can be forwarded as required. The FortiGate units must be able to communicate with each other, routing on the client network must allow packets destined for



the web server network to be received by the client-side FortiGate unit, and packets from the server-side FortiGate unit must be able to reach the web servers etc.

You can use the following `get` and `diagnose` commands to display information about how WAN optimization is operating

Enter the following command to display WAN optimization tunnel protocol statistics. The `http` tunnel and `tcp` tunnel parts of the command output below shows that WAN optimization has been processing HTTP and TCP packets.

```
get test wad 11
wad tunnel protocol stats:
 http tunnel
 bytes_in=1751767 bytes_out=325468
 ftp tunnel
 bytes_in=0 bytes_out=0
 cifs tunnel
 bytes_in=0 bytes_out=0
 mapi tunnel
 bytes_in=0 bytes_out=0
 tcp tunnel
 bytes_in=3182253 bytes_out=200702
 maintenance tunnel
 bytes_in=11800 bytes_out=15052
```

Enter the following command to display the current WAN optimization peers. You can use this command to make sure all peers are configured correctly. The command output for the client-side FortiGate unit shows one peer with IP address 192.168.20.1, peer name `Web-servers`, and with 10 active tunnels.

```
get test wad 26
peer name=Web-servers ip=192.168.20.1 vd=0 version=1
 tunnels(active/connecting/failover)=10/0/0
 sessions=0 n_retries=0 version_valid=true
```

Enter the following command to list all of the running WAN optimization tunnels and display information about each one. The command output shows 3 tunnels all created by peer-to-peer WAN optimization rules (auto-detect set to on).

```
diagnose wad tunnel list

Tunnel: id=139 type=auto
 vd=0 shared=no uses=0 state=1
 peer name= id=0 ip=unknown
 SSL-secured-tunnel=no auth-grp=test
 bytes_in=744 bytes_out=76

Tunnel: id=141 type=auto
 vd=0 shared=no uses=0 state=1
 peer name= id=0 ip=unknown
 SSL-secured-tunnel=no auth-grp=test
 bytes_in=727 bytes_out=76

Tunnel: id=142 type=auto
 vd=0 shared=no uses=0 state=1
 peer name= id=0 ip=unknown
 SSL-secured-tunnel=no auth-grp=test
 bytes_in=727 bytes_out=76

Tunnels total=3 manual=0 auto=3
```

## Example: Adding secure tunneling to an active-passive WAN optimization configuration


This example shows how to configure two FortiGate units for active-passive WAN optimization with secure tunneling. The same authentication group is added to both FortiGate units. The authentication group includes a password (or pre-shared key) and has *Peer Acceptance* set to *Accept any Peer*. An active policy is added to the client-side FortiGate unit and a passive policy to the server-side FortiGate unit. The active policy includes a profile that performs secure tunneling, optimizes HTTP traffic, and uses Transparent Mode and byte caching.

The authentication group is named *Auth-Secure-Tunnel* and the password for the pre-shared key is *2345678*. The topology for this example is shown in [Figure 400](#). This example includes web-based manager configuration steps followed by equivalent CLI configuration steps. For information about secure tunneling, see [“Secure tunneling” on page 2615](#).

### Network topology and assumptions

This example configuration includes a client-side FortiGate unit called Client-net with a WAN IP address of 172.30.120.1. This unit is in front of a network with IP address 172.20.120.0. The server-side FortiGate unit is called Web-servers and has a WAN IP address of 192.168.20.1. This unit is in front of a web server network with IP address 192.168.10.0.

**Figure 400:**Example active-passive WAN optimization and secure tunneling topology



## General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Configure the client-side FortiGate unit:
  - Add peers.
  - Add an authentication group.
  - Add an active WAN optimization policy.
2. Configure the server-side FortiGate unit.
  - Add peers.
  - Add the same authentication group
  - Add a WAN optimization tunnel policy.
  - Add a passive WAN optimization policy that applies application control.

Also note that if you perform any additional actions between procedures, your configuration may have different results.

## Configuring WAN optimization with secure tunneling - web-based manager

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit web-based manager. (CLI steps follow.)

### To configure the client-side FortiGate unit

1. Go to *WAN Opt. & Cache > WAN Opt. Peer > Peer* and enter a *Local Host ID* for the client-side FortiGate unit:

<b>Local Host ID</b>	Client-Fgt
----------------------	------------

2. Select *Apply* to save your setting.
3. Select *Create New* and add a *Peer Host ID* and the *IP Address* for the server-side FortiGate unit:

<b>Peer Host ID</b>	Server-Fgt
<b>IP Address</b>	192.168.20.1

4. Select *OK*.
5. Go to *Wan Opt. & Cache > WAN Opt. Peer > Authentication Group* and select *Create New* to add the authentication group to be used for secure tunneling:

<b>Name</b>	Auth-Secure-Tunnel
<b>Authentication Method</b>	Pre-shared key
<b>Password</b>	2345678
<b>Peer Acceptance</b>	Accept Any Peer

6. Select *OK*.
7. Go to *Wan Opt. & Cache > WAN Opt. Profile > Profile* and select *Create New* to add a WAN optimization profile that enables secure tunneling and includes the authentication group:

<b>Name</b>	Secure-wan-op-pro
<b>Transparent Mode</b>	Select
<b>Authentication Group</b>	Auth-Secure-tunnel

8. Select the *HTTP* protocol, select *Secure Tunneling* and *Byte Caching* and set the *Port* to 80.
9. Select *OK*.
10. Go to *Firewall Objects > Address > Address* and select *Create New* to add a firewall address for the client network.

<b>Category</b>	Address
<b>Name</b>	Client-Net
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	172.20.120.0/24
<b>Interface</b>	port1

11. Select *Create New* to add a firewall address for the web server network.

<b>Category</b>	Address
<b>Address Name</b>	Web-Server-Net
<b>Type</b>	Subnet

<b>Subnet / IP Range</b>	192.168.10.0/24
<b>Interface</b>	port2

- Go to *Policy > Policy > Policy* and select *Create New* to add an active WAN optimization security policy:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	port1
<b>Source Address</b>	Client-Net
<b>Outgoing Interface</b>	port2
<b>Destination Address</b>	Web-Server-Net
<b>Schedule</b>	always
<b>Service</b>	HTTP
<b>Action</b>	ACCEPT

- Turn on Antivirus and select the *default* antivirus profile.

- Select *Enable WAN Optimization* and configure the following settings:

<b>Enable WAN Optimization</b>	active
<b>Profile</b>	Secure-wan-opt-pro

- Select *OK*.

#### To configure the server-side FortiGate unit

- Go to *WAN Opt. & Cache > WAN Opt. Peer > Peer* and enter a *Local Host ID* for the server-side FortiGate unit:

<b>Local Host ID</b>	Server-Fgt
----------------------	------------

- Select *Apply* to save your setting.
- Select *Create New* and add a *Peer Host ID* and the *IP Address* for the client-side FortiGate unit:

<b>Peer Host ID</b>	Client-Fgt
<b>IP Address</b>	172.30.120.1

- Select *OK*.

- Go to *Wan Opt. & Cache > WAN Opt. Peer > Authentication Group* and select *Create New* and add an authentication group to be used for secure tunneling:

<b>Name</b>	Auth-Secure-Tunnel
<b>Authentication Method</b>	Pre-shared key
<b>Password</b>	2345678
<b>Peer Acceptance</b>	Accept Any Peer

- Select *OK*.
- Go to *Firewall Objects > Address > Address* and select *Create New* to add a firewall address for the client network.

<b>Category</b>	Address
<b>Name</b>	Client-Net
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	172.20.120.0/24
<b>Interface</b>	port1

- Select *Create New* to add a firewall address for the web server network.

<b>Category</b>	Address
<b>Address Name</b>	Web-Server-Net
<b>Type</b>	Subnet
<b>Subnet / IP Range</b>	192.168.10.0/24
<b>Interface</b>	port2

- Go to *Policy > Policy > Policy* and select *Create New* to add a WAN optimization tunnel policy.

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	wanopt
<b>Source Address</b>	all
<b>Outgoing Interface</b>	port1
<b>Destination Address</b>	all
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT

10. Select *OK*.

11. Select *Create New* to add a passive WAN optimization policy that applies application control.

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	port2
<b>Source Address</b>	Client-Net
<b>Outgoing Interface</b>	port1
<b>Destination Address</b>	Web-Server-Net
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT

12. Turn on *Application Control* and select the *default* application control sensor.

13. Select *Enable WAN Optimization* and configure the following settings:

<b>Enable WAN Optimization</b>	passive
<b>Passive Option</b>	default

14. Select *OK*.

## Configuring WAN optimization with secure tunneling - CLI

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit CLI.

### To the client-side FortiGate unit

1. Add the Local Host ID to the client-side FortiGate configuration:

```
config wanopt settings
 set host-id Client-Fgt
end
```

2. Add the server-side Local Host ID to the client-side peer list:

```
config wanopt peer
 edit Server-Fgt
 set ip 192.168.20.1
 end
```

3. Add a new authentication group to be used for secure tunneling:

```
config wanopt auth-group
 edit Auth-Secure-Tunnel
 set auth-method psk
 set psk 2345678
 end
```

Leave *peer-accept* at its default value.

4. Add a WAN optimization profile that enables secure tunneling and includes the authentication group, enables HTTP protocol optimization, and enables secure tunneling and byte caching for HTTP traffic:

```
config wanopt profile
 edit Secure-wan-op-pro
 set auth-group Auth-Secure-Tunnel
 config http
 set status enable
 set secure-tunnel enable
 set byte-caching enable
 set port 80
 end
 end
```

5. Add a firewall address for the client network.

```
config firewall address
 edit Client-Net
 set type ipmask
 set subnet 172.20.120.0 255.255.255.0
 set associated-interface port1
 end
```

6. Add a firewall address for the web server network.

```
config firewall address
 edit Web-Server-Net
 set type ipmask
 set subnet 192.168.10.0 255.255.255.0
 set associated-interface port2
 end
```

7. Add an active WAN optimization security policy that includes the WAN optimization profile that enables secure tunneling and that applies virus scanning:

```
config firewall policy
 edit 0
 set srcintf port1
 set dstintf port2
 set srcaddr Client-Net
 set dstaddr Web-Server-Net
 set action accept
 set service HTTP
 set schedule always
 set wanopt enable
 set wanopt-detection active
 set wanopt-profile Secure-wan-opt-pro
 set utm-status enable
 set av-profile default
 end
```



## To configure the server-side FortiGate unit

1. Add the Local Host ID to the server-side FortiGate configuration:

```
config wanopt settings
 set host-id Server-Fgt
end
```

2. Add the client-side Local Host ID to the server-side peer list:

```
config wanopt peer
 edit Client-Fgt
 set ip 172.20.120.1
end
```

3. Add an authentication group to be used for secure tunneling:

```
config wanopt auth-group
 edit Auth-Secure-Tunnel
 set auth-method psk
 set psk 2345678
end
```

Leave `peer-accept` at its default value.

4. Add a firewall address for the client network.

```
config firewall address
 edit Client-Net
 set type ipmask
 set subnet 172.20.120.0 255.255.255.0
 set associated-interface port1
end
```

5. Add a firewall address for the web server network.

```
config firewall address
 edit Web-Server-Net
 set type ipmask
 set subnet 192.168.10.0 255.255.255.0
 set associated-interface port2
end
```

6. Add a WAN optimization tunnel policy.

```
config firewall policy
 edit 0
 set srcintf wanopt
 set dstintf port1
 set srcaddr all
 set dstaddr all
 set action accept
 set service ALL
 set schedule always
end
```

**7. Add a passive WAN optimization policy.**

```
config firewall policy
 edit 0
 set srcintf port1
 set dstintf port2
 set srcaddr Client-Net
 set dstaddr Web-Server-Net
 set action accept
 set service ALL
 set schedule always
 set wanopt enable
 set wanopt-detection passive
 set wanopt-passive-opt default
 set utm-status enable
 set application-list default
 end
```

# Web caching and SSL offloading

FortiGate web caching is a form of object caching that accelerates web applications and web servers by reducing bandwidth usage, server load, and perceived latency. Web caching supports caching of HTTP 1.0 and HTTP 1.1 web sites. See [RFC 2616](#) for information about web caching for HTTP 1.1.



Web caching does not cache audio and video streams including Flash videos and streaming content.

Web caching caches compressed and non-compressed versions of the same file separately. If the HTTP protocol considers the compressed and uncompressed versions of a file the same object, only the compressed or uncompressed file will be cached.

Web caching involves storing HTML pages, images, servlet responses and other web-based objects for later retrieval. You can also go to *System > Config > Advanced > Disk Management* to view the storage locations on the FortiGate unit hard disks. You can change the default storage configuration using the `config wanopt storage` command.

There are three significant advantages to using web caching to improve HTTP and WAN performance:

- reduced bandwidth consumption because fewer requests and responses go over the WAN or Internet.
- reduced web server load because there are fewer requests for web servers to handle.
- reduced latency because responses for cached requests are available from a local FortiGate unit instead of from across the WAN or Internet.

You can use web caching to cache any web traffic that passes through the FortiGate unit, including web pages from web servers on a LAN, WAN or on the Internet. You apply web caching by enabling the web caching option in any security policy. When enabled in a security policy, web caching is applied to all HTTP sessions accepted by the security policy. If the security policy is an explicit web proxy security policy, the FortiGate unit caches explicit web proxy sessions.

This section contains the following topics:

- [Turning on web caching for HTTP and HTTPS traffic](#)
- [Turning on web caching and SSL offloading for HTTPS traffic](#)
- [Changing the ports on which to look for HTTP and HTTPS traffic to cache](#)
- [Web caching and HA](#)
- [Web caching and memory usage](#)
- [Changing web cache settings](#)
- [Forwarding URLs to forwarding servers and exempting web sites from web caching](#)
- [Monitoring Web caching performance](#)
- [Example: Web caching of HTTP and HTTPS Internet content for users on an internal network](#)
- [Example: reverse proxy web caching and SSL offloading for an Internet web server using a static one-to-one virtual IP](#)

## Turning on web caching for HTTP and HTTPS traffic

Web caching can be applied to any HTTP or HTTPS traffic by enabling web caching in a security policy that accepts the traffic. This includes WAN optimization and explicit web proxy traffic. Web caching caches all HTTP traffic accepted by a policy on TCP port 80.

You can add web caching to a security policy to:

- Cache Internet HTTP traffic for users on an internal network to reduce Internet bandwidth use. Do this by selecting the web cache option for security policies that allow users on the internal network to browse web sites on the Internet.
- Reduce the load on a public facing web server by caching objects on the FortiGate unit. This is a reverse proxy with web caching configuration. Do this by selecting the web cache option for a security policy that allows users on the Internet to connect to the web server.
- Cache outgoing explicit web proxy traffic when the explicit proxy is used to proxy users in an internal network who are connecting to the web servers on the Internet. Do this by selecting the web cache option for explicit web proxy security policies that allow users on the internal network to browse web sites on the Internet.
- Combine web caching with WAN optimization. You can enable web caching in any WAN optimization security policy. This includes manual, active, and passive WAN optimization policies and WAN optimization tunnel policies. You can enable web caching on both the client-side and the server-side FortiGate units or on just one or the other. For optimum performance you can enable web caching on both the client-side and server-side FortiGate units. In this way only uncached content is transmitted through the WAN optimization tunnel. All cached content is access locally by clients from the client side FortiGate unit.



One important use for web caching is to cache software updates (for example, Windows Updates or iOS updates). When updates occur a large number of users may all be trying to download these updates at the same time. Caching these updates will be a major performance improvement and also have a potentially large impact on reducing Internet bandwidth use. You may want to adjust the maximum cache object size to make sure these updates are cached. See [“Max cache object size” on page 2648](#).

## Turning on web caching and SSL offloading for HTTPS traffic

Web caching can also cache the content of HTTPS traffic on TCP port 443. With HTTPS web caching, the FortiGate unit receives the HTTPS traffic on behalf of the client, opens up the encrypted traffic and extracts content to be cached. Then FortiGate unit re-encrypts the traffic and sends it on to its intended recipient. It is very similar to a man-in-the-middle attack. You enable HTTPS web caching from the CLI in a security policy that accepts the traffic to be cached using `webcache-https`:

```
config firewall policy
 edit 0
 .
 .
 .
 set webcache enable
 set webcache-https any
 .
 .
 .
 end
```

The any setting causes the FortiGate unit to re-encrypt the traffic with the FortiGate unit's certificate rather than the original certificate. This configuration can cause errors for HTTPS clients because the name on the certificate does not match the name on the web site.

You can stop these errors from happening by configuring HTTPS web caching to use the web server's certificate by setting `webcache-https` to `ssl-server`:

```
config firewall policy
 edit 0
 .
 .
 .
 set webcache enable
 set webcache-https ssl-server
 .
 .
 .
 end
```

The `ssl-server` option causes the FortiGate unit to re-encrypt the traffic with the certificate that you imported into the FortiGate unit. The certificate is added to an SSL server configuration using the following command:

```
config wanopt ssl-server
 edit example_server
 set ip <Web-Server-IP>
 set port 443
 set ssl-mode { full | half}
 set ssl-cert <Web-Server-Cert>
 end
```

Where:

`Web-Server-IP` is the web server's IP address.

`Web-Server-Cert` is the original web server certificate imported into the FortiGate unit.


The SSL server configuration also determines whether the SSL server is operating in half or full mode and the port used for the HTTPS traffic.

Using the SSL server configuration, web caching also supports SSL offloading that uses the FortiGate unit's FortiASIC SSL encryption/decryption engine to accelerate SSL performance.

## Full mode SSL server configuration

The `ssl-mode` option determines whether the SSL server operates in half or full mode. In full mode the FortiGate unit performs both decryption and encryption of the HTTPS traffic. The full mode sequence is shown in [Figure 401](#).

**Figure 401:**Full mode SSL server configuration



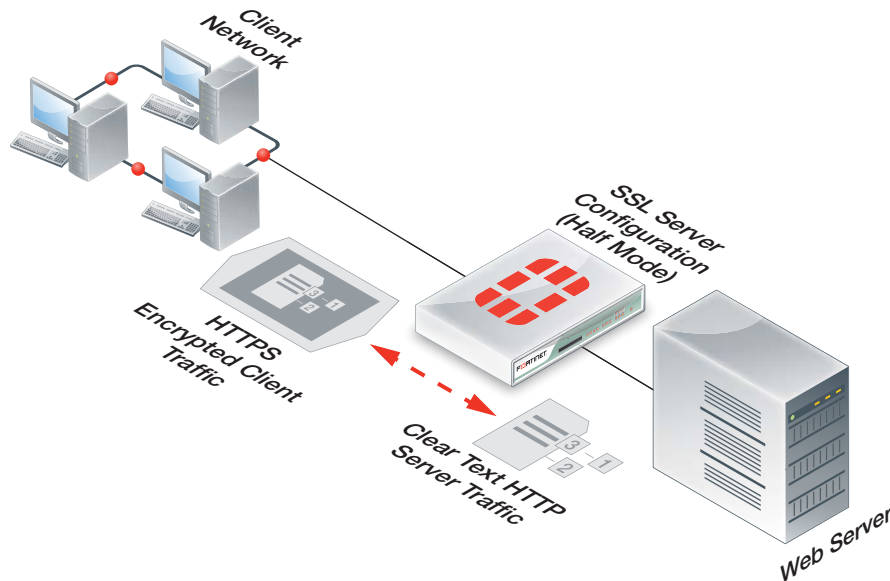
In full mode the FortiGate unit is acting as a man in the middle, decrypting and encrypting the traffic. So both the client and the web server see encrypted packets.

Usually the port of the encrypted HTTPS traffic is always 443. However, in the SSL server configuration you can set the port used for HTTPS traffic. This port is not altered by the SSL Server. So for example, if the SSL Server receives HTTPS traffic on port 443, the re-encrypted traffic forwarded to the FortiGate unit to the server or client will still use port 443.

### Half mode SSL server configuration

In half mode, the FortiGate unit only performs one encryption or decryption action. If HTTP packets are received, the half mode SSL server encrypts them and converts them to HTTPS packets. If HTTPS packets are received, the SSL server decrypts them and converts them to HTTP packets. The half mode sequence is shown in [Figure 402](#):

**Figure 402:**Half mode SSL server configuration



In half mode, the FortiGate unit is acting like an SSL accelerator, offloading HTTPS decryption from the web server to the FortiGate unit. Since FortiGate units can accelerate SSL processing, the end result could be improved web site performance.

Usually the port of the encrypted traffic is always 443. However, in the SSL server configuration you can set the port used for HTTPS traffic. No matter what port is used for the HTTPS traffic, the decrypted HTTP traffic uses port 80.

## Changing the ports on which to look for HTTP and HTTPS traffic to cache

By default FortiOS assumes HTTP traffic uses TCP port 80 and HTTPS traffic uses port 443. So web caching caches all HTTP traffic accepted by a policy on TCP port 80 and all HTTPS traffic on TCP port 443. If you want to cache HTTP or HTTPS traffic on other ports, you can enable security features for the security policy and configure a proxy options profile to that looks for HTTP and HTTPS traffic on other TCP ports.

Setting the HTTP port to *Any* in a proxy options profile is not compatible with web caching. If you set the HTTP port to any, web caching only caches HTTP traffic on port 80.

## Web caching and HA

You can configure web caching on a FortiGate HA cluster. The recommended best practice HA configuration for web caching is active-passive mode. When the cluster is operating, all web caching sessions are processed by the primary unit only. Even if the cluster is operating in active-active mode, HA does not load-balance web caching sessions.

In a cluster, only the primary unit stores the web cache database. The databases is not synchronized to the subordinate units. So, after a failover, the new primary unit must build its web cache.

## Web caching and memory usage

To accelerate and optimize disk access and to provide better throughput and less latency, web caching uses provisioned memory to reduce disk I/O and increase disk I/O efficiency. In addition, web caching requires a small amount of additional memory per session for comprehensive flow control logic and efficient traffic forwarding.

When web caching is enabled you will see a reduction in available memory. The reduction increases when more web caching sessions are being processed. If you are thinking of enabling web caching on an operating FortiGate unit, make sure its memory usage is not maxed out during high traffic periods.

In addition to using the system dashboard to see the current memory usage you can use the `get test wad 1` command to see how much memory is currently being used by web caching. See “`get test {wa_cs | wa_dbd | wad | wad_diskd | wccpd} <test_level>`” on page 2719 for more information.

## Changing web cache settings

In most cases, the default settings for the WAN optimization web cache are acceptable. However, you may want to change them to improve performance or optimize the cache for your configuration. To change these settings, go to *WAN Opt. & Cache > Cache > Settings*.

From the FortiGate CLI, you can use the `config wanopt webcache` command to change these WAN optimization web cache settings.



For more information about many of these web cache settings, see [RFC 2616](#).

---

### Always revalidate

Select to always revalidate requested cached objects with content on the server before serving them to the client.

### Max cache object size

Set the maximum size of objects (files) that are cached. The default size is 512000 KB and the range is 1 to 4294967 KB. This setting determines the maximum object size to store in the web cache. Objects that are larger than this size are still delivered to the client but are not stored in the FortiGate web cache.

For most web traffic the default maximum cache object size is recommended. However, since web caching can also cache larger objects such as Windows updates, Mac OS updates, iOS updates or other updates delivered using HTTP you might want to increase the object size to make sure these updates are cached. Caching these updates can save a lot of Internet bandwidth and improve performance when major updates are released by these vendors.

### Negative response duration

Set how long in minutes that the FortiGate unit caches error responses from web servers. If error responses are cached, then subsequent requests to the web cache from users will receive the error responses regardless of the actual object status.



The default is 0, meaning error responses are not cached. The content server might send a client error code (4xx HTTP response) or a server error code (5xx HTTP response) as a response to some requests. If the web cache is configured to cache these negative responses, it returns that response in subsequent requests for that page or image for the specified number of minutes.

### Fresh factor

Set the fresh factor as a percentage. The default is 100, and the range is 1 to 100%. For cached objects that do not have an expiry time, the web cache periodically checks the server to see if the objects have expired. The higher the *Fresh Factor* the less often the checks occur.

For example, if you set the *Max TTL* value and *Default TTL* to 7200 minutes (5 days) and set the *Fresh Factor* to 20, the web cache check the cached objects 5 times before they expire, but if you set the *Fresh Factor* to 100, the web cache will check once.

### Max TTL

The maximum amount of time (Time to Live) an object can stay in the web cache without the cache checking to see if it has expired on the server. The default is 7200 minutes (120 hours or 5 days) and the range is 1 to 5256000 minutes (5256000 minutes in a year).

### Min TTL

The minimum amount of time an object can stay in the web cache before the web cache checks to see if it has expired on the server. The default is 5 minutes and the range is 1 to 5256000 minutes (5256000 minutes in a year).

### Default TTL

The default expiry time for objects that do not have an expiry time set by the web server. The default expiry time is 1440 minutes (24 hours) and the range is 1 to 5256000 minutes (5256000 minutes in a year).

### Proxy FQDN

The fully qualified domain name (FQDN) for the proxy server. This is the domain name to enter into browsers to access the proxy server. This field is for information only can be changed from the explicit web proxy configuration.

### Max HTTP request length

The maximum length of an HTTP request that can be cached. Larger requests will be rejected. This field is for information only can be changed from the explicit web proxy configuration.

### Max HTTP message length

The maximum length of an HTTP message that can be cached. Larger messages will be rejected. This field is for information only can be changed from the explicit web proxy configuration.

### Ignore

Select the following options to ignore some web caching features.

- If-modified-since
- By default, if the time specified by the if-modified-since (IMS) header in the client's conditional request is greater than the last modified time of the object in the cache, it is a

strong indication that the copy in the cache is stale. If so, HTTP does a conditional GET to the Overlay Caching Scheme (OCS), based on the last modified time of the cached object.

- Enable ignoring if-modified-since to override this behavior.
- HTTP 1.1 conditionals
- HTTP 1.1 provides additional controls to the client over the behavior of caches toward stale objects. Depending on various cache-control headers, the FortiGate unit can be forced to consult the OCS before serving the object from the cache. For more information about the behavior of cache-control header values, see [RFC 2616](#).
- Enable ignoring HTTP 1.1 Conditionals to override this behavior.
- Pragma-no-cache
- Typically, if a client sends an HTTP GET request with a pragma no-cache (PNC) or cache-control no-cache header, a cache must consult the OCS before serving the content. This means that the FortiGate unit always re-fetches the entire object from the OCS, even if the cached copy of the object is fresh.
- Because of this behavior, PNC requests can degrade performance and increase server-side bandwidth utilization. However, if you enable ignoring Pragma-no-cache, then the PNC header from the client request is ignored. The FortiGate unit treats the request as if the PNC header is not present.
- IE Reload
- Some versions of Internet Explorer issue Accept / header instead of Pragma no-cache header when you select *Refresh*. When an Accept header has only the / value, the FortiGate unit treats it as a PNC header if it is a type-N object.
- Enable ignoring IE reload to cause the FortiGate unit to ignore the PNC interpretation of the Accept / header.

### Cache Expired Objects

Applies only to type-1 objects. When this option is selected, expired type-1 objects are cached (if all other conditions make the object cacheable).

### Revalidated Pragma-no-cache

The pragma-no-cache (PNC) header in a client's request can affect how efficiently the FortiGate unit uses bandwidth. If you do not want to completely ignore PNC in client requests (which you can do by selecting to ignore Pragma-no-cache, above), you can nonetheless lower the impact on bandwidth usage by selecting *Revalidate Pragma-no-cache*.

When you select *Revalidate Pragma-no-cache*, a client's non-conditional PNC-GET request results in a conditional GET request sent to the OCS if the object is already in the cache. This gives the OCS a chance to return the 304 Not Modified response, which consumes less server-side bandwidth, because the OCS has not been forced to otherwise return full content.

By default, *Revalidate Pragma-no-cache* is disabled and is not affected by changes in the top-level profile.

Most download managers make byte-range requests with a PNC header. To serve such requests from the cache, you should also configure byte-range support when you configure the *Revalidate pragma-no-cache* option.

## Forwarding URLs to forwarding servers and exempting web sites from web caching

You can go to *WAN Opt. & Cache > Cache > URL Match List* and use the URL match list to forward URL patterns to forwarding server and create a list of URLs that are exempt from web caching.

- [Forwarding URLs and URL patterns to forwarding servers](#)
- [Exempting web sites from web caching](#)

### Forwarding URLs and URL patterns to forwarding servers

As part of configuring the explicit web proxy you can configure proxy chaining by adding web proxy forwarding servers. See [“Proxy chaining \(web proxy forwarding servers\)” on page 2578](#).

You can then use the URL match list to always forward explicit web proxy traffic destined for configured URLs or URL patterns to one of these forwarding servers. For example, you might want to forward all traffic for a specific country to a proxy server located in that country.

To forward traffic destined for a URL to a forwarding server that you have already added, go to *WAN Opt. & Cache > Cache > URL Match List* and select *Create New*. Add a name for the URL match entry and enter the URL or URL pattern. You can use wildcards such as \* and ? and you can use a numeric IP address. Select *Forward to Server* and select a web proxy forwarding server from the list.

You can also exempt the URL or URL pattern from web caching.

Use the following command to forward all .ca traffic to a proxy server and all .com traffic to another proxy server.

```
config web-proxy url-match
 edit "com"
 set forward-server "server-commercial"
 set url-pattern "com"
 next
 edit "ca"
 set forward-server "server-canada"
 set url-pattern "ca"
 next
 edit "www.google.ca"
 set cache-exemption enable
 set url-pattern "www.google.ca"
 next
end
```

### Exempting web sites from web caching

You may want to exempt some URLs from web caching for a number of reasons. For example, if your users access websites that are not compatible with FortiGate web caching you can add the URLs of these web sites to the web caching exempt list. You can add URLs and numeric IP addresses to the web cache exempt list.

You can also add URLs to the web cache exempt list by going to *WAN Opt. & Cache > Cache > URL Match List* and selecting *Create New*. Add a URL pattern to be exempt and select *Exempt from Cache*.

You can also add URLs and addresses to be exempt from the CLI. Enter the following command to add `www.example.com` to the web cache exempt list.

```
config web-proxy url-match
 set cache-exemption enable
 set url-pattern www.example.com
end
```


## Monitoring Web caching performance

The web cache monitor shows the percentage of web cache requests that retrieved content from the cache (hits) and the percentage that did not receive content from the cache (misses). A higher the number of hits usually indicates that the web cache is being more effective at reducing WAN traffic.

The web cache monitor also shows a graph of web traffic on the WAN and LAN. A lower WAN line on the graph indicates the web cache is reducing traffic on the WAN. The web cache monitor also displays the total number of web requests processed by the web cache.

To view the web cache monitor, go to *WAN Opt. & Cache > Monitor > Cache Monitor*.

**Figure 403:**Web cache monitor



## Example: Web caching of HTTP and HTTPS Internet content for users on an internal network

This example describes how to configure web caching of HTTP and HTTPS for users on a private network connecting to the Internet.

### Network topology and assumptions


This example includes a client network with subnet address 10.31.101.0 connecting to web servers on the Internet (Figure 404). All of the users on the private network access the Internet through a single general security policy on the FortiGate unit that accepts all sessions connecting to the Internet. Web caching for HTTP and HTTPS traffic is just added to this security policy.

Since users on the private network have unrestricted access to the Internet and can be accessing many web servers the `webcache-https` is set to `any` and users may see error messages on their web browsers when accessing HTTPS content.

Initially, security profiles are not selected so the example caches all HTTP traffic on TCP port 80 and HTTPS traffic on port 443. The example also describes how to configure the security policy

to cache HTTP traffic on port 80 and 8080 by added a proxy options profile that looks for HTTP traffic on TCP ports 80 and 8080. The example also describes how to configure the security policy to cache HTTPS traffic on port 443 and 8443 using the same proxy options profile.

**Figure 404:**Example web caching topology



### General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Add HTTP web caching to the security policy that all users on the private network use to connect to the Internet.
2. Add HTTPS web caching.
3. Add a protocol options profile to look for HTTP traffic on ports 80 and 8080 and HTTPS traffic on ports 443 and 8443 and add this protocol options profile to the security policy.

If you perform any additional actions between procedures, your configuration may have different results.

### Configuration Steps - web-based manager

Use the following steps to configure the example configuration from the FortiGate web-based manager.

#### To add HTTP web caching to a security policy

1. Go to *Policy > Policy > Policy* and add a security policy that allows all users on the internal network to access the Internet.

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	Internal
<b>Source Address</b>	all
<b>Outgoing Interface</b>	wan1
<b>Destination Address</b>	all
<b>Schedule</b>	always

<b>Service</b>	ALL
<b>Action</b>	ACCEPT

2. Select *Enable NAT* and select *Use Destination Interface Address*.
3. Select *Enable Web cache*.
4. Select *OK*.

#### To add HTTPS web caching

1. From the CLI enter the following command to add HTTPS web caching to the security policy. Assume the index number of the policy is 5.

```
config firewall policy
 edit 5
 set webcache-https any
 end
```

#### To cache HTTP traffic on port 80 and 8080 and HTTPS traffic on ports 443 and 8443

1. Go to *Policy > Policy > Proxy Options* and edit the *default* proxy options profile. You could also add a new profile.
2. Under *Protocol Port Mapping* enable *HTTP* and under *Inspection Ports* enter *80,8080*.
3. Under *SSL Inspection Options* enable *HTTPS* and under *Inspection Ports* enter *443,8443*.
4. From the CLI, enter the following command to add the default proxy options profile to the firewall policy.

```
config firewall policy
 edit 5
 set utm-status enable
 set profile-protocol-options default
 end
```

### Configuration Steps - CLI

Use the following steps to configure the example configuration from the FortiGate CLI.

#### To add HTTP and HTTPS web caching to a security policy

1. Enter the following command to add a security policy that allows all users on the internal network to access the Internet and that includes web caching of HTTP and HTTPS traffic.

```
config firewall policy
 edit 0
 set srcintf internal
 set srcaddr all
 set dstintf wan1
 set distinf all
 set schedule always
 set service ANY
 set action accept
 set nat enable
 set webcache enable
 set webcache-https any
 end
```

### To cache HTTP traffic on port 80 and 8080 and HTTPS traffic on ports 443 and 8443

1. Enter the following command to edit the *default* proxy options profile to configure it to look for HTTP traffic on ports 80 and 8080 and to look for HTTPS traffic on ports 443 and 8443:

```
config firewall profile-protocol-options
 edit default
 config http
 set status enable
 set ports 80 8080
 end
 config https
 set status enable
 set ports 443 8443
 end
 end
end
```

2. Enter the following command to add the protocol options profile to the security policy:

```
config firewall policy
 edit 1
 set utm-status enable
 set profile-protocol-options default
 end
```

## Example: reverse proxy web caching and SSL offloading for an Internet web server using a static one-to-one virtual IP

This section describes configuring SSL offloading for a reverse proxy web caching configuration using a static one-to-one firewall virtual IP (VIP). While the static one-to-one configuration described in this example is valid, it's also common to change the destination port of the unencrypted HTTPS traffic to a commonly used HTTP port such as 8080 using a port forwarding virtual IP.

### Network topology and assumptions

In this configuration, clients on the Internet use HTTP and HTTPS to browse to a web server that is behind a FortiGate unit. A policy added to the FortiGate unit forwards the HTTP traffic to the web server. The policy also offloads HTTPS decryption and encryption from the web server so the web server only sees HTTP traffic.

The FortiGate unit also caches HTTP and HTTPS pages from the web server so when users access cached pages the web server does not see the traffic. Replies to HTTPS sessions are encrypted by the FortiGate unit before returning to the clients.

In this configuration, the FortiGate unit is operating as a web cache in reverse proxy mode. Reverse proxy caches can be placed directly in front of a web server. Web caching on the FortiGate unit reduces the number of requests that the web server must handle, therefore leaving it free to process new requests that it has not serviced before.

Using a reverse proxy configuration:

- avoids the capital expense of additional web servers by increasing the capacity of existing servers
- serves more requests for static content from web servers
- serves more requests for dynamic content from web servers
- reduces operating expenses including the cost of bandwidth required to serve content
- accelerates the response time of web servers and of page download times to end users.

When planning a reverse proxy implementation, the web server's content should be written so that it is "cache aware" to take full advantage of the reverse proxy cache.

In reverse proxy mode, the FortiGate unit functions more like a web server for clients on the Internet. Replicated content is delivered from the proxy cache to the external client without exposing the web server or the private network residing safely behind the firewall.

In this example, the site URL translates to IP address 192.168.10.1, which is the port2 IP address of the FortiGate unit. The port2 interface is connected to the Internet.

This example assumes that all HTTP traffic uses port 80 and all HTTPS traffic uses port 443.

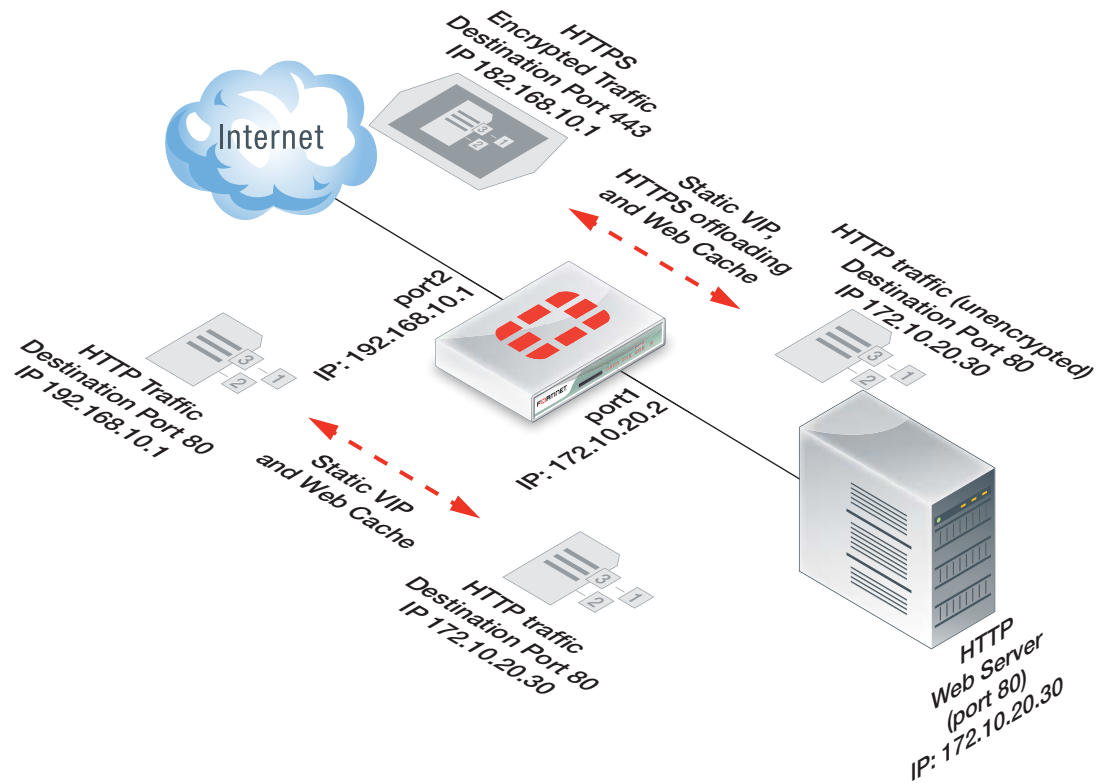
The FortiGate unit includes the web server CA and an SSL server configuration for IP address 172.10.20.30 and port to 443. The name of the file containing the CA is Rev\_Proxy\_Cert\_1.crt.

The destination address of incoming HTTP and HTTPS sessions is translated to the IP address of the web server using a static one-to-one virtual IP that performs destination address translation (DNAT) for the HTTP packets. The DNAT translates the destination address of the packets from 192.168.10.1 to 172.10.20.30 but does not change the destination port number.

When the SSL server on the FortiGate unit decrypts the HTTPS packets their destination port is changed to port 80.



**Figure 405:** Reverse proxy web caching and SSL offloading for an Internet web server using static one-to-one virtual IPs



## General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Configure the FortiGate unit as a reverse proxy web cache server.
2. Configure the FortiGate unit for SSL offloading of HTTPS traffic.
3. Add an SSL server to offload SSL encryption and decryption for the web server.

Also note that if you perform any additional actions between procedures, your configuration may have different results.

## Configuration steps - web-based manager

### To configure the FortiGate unit as a reverse proxy web cache server

1. Go to *Firewall Objects > Virtual IP > Virtual IP* and select *Create New* to add a static NAT virtual IP that translates destination IP addresses from 192.168.10.1 to 172.10.20.30 (and does not translate destination ports):

<b>Name</b>	Reverse_proxy_VIP
<b>External Interface</b>	port2
<b>Type</b>	Static NAT
<b>Source Address Filter</b>	Do not select.
<b>External IP Address/Range</b>	192.168.10.1

<b>Mapped IP Address/Range</b>	172.10.20.30
<b>Port Forwarding</b>	Do not select.

2. Select *OK*.
3. Go to *Policy > Policy > Policy* and select *Create New* to add a port2 to port1 security policy that accepts HTTP and HTTPS traffic from the Internet.

Do not select security profiles. Set the destination address to the virtual IP. You do not have to enable NAT.

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	port2
<b>Source Address</b>	all
<b>Outgoing Interface</b>	port1
<b>Destination Address</b>	Reverse_proxy_VIP
<b>Schedule</b>	always
<b>Service</b>	HTTP HTTPS
<b>Action</b>	ACCEPT

4. Select *Enable Web cache*.
5. Select *OK*.

#### To configure the FortiGate unit to offload SSL encryption and cache HTTPS content

1. Go to *System > Certificates > Local Certificates* and select *Import* to import the web server's CA.

For *Type*, select *Local Certificate*. Select the *Browse* button to locate the file `Rev_Proxy_Cert_1.crt`.

The certificate key size must be 1024 or 2048 bits. 4096-bit keys are not supported.

2. From the CLI enter the following command to add HTTPS web caching to the security policy. Assume the index number of the policy is 5.

```
config firewall policy
 edit 5
 set webcache-https ssl-server
 end
```

3. From the CLI, enter the following command to add the SSL server.

The SSL server `ip` must match the destination address of the SSL traffic after being translated by the virtual IP (172.10.20.30) and the SSL server `port` must match the

destination port of the SSL traffic (443). The SSL server operates in half mode since it performs a single-step conversion (HTTPS to HTTP or HTTP to HTTPS).

```
config wanopt ssl-server
 edit rev_proxy_server
 set ip 172.10.20.30
 set port 443
 set ssl-mode half
 set ssl-cert Rev_Proxy_Cert_1
 end
```

## Configuration steps - CLI

### To configure the FortiGate unit as a reverse proxy web cache server

1. Enter the following command to add a static NAT virtual IP that translates destination IP addresses from 192.168.10.1 to 172.10.20.30 (and does not translate destination ports):

```
config firewall vip
 edit Reverse_proxy_VIP
 set extintf port2
 set type static-nat
 set extip 192.168.10.1
 set mappedip 172.10.20.30
 end
```

2. Enter the following command to add a port2 to port1 security policy that accepts HTTP and HTTPS traffic from the Internet. Enable web caching and HTTPS web caching.

Do not select security profiles. Set the destination address to the virtual IP. You do not have to enable NAT.

```
config firewall policy
 edit 0
 set srcintf port2
 set srcaddr all
 set dstintf port1
 set dstaddr Reverse_proxy_VIP
 set schedule always
 set service HTTP HTTPS
 set action accept
 set webcache enable
 set webcache-https ssl-server
 end
```

### To add an SSL server to offload SSL encryption and decryption for the web server

1. Place a copy of the web server's CA (file name Rev\_Proxy\_Cert\_1.crt) in the root folder of a TFTP server.
2. Enter the following command to import the web server's CA from a TFTP server. The IP address of the TFTP server is 10.31.101.30:

```
execute vpn certificate local import tftp Rev_Proxy_Cert_1.crt
 10.31.101.30
```

The certificate key size must be 1024 or 2048 bits. 4096-bit keys are not supported.

3. From the CLI, enter the following command to add the SSL server.

The SSL server `ip` must match the destination address of the SSL traffic after being translated by the virtual IP (172.10.20.30) and the SSL server `port` must match the destination port of the SSL traffic (443). The SSL server operates in half mode since it performs a single-step conversion (HTTPS to HTTP or HTTP to HTTPS).

```
config wanopt ssl-server
 edit rev_proxy_server
 set ip 172.10.20.30
 set port 443
 set ssl-mode half
 set ssl-cert Rev_Proxy_Cert_1
 end
```


4. Configure other `ssl-server` settings that you may require for your configuration.

# FortiClient WAN optimization

FortiClient WAN optimization supports protocol optimization and byte caching in IPsec VPN and SSL VPN tunnels between FortiClient and a FortiGate unit. To add WAN optimization to FortiClient, configure FortiClient Advanced settings and enable WAN optimization. This setting can then apply WAN optimization to any IPsec or SSL VPN tunnel between FortiClient and FortiGate, if the FortiGate IPsec or SSL VPN configuration also includes WAN optimization.

When FortiClient with WAN optimization enabled attempts to connect a server-side FortiGate unit, FortiClient automatically detects if WAN optimization has been added to the FortiGate tunnel configuration. If WAN optimization is detected and FortiClient can successfully negotiate with the FortiGate unit, WAN optimization starts.


**Figure 406:**FortiClient WAN optimization topology



## FortiClient WAN optimization over SSL VPN configuration example

This example shows how to add WAN optimization to a FortiClient SSL VPN. The SSL VPN tunnel allows remote FortiClient users to connect to the internal network behind the FortiGate unit as shown in [Figure 407](#).

**Figure 407:**Example FortiClient WAN optimization configuration



### To configure the FortiGate unit

Because computers running FortiClient can have IP addresses that change often, it is usually not practical to add FortiClient peers to the FortiGate WAN optimization peer list. Instead, a FortiGate unit that accepts WAN optimization tunnel requests from FortiClient is usually configured to accept any peer (see [“Accepting any peers” on page 2610](#)). This example does this by adding a WAN optimization authentication group with *Peer acceptance* set to *Accept Any Peer*.

In addition this example includes a *wanopt* to *internal* policy to allow WAN optimization traffic reach the internal network. Finally passive WAN optimization is added to the *ssl.root* policy because WAN optimization is accepting traffic from the SSL VPN tunnel.

1. Go to *WAN Opt. & Cache > WAN Opt. Peer > Authentication Group* and select *Create New*.
2. Configure the WAN optimization authentication group:

<b>Name</b>	auth-fc
<b>Authentication Method</b>	Certificate
<b>Certificate</b>	Fortinet_Firmware
<b>Peer Acceptance</b>	Accept Any Peer

3. Select *OK*.
4. Go to *WAN Opt. & Cache > WAN Opt. Profiles > Profiles* and select *Create New* (select the + button).

5. Add a profile for FortiClient WAN optimization sessions:

<b>Name</b>	Fclient_Pro
<b>Transparent Mode</b>	Select
<b>Authentication Group</b>	auth-fc

6. Select any Protocols and any settings for each protocol.  
7. Select *OK*.  
8. Go to *Firewall Objects > Address > Addresses* and select *Create New* to add a firewall address for the internal network that FortiClient users can access.

<b>Category</b>	Address
<b>Address Name</b>	Internal-Server-Net
<b>Type</b>	IP Range
<b>Subnet / IP Range</b>	192.168.10.0/24
<b>Interface</b>	internal

9. Go to *Policy > Policy > Policy* and select *Create New* to add a WAN optimization tunnel policy.

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	wanopt
<b>Source Address</b>	all
<b>Outgoing Interface</b>	internal
<b>Destination Address</b>	all
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT

10. Select *OK*.

11. Select *Create New* to add an ssl.root policy with passive WAN optimization.



If you already have an ssl.root to internal policy you can edit it and enable passive WAN optimization as shown in [Step 13](#).

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address

<b>Incoming Interface</b>	ssl.root
<b>Source Address</b>	all
<b>Outgoing Interface</b>	internal
<b>Destination Address</b>	Internal-Server-Net
<b>Schedule</b>	always
<b>Service</b>	ALL
<b>Action</b>	ACCEPT

12. Optionally, select the *Security Profiles* to apply to the WAN optimization traffic.

13. Select *Enable WAN Optimization* and configure the following settings:

<b>Enable WAN Optimization</b>	passive
<b>Passive Option</b>	default

14. Select *OK*.

**To configure FortiClient and start the WAN optimization SSL VPN connection**

1. Open FortiClient, configure *Advanced* settings, and select *Enable WAN optimization*.
2. Add a new SSL VPN connection.

Set the Server to the WAN1 IP address of the FortiGate unit (172.20.120.30 in this example) and the correct SSL VPN port number (usually 10443 or 443).

No other settings are required for this example. You can add authentication in the form of a user name and password if required by the FortiGate unit.

3. Start the SSL VPN tunnel.
4. Accept the certificate.

You should be connected to the SSL VPN tunnel and traffic in it should be optimized.



# The FortiGate explicit web proxy

You can use the FortiGate explicit web proxy to enable explicit HTTP, and HTTPS proxying on one or more FortiGate interfaces. The explicit web proxy also supports proxying FTP sessions from a web browser and proxy auto-config (PAC) to provide automatic proxy configurations for explicit web proxy users. From the CLI you can also configure the explicit web proxy to support IPv6 traffic and SOCKS sessions from a web browser.

The explicit web and FTP proxies can be operating at the same time on the same or on different FortiGate interfaces.



The explicit web proxy is configured for each VDOM when multiple VDOMs are enabled.

---

In most cases you would configure the explicit web proxy for users on a network by enabling the explicit web proxy on the FortiGate interface connected to that network. Users on the network would configure their web browsers to use a proxy server for HTTP and HTTPS, FTP, or SOCKS and set the proxy server IP address to the IP address of the FortiGate interface connected to their network. Users could also enter the PAC URL into their web browser PAC configuration to automate their web proxy configuration using a PAC file stored on the FortiGate unit.




Enabling the explicit web proxy on an interface connected to the Internet is a security risk because anyone on the Internet who finds the proxy could use it to hide their source address.

---

If the FortiGate unit is operating in Transparent mode, users would configure their browsers to use a proxy server with the FortiGate management IP address.

The web proxy receives web browser sessions to be proxied at FortiGate interfaces with the explicit web proxy enabled. The web proxy uses FortiGate routing to route sessions through the FortiGate unit to a destination interface. Before a session leaves the exiting interface, the explicit web proxy changes the source addresses of the session packets to the IP address of the exiting interface. When the FortiGate unit is operating in Transparent mode the explicit web proxy changes the source addresses to the management IP address. For more information about explicit web proxy sessions, see [“Explicit proxy sessions and user limits” on page 2686](#).

**Figure 408:**Example explicit web proxy topology



To allow all explicit web proxy traffic to pass through the FortiGate unit you can set the explicit web proxy default firewall proxy action to accept. However, in most cases you would want to use security policies to control explicit web proxy traffic and apply security features such as access control/authentication, security profiles such as antivirus and application control, and traffic logging. You can do this by keeping the default explicit web proxy security policy action to deny and then adding web-proxy security policies.

You can also change the explicit web proxy default security policy action to accept and add explicit web proxy security policies. If you do this, sessions that match web-proxy security policies are processed according to the security policy settings. Connections to the explicit web proxy that do not match a web-proxy security policy are allowed with no restrictions or additional security processing. This configuration is not recommended and is not a best practice.

Web-proxy security policies can selectively allow or deny traffic, apply authentication using user identity-based policies, enable traffic logging, and use security profiles to apply virus scanning, web filtering, IPS, application control, and DLP to explicit web proxy traffic.

You cannot configure Traffic shaping for explicit web proxy traffic. Security policies for the web proxy can only include firewall addresses not assigned to a FortiGate unit interface or with interface set to *Any*. (On the web-based manager you must set the interface to *Any*. In the CLI you must `unset` the associated-interface.)

Authentication of explicit web proxy sessions uses HTTP authentication and can be based on the user's source IP address or on cookies from the user's web browser. For more information, see [“Explicit web proxy authentication” on page 2677](#).

To use the explicit web proxy, users must add the IP address of a FortiGate interface on which the explicit web proxy is enabled and the explicit web proxy port number (default 8080) to the proxy configuration settings of their web browsers.

On FortiGate units that support it, you can also enable web caching for explicit web proxy sessions.

This section describes:

- [Explicit web proxy configuration overview](#)
- [IPv6 Explicit web proxy](#)
- [Proxy chaining \(web proxy forwarding servers\)](#)
- [Explicit web proxy authentication](#)
- [Security profiles, client reputation, device identification, and the explicit web proxy](#)
- [Web Proxy firewall services and service groups](#)
- [Example: users on an internal network browsing the Internet through the explicit web proxy with web caching, RADIUS authentication, web filtering and virus scanning](#)
- [Explicit proxy sessions and user limits](#)

## Explicit web proxy configuration overview

This section describes:

- [General configuration steps](#)
- [Proxy auto-config \(PAC\) configuration](#)
- [Unknown HTTP version](#)
- [Authentication realm](#)
- [Other explicit web proxy options](#)
- [Restricting the IP address of the explicit web proxy](#)
- [Restricting the outgoing source IP address of the explicit web proxy](#)



For explicit FTP proxy options, see [“Explicit FTP proxy configuration overview”](#) on page 2692.



For web proxy forwarding server options, see [“Proxy chaining \(web proxy forwarding servers\)”](#) on page 2674.

---

### General configuration steps

You can use the following general steps to configure the explicit web proxy.

#### To enable the explicit web proxy - web-based manager

1. Go to *System > Network > Explicit Proxy*. Select *Enable Explicit Web Proxy* to turn on the explicit web proxy for HTTP and HTTPS traffic.

You can also select FTP to enable the web proxy for FTP over HTTP sessions in a web browser (not an FTP client) and PAC to enable automatic proxy configuration.

You can also optionally change the HTTP port that the proxy listens on (the default is 8080) and optionally specify different ports for HTTPS, FTP, and PAC.

2. Select *Apply*.

The default explicit web proxy configuration has *Default Firewall Policy Action* set to Deny and requires you to add a security policy to allow access to the explicit web proxy. This configuration is recommended as a best practice because you can use security policies to control access to the explicit web proxy and also apply security features such as logging, security profiles, and authentication (by adding identity-based policies).

3. Go to *System > Network > Interface* and select one or more interfaces for which to enable the explicit web proxy. Edit the interface configuration and select *Enable Explicit Web Proxy*.



Enabling the explicit web proxy on an interface connected to the Internet is a security risk because anyone on the Internet who finds the proxy could use it to hide their source address. If you enable the proxy on such an interface make sure authentication is required to use the proxy.

4. Go to *Firewall Objects > Address > Addresses* and select *Create New* to add a firewall address that matches the source address of packets to be accepted by the explicit proxy.

<b>Name</b>	Internal_subnet
<b>Type</b>	IP Range
<b>Subnet / IP Range</b>	10.31.101.1 - 10.31.101.255
<b>Source Address</b>	Internal_subnet
<b>Interface</b>	Any*

\*The *Interface* must be set to *Any*.

5. Go to *Policy > Policy > Policy* and select *Create New* and set the *Incoming Interface* to *web-proxy*. Configure the security policy as required to accept the traffic that you want to be allowed to use the explicit web proxy.

The source address of the policy must match client source IP addresses. The interface of this firewall address must be set to *Any*.

The destination address of the policy should match the IP addresses of web sites that clients are connecting to. Usually the destination address would be *all* if proxying Internet web browsing.

If *Default Firewall Policy Action* is set to *Deny*, traffic sent to the explicit web proxy that is not accepted by a web-proxy security policy is dropped. If *Default Firewall Policy Action* is set to *Allow* then all web-proxy sessions that don't match with a security policy are allowed.

For example, the following security policy allows users on an internal network to access the Internet through the wan1 interface of a FortiGate unit.

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	web-proxy
<b>Source Address</b>	Internal_subnet
<b>Outgoing Interface</b>	wan1
<b>Destination Address</b>	all

<b>Service</b>	webproxy
<b>Action</b>	ACCEPT

Set the *Policy Subtype* to *User Identity* to require authentication to access the explicit web proxy. For example:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	User Identity
<b>Incoming Interface</b>	web-proxy
<b>Source Address</b>	Internal_subnet
<b>Outgoing Interface</b>	wan1
<b>Destination Address</b>	all
<b>Service</b>	webproxy

Select *Create New* to add an *Authentication Rule* and configure the rule as follows:

<b>Groups</b>	Proxy-Group
<b>Users</b>	(optional)
<b>Schedule</b>	always
<b>Action</b>	ACCEPT

Add security profiles as required.

You can add multiple user identity policies to apply different authentication for different user groups and users and also apply different security profile and logging settings for different user groups.

6. Select *OK*.

### To enable the explicit web proxy - CLI

1. Enter the following command to turn on the explicit web proxy for HTTP and HTTPS traffic.

```
config web-proxy explicit
 set status enable
end
```

You can also enter the following command to enable the web proxy for FTP sessions in a web browser.

```
config web-proxy explicit
 set ftp-over-http enable
end
```

The default explicit web proxy configuration has *sec-default-action* set to *deny* and requires you to add a security policy to allow access to the explicit web proxy.

2. Enter the following command to enable the explicit web proxy for the internal interface.

```
config system interface
edit internal
set explicit-web-proxy enable
end
end
```

3. Use the following command to add a firewall address that matches the source address of users who connect to the explicit web proxy.

```
config firewall address
edit Internal_subnet
set type iprange
set start-ip 10.31.101.1
set end-ip 10.31.101.255
end
```

The source address for a web-proxy security policy cannot be assigned to a FortiGate interface.

4. Use the following command to add a security policy that allows all users on the 10.31.101.0 subnet to use the explicit web proxy for connections through the wan1 interface to the Internet.

```
config firewall policy
edit 0
set srcintf web-proxy
set dstintf wan1
set scraddr Internal_subnet
set dstaddr all
set action accept
set service webproxy
set schedule always
end
```

You can also add authentication to this policy.

```
config firewall policy
edit 0
set srcintf web-proxy
set dstintf wan1
set scraddr Internal_subnet
set dstaddr all
set action accept
set service webproxy
set identity-based enable
config identity-based-policy
edit 1
set groups Internal_users
set utm-status enable
set profile-protocol-options default
set av-profile default
set logtraffic enable
set schedule always
end
end
```

5. Use the following command to change global web proxy settings, for example to set the maximum request length for the explicit web proxy to 10:

```
config web-proxy global
 set max-request-length 10
end
```

## Proxy auto-config (PAC) configuration

A proxy auto-config (PAC) file defines how web browsers can choose a proxy server for receiving HTTP content. PAC files include the FindProxyForURL(url, host) JavaScript function that returns a string with one or more access method specifications. These specifications cause the web browser to use a particular proxy server or to connect directly.

To configure PAC for explicit web proxy users, you can use the port that PAC traffic from client web browsers use to connect to the explicit web proxy. explicit web proxy users must configure their web browser's PAC proxy settings to use the PAC port.

### PAC File Content

You can edit the default PAC file from the web-based manager or use the following command to upload a custom PAC file:

```
config web-proxy explicit
 set pac-file-server-status enable
 set pac-file-data <pac_file_str>
end
```

Where <pac\_file\_str> is the contents of the PAC file. Enter the PAC file text in quotes. You can copy the contents of a PAC text file and paste the contents into the CLI using this option. Enter the command followed by two sets of quotes then place the cursor between the quotes and paste the file content.

The maximum PAC file size is 256 kbytes. If your FortiGate unit is operating with multiple VDOMs each VDOM has its own PAC file. The total amount of FortiGate memory available to store all of these PAC files 2 MBytes. If this limit is reached you will not be able to load any additional PAC files.

You can use any PAC file syntax that is supported by your users's browsers. The FortiGate unit does not parse the PAC file.

To use PAC, users must add an automatic proxy configuration URL (or PAC URL) to their web browser proxy configuration. The default FortiGate PAC file URL is:

```
http://<interface_ip>:<PAC_port_int>/<pac_file_str>
```

For example, if the interface with the explicit web proxy has IP address 172.20.120.122, the PAC port is the same as the default HTTP explicit web proxy port (8080) and the PAC file name is proxy.pac the PAC file URL would be:

```
http://172.20.120.122:8080/proxy.pac
```

From the CLI you can use the following command to display the PAC file urls:

```
get web-proxy explicit
```

## Unknown HTTP version

You can select the action to take when the proxy server must handle an unknown HTTP version request or message. Set unknown HTTP version to Reject or Best Effort. Best Effort attempts to handle the HTTP traffic as best as it can. Reject treats known HTTP traffic as malformed and drops it. The Reject option is more secure.

## Authentication realm

You can enter an authentication realm to identify the explicit web proxy. The realm can be any text string of up to 63 characters. If the realm includes spaces enclose it in quotes. When a user authenticates with the explicit web proxy the HTTP authentication dialog includes the realm so you can use the realm to identify the explicitly web proxy for your users.

## Other explicit web proxy options

You can change the following explicit web proxy options as required by your configuration.

<b>HTTP port, HTTPS port, FTP port, PAC port</b>	The TCP port that web browsers use to connect to the explicit proxy for HTTP, HTTPS, FTP and PAC services. The default port is 8080 for all services. By default HTTPS, FTP, and PAC use the same port as HTTP. You can change any of these ports as required. Users configuring their web browsers to use the explicit web proxy should add the same port numbers to their browser configurations.
<b>Proxy FQDN</b>	Enter the fully qualified domain name (FQDN) for the proxy server. This is the domain name to enter into browsers to access the proxy server.
<b>Max HTTP request length</b>	Enter the maximum length of an HTTP request in Kbytes. Larger requests will be rejected.
<b>Max HTTP message length</b>	Enter the maximum length of an HTTP message in Kbytes. Larger messages will be rejected.

## Restricting the IP address of the explicit web proxy

You can use the following command to restrict access to the explicit web proxy using only one IP address. The IP address that you specify must be the IP address of an interface that the explicit HTTP proxy is enabled on. You might want to use this option if the explicit FTP proxy is enabled on an interface with multiple IP addresses.

For example, to require users to connect to the IP address 10.31.101.100 to connect to the explicit HTTP proxy:

```
config web-proxy explicit
 set incoming-ip 10.31.101.100
end
```

## Restricting the outgoing source IP address of the explicit web proxy

You can use the following command to restrict the source address of outgoing web proxy packets to a single IP address. The IP address that you specify must be the IP address of an interface that the explicit HTTP proxy is enabled on. You might want to use this option if the explicit HTTP proxy is enabled on an interface with multiple IP addresses.

For example, to restrict the outgoing packet source address to 172.20.120.100:

```
config http-proxy explicit
 set outgoing-ip 172.20.120.100
end
```



## IPv6 Explicit web proxy

You can use the explicit web proxy for IPv6 web traffic. To do this you need to:

- Enable the IPv6 explicit web proxy from the CLI
- Enable the explicit web proxy for one or more FortiGate interfaces. These interfaces also need an IPv6 address
- Add web proxy security policies and add IPv6 firewall addresses to allow the explicit web proxy to accept IPv6 traffic.



If you have enabled both the IPv4 and the IPv6 explicit web proxy you can combine IPv4 and IPv6 addresses in a single explicit web proxy policy to allow both IPv4 and IPv6 traffic through the proxy.

Use the following steps to set up a FortiGate unit to accept IPv4 and IPv6 traffic for the explicit web proxy at the Internal interface and forward IPv4 and IPv6 explicit proxy traffic out the wan1 interface to the Internet.

1. Enter the following CLI command to enable the IPv6 explicit web proxy:

```
config web-proxy explicit
 set status enable
 set ipv6-status enable
end
```

2. Go to *System > Network > Interface* and edit the *internal* interface, select *Enable Explicit Web Proxy* and select *OK*.
3. Go to *Policy > Policy > Policy* and select *Create New* to add an IPv6 explicit web proxy security policy:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	web-proxy
<b>Source Address</b>	Internal-IPv4-subnet
<b>Source IPv6 Address</b>	Internal-IPv6-subnet
<b>Outgoing Interface</b>	wan1
<b>Destination Address</b>	all
<b>Destination IPv6 Address</b>	all
<b>Service</b>	webproxy
<b>Action</b>	ACCEPT

This IPv6 explicit web proxy policy allows traffic from all IPv6 IP addresses to connect through the explicit web proxy and through the wan1 interface to any IPv6 addresses that are accessible from the wan1 interface.

## Restricting the IP address of the explicit IPv6 web proxy

You can use the following command to restrict access to the IPv6 explicit web proxy to use only one IPv6 IP address. The IPv6 address that you specify must be the IPv6 address of an interface that the explicit HTTP proxy is enabled on. You might want to use this option if the explicit web proxy is enabled on an interface with multiple IPv6 addresses.

For example, to require users to connect to the IPv6 address 2001:db8:0:2::30 to connect to the explicit IPv6 HTTP proxy:

```
config web-proxy explicit
 set incoming-ipv6 2001:db8:0:2::30
end
```

## Restricting the outgoing source IP address of the IPv6 explicit web proxy

You can use the following command to restrict the source address of outgoing web proxy packets to a single IPv6 address. The IP address that you specify must be the IPv6 address of an interface that the explicit HTTP proxy is enabled on. You might want to use this option if the explicit HTTP proxy is enabled on an interface with multiple IPv6 addresses.

For example, to restrict the outgoing packet source address to 2001:db8:0:2::50:

```
config http-proxy explicit
 set outgoing-ip6 2001:db8:0:2::50
end
```

## Proxy chaining (web proxy forwarding servers)

For the explicit web proxy you can configure web proxy forwarding servers to use proxy chaining to redirect web proxy sessions to other proxy servers. Proxy chaining can be used to forward web proxy sessions from the FortiGate unit to one or more other proxy servers on your network or on a remote network. You can use proxy chaining to integrate the FortiGate explicit web proxy with an already existing web proxy solution.

A FortiGate unit can forward sessions to most web proxy servers including a remote FortiGate unit with the explicit web proxy enabled. No special configuration of the explicit web proxy on the remote FortiGate unit is required.

You can deploy the explicit web proxy with proxy chaining in an enterprise environment consisting of small satellite offices and a main office. If each office has a FortiGate unit, users at each of the satellite offices can use their local FortiGate unit as an explicit web proxy server. The satellite office FortiGate units can forward explicit web proxy sessions to an explicit web proxy server at the central office. From here the sessions can connect to web servers on the Internet.

FortiGate proxy chaining does not support authenticating with the remote forwarding server.

This section also describes:

- [Adding a web proxy forwarding server](#)
- [Web proxy forwarding server monitoring and health checking](#)
- [Adding proxy chaining to an explicit web proxy security policy](#)

## Adding a web proxy forwarding server

To add a forwarding server, select *Create New* in the *Web Proxy Forwarding Servers* section of the *Explicit Proxy* page by going to *System > Network > Explicit Proxy*.

<b>Server Name</b>	Enter the name of the forwarding server.
<b>Proxy Address</b>	Enter the IP address of the forwarding server.
<b>Proxy Address Type</b>	Select the type of IP address of the forwarding server. A forwarding server can have an FQDN or IP address.
<b>Port</b>	Enter the port number.
<b>Server Down action</b>	Select what action the FortiGate unit will take if the forwarding server is down.
<b>Enable Health Monitor</b>	Select to enable health check monitoring.
<b>Health Check Monitor Site</b>	Enter the URL address of the health check monitoring site.

Use the following CLI command to add a web proxy forwarding server named `fwd-srv` at address `proxy.example.com` and port 8080.

```
config web-proxy forward-server
 edit fwd-srv
 set addr-type fqdn
 set fqdn proxy.example.com
 set port 8080
 end
```

## Web proxy forwarding server monitoring and health checking

By default, a FortiGate unit monitors web proxy forwarding server by forwarding a connection to the remote server every 10 seconds. If the remote server does not respond it is assumed to be down. Checking continues and when the server does send a response the server is assumed to be back up. If you configure health checking, every 10 seconds the FortiGate unit attempts to get a response from a web server by connecting through the remote forwarding server.

You can configure health checking for each remote server and specify a different website to check for each one.

If the remote server is found to be down you can configure the FortiGate unit to block sessions until the server comes back up or to allow sessions to connect to their destination, bypassing the remote forwarding server. You cannot configure the FortiGate unit to fail over to another remote forwarding server.

Configure the server down action and enable health monitoring from the web-based manager by going to *System > Network > Explicit Proxy*, selecting a forwarding server, and changing the server down action and changing the health monitor settings.

Use the following CLI command to enable health checking for a web proxy forwarding server and set the server down option to bypass the forwarding server if it is down.

```
config web-proxy forward-server
 edit fwd-srv
 set healthcheck enable
 set monitor http://example.com
 set server-down-option pass
 end
```

## Adding proxy chaining to an explicit web proxy security policy

You enable proxy chaining for web proxy sessions by adding a web proxy forwarding server to an explicit web proxy security policy. In a policy you can select one web proxy forwarding server. All explicit web proxy traffic accepted by this security policy is forwarded to the specified web proxy forwarding server.

### To add an explicit web proxy forwarding server - web-based manager

1. Go to *Policy > Policy > Policy* and select Create New.
2. Configure the security policy:

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	web-proxy
<b>Source Address</b>	Internal_subnet
<b>Outgoing Interface</b>	wan1
<b>Destination Address</b>	all
<b>Schedule</b>	always
<b>Service</b>	webproxy
<b>Action</b>	ACCEPT
<b>Web Proxy Forwarding Server</b>	Select, fwd-srv

3. Select OK to save the security policy.

### To add an explicit web proxy forwarding server - CLI

1. Use the following command to add a security policy that allows all users on the 10.31.101.0 subnet to use the explicit web proxy for connections through the wan1 interface to the

```
Internet. The policy forwards web proxy sessions to a remote forwarding server named fwd-srv
config firewall policy
edit 2
set srcintf web-proxy
set dstintf wan1
set scraddr Internal_subnet
set dstaddr all
set action accept
set schedule always
set service webproxy
set webproxy-forward-server fwd-srv
end
```

## Explicit web proxy authentication

You can add identity-based policies to apply authentication to explicit web proxy sessions. You can use authentication to control access to the explicit web proxy. You can also use identity-based policies to identify users and apply different security profiles to different users.

Authentication of web proxy sessions uses HTTP basic and digest authentication as described in [RFC 2617 \(HTTP Authentication: Basic and Digest Access Authentication\)](#) and prompts the user for credentials from the browser allowing individual users to be identified by their web browser instead of IP address. HTTP authentication allows the FortiGate unit to identify multiple users accessing services from a shared IP address. You can also select IP-based authentication to authenticate users according to their source IP address.

### IP-Based authentication

IP-based authentication applies authentication by source IP address. For explicit web proxy, IP authentication is compatible with basic, digest, NTLM, form or FSSO authentication methods. Once a user authenticates, all sessions to the explicit web proxy from that IP address are assumed to be from that user and are accepted until the authentication timeout ends or the session times out.

This method of authentication is similar to standard (non-web proxy) firewall authentication and may not produce the desired results if multiple users share IP addresses (such as in a network that uses virtualization solutions or includes a NAT device between the users and the explicit web proxy).

To configure IP-based authentication, add a security policy for the explicit web proxy, set the *Policy Subtype* to *User Identity*, set the *Incoming Interface* to *web-proxy*, and make sure *IP Based* is selected before adding identity-based policies. You can also set the authentication method to basic, digest, NTLM, form or FSSO.

Use the following CLI command to add IP-based authentication to a web proxy security policy. IP-based authentication is selected by setting `ip-based` to `enable`.

```
config firewall policy
 edit 3
 set srcintf web-proxy
 set dstintf port1
 set scraddr User_network
 set dstaddr all
 set action accept
 set identity-based enable
 set ip-based enable
 config identity-based-policy
 edit 1
 set groups Internal_users
 set service ANY
 set schedule always
 end
 end
end
```

## Per session authentication

If you don't select *IP Based* the FortiGate unit applies HTTP authentication per session. This authentication is browser-based (see [Figure 409 on page 2680](#)). When a user enters a user name and password in their browser to authenticate with the explicit web proxy, this information is stored by the browser in a session cookie. Each new session started by the same web browser uses the session cookie for authentication. When the session cookie expires the user has to re-authenticate. If the user starts another browser on the same PC or closes and then re-opens their browser they have to authenticate again.

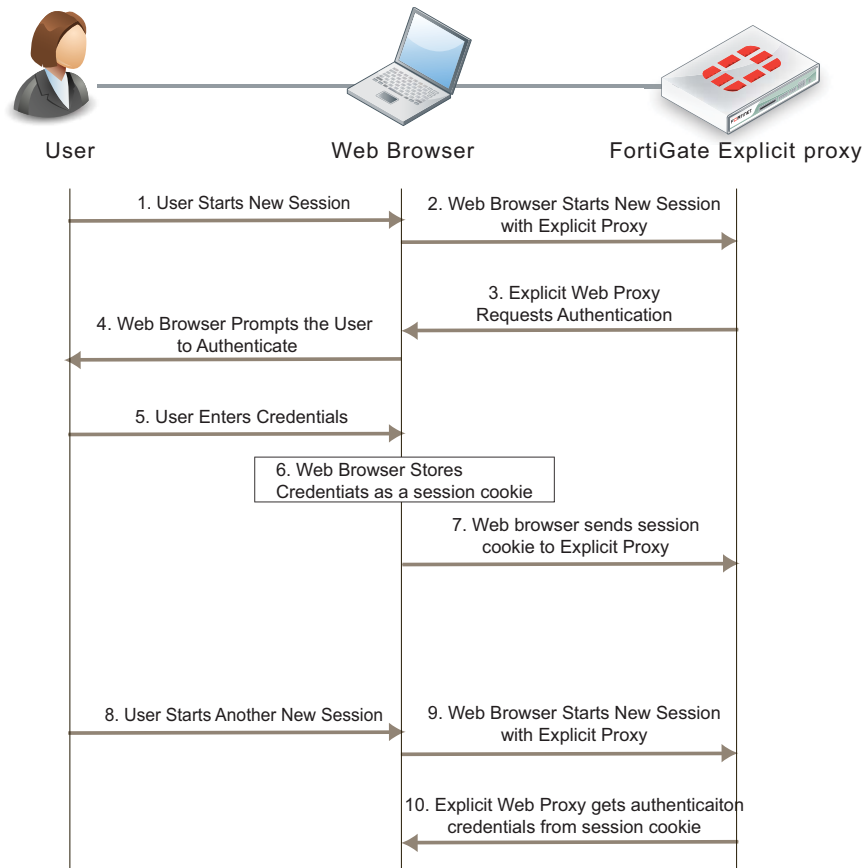
Since the authentication is browser-based, multiple clients with the same IP address can authenticate with the proxy using their own credentials. HTTP authentication provides authentication for multiple user sessions from the same source IP address. This can happen if there is a NAT device between the users and the FortiGate unit. HTTP authentication also supports authentication for other configurations that share one IP address among multiple users. These includes Citrix products and Windows Terminal Server and other similar virtualization solutions.

To configure per session authentication, add a security policy for the explicit web proxy, set the *Policy Subtype* to *User Identity*, set the *Incoming Interface* to *web-proxy*, and make sure *IP Based* is not selected before adding identity-based policies.

Use the following CLI command to add per session authentication to a security policy. Per session authentication is selected by setting `ip-based` to `disable`.

```
config firewall policy
 edit 5
 set srcintf web-proxy
 set dstintf port1
 set scraddr User_network
 set dstaddr all
 set action accept
 set identity-based enable
 set ip-based disable
 config identity-based-policy
 edit 1
 set groups Internal_users
 set service ANY
 set schedule always
 end
 end
end
```

**Figure 409:**Per session HTTP authentication



## Security profiles, client reputation, device identification, and the explicit web proxy

You can apply all security profiles to explicit web proxy sessions. This includes antivirus, web filtering, intrusion protection (IPS), application control, and data leak prevention (DLP) including DLP archiving features to explicit web proxy sessions. Security profiles are applied by selecting them in a web proxy security policy or a user identity policy in a web proxy security policy.

You can also enable client reputation for explicit web proxy policies.

The explicit web proxy is not compatible with device identification.

Since the traffic accepted by the explicit web proxy is known to be either HTTP, HTTPS, or FTP over HTTP and since the ports are already known by the proxy, the explicit web proxy does not use the HTTP or HTTPS proxy options settings. The explicit web proxy does support the following proxy options:

- Enable chunked bypass
- HTTP oversized file action and threshold

The explicit web proxy does not support the following proxy options:

- Client comforting
- Server comforting
- Monitor content information from dashboard. URLs visited by explicit web proxy users are not added to dashboard usage and log and archive statistics widgets.



For explicit web proxy sessions, the FortiGate unit applies antivirus scanning to HTTP POST requests and HTTP responses. The FortiGate unit starts virus scanning a file in an HTTP session when it receives a file in the body of an HTML request. The explicit web proxy can receive HTTP responses from either the originating web server or the FortiGate web cache module.

Flow-based virus scanning is not available for explicit web proxy sessions. Even if the FortiGate unit is configured to use flow-based antivirus, explicit web proxy sessions use the regular virus database.

## Web Proxy firewall services and service groups


Configure web proxy services by selecting *Explicit Proxy* when configuring a service. Web proxy services can only be selected in a security policy when *web-proxy* is selected as the source interface.

Web proxy services are similar to standard firewall services. You can configure web proxy services to define one or more protocols and port numbers that are associated with each web proxy service. Web proxy services can also be grouped into web proxy service groups.

### Example: users on an internal network browsing the Internet through the explicit web proxy with web caching, RADIUS authentication, web filtering and virus scanning

This example describes how to configure the explicit web proxy for the example network shown in [Figure 410](#). In this example, users on the internal network connect to the explicit web proxy through the Internal interface of the FortiGate unit. The explicit web proxy is configured to use port 8888 so users must configure their web browser proxy settings to use port 8888 and IP address 10.31.101.100.

**Figure 410:**Example explicit web proxy network topology



Explicit web proxy users must authenticate with a RADIUS server before getting access to the proxy. The security policy that accepts explicit web proxy traffic applies per session authentication and includes a RADIUS server user group. The identity based policy also applies web filtering and virus scanning.

## General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Enable the explicit web proxy for HTTP and HTTPS and change the HTTP and HTTPS ports to 8888.
2. Enable the explicit web proxy on the internal interface.
3. Add a RADIUS server and user group for the explicit web proxy.
4. Add a user identity security policy for the explicit web proxy.  
Enable web caching  
Add an authentication rule and enable antivirus and web filtering.

## Configuring the explicit web proxy - web-based manager

Use the following steps to configure the explicit web proxy from FortiGate web-based manager.

### To enable and configure the explicit web proxy

1. Go to *System > Network > Explicit Proxy* and change the following settings:

<b>Enable Explicit Web Proxy</b>	Select <i>HTTP/HTTPS</i> .
<b>Listen on Interfaces</b>	No change. This field will eventually show that the explicit web proxy is enabled for the Internal interface.
<b>HTTP Port</b>	8888
<b>HTTPS Port</b>	8888
<b>Realm</b>	You are authenticating with the explicit web proxy.
<b>Default Firewall Policy Action</b>	Deny

2. Select *Apply*.

### To enable the explicit web proxy on the Internal interface

1. Go to *System > Network > Interface*.
2. Edit the internal interface.
3. Select *Enable Explicit Web Proxy*.
4. Select *OK*.

### To add a RADIUS server and user group for the explicit web proxy

1. Go to *User & Device > Authentication > RADIUS Server* and select *Create New* to add a new RADIUS server:

<b>Name</b>	RADIUS_1
<b>Primary Server Name/IP</b>	10.31.101.200
<b>Primary Server Secret</b>	RADIUS_server_secret

2. Select *OK*.

3. Go to *User & Device > User > User Group* and select *Create New* to add a new user group.

<b>Name</b>	Explicit_proxy_user_group
<b>Type</b>	Firewall
<b>Remote authentication servers</b>	RADIUS_1
<b>Group Name</b>	Any

4. Select OK.

#### To add a security policy for the explicit web proxy

1. Go to *Firewall Objects > Address > Addresses* and select *Create New*.
2. Add a firewall address for the internal network:

<b>Address Name</b>	Internal_subnet
<b>Type</b>	Subnet / IP Range
<b>Subnet / IP Range</b>	10.31.101.[1-255]
<b>Interface</b>	Any

3. Go to *Policy > Policy > Policy* and select *Create New*.
4. Configure the explicit web proxy security policy.

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	User Identity
<b>Incoming Interface</b>	web-proxy
<b>Source Address</b>	Internal_subnet
<b>Outgoing Interface</b>	wan1
<b>Destination Address</b>	all
<b>Service</b>	webproxy

5. Select *Enable Web cache*.
6. Under *Configure Authentication Rules* select *Create New* to add an authentication rule:

<b>Groups</b>	Explicit_policy
<b>Users</b>	Leave blank
<b>Schedule</b>	always
<b>Action</b>	ACCEPT

7. Turn on *Antivirus* and *Web Filter* and select the *default* profiles for both.
8. Select the *default proxy options* profile.
9. Select OK.

10. Make sure *IP Based* is not selected.

11. Select OK.

## Configuring the explicit web proxy - CLI

Use the following steps to configure the example explicit web proxy configuration from the CLI.

### To enable the explicit web proxy on the Internal interface

1. Enter the following command to enable the explicit web proxy on the internal interface.

```
config system interface
 edit internal
 set explicit-web-proxy enable
 end
```

### To enable and configure the explicit web proxy

1. Enter the following command to enable the explicit web proxy and set the TCP port that proxy accepts HTTP and HTTPS connections on to 8888.

```
config web-proxy explicit
 set status enable
 set http-incoming-port 8888
 set https-incoming-port 8888
 set realm "You are authenticating with the explicit web proxy"
 set sec-default-action deny
end
```

### To add a RADIUS server and user group for the explicit web proxy

1. Enter the following command to add a RADIUS server:

```
config user radius
 edit RADIUS_1
 set server 10.31.101.200
 set secret RADIUS_server_secret
 end
```

2. Enter the following command to add a user group for the RADIUS server.

```
config user group
 edit Explicit_proxy_user_group
 set group-type firewall
 set member RADIUS_1
 end
```

### To add a security policy for the explicit web proxy

1. Enter the following command to add a firewall address for the internal subnet:

```
config firewall address
 edit Internal_subnet
 set type iprange
 set start-ip 10.31.101.1
 set end-ip 10.31.101.255
 end
```

2. Enter the following command to add the explicit web proxy security policy:

```
config firewall policy
 edit 0
 set srcintf web-proxy
 set dstintf wan1
 set srcaddr Internal_subnet
 set dstaddr all
 set action accept
 set service webproxy
 set webcache enable
 set identity-based enable
 set ipbased disable
 set auth-method basic
 config identity-based-policy
 edit 1
 set groups Explicit_Proxy_user_group
 set schedule always
 set utm-status enable
 set av-profile default
 set webfilter-profile default
 set profile-protocol-options default
 end
 end
 end
```

## Testing and troubleshooting the configuration

You can use the following steps to verify that the explicit web proxy configuration is working as expected:

### To test the explicit web proxy configuration

1. Configure a web browser on the internal subnet to use a web proxy server at IP address 10.31.101.100 and port 8888.
2. Browse to an Internet web page.  
The web browser should pop up an authentication window that includes the phrase that you added to the Realm option.
3. Enter the username and password for an account on the RADIUS server.  
If the account is valid you should be allowed to browse web pages on the Internet.
4. Close the browser and clear its cache and cookies.
5. Restart the browser and connect to the Internet.  
You could also start a second web browser on the same PC. Or you could start a new instance of the same browser as long as the browser asks for a user name and password again.  
You should have to authenticate again because identity-based policies are set to session-based authentication.
6. If this basic functionality does not work, check your FortiGate and web browser configuration settings.
7. Browse to a URL on the URL filter list and confirm that the web page is blocked.
8. Browse to <http://eicar.org> and attempt to download an anti-malware test file.  
The antivirus configuration should block the file.

Sessions for web-proxy security policies do not appear on the Top Sessions dashboard widget and the count column for security policies does not display a count for explicit web proxy security policies.

9. You can use the following command to display explicit web proxy sessions

```
get test wad 60
```

```
IP based users:
```

```
Session based users:
```

```
user:0x9c20778, username:User1, vf_id:0, ref_cnt:9
```

```
Total allocated user:1
```

```
Total user count:3, shared user quota:50, shared user count:3
```

This command output shows one explicit proxy user with user name `User1` authenticated using session-based authentication.

## Explicit proxy sessions and user limits

Web browsers and web servers open and close multiple sessions with the explicit web proxy. Some sessions open and close very quickly. HTTP 1.1 keepalive sessions are persistent and can remain open for long periods of time. Sessions can remain on the explicit web proxy session list after a user has stopped using the proxy (and has, for example, closed their browser). If an explicit web proxy session is idle for more than 3600 seconds it is torn down by the explicit web proxy. See [RFC 2616](#) for information about HTTP keepalive/persistent HTTP sessions.

This section describes proxy sessions and user limits for both the explicit web proxy and the explicit FTP proxy. Session and user limits for the two proxies are counted and calculated together. However, in most cases if both proxies are active there will be many more web proxy sessions than FTP proxy sessions.

The FortiGate unit adds two sessions to its session table for every explicit proxy session started by a web browser and every FTP session started by an FTP client. An entry is added to the session table for the session from the web browser or client to the explicit proxy. All of these sessions have the same destination port as the explicit web proxy port (usually 8080 for HTTP and 21 for FTP). An entry is also added to the session table for the session between the exiting FortiGate interface and the web or FTP server destination of the session. All of these sessions have a FortiGate interface IP address and the source address of the session and usually have a destination port of 80 for HTTP and 21 for FTP.

Proxy sessions that appear in the Top sessions dashboard widget do not include the Policy ID of the web-proxy or ftp-proxy security policy that accepted them. However, the explicit proxy sessions appear in the Top Sessions dashboard widget with a destination port that matches the explicit proxy port number (usually 8080 for the web proxy and 21 for the FTP proxy). The proxied sessions from the FortiGate unit have their source address set to the IP address of the FortiGate unit interface that the sessions use to connect to their destinations (for example, for connections to the Internet the source address would be the IP address of the FortiGate interface connected to the Internet).

FortiOS limits the number of explicit proxy users. This includes both explicit FTP proxy and explicit web proxy users. The number of users varies by FortiGate model from as low as 10 to up to 18000 for high end models. You can use the following command to display the limit on the number of explicit web proxy users for a FortiGate unit:

```
get test wad 62
```

```
Total user count:1, shared user quota:500, shared user count:1
form_auth_keepalive=0 vd=root max=0 guarantee=0 used=1
```

This command output shows that the explicit proxy user limit (the `shared user quota`) for this FortiGate unit is 500 users.

You cannot change this limit. If your FortiGate unit is configured for multiple VDOMs this limit must be shared by all VDOMs. You can also use VDOM resource limiting to limit the number of explicit proxy users for the FortiGate unit and for each VDOM. To limit the number of explicit proxy users for the FortiGate unit from the web-based manager enable multiple VDOMs and go to *System > VDOM > Global Resources* set the number of Concurrent explicit proxy users or use the following command:

```
config global
 config system resource-limits
 set proxy 50
 end
end
```

To limit the number of explicit proxy users for a VDOM, from the web-based manager enable multiple VDOMs and go to *System > VDOM > VDOM* and edit a VDOM or use the following command to change the number of explicit web proxy users for VDOM\_1:

```
config global
 config system vdom-property
 edit VDOM_1
 set proxy 25
 end
end
```

The VDOM resource limit pages on the web-based manager also display the current number of explicit web proxy users. You can also use the `get test wad 60` CLI command to view the number of explicit web proxy users. For example:

```
get test wad 60
IP based users:
 user:0x9ab8350 username:User1, vf_id:0, ip_addr:10.31.101.10,
 ref_cnt:9
```

```
Session based users:
 user:0x9ac3c40, username:User2, vf_id:0, ref_cnt:3
 user:0x9ab94f0, username:User3, vf_id:0, ref_cnt:1
```

```
Total allocated user:3
```

```
Total user count:3, shared user quota:50, shared user count:3
```

Users may be displayed with this command even if they are no longer actively using the proxy. All idle sessions time out after 3600 seconds.

The command output shows three explicit proxy users. The user named User1 has authenticated with a security policy that includes IP-based authentication and the user's source IP address is 10.31.101.10. The users named User2 and User3 have authenticated with a security policy that includes session-based authentication.

You can use the following command to flush all current explicit proxy users. This means delete information about all users and force them re-authenticate.

```
get test wad 61
```



Users that authenticate with explicit web-proxy or ftp-proxy security policies do not appear in the *User & Device > Monitor > Firewall* list and selecting *De-authenticate All Users* has no effect on explicit proxy users.

---

How the number of concurrent explicit proxy users is determined depends on their authentication method:

- For session-based authenticated users, each authenticated user is counted as a single user. Since multiple users can have the same user name, the proxy attempts to identify users according to their authentication membership (based upon whether they were authenticated using RADIUS, LADAP, FSAE, local database etc.). If a user of one session has the same name and membership as a user of another session, the explicit proxy assumes this is one user.
- For IP Based authentication, or no authentication, or if no web-proxy security policy has been added, the source IP address is used to determine a user. All sessions from a single source address are assumed to be from the same user.

The explicit proxy does not limit the number of active sessions for each user. As a result the actual explicit proxy session count is usually much higher than the number of explicit web proxy users. If an excessive number of explicit web proxy sessions is compromising system performance you can limit the amount of users if the FortiGate unit is operating with multiple VDOMs.



# The FortiGate explicit FTP proxy

You can use the FortiGate explicit FTP proxy to enable explicit FTP proxying on one or more FortiGate interfaces. The explicit web and FTP proxies can be operating at the same time on the same or on different FortiGate interfaces.



Explicit FTP proxies are configured for each VDOM when multiple VDOMs are enabled.

In most cases you would configure the explicit FTP proxy for users on a network by enabling the explicit FTP proxy on the FortiGate interface connected to that network. Users on the network would connect to and authenticate with the explicit FTP proxy before connecting to an FTP server. In this case the IP address of the explicit FTP proxy is the IP address of the FortiGate interface on which the explicit FTP proxy is enabled.




Enabling the explicit FTP proxy on an interface connected to the Internet is a security risk because anyone on the Internet who finds the proxy could use it to hide their source address.

If the FortiGate unit is operating in Transparent mode, users would configure their browsers to use a proxy server with the FortiGate unit management IP address.

The FTP proxy receives FTP sessions to be proxied at FortiGate interfaces with the explicit FTP proxy enabled. The FTP proxy uses FortiGate routing to route sessions through the FortiGate unit to a destination interface. Before a session leaves the exiting interface, the explicit FTP proxy changes the source addresses of the session packets to the IP address of the exiting interface. When the FortiGate unit is operating in Transparent mode the explicit web proxy changes the source addresses to the management IP address.

**Figure 411:**Example explicit FTP proxy topology



To allow anyone to anonymously log into explicit FTP proxy and connect to any FTP server you can set the explicit FTP proxy default firewall proxy action to accept. When you do this, users can log into the explicit FTP proxy with any username and password.

In most cases you would want to use security policies to control explicit FTP proxy traffic and apply security features such as access control/authentication, security profiles, and traffic logging. You can do this by keeping the default explicit FTP proxy firewall policy action to deny and then adding ftp-proxy security policies. In most cases you would also want users to authenticate with the explicit FTP proxy. By default an anonymous FTP login is required. Usually you would add authentication, in the form of identity based policies, to ftp-proxy security policies. Users can then authenticate with the explicit FTP proxy according to user groups added to the identity based policies. User groups added to FTP proxy identity based policies can use any authentication method supported by FortiOS including the local user database and RADIUS and other remote servers.

If you leave the default firewall policy action set to deny and add ftp-proxy security policies, all connections to the explicit FTP proxy must match an ftp-proxy security policy or else they will be dropped. Sessions that are accepted are processed according to the ftp-proxy security policy settings.

You can also change the explicit FTP proxy default firewall policy action to accept and add explicit FTP proxy security policies. If you do this, sessions that match ftp-proxy security policies are processed according to the security policy settings. Connections to the explicit FTP proxy that do not match an ftp-proxy security policy are allowed and the users can authenticate with the proxy anonymously user any username and password.

There are some limitations to the security profile features that can be applied to explicit web proxy sessions. See [“Security profiles, client reputation, device identification, and the explicit FTP proxy” on page 2697](#).

You cannot configure IPsec, SSL VPN, or Traffic shaping for explicit FTP proxy traffic. Security policies for the FTP proxy can only include firewall addresses not assigned to a FortiGate unit interface or with interface set to *any*. (On the web-based manager you must set the interface to *Any*. In the CLI you must unset the associated-interface.)

This section describes:

- [How to use the explicit FTP proxy to connect to an FTP server](#)
- [Explicit FTP proxy configuration overview](#)
- [Security profiles, client reputation, device identification, and the explicit FTP proxy](#)
- [Example: users on an internal network connecting to FTP servers on the Internet through the explicit FTP with RADIUS authentication and virus scanning](#)
- [Explicit FTP proxy sessions and user limits](#)

## How to use the explicit FTP proxy to connect to an FTP server

To connect to an FTP server using the explicit FTP proxy, users must run an FTP client and connect to the IP address of a FortiGate interface on which the explicit FTP proxy is enabled. This connection attempt must use the configured explicit FTP proxy port number (default 21).

The explicit FTP proxy is not compatible with using a web browser as an FTP client. To use web browsers as FTP clients configure the explicit web proxy to accept FTP sessions.

The following steps occur when a user starts an FTP client to connect to an FTP server using the explicit FTP proxy. Any RFC-compliant FTP client can be used. This example describes using a command-line FTP client. Some FTP clients may require a custom FTP proxy connection script.

1. The user enters a command on the FTP client to connect to the explicit FTP proxy.

For example, if the IP address of the FortiGate interface on which the explicit FTP proxy is enabled is 10.31.101.100, enter:

```
ftp 10.31.101.100
```

2. The explicit FTP proxy responds with a welcome message and requests the user's FTP proxy user name and password and a username and address of the FTP server to connect to:

```
Connected to 10.31.101.100.
220 Welcome to Fortigate FTP proxy
Name (10.31.101.100:user):
```

You can change the message by editing the FTP Explicit Banner Message replacement message.

3. At the prompt the user enters their FTP proxy username and password and a username and address for the FTP server. The FTP server address can be a domain name or numeric IP address. This information is entered using the following syntax:

```
<proxy-user>:<proxy-password>:<server-user>@<server-address>
```

For example, if the proxy username and password are `p-name` and `p-pass` and a valid username for the FTP server is `s-name` and the server's IP address is `ftp.example.com` the syntax would be:

```
p-name:p-pass:s-name@ftp.example.com
```



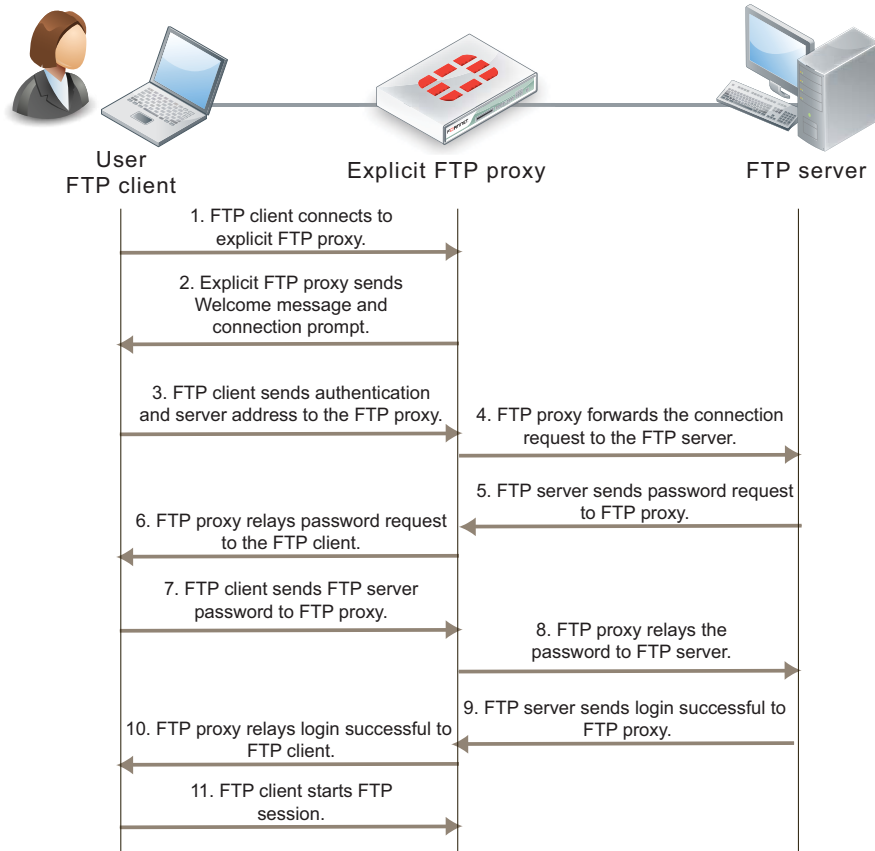
If the FTP proxy accepts anonymous logins `p-name` and `p-pass` can be any characters.

---

4. The FTP proxy forwards the connection request, including the user name, to the FTP server.
5. If the user name is valid for the FTP server it responds with a password request prompt.
6. The FTP proxy relays the password request to the FTP client.
7. The user enters the FTP server password and the client sends the password to the FTP proxy.
8. The FTP proxy relays the password to the FTP server.
9. The FTP server sends a login successful message to the FTP proxy.
10. The FTP proxy relays the login successful message to the FTP client.
11. The FTP client starts the FTP session.

All commands entered by the client are relayed by the proxy to the server. Replies from the server are relayed back to the FTP client.

**Figure 412:**Explicit FTP proxy session



From a simple command line FTP client connecting to an the previous sequence could appear as follows:

```
ftp 10.31.101.100 21
Connected to 10.31.101.100.
220 Welcome to Fortigate FTP proxy
Name (10.31.101.100:user): p-name:p-pass:s-name@ftp.example.com
331 Please specify the password.
Password: s-pass
230 Login successful.
Remote system type is UNIX
Using binary mode to transfer files.
ftp>
```

## Explicit FTP proxy configuration overview

This section describes:

- [General configuration steps](#)
- [Restricting the IP address of the explicit FTP proxy](#)
- [Restricting the outgoing source IP address of the explicit FTP proxy](#)

### General configuration steps

You can use the following general steps to configure the explicit FTP proxy.

### To enable the explicit FTP proxy - web-based manager

1. Go to *System > Network > Explicit Proxy > Explicit FTP Proxy Options*. Select *Enable Explicit FTP Proxy* to turn on the explicit FTP proxy.

2. Select *Apply*.

The default explicit FTP proxy configuration has *Default Firewall Policy Action* set to *Deny* and requires you to add a security policy to allow access to the explicit FTP proxy. This configuration is recommended and is a best practice because you can use security policies to control access to the explicit web proxy and also apply security features such as logging, security profiles, and authentication (by adding identity-based policies).

3. Go to *System > Network > Interface* and select one or more interfaces for which to enable the explicit web proxy. Edit the interface configuration and select *Enable Explicit FTP Proxy*.



Enabling the explicit FTP proxy on an interface connected to the Internet is a security risk because anyone on the Internet who finds the proxy could use it to hide their source address. If you enable the proxy on such an interface make sure authentication is required to use the proxy.

---

4. Go to *Policy > Policy > Policy* and select *Create New* and set the *Source Interface/Zone* to *ftp-proxy*.

You can add multiple ftp-proxy security policies.

- Configure the security policy as required to accept the traffic that you want to be processed by the explicit web proxy.

The source address of the policy should match client source IP addresses. The firewall address selected as the source address cannot be assigned to a FortiGate interface. The Interface field of the firewall address must be blank or it must be set to *Any*.

The destination address of the policy should match the IP addresses of FTP servers that clients are connecting to. The destination address could be *all* to allow connections to any FTP server.

If *Default Firewall Policy Action* is set to Deny, traffic sent to the explicit FTP proxy that is not accepted by an ftp-proxy security policy is dropped. If *Default Firewall Policy Action* is set to Allow then all web-proxy sessions that don't match with a security policy are allowed.

For example the following security policy allows users on an internal network to access FTP servers on the Internet through the wan1 interface of a FortiGate unit.

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	Address
<b>Incoming Interface</b>	ftp-proxy
<b>Source Address</b>	Internal_subnet
<b>Outgoing Interface</b>	wan1
<b>Destination Address</b>	all
<b>Schedule</b>	always
<b>Action</b>	ACCEPT

The following security policy requires users on an internal network to authenticate with the FortiGate unit before accessing FTP servers on the Internet through the wan1 interface.

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	User Identity
<b>Incoming Interface</b>	ftp-proxy
<b>Source Address</b>	Internal_subnet
<b>Outgoing Interface</b>	wan1
<b>Destination Address</b>	all

Select *Create New* to add an *Authentication Rule* and configure the rule as follows:

<b>Groups</b>	Proxy-Group
<b>Users</b>	(optional)
<b>Schedule</b>	always
<b>Action</b>	ACCEPT

Add security profiles as required and select OK.

You can add multiple user identity policies to apply different authentication for different user groups and users and also apply different security profiles and logging settings for different user groups.

6. You can select other security policy options as required.

For example, you can apply security profiles to web proxy sessions and log allowed ftp proxy traffic.

7. Select OK.

### To enable the explicit web proxy - CLI

1. Enter the following command to turn on the explicit FTP proxy. This command also changes the explicit FTP proxy port to 2121.

```
config ftp-proxy explicit
 set status enable
 set incoming-port 2121
end
```

The default explicit FTP proxy configuration has `sec-default-action` set to `deny` and requires you to add a security policy to allow access to the explicit FTP proxy.

2. Enter the following command to enable the explicit FTP proxy for the internal interface.

```
config system interface
 edit internal
 set explicit-ftp-proxy enable
 end
end
```

3. Use the following command to add a firewall address that matches the source address of users who connect to the explicit FTP proxy.

```
config firewall address
 edit Internal_subnet
 set type iprange
 set start-ip 10.31.101.1
 set end-ip 10.31.101.255
 end
```

The source address for a ftp-proxy security policy cannot be assigned to a FortiGate unit interface.

4. Use the following command to add a security policy that allows all FTP proxy users on the 10.31.101.0 subnet to use the explicit FTP proxy for connections through the wan1 interface to the Internet. This policy also applies virus scanning to FTP proxy traffic.

```
config firewall policy
 edit 2
 set srcintf ftp-proxy
 set dstintf wan1
 set scraddr Internal_subnet
 set dstaddr all
 set action accept
 set schedule always
 set utm-status enable
 set profile-protocol-options default
 set av-profile default
 end
```

The following command requires FTP proxy users to authenticate with the FortiGate unit before accessing FTP servers on the Internet.

```
config firewall policy
 edit 2
 set srcintf ftp-proxy
 set dstintf wan1
 set scraddr Internal_subnet
 set dstaddr all
 set action accept
 set identity-based enable
 config identity-based-policy
 edit 1
 set groups Internal_users
 set schedule always
 set utm-status enable
 set profile-protocol-options default
 set av-profile default
 end
 end
 end
```

## Restricting the IP address of the explicit FTP proxy

You can use the following command to restrict access to the explicit FTP proxy using only one IP address. The IP address that you specify must be the IP address of an interface that the explicit FTP proxy is enabled on. You might want to use this option if the explicit FTP proxy is enabled on an interface with multiple IP addresses.

For example, to require users to connect to the IP address 10.31.101.100 to connect to the explicit FTP proxy:

```
config ftp-proxy explicit
 set incoming-ip 10.31.101.100
end
```



## Restricting the outgoing source IP address of the explicit FTP proxy

You can use the following command to restrict the source address of outgoing FTP proxy packets to a single IP address. The IP address that you specify must be the IP address of an interface that the explicit FTP proxy is enabled on. You might want to use this option if the explicit FTP proxy is enabled on an interface with multiple IP addresses.

For example, to restrict the outgoing packet source address to 172.20.120.100:

```
config ftp-proxy explicit
 set outgoing-ip 172.20.120.100
end
```

## Security profiles, client reputation, device identification, and the explicit FTP proxy

You can apply security profiles to explicit FTP proxy sessions. This includes antivirus, intrusion protection (IPS), application control, and data leak prevention (DLP) including DLP archiving. Security profiles are applied by selecting them in a ftp proxy security policy or an identity based policy in a FTP proxy security policy.

You can also enable client reputation for explicit FTP proxy policies.

The explicit FTP proxy is not compatible with device identification.

## Explicit FTP proxy sessions and protocol options

Since the traffic accepted by the explicit FTP proxy is known to be FTP and since the ports are already known by the proxy, the explicit FTP proxy does not use the FTP port protocol options settings.

When adding security profiles to an FTP proxy security policy, you must select a protocol options profile. In most cases you can select the default protocol options profile. You could also create a custom protocol options profile.

The explicit FTP proxy supports the following protocol options:

- FTP oversized file action and threshold

The explicit FTP proxy does not support the following protocol options:

- Client comforting
- Server comforting
- Monitor content information from dashboard. URLs visited by explicit FTP proxy users are not added to dashboard usage and log and archive statistics widgets.

## Explicit FTP proxy sessions and antivirus


For explicit FTP proxy sessions, the FortiGate unit applies antivirus scanning to FTP file GET and PUT requests. The FortiGate unit starts virus scanning a file in an FTP session when it receives a file in the body of an FTP request.

Flow-based virus scanning is not available for explicit FTP proxy sessions. Even if the FortiGate unit is configured to use flow-based antivirus, explicit FTP proxy sessions use the regular virus database.

## Example: users on an internal network connecting to FTP servers on the Internet through the explicit FTP with RADIUS authentication and virus scanning

This example describes how to configure the explicit FTP proxy for the example network shown in Figure 413. In this example, users on the internal network connect to the explicit FTP proxy through the Internal interface with IP address 10.31.101.100 of the FortiGate-51B unit. The explicit web proxy is configured to use port 2121 so to connect to an FTP server on the Internet users must first connect to the explicit FTP proxy using IP address 10.31.101.100 and port 2121.

**Figure 413:**Example explicit FTP proxy network topology



In this example, explicit FTP proxy users must authenticate with a RADIUS server before getting access to the proxy. To apply authentication, the security policy that accepts explicit FTP proxy traffic includes an identity based policy that applies per session authentication to explicit FTP proxy users and includes a user group with the RADIUS server in it. The identity based policy also applies virus scanning and DLP.

### General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Enable the explicit FTP proxy and change the FTP port to 2121.
2. Enable the explicit FTP proxy on the internal interface.
3. Add a RADIUS server and user group for the explicit FTP proxy.
4. Add a user identity security policy for the explicit FTP proxy.  
Enable antivirus and DLP features for the identity-based policy.

### Configuring the explicit FTP proxy - web-based manager

Use the following steps to configure the explicit FTP proxy from FortiGate web-based manager.

### To enable and configure the explicit FTP proxy

1. Go to *System > Network > Explicit Proxy > Explicit FTP Proxy Options* and change the following settings:

<b>Enable Explicit FTP Proxy</b>	Select.
<b>Listen on Interface</b>	No change. This field will eventually show that the explicit web proxy is enabled for the Internal interface.
<b>FTP Port</b>	2121
<b>Default Firewall Policy Action</b>	Deny

2. Select *Apply*.

### To enable the explicit FTP proxy on the Internal interface

1. Go to *System > Network > Interface*.
2. Edit the internal interface.
3. Select *Enable Explicit FTP Proxy*.
4. Select *OK*.

### To add a RADIUS server and user group for the explicit FTP proxy

1. Go to *User > Remote > RADIUS*.
2. Select *Create New* to add a new RADIUS server:

<b>Name</b>	RADIUS_1
<b>Primary Server Name/IP</b>	10.31.101.200
<b>Primary Server Secret</b>	RADIUS_server_secret

3. Go to *User > User Group > User Group* and select *Create New*.

<b>Name</b>	Explicit_proxy_user_group
<b>Type</b>	Firewall
<b>Remote authentication servers</b>	RADIUS_1
<b>Members</b>	RADIUS_1

4. Select *OK*.

### To add a security policy for the explicit FTP proxy

1. Go to *Firewall Objects > Address > Address* and select *Create New*.
2. Add a firewall address for the internal network:

<b>Address Name</b>	Internal_subnet
<b>Type</b>	Subnet / IP Range

<b>Subnet / IP Range</b>	10.31.101.[1-255]
<b>Interface</b>	Any

3. Go to *Policy > Policy > Policy* and select *Create New*.
4. Configure the explicit FTP proxy security policy.

<b>Policy Type</b>	Firewall
<b>Policy Subtype</b>	User Identity
<b>Incoming Interface</b>	ftp-proxy
<b>Source Address</b>	Internal_subnet
<b>Outgoing Interface</b>	wan1
<b>Destination Address</b>	all

5. Select *Enable Identity Based Policy*, make sure *IP Based* is not selected and *Auth Method* is set to *Basic*.
6. Under *Configure Authentication Rules* select *Create New* to add an authentication rule:

<b>Groups</b>	Explicit_policy
<b>Users</b>	Leave blank
<b>Schedule</b>	always
<b>Action</b>	ACCEPT

7. Select *Antivirus* and *Web Filter* and select the *default* profiles for both.
8. Select the *default* proxy options profile.
9. Select OK.
10. Select OK.

## Configuring the explicit FTP proxy - CLI

Use the following steps to configure the example explicit web proxy configuration from the CLI.

### To enable and configure the explicit FTP proxy

1. Enter the following command to enable the explicit FTP proxy and set the TCP port that proxy accepts FTP connections on to 2121.

```
config ftp-proxy explicit
 set status enable
 set incoming-port 2121
 set sec-default-action deny
end
```

### To enable the explicit FTP proxy on the Internal interface

1. Enter the following command to enable the explicit FTP proxy on the internal interface.

```
config system interface
 edit internal
 set explicit-ftp-proxy enable
 end
```

### To add a RADIUS server and user group for the explicit FTP proxy

1. Enter the following command to add a RADIUS server:

```
config user radius
 edit RADIUS_1
 set server 10.31.101.200
 set secret RADIUS_server_secret
 end
```

2. Enter the following command to add a user group for the RADIUS server.

```
config user group
 edit Explicit_proxy_user_group
 set group-type firewall
 set member RADIUS_1
 end
```

### To add a security policy for the explicit FTP proxy

1. Enter the following command to add a firewall address for the internal subnet:

```
config firewall address
 edit Internal_subnet
 set type iprange
 set start-ip 10.31.101.1
 set end-ip 10.31.101.255
 end
```

2. Enter the following command to add the explicit FTP proxy security policy:

```
config firewall policy
 edit 0
 set srcintf web-proxy
 set dstintf wan1
 set srcaddr Internal_subnet
 set dstaddr all
 set action accept
 set identity-based enable
 set auth-method basic
 config identity-based-policy
 edit 1
 set groups Explicit_Proxy_user_group
 set schedule always
 set utm-status enable
 set av-profile default
 set dlp-sensor default
 set profile-protocol-options default
 end
 end
 end
```

## Testing and troubleshooting the configuration

You can use the following steps to verify that the explicit FTP proxy configuration is working as expected. These steps use a command line FTP client.

### To test the explicit web proxy configuration

1. From a system on the internal network start an FTP client and enter the following command to connect to the FTP proxy:

```
ftp 10.31.101.100
```

The explicit FTP proxy should respond with a message similar to the following:

```
Connected to 10.31.101.100.
```

```
220 Welcome to Fortigate FTP proxy
```

```
Name (10.31.101.100:user) :
```

2. At the prompt enter a valid username and password for the RADIUS server followed by a user name for an FTP server on the Internet and the address of the FTP server. For example, if a valid username and password on the RADIUS server is `ex_name` and `ex_pass` and you attempt to connect to an FTP server at `ftp.example.com` with user name `s_name`, enter the following at the prompt:

```
Name (10.31.101.100:user) :ex_name:ex_pass:s_name@ftp.example.com
```

3. You should be prompted for the password for the account on the FTP server.
4. Enter the password and you should be able to connect to the FTP server.
5. Attempt to explore the FTP server file system and download or upload files.
6. To test security profiles functionality, attempt to upload or download an ECAR test file. Or upload or download a tex file containing text that would be matched by the DLP sensor. For eicar test files, go to <http://eicar.org>.

## Explicit FTP proxy sessions and user limits

FTP clients do not open large numbers of sessions with the explicit FTP proxy. Most sessions stay open for a short while depending on how long a user is connected to an FTP server and how large the file uploads or downloads are. So unless you have large numbers of FTP users, the explicit FTP proxy should not be adding large numbers of sessions to the session table.

Explicit FTP proxy sessions and user limits are combined with explicit web proxy session and user limits. For information about explicit proxy session and user limits, see [“Explicit proxy sessions and user limits” on page 2686](#).

# FortiGate WCCP

The Web Cache Communication Protocol (WCCP) can be used to provide web caching with load balancing and fault tolerance. In a WCCP configuration, a WCCP server receives HTTP requests from user's web browsers and redirects the requests to one or more WCCP clients. The clients either return cached content or request new content from the destination web servers before caching it and returning it to the server which in turn returns the content to the original requestor. If a WCCP configuration includes multiple WCCP clients, the WCCP server load balances traffic among the clients and can detect when a client fails and failover sessions to still operating clients. WCCP is described by the [Web Cache Communication Protocol internet draft](#).

The sessions that are cached by WCCP depend on the configuration of the WCCP clients. If the client is a FortiGate unit, you can configure the port numbers and protocol number of the sessions to be cached. For example, to cache HTTPS traffic on port 443 the WCCP client port must be set to 443 and protocol must be set to 6. If the WCCP client should also cache HTTPS traffic on port 993 the client ports option should include both port 443 and 993.

On a FortiGate unit, WCCP sessions are accepted by a security policy before being cached. If the security policy that accepts sessions that do not match the port and protocol settings in the WCCP clients the traffic is dropped.

WCCP is configured per-VDOM. A single VDOM can operate as a WCCP server or client (not both at the same time). FortiGate units are compatible with third-party WCCP clients and servers. If a FortiGate unit is operating as an Internet firewall for a private network, you can configure it to cache and serve some or all of the web traffic on the private network using WCCP by adding one or more WCCP clients, configuring WCCP server settings on the FortiGate unit and adding WCCP security policies that accept HTTP session from the private network.

FortiGate units support WCCPv1 and WCCPv2. A FortiGate unit in NAT/Route or transparent mode can operate as a WCCP server. To operate as a WCCP client a FortiGate unit must be in NAT/Route mode. FortiGate units communicate between WCCP servers and clients over UDP port 2048. This communication can be encapsulated in a GRE tunnel or just use layer 2 forwarding.



A WCCP server can also be called a WCCP router. A WCCP client can also be called a WCCP cache engine.

---

This section describes:

- [WCCP service groups, service numbers, service IDs and well known services](#)
- [WCCP configuration overview](#)
- [Example: caching HTTP sessions on port 80 using WCCP](#)
- [Example: caching HTTP sessions on port 80 and HTTPS sessions on port 443 using WCCP](#)
- [WCCP packet flow](#)
- [Configuring the forward and return methods and adding authentication](#)
- [WCCP Messages](#)
- [Troubleshooting WCCP](#)

## WCCP service groups, service numbers, service IDs and well known services

A FortiGate unit configured as a WCCP server or client can include multiple server or client configurations. Each of these configurations is called a WCCP service group. A service group consists of one or more WCCP servers (or routers) and one or more WCCP clients working together to cache a specific type of traffic. The service group configuration includes information about the type of traffic to be cached, the addresses of the WCCP clients and servers and other information about the service.

A service group is identified with a numeric WCCP service ID (or service number) in the range 0 to 255. All of the servers and clients in the same WCCP service group must have service group configurations with the same WCCP service ID.

The value of the service ID provides some information about the type of traffic to be cached by the service group. Service IDs in the range 0 to 50 are reserved for well known services. A well known service is any service that is defined by the WCCP standard as being well known. Since the service is well known, just the service ID is required to identify the traffic to be cached.

Even though the well known service ID range is 0 to 50, at this time only one well known service has been defined. Its service ID 0, which is used for caching HTTP (web) traffic.

So to configure WCCP to cache HTTP sessions you can add a service group to the WCCP router and WCCP clients with a service ID of 0. No other information about the type of traffic to cache needs to be added to the service group.

Since service IDs 1 to 50 are reserved for well know services and since these services are not defined yet, you should not add service groups with IDs in the range 1 to 50.



FortiOS does allow you to add service groups with IDs between 1 and 50. Since these service groups have not been assigned well known services; however, they will not cache any sessions. Service groups with IDs 51 to 255 allow you to set the port numbers and protocol number of the traffic to be cached. So you can use service groups with IDs 51 to 255 to cache different kinds of traffic based on port numbers and protocol number of the traffic. Service groups 1 to 50; however, do not allow you to set port numbers or protocol numbers so cannot be used to cache any traffic.

---

To cache traffic other than HTTP traffic you must add service groups with IDs in the range 51 to 255. These service group configurations must include the port numbers and protocol number of the traffic to be cached. It is the port and protocol number configuration in the service group that determines what traffic will be cached by WCCP.

### Example WCCP server and client configuration for caching HTTP sessions (service ID = 0)

Enter the following command to add a WCCP service group to a WCCP server that caches HTTP sessions. The IP address of the server is 10.31.101.100 and the WCCP clients are on the 10.31.101.0 subnet. The service

ID of this service group is 0.

```
config system wccp
 edit 0
 set router-id 10.31.101.100
 set server-list 10.31.101.0 255.255.255.0
 end
```



Enter the following commands to configure a FortiGate unit to operate as a WCCP client and add a service group that configures the client to cache HTTP sessions. The IP address of the server is 10.31.101.100 and IP address of this WCCP clients is 10.31.101.1 subnet. The service ID of this service group is 0.

```
config system settings
 set wccp-cache-engine enable
end

config system wccp
 edit 0
 set cache-id 10.31.101.1
 set router-list 10.31.101.100
 end
```



You cannot enter the `wccp-cache-engine enable` command if you have already added a WCCP service group. When you enter this command an interface named `w.<vdom_name>` is added to the FortiGate configuration (for example `w.root`). All traffic redirected from a WCCP router is considered to be received at this interface of the FortiGate unit operating as a WCCP client. A default route to this interface with lowest priority is added.

## Example WCCP server and client configuration for caching HTTPS sessions

Enter the following command to add a service group to a WCCP server that caches HTTPS content on port 443 and protocol 6. The IP address of the server is 10.31.101.100 and the WCCP clients are on the 10.31.101.0 subnet. The service ID of this service group is 80.

```
config system wccp
 edit 80
 set router-id 10.31.101.100
 set server-list 10.31.101.0 255.255.255.0
 set ports 443
 set protocol 6
 end
```

Enter the following commands to configure a FortiGate unit to operate as a WCCP client and add a service group that configures client to cache HTTPS sessions on port 443 and protocol 6. The IP address of the server is 10.31.101.100 and IP address of this WCCP clients is 10.31.101.1 subnet. The service ID of this service group must be 80 to match the service ID added to the server.

```
config system settings
 set wccp-cache-engine enable
end

config system wccp
 edit 80
 set cache-id 10.31.101.1
 set router-list 10.31.101.100
 set ports 443
 set protocol 6
 end
```

## Example WCCP server and client configuration for caching HTTP and HTTPS sessions

You could do this by configuring two WCCP service groups as described in the previous examples. Or you could use the following commands to configure one service group for both types of traffic. The example also caches HTTP sessions on port 8080.

Enter the following command to add a service group to a WCCP server that caches HTTP sessions on ports 80 and 8080 and HTTPS sessions on port 443. Both of these protocols use protocol number 6. The IP address of the server is 10.31.101.100 and the WCCP clients are on the 10.31.101.0 subnet. The service ID of this service group is 90.

```
config system wccp
 edit 90
 set router-id 10.31.101.100
 set server-list 10.31.101.0 255.255.255.0
 set ports 443 80 8080
 set protocol 6
 end
```

Enter the following commands to configure a FortiGate unit to operate as a WCCP client and add a service group that configures client to cache HTTP sessions on port 80 and 8080 and HTTPS sessions on port 443. The IP address of the server is 10.31.101.100 and IP address of this WCCP clients is 10.31.101.1 subnet. The service ID of this service group must be 90 to match the service ID added to the server.

```
config system settings
 set wccp-cache-engine enable
end

config system wccp
 edit 90
 set cache-id 10.31.101.1
 set router-list 10.31.101.100
 set ports 443 80 8080
 set protocol 6
 end
```

## Other WCCP service group options

In addition to using WCCP service groups to define the types of traffic to be cached by WCCP the following options are available for servers and clients.

### Server configuration options

The server configuration must include the `router-id`, which is the WCCP server IP address. This is the IP address of the interface that the server uses to communicate with WCCP clients.

The `group-address` is used for multicast WCCP configurations to specify the multicast addresses of the clients.

The `server-list` defines the IP addresses of the WCCP clients that the server can connect to. Often the server list can be the address of the subnet that contains the WCCP clients.

The `authentication` option enables or disables authentication for the WCCP service group. Authentication must be enabled on all servers and clients in a service group and members of the group must have the same `password`.

The `forward-method` option specifies the protocol used for communication between the server and clients. The default forwarding method is GRE encapsulation. If required by your network you can also select to use unencapsulated layer-2 packets instead of GRE or select any to allow both. The `return-method` allows you to specify the communication method from the client to the server. Both GRE and layer-2 are supported.

The `assignment-method` determines how the server load balances sessions to the clients if there are multiple clients. Load balancing can be done using hashing or masking.

### Client configuration options

The client configuration includes the `cache-id` which is the IP address of the FortiGate interface of the client that communicates with WCCP server. The `router-list` option is the list of IP addresses of the WCCP servers in the WCCP service group.

The `ports` option lists the port numbers of the sessions to be cached by the client and the `protocol` sets the protocol number of the sessions to be cached. For TCP sessions the protocol is 6.

The `service-type` option can be auto, dynamic or standard. Usually you would not change this setting.

The client configuration also includes options to influence load balancing including the `primary-hash`, `priority`, `assignment-weight` and `assignment-bucket-format`.

## WCCP configuration overview

To configure WCCP you must create a service group that includes WCCP servers and clients. WCCP servers intercept sessions to be cached (for example, sessions from users browsing the web from a private network). To intercept sessions to be cached the WCCP server must include a security policy that accepts sessions to be cached and WCCP must be enabled in this security policy.

The server must have an interface configured for WCCP communication with WCCP clients. That interface sends and receives encapsulated GRE traffic to and from WCCP clients. The server must also include a WCCP service group that includes a service ID and the addresses of the WCCP clients as well as other WCCP configuration options.

To use a FortiGate unit as a WCCP client, the FortiGate unit must be set to be a WCCP client (or cache engine). You must also configure an interface on the client for WCCP communication. The client sends and receives encapsulated GRE traffic to and from the WCCP server using this interface.

The client must also include a WCCP service group with a service ID that matches a service ID on the server. The client service group also includes the IP address of the servers in the service group and specifies the port numbers and protocol number of the sessions that will be cached on the client.

When the client receives sessions from the server on its WCCP interface, it either returns cached content over the WCCP interface or connects to the destination web servers using the appropriate interface depending on the client routing configuration. Content received from web servers is then cached by the client and returned to the WCCP server over the WCCP link. The server then returns the received content to the initial requesting user web browser.

Finally you may also need to configure routing on the server and client FortiGate units and additional security policies may have to be added to the server to accept sessions not cached by WCCP.

## Example: caching HTTP sessions on port 80 using WCCP

In this example configuration (shown in [Figure 414](#)), a FortiGate unit with host name WCCP\_srv is operating as an Internet firewall for a private network is also configured as a WCCP server. The port1 interface of WCCP\_srv is connected to the Internet and the port2 interface is connected to the internal network.

All HTTP traffic on port 80 that is received at the port2 interface of WCCP\_srv is accepted by a port2 to port1 security policy with WCCP enabled. All other traffic received at the port2 interface is allowed to connect to the Internet by adding a general port2 to port1 security policy below the HTTP on port 80 security policy.


A WCCP service group is added to WCCP\_srv with a service ID of 0 for caching HTTP traffic on port 80. The port5 interface of WCCP\_srv is configured for WCCP communication.

A second FortiGate unit with host name WCCP\_client is operating as a WCCP client. The port1 interface of WCCP\_client is connected to port5 of WCCP\_srv and is configured for WCCP communication.

WCCP\_client is configured to cache HTTP traffic because it also has a WCCP service group with a service ID of 0.

WCCP\_client connects to the Internet through WCCP\_srv. To allow this, a port5 to port1 security policy is added to WCCP\_srv.

**Figure 414:**FortiGate WCCP server and client configuration



### Configuring the WCCP server (WCCP\_srv)

Use the following steps to configure WCCP\_srv as the WCCP server for the example network. The example steps only describe the WCCP-related configuration.

## To configure WCCP\_srv as a WCCP server

1. Add a port2 to port1 security policy that accepts HTTP traffic on port 80 and is configured for WCCP:

```
config firewall policy
 edit 0
 set srtintf port2
 set dstintf port1
 set srcaddr all
 set dstaddr all
 set action accept
 set schedule always
 set service HTTP
 set wccp enable
 set nat enable
 end
```

2. Add another port2 to port1 security policy to allow all other traffic to connect to the Internet.

```
.config firewall policy
 edit 0
 set srtintf port2
 set dstintf port1
 set srcaddr all
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
 set nat enable
 end
```

3. Move this policy below the WCCP policy in the port2 to port1 policy list.
4. Enable WCCP on the port5 interface.

```
config system interface
 edit port5
 set wccp enable
 end
```

5. Add a WCCP service group with service ID 0.

```
config system wccp
 edit 0
 set router-id 10.51.101.100
 set server-list 10.51.101.0 255.255.255.0
 end
```

- 6 Add a firewall address and security policy to allow the WCCP\_client to connect to the internet.

```
config firewall address
 edit WCCP_client_addr
 set subnet 10.51.101.10
 end

config firewall policy
 edit 0
 set srtintf port5
 set dstintf port1
 set srcaddr WCCP_client_addr
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
 set nat enable
 end
```

## Configuring the WCCP client (WCCP\_client)

Use the following steps to configure WCCP\_client as the WCCP client for the example network. The example steps only describe the WCCP-related configuration.

### To configure WCCP\_client as a WCCP client

1. Configure WCCP\_client to operate as a WCCP client.

```
config system settings
 set wccp-cache-engine enable
end
```



You cannot enter the `wccp-cache-engine enable` command if you have already added a WCCP service group. When you enter this command an interface named `w.<vdom_name>` is added to the FortiGate configuration (for example `w.root`). All traffic redirected from a WCCP router is considered to be received at this interface of the FortiGate unit operating as a WCCP client. A default route to this interface with lowest priority is added.

- 2 Enable WCCP on the port1 interface.

```
config system interface
 edit port1
 set wccp enable
 end
```

- 3 Add a WCCP service group with service ID 0.

```
config system wccp
 edit 0
 set cache-id 10.51.101.10
 set router-list 10.51.101.100
 end
```

## Example: caching HTTP sessions on port 80 and HTTPS sessions on port 443 using WCCP

This example configuration is the same as that shown in [Figure 414](#) and described in “[Example: caching HTTP sessions on port 80 using WCCP](#)” on [page 2708](#) except that WCCP now also caches HTTPS traffic on port 443. To cache HTTP and HTTPS traffic the WCCP service group must have a service ID in the range 51 to 255 and you must specify port 80 and 443 and protocol 6 in the service group configuration of the WCCP client.

Also the security policy on the WCCP\_srv that accepts sessions from the internal network to be cached must accept HTTP and HTTPS sessions.

### Configuring the WCCP server (WCCP\_srv)

Use the following steps to configure WCCP\_srv as the WCCP server for the example network. The example steps only describe the WCCP-related configuration.

#### To configure WCCP\_srv as a WCCP server

1. Add a port2 to port1 security policy that accepts HTTP traffic on port 80 and HTTPS traffic on port 443 and is configured for WCCP:

```
config firewall policy
 edit 0
 set srtintf port2
 set dstintf port1
 set srcaddr all
 set dstaddr all
 set action accept
 set schedule always
 set service HTTP HTTPS
 set wccp enable
 set nat enable
 end
```

2. Add another port2 to port1 security policy to allow all other traffic to connect to the Internet.

```
.config firewall policy
 edit 0
 set srtintf port2
 set dstintf port1
 set srcaddr all
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
 set nat enable
 end
```

3. Move this policy below the WCCP policy in the port2 to port1 policy list.

4. Enable WCCP on the port5 interface.

```
config system interface
 edit port5
 set wccp enable
 end
```

- 5 Add a WCCP service group with service ID 90 (can be any number between 51 and 255).

```
config system wccp
 edit 90
 set router-id 10.51.101.100
 set server-list 10.51.101.0 255.255.255.0
 end
```

- 6 Add a firewall address and security policy to allow the WCCP\_client to connect to the internet.

```
config firewall address
 edit WCCP_client_addr
 set subnet 10.51.101.10
 end

.config firewall policy
 edit 0
 set srtintf port5
 set dstintf port1
 set srcaddr WCCP_client_addr
 set dstaddr all
 set action accept
 set schedule always
 set service ANY
 set nat enable
 end
```

## Configuring the WCCP client (WCCP\_client)

Use the following steps to configure WCCP\_client as the WCCP client for the example network. The example steps only describe the WCCP-related configuration.

### To configure WCCP\_client as a WCCP client

1. Configure WCCP\_client to operate as a WCCP client.

```
config system settings
 set wccp-cache-engine enable
end
```



You cannot enter the `wccp-cache-engine enable` command if you have already added a WCCP service group. When you enter this command an interface named `w.<vdom_name>` is added to the FortiGate configuration (for example `w.root`). All traffic redirected from a WCCP router is considered to be received at this interface of the FortiGate unit operating as a WCCP client. A default route to this interface with lowest priority is added.

---

- 2 Enable WCCP on the port1 interface.

```
config system interface
 edit port1
 set wccp enable
 end
```



- 3 Add a WCCP service group with service ID 90. This service group also specifies to cache sessions on ports 80 and 443 (for HTTP and HTTPS) and protocol number 6.

```
config system wccp
 edit 90
 set cache-id 10.51.101.10
 set router-list 10.51.101.100
 ports 80 443
 set protocol 6
end
```

## WCCP packet flow

The following packet flow sequence assumes you have configured a FortiGate unit to be a WCCP server and one or more FortiGate units to be WCCP clients.

1. A user's web browser sends a request for web content.
2. The FortiGate unit configured as a WCCP server includes a security policy that intercepts the request and forwards it to a WCCP client.  
The security policy can apply security profiles to traffic accepted by the policy.
3. The WCCP client receives the WCCP session.
4. The client either returns requested content to the WCCP server if it is already cached, or connects to the destination web server, receives and caches the content and then returns it to the WCCP server.
5. The WCCP server returns the requested content to the user's web browser.
6. The WCCP router returns the request to the client web browser.  
The client web browser is not aware that all this is taking place and does not have to be configured to use a web proxy.

## Configuring the forward and return methods and adding authentication

The WCCP forwarding method determines how intercepted traffic is transmitted from the WCCP router to the WCCP cache engine. There are two different forwarding methods:

- GRE forwarding (the default) encapsulates the intercepted packet in an IP GRE header with a source IP address of the WCCP router and a destination IP address of the target WCCP cache engine. The result is a tunnel that allows the WCCP router to be multiple hops away from the WCCP cache server.
- L2 forwarding rewrites the destination MAC address of the intercepted packet to match the MAC address of the target WCCP cache engine. L2 forwarding requires that the WCCP router is Layer 2 adjacent to the WCCP client.

You can use the following command on a FortiGate unit configured as a WCCP router to change the forward and return methods to L2:

```
config system wccp
 edit 1
 set forward-method L2
 set return-method L2
end
```

You can also set the forward and return methods to any in order to match the cache server configuration.

By default the WCCP communication between the router and cache servers is unencrypted. If you are concerned about attackers sniffing the information in the WCCP stream you can use the following command to enable hash-based authentication of the WCCP traffic. You must enable authentication on the router and the cache engines and all must have the same password.

```
config system wccp
 edit 1
 set authentication enable
 set password <password>
 end
```

## WCCP Messages

When the WCCP service is active on a web cache server it periodically sends a WCCP HERE I AM broadcast or unicast message to the FortiGate unit operating as a WCCP router. This message contains the following information:

- Web cache identity (the IP address of the web cache server).
- Service info (the service group to join).

If the information received in the previous message matches what is expected, the FortiGate unit replies with a WCCP I SEE YOU message that contains the following details:

- Router identity (the FortiGate unit's IP address).
- Sent to IP (the web cache IP addresses to which the packets are addressed)

When both ends receive these two messages the connection is established, the service group is formed and the designated web cache is elected.

## Troubleshooting WCCP

Two types of debug commands are available for debugging or troubleshooting a WCCP connection between a FortiGate unit operating as a WCCP router and its WCCP cache engines.

### Real time debugging

The following commands can capture live WCCP messages:

```
diag debug en
diag debug application wccpd <debug level>
```

### Application debugging

The following commands display information about WCCP operations:

```
get test wccpd <integer>
diag test application wccpd <integer>
```

Where <integer> is a value between 1 and 5:

1. Display WCCP stats
2. Display WCCP config
3. Display WCCP cache servers
4. Display WCCP services
5. Display WCCP assignment

Enter the following command to view debugging output:

```
diag test application wccpd 3
```

Sample output from a successful WCCP connection:

```
service-0 in vdom-root: num=1, usable=1
cache server ID:
len=44, addr=172.16.78.8, weight=4135, status=0
rcv_id=6547, usable=1, fm=1, nq=0, dev=3(k3),
to=192.168.11.55
ch_no=0, num_router=1:
192.168.11.55
```

Sample output from the same command from an unsuccessful WCCP connection (because of a service group password mismatch):

```
service-0 in vdom-root: num=0, usable=0
diag debug application wccpd -1
Sample output:
wccp_on_recv()-98: vdom-root recv: num=160, dev=3(3),
172.16.78.8->192.168.11.55
wccp2_receive_pkt()-1124: len=160, type=10, ver=0200,
length=152
wccp2_receive_pkt()-1150: found component:t=0, len=20
wccp2_receive_pkt()-1150: found component:t=1, len=24
wccp2_receive_pkt()-1150: found component:t=3, len=44
wccp2_receive_pkt()-1150: found component:t=5, len=20
wccp2_receive_pkt()-1150: found component:t=8, len=24
wccp2_check_security_info()-326: MD5 check failed
```

# Storage

WAN optimization storage is used for storing the byte cache and web cache databases. In most cases, you can accept the default WAN optimization storage configuration because all of the disk space available on the FortiGate unit is in one partition. By default WAN optimization and logging and archiving are configured to use this partition.

You only have to configure WAN optimization storage if you have more than one possible storage location. This can happen if you have multiple partitions that you can use for storage locations. If you have more than one storage location you can move WAN optimization storage to it. You can also configure WAN optimization to use multiple storage locations.

You can also optionally configure WAN optimization storage if you want to adjust the relative amounts of disk space available for byte caching and web caching.

This chapter contains the following topics:

- [Formatting the hard disk](#)
- [Configuring WAN optimization and Web cache storage](#)

## Formatting the hard disk

In most cases the hard disks on your FortiGate unit should be formatted with one partition that is used for WAN optimization and Logging and Archiving. If for some reason the hard disk is not formatted you can use the following information to format it. In some cases you might also want to use the following options to erase all data from the hard disk by reformatting it.

From the web-based manager go to *System > Config > Advanced > Disk Management* to display information about the hard disk or disks available to the FortiGate unit. To format a hard disk, select the format icon. The hard disk format takes a few minutes and the FortiGate unit restarts after formatting is complete.

From this web-based manager page you can also view and change the WAN optimization and Web Cache Storage size and view how much of the WAN optimization and web cache storage has been used.

From the CLI you can use the following command to view the current disk format and partition status. See the following example for a FortiGate-51B unit.

```
execute disk list
```

```
Device I1 29.9 GB ref: 256 SUPER TALENT (IDE)
partition 1 29.9 GB ref: 257 label: 2B6375792136C707
```

You can use the following command to reformat the hard disk. Use this command if for some reason the disk is not formatted correctly. The command includes the device partition reference number (256) so formats the entire disk and not just the partition.

```
execute disk format 256
```

You can use the following command to reformat the partition. The command includes the partition reference number so formats the partition, removing all data from it. You can use this command to delete all data from the partition and to fix partition errors.

```
execute disk format 257
```

## Configuring WAN optimization and Web cache storage

You can use the following command to add multiple WAN optimization storage locations if your FortiGate unit has multiple disk partitions and you want to use more than one for WAN optimization storage:

```
config system storage
```

Enter `get` to see the name of the default storage location. You cannot edit this storage location, but you can add new ones:

```
config system storage
 edit new_storage
 set partition <partition_number>
 end
```

Where `<partition_number>` is the number of the partition to create a storage location in. This cannot be the same as the partition added to the default storage location. This command automatically adds a WAN optimization storage location with the name `new_storage`.

## Changing the amount of space allocated for WAN optimization and Web cache storage

From the web-based manager you can go to *System > Config > Advanced > Disk Management* to edit the WAN optimization & Web Cache storage and change the allocation size to limit the amount of storage available for WAN optimization byte caching and web caching. The size is in Mbytes.

You can use the following command to change the size of any WAN optimization storage location. For example, in the FortiGate-51B the default WAN optimization storage is `Internal`. Use the following command to limit the amount of space allocated for WAN optimization to 20 Gbytes

```
config wanopt storage
 edit Internal
 set size 20000
 end
```

## Adjusting the relative amount of disk space available for byte caching and web caching

By default the `config wanopt storage` command allocates the same amount disk for byte caching and for web caching. In some cases you may want to adjust the relative amounts of disk space available for these two uses. For example, if you have not implemented web caching you may want to reduce the amount of disk space used for web caching and increase the amount of space used for byte caching.

You can adjust the relative amount of disk space used for byte caching using the `webcache-storage-percentage` option of the `config wanopt storage` command. This option adjusts the percentage in the range of 0 to 100. The default percentage is 50.

To reduce the percentage of space allocated on the Internal disk for web caching to 10% (resulting in the amount of space for byte caching increasing to 90%) enter:

```
config wanopt storage
 edit Internal
 set webcache-storage-percentage 10
 end
```

You can enter this command at any time without disrupting web caching or byte caching performance. Data may be lost from the cache that is reduced in size.

# Diagnose commands

The following get and diagnose commands are available for troubleshooting WAN optimization, web cache, explicit proxy and WCCP.

- `get test {wa_cs | wa_dbd | wad | wad_diskd | wccpd} <test_level>`
- `diagnose wad`
- `diagnose wacs`
- `diagnose wadb`
- `diagnose debug application {wa_cs | wa_dbd | wad | wad_diskd | wccpd} [<debug_level>]`

## `get test {wa_cs | wa_dbd | wad | wad_diskd | wccpd} <test_level>`

Display usage information about WAN optimization and web-cache-related applications. Use `<test_level>` to display different information.

```
get test wa_cs <test_level>
get test wa_dbd <test_level>
get test wad <test_level>
get test wad_diskd <test_level>
get test wccpd <test_level>
```

Variable	Description
wad	Display information about WAN optimization, web caching, the explicit web proxy, and the explicit FTP proxy.
wa_cs	Display information about the WAN optimization web cache server.
wa_dbd	Display information about the WAN optimization storage server application.
wad_diskd	Display information about the WAN optimization disk access daemon application.
wccpd	Display information about the WCCP application.

## Examples

Enter the following command to display WAN optimization tunnel protocol statistics. The http tunnel and tcp tunnel parts of the command output below shows that WAN optimization has been processing HTTP and TCP packets.

```
get test wad 11
wad tunnel protocol stats:
 http tunnel
 bytes_in=1751767 bytes_out=325468
 ftp tunnel
 bytes_in=0 bytes_out=0
 cifs tunnel
 bytes_in=0 bytes_out=0
```

```
mapi tunnel
 bytes_in=0 bytes_out=0
tcp tunnel
 bytes_in=3182253 bytes_out=200702
maintenance tunnel
 bytes_in=11800 bytes_out=15052
```

Enter the following command to display the current WAN optimization peers. You can use this command to make sure all peers are configured correctly. The command output shows one peer with IP address 172.20.120.141, peer name Web\_servers, with 10 active tunnels.

```
get test wad 26
peer name=Web_servers ip=172.20.120.141 vd=0 version=1
 tunnels(active/connecting/failover)=10/0/0
 sessions=0 n_retries=0 version_valid=true
```

Enter the following command to restart the WAN optimization web cache server.

```
get test wa_cs 99
```

Enter the following command to display all test options:

```
get test wad
```

WAD Test Usage

```
1: display total memory usage
3: display proxy status
4: display all stats and connections
5: toggle AV conserve mode(for debug purpose).
8: display all fix-sized advanced memory stats
10: toggle cifs read-ahead
11: display tunnel protocol stats
12: flush tunnel protocol stats
13: display http protocol stats
14: flush http protocol stats
15: display cifs protocol stats
16: flush cifs protocol stats
17: display ftp protocol stats
18: flush ftp protocol stats
19: display mapi protocol stats
20: flush mapi protocol stats
21: display tcp protocol stats
22: flush tcp protocol stats
23: display all protocols stats
24: flush all protocols stats
25: display all listeners
26: display all peers
27: display DNS stats
28: display security profile mapping for regular firewall policy
30: display Byte Cache DB state
31: flush Byte Cache DB stats
32: display Web Cache DB state
33: flush Web Cache DB stats
```



35: display tunnel compressor state  
36: flush tunnel compressor stats  
37: discard all wad debug info that is currently pending  
38: display rules  
39: display video cache rules (patterns)  
40: display cache state  
41: flush cache stats  
42: display all fix-sized advanced memory stats in details  
45: display memory cache state  
46: flush memory cache stats  
47: display SSL stats  
48: flush SSL stats  
49: display SSL mem stats  
50: display Web Cache stats  
51: flush Web Cache stats  
52: flush idle Web cache objects  
53: display firewall policies  
54: display WAD tunnel stats.  
55: display WAD fsae state.  
56yxxx: set xxx concurrent Web Cache session for object storage y.  
57yxxx: set xxxK(32K, 64K,...) unconfirmed write/read size per Web  
Cache object for object storage y.  
58yxxxx: set xxxxK maximum output buffer size for object storage y.  
59yxx: set lookup lowmark(only if more to define busy status) to be  
xx for object storage y.  
60: display current web proxy users  
61: flush current web proxy users  
62: display current web proxy user summary  
63: display web cache cache sessions  
65: display cache exemption patterns  
66: toggle dumping URL when daemon crashes.  
67: list all used fqdns.  
68: list all current ftpproxy sessions.  
69: display ftpproxy stats.  
70: clear ftpproxy stats.  
600000..699999 cmem bucket stats (699999 for usage)  
70yxxx: set xxxK maximum output buffer size for byte storage y.  
71yxxx: set number of buffered add requests to be xxx for byte  
storage y.  
72yxxxx: set number of buffered query requests to be xxxx for byte  
storage y.  
73yxxxxx: set number of concurrent query requests to be xxxxx for  
byte storage y.  
79xxxx: set xxxxMiB maximum AV memory.(0: set to default.  
80: display av memory usage  
81: toggle av memory protection  
800..899: mem\_diag commands (800 for help & usage)  
800000..899999: mem\_diag commands with 1 arg (800 for help & usage)  
80000000..899999999: mem\_diag commands with 2 args (800 for help &  
usage)

```

90: set to test disk failure
91: unset to test disk failure
92: trigger a disk failure event
98: gracefully stopping wad proxy
99: restart proxy

```

## diagnose wad

Display diagnostic information about the WAN optimization daemon (wad).

```

diagnose wad console-log {disable | enable}
diagnose wad debug-url {disable | enable}
diagnose wad filter {clear | dport | dst | list | negate | protocol |
 sport | src | vd}
diagnose wad history {clear | list}
diagnose wad session {clear | list}
diagnose wad stats {cache | cifs | clear | crypto | ftp | http | list
 | mapi | mem | scan | scripts | summary | tcp | tunnel}
diagnose wad user {clear | list}
diagnose wad webcache {clear | list}

```

Variable	Description
console-log	Enable or disable displaying WAN optimization log messages on the CLI console.
filter	Set a filter for listing WAN optimization daemon sessions or tunnels. clear reset or clear the current log filter settings. dport enter the destination port range to filter by. dst enter the destination address range to filter by. list display the current log filter settings
history	Display statistics for one or more WAN optimization protocols for a specified period of time (the last 10 minutes, hour, day or 30 days).
session	Display diagnostics for WAN optimization sessions or clear active sessions.
stats	Display statistics for various parts of WAN optimization such as cache statistics, CIFS statistics, MAPI statistics, HTTP statistics, tunnel statistics etc. You can also clear WAN optimization statistics and display a summary.
tunnel	Display diagnostic information for one or all active WAN optimization tunnels. Clear all active tunnels. Clear all active tunnels.

### Example: diagnose wad tunnel list

Enter the following command to list all of the running WAN optimization tunnels and display information about each one. The command output shows 10 tunnels all created by peer-to-peer WAN optimization rules (auto-detect set to off).

```
diagnose wad tunnel list
```

```
Tunnel: id=100 type=manual
 vd=0 shared=no uses=0 state=3
 peer name=Web_servers id=100 ip=172.20.120.141
 SSL-secured-tunnel=no auth-grp=
 bytes_in=348 bytes_out=384
```

```
Tunnel: id=99 type=manual
 vd=0 shared=no uses=0 state=3
 peer name=Web_servers id=99 ip=172.20.120.141
 SSL-secured-tunnel=no auth-grp=
 bytes_in=348 bytes_out=384
```

```
Tunnel: id=98 type=manual
 vd=0 shared=no uses=0 state=3
 peer name=Web_servers id=98 ip=172.20.120.141
 SSL-secured-tunnel=no auth-grp=
 bytes_in=348 bytes_out=384
```

```
Tunnel: id=39 type=manual
 vd=0 shared=no uses=0 state=3
 peer name=Web_servers id=39 ip=172.20.120.141
 SSL-secured-tunnel=no auth-grp=
 bytes_in=1068 bytes_out=1104
```

```
Tunnel: id=7 type=manual
 vd=0 shared=no uses=0 state=3
 peer name=Web_servers id=7 ip=172.20.120.141
 SSL-secured-tunnel=no auth-grp=
 bytes_in=1228 bytes_out=1264
```

```
Tunnel: id=8 type=manual
 vd=0 shared=no uses=0 state=3
 peer name=Web_servers id=8 ip=172.20.120.141
 SSL-secured-tunnel=no auth-grp=
 bytes_in=1228 bytes_out=1264
```

```
Tunnel: id=5 type=manual
 vd=0 shared=no uses=0 state=3
 peer name=Web_servers id=5 ip=172.20.120.141
 SSL-secured-tunnel=no auth-grp=
 bytes_in=1228 bytes_out=1264
```

```
Tunnel: id=4 type=manual
 vd=0 shared=no uses=0 state=3
 peer name=Web_servers id=4 ip=172.20.120.141
 SSL-secured-tunnel=no auth-grp=
 bytes_in=1228 bytes_out=1264
```

```

Tunnel: id=1 type=manual
 vd=0 shared=no uses=0 state=3
 peer name=Web_servers id=1 ip=172.20.120.141
 SSL-secured-tunnel=no auth-grp=
 bytes_in=1228 bytes_out=1264

Tunnel: id=2 type=manual
 vd=0 shared=no uses=0 state=3
 peer name=Web_servers id=2 ip=172.20.120.141
 SSL-secured-tunnel=no auth-grp=
 bytes_in=1228 bytes_out=1264

Tunnels total=10 manual=10 auto=0

```

### Example: diagnose wad webcache list

This following command displays the web caching stats for the last 10 minutes of activity. The information displayed is divided into 20 slots and each slot contains stats for 30 seconds

```

20 * 30 seconds = 600 seconds = 10 minutes
diagnose wad webcache list 10min
web cache history vd=0 period=last 10min

```

The first 20 slots are for HTTP requests in the last 10 minutes. Each slot of stats has four numbers, which is the total number of HTTP requests, the number of cacheable HTTP requests, the number of HTTP requests that are processed by the web cache (hits), and the number of HTTP requests that are processed without checking the web cache (bypass). There are many reasons that a HTTP request may bypass web cache.

```

total cacheable hits bypass

36 10 3 1
128 92 1 10
168 97 2 3
79 56 0 3
106 64 5 3
180 118 6 11
88 53 7 3
80 43 4 4
107 44 9 2
84 12 0 2
228 139 52 10
32 2 0 5
191 88 13 7
135 25 40 3
48 10 0 8
193 13 7 7
67 31 1 2
109 35 24 6
117 36 10 5
22 0 0 4

```

The next 20 slots are for video requests in the last 10 minutes. Each slot has two numbers for each 30 seconds: total number of video requests, and the number of video requests that are processing using cached data.

```

video total video hit

0 0
0 0
0 0
0 0
0 0
0 0
0 0
0 0
0 0
0 0
0 0
0 0
0 0
0 0
0 0
0 0
0 0
0 0
0 0
0 0

```

These 20 slots are for traffic details in last 10 minutes. Each slot has four numbers for 30 seconds each.

```

--- LAN --- --- WAN ---
bytes_in bytes_out bytes_in bytes_out

34360 150261 141086 32347
105408 861863 858501 100670
128359 1365919 1411849 127341
60103 602813 818075 59967
105867 1213192 1463736 97489
154961 1434784 1344911 158667
73967 370275 369847 70626
129327 602834 592399 123676
115719 663446 799445 111262
58151 724993 631721 59989
175681 2092925 1092556 166212
37805 33042 41528 37779
183686 1255118 1114646 172371
106125 904178 807152 81520
66147 473983 543507 66782
170451 1289530 1201639 165540
69196 544559 865370 68446
134142 579605 821430 132113

```

```
96895 668037 730633 89872
59576 248734 164002 59448
```

## diagnose wacs

Display diagnostic information for the web cache database daemon (wacs).

```
diagnose wacs clear
diagnose wacs reents
diagnose wacs restart
diagnose wacs stats
```

Variable	Description
clear	Remove all entries from the web cache database.
reents	Display recent web cache database activity.
restart	Restart the web cache daemon and reset statistics.
stats	Display web cache statistics.

## diagnose waddb

Display diagnostic information for the WAN optimization database daemon (waddb).

```
diagnose waddb {check | clear | reents | restart | stats}
```

Variable	Description
check	Check WAN optimization database integrity.
clear	Remove all entries from the WAN optimization database.
reents	Display recent WAN optimization database activity.
restart	Restart the WAN optimization daemon and reset statistics.
stats	Display WAN optimization statistics.

## diagnose debug application {wa\_cs | wa\_dbd | wad | wad\_diskd | wccpd} [<debug\_level>]

View or set the debug level for displaying WAN optimization and web cache-related daemon debug messages. Include a <debug\_level> to change the debug level. Leave the <debug\_level> out to display the current debug level. Default debug level is 0.

```
diagnose debug application wa_cs [<debug_level>]
diagnose debug application wa_dbd [<debug_level>]
diagnose debug application wad [<debug_level>]
```

```
diagnose debug application wccpd [<debug_level>]
```

<b>Variable</b>	<b>Description</b>
wa_cs	Set the debug level for the web cache server.
wa_dbd	Set the debug level for the WAN optimization database server.
wad	Set the debug level for the WAN optimization daemon.
wccpd	Set the debug level for the WCCP daemon.

# Index

## Numerics

- 3DES 1096, 1407
- 3DES-Triple-DES 1648
- 802.11 wireless protocols 784
- 802.1Q 1523, 1527, 1530
- 802.3ad 1491
  - aggregate interface 1188
- 802.3ad aggregate interface
  - full mesh HA 1240
  - HA MAC addresses 1188
  - port monitoring 1188

## A

- A 518
- a-a
  - load balance schedule 1347
- abort 1414
- accelerated interfaces 2443
- ACCEPT 952
- accept any peer 2610
- Accept peer certificate 1643
- Accept this peer certificate
  - group only 1643
- access control list (ACL) 277
- access controls 1415
- access point
  - adding 815
  - enabling 817
- Access Point Number (APN) 753
- Active Directory - see Directory Service
- Active Directory (AD) 165, 460, 481, 554
  - access mode 581
- active mode
  - real server 1886
- active sessions
  - HA statistics 1272
- active-active
  - best practice 1143
  - device failover 1344
  - link failover 1344
  - load balancing 1129, 1343
  - network processor accelerated interfaces 1346
  - operation mode 1129
  - redundant interfaces 1201
  - session failover 1344
  - UTM sessions continue after a failover 1335
- active-active HA 1075, 1095
- active-passive
  - device failover 1128
  - failover 1290
  - LACP 1189
  - link failover 1128
  - operating mode 1128

- active-passive mode
  - redundant interfaces 1201
- adding configuring defining
  - attack logging 1935
  - logging practices 1930, 1955, 1960, 1967
- adding or configuring
  - authenticated access 514
  - local users 489
  - logging of events 559, 594
  - peer user groups 507
  - peer users 493
  - user groups 504
- adding, configuring defining
  - administrator settings 1441
  - authentication settings 464
  - backing up configuration 1395
  - changing administrator's password 1397
  - dashboards 1391
  - DHCP interface settings 1492
  - DHCP relay agent 1578
  - DHCP server 1577
  - endpoint profile 2006
  - firmware version 1394
  - formatting USB disks 1396
  - FSSO 482
  - general system settings 1441
  - HA 1149
  - health check monitor 1888
  - interface 1483
  - IPSec VPN phase 1 1625
  - IPSec VPN phase 1 advanced options 1626
  - IPSec VPN phase 2 1629
  - IPSec VPN phase 2 advanced options 1629
  - LDAP authentication for administrators 1440
  - password authentication 1435
  - per-IP traffic shaping 2247
  - PKI authentication, administrators 1441
  - port monitoring (HA) 1150
  - RADIUS authentication, administrators 1440
  - RAID disk 1401
  - replacement message images 1591
  - replacement messages 1592
  - restoring configuration 1397
  - secondary IP address 1497
  - shared traffic shapers 2245
  - SNMP community 1518
  - synchronizing with NTP server 1394
  - system configuration backup and restore,
    - FortiManager 1396
  - system time 1394
  - TACACS+ authentication 1440
  - text strings (names) 1389
  - uploading scripts 1600
  - URL filter 2124
  - user authentication settings 464



- web filter profile 2119
- adding, configuring or defining
  - anti-overbilling protection 697
  - APN filtering 687
  - carrier duplicate message 680
  - carrier end point MMS filter list 705
  - carrier endpoint filter list 681
  - carrier GTP profile, advanced filtering rule 693
  - carrier utm, message flood 679
  - carrier utm, notification list 676
  - GTP log settings 697
  - GTP message type filtering 687
  - GTP profile 683
  - MMS profile 662
  - server load balance port forwarding virtual IP 1917
  - server load balance virtual IP 1911
- address
  - MAC 2196
- Address Resolution Protocol (ARP) 1551, 2380
- address spoofing 739
- address translation
  - MMS 637
- addresses 909
- ADM 1346
- admin
  - administrator account 1402
  - concurrent sessions 1436
  - disclaimer, login
    - disclaimer 1439
  - password 1429
- administration
  - schools 1585
- administrative access 1495
  - changing 1403
- administrative distance 279, 280
- administrative interface. *See* web-based manager
- administrator
  - account 1402
  - lockout 1437
  - password 1402
  - settings 1441
- administrator profiles
  - global 1435
  - vdom 1435
- administrators
  - LDAP authentication 1440
  - management access 1435
  - monitoring *See also* widgets 1397
- ADM-XD4
  - security processing module 1115
- ADSL 1484
- advanced mode (FSSO), nested groups 567
- advertisement message
  - adjusting the interval 1369
- advertisement messages
  - VRRP 1364
- AES-128 1096
- AES128,192 ,256 1648
- AES-192 1096
- AES-256 1096
- age
  - age difference margin 1133
    - changing the age difference margin 1133
    - displaying cluster unit age 1134
    - primary unit selection 1133
    - reset the cluster age 1135
    - resetting cluster unit age 1135
  - age difference margin 1133
    - changing 1133
  - aggregate interface
    - best practice 1144
    - HA MAC addresses 1188
    - interface monitoring 1188
  - aggregate interfaces 1491
  - aggregated subnets
    - for hub-and-spoke VPN 1680
  - aggregation, link 1094
  - aggressive mode 1637
  - air flow 1456
  - alert email 2361
    - HA 1277
  - alert message console 1399
  - alert notification 714, 728
- ALG
  - changing the port numbers that the SIP ALG listens on 2516
- all policies 2244, 2245
- allow access 1495
- Allow inbound, encryption policy 1661
- Allow Invalid SSL Certificate 950
- Allow outbound, encryption policy 1661
- always revalidate 2648
- always up 1633, 2201
- ambient temperature 1456
- ambiguous routing
  - resolving in FortiGate dialup-client configuration 1725
  - VPN routing 1672
- AMC
  - hard disk 1143
- antenna 785
- anti-overbilling 697
- antireplay 1078, 1096, 1097, 1816
- antispam, *see* email filtering *and* FortiGuard, AntiSpam
- anti-spoofing 279, 764, 2442
- AntiVirus 945
- antivirus 1095, 2032
  - and* PCI DSS 777
  - archive scan depth 2039
  - change default database 2038
  - concepts 2032
  - databases 2036
  - enabling scanning 2037
  - example 2046
  - explicit FTP proxy 2697
  - explicit web proxy 2680
  - file filtering 2022
  - flow-based scanning 2033
  - FortiAnalyzer 2022

- HTTPS, IMAPS, POP3S, SMTPS 2176
- maximum file size 2039
- proxy-based scanning 2032
- scan buffer size 2038
- scanning order 2033
- SymbOS/Commwar.A!worm 702
- SymbOS/Commwar.B!wm 701
- SymbOS/Commwarrie.C-wm 702
- antivirus monitor 2184
- antivirus quarantine
  - HTTPS, IMAPS, POP3S, SMTPS 2176
- antivirus updates 1465
- AP profile
  - creating 796
  - described 793
- APN filtering 687
- Application Control 945
- application control 2023, 2248
  - explicit FTP proxy 2697
  - explicit web proxy 2680
  - monitor 2155
- application layer 1095
- Application Layer Firewalls 892
- application monitor 2155, 2187
- archive and data leak monitor 2186
- archive antivirus scan depth 2039
- archiving
  - DLP 2144
- area 403
- ARP 1884
  - cache 2316
  - gratuitous 1300
  - proxy ARP 1884
  - request 2383
  - resolution 2443
- arp table 1323, 1361
- arp-reply
  - load balance virtual server 1885
- arps
  - CLI command 1301
  - gratuitous 1301
- arps-interval
  - CLI command 1301
- AS
  - multihomed 314
  - number (ASN) 314
  - stub 314
- ASCII 480, 1420
- asymmetric routing 764, 1554, 2442
- attached network equipment
  - failover 1335
- attack updates
  - scheduling 1465
- attributes, RADIUS 468
- authenticating
  - based on peer IDs 1644
  - FortiGate unit pre-shared key 1641
  - IPsec VPN peers and clients 1642
  - L2TP clients 1565
  - PPTP clients 1557
    - through IPsec certificate 1639
    - through XAuth settings 1651
- authenticating users
  - FortiGate, local 488
    - with LDAP servers 489
    - with RADIUS servers 489
    - with TACACS+ servers 489
- authentication 787, 2602
  - authentication method 2613
    - certificate-based 517
      - Citrix 2678
      - client certificates 2200
      - configuring access 514
      - defining settings 464
      - disclaimer 521
      - explicit web proxy 2677, 2678
      - firewall policy 517, 521
      - guest users 559, 594
      - heartbeat 1299
      - HTTP 2678
      - IP Based 2349
      - IPsec VPN 528
      - L2TP 462, 531
      - NAT device 2678
      - overview 533
      - peer 2611
      - PKI certificate, administrators 1441
      - PPTP 462, 530
      - protocols 517
      - proxy 2678
      - RADIUS for administrators 1440
      - replacement messages 518
      - SCP 1445
      - server certificate and SSL VPN 2200
      - SHA-1, 256, 384, 512 1648
      - SSL VPN 528
      - SSL VPN timeout 528
      - timeout 514
      - timeout setting 2197
      - VPN 528
      - VPN client-based 463
      - WAN optimization peer authentication 2610
      - web proxy 2678
      - web-based user 463
      - Windows Terminal Server 2678
      - XAuth 529
  - authentication group
    - authentication method 2613
      - certificate 2613
      - password 2613
      - pre-shared key 2613
- authentication protocols 518
  - ASCII 480
  - CHAP 480
  - MS-CHAP 480
  - PAP 480
  - setting 517
  - TACACS+ servers 480
- authentication realm
  - explicit web proxy 2672

- authentication server, external
  - for L2TP 1565
  - for PPTP 1557
- authentication servers
  - Directory Service 481
  - LDAP 471
  - RADIUS 466
  - TACACS+ 480
- authorization, LDAP 1432
- auto connect 1633, 2201
- Autokey
  - IPSec VPN 1624
  - keep alive 1654

## B

- back to HA monitor
  - HA statistics 1272
- backing up configuration
  - See widgets, system information
- backup
  - cluster configuration 1276
- backup and restore configuration, central management 1396
- backup configuration
  - SCP 1443
  - USB 1453
- backup password 476
- backup unit 1122
  - See Also subordinate unit 1122
- backup VPN 1765
- band
  - radio bands for wireless LANs 784
- bandwidth 789, 2238
  - guaranteed 2236, 2244, 2246, 2247
  - maximum 2244, 2246, 2247
  - zero 2244
- banned users
  - cause or rule 502
- baud rate 1423
- Berkeley Packet Filtering (BPF) 2311
- best practice 1145
- BGP
  - graceful restart 1316
- bgp
  - attribute
    - AS\_PATH 365
    - ATOMIC\_AGGREGATE 367
    - COMMUNITY 366
    - MULTI\_EXIT\_DESC 366
    - NEXT\_HOP 366
  - BGP-4+ 358
  - clearing routes 360, 371
  - control plane 373
  - flap 373
  - graceful restart 373
  - MED 366
  - neighbors 360
  - password, MD5 360
  - RFC 1997 366
  - route reflectors (RR) 362

- stabilizing the network 373
- Bi-directional Forwarding Detection (BFD) 374
- binding
  - LDAP servers 472
- bits per second (bps) 1406
- black list 2023
- blackhole route 278
- block 2563
- block traffic 2244
- blocking of users
  - Endpoint Control 2005
- Blowfish 1407
- body
  - SIP message 2490
- bookmarks, personal 2203
- boot interrupt 1405
- Boot Strap Router (BSR) 1030
- border gateway protocol (BGP). See routing, BGP
- BPDU
  - message exchange 1362
- branch 2571
- bridge protocol data unit 1362
- broadcast
  - domains 1523
  - storm 1551
- broken cluster unit
  - replacing 1187
- buffer size
  - IPS 2071
- busy
  - load balance 1350
- byte cache 2580
  - changing the relative amount of disk space 2717
- byte caching
  - dynamic data chunking 2599

## C

- CA certificate 2174
- cache
  - exempting from web caching 2651
  - iOS updates 2644, 2648
  - Windows updates 2644, 2648
- cache cleaner 2194
- cache engine
  - WCCP 2703
- cache expired objects 2650
- call-keepalive 2519
- captive portal 786
- carrier
  - duplicate message, configuring 680
  - GTP profile 683
  - HA 739
  - message flood 679
  - MMS profile 662
  - notification list 676
- carrier end point
  - filter list 705
  - filtering 636, 735
  - IP filtering 706

- MMS filtering 664
- pattern, Perl 681
- pattern,regular expression 681
- pattern,wildcard 681
- carrier endpoint
  - filter list 681
- case sensitivity
  - Perl regular expressions 1424
- CB2 1346
- CDE
  - defined 768
- central management
  - backup and restore configuration 1396
- Central NAT Table 902
- certificate
  - authentication group 2613
  - key size 2175
  - SSL 2174
- certificate authority (CA) 533
- Certificate Name, Phase 1 1640
- certificate request 536, 546
  - generating 536, 545
  - key size 537
- certificate revocation list (CRL)
  - importing 539
- certificate signing request (CSR) 536
- certificate, IPsec
  - group 1643
  - Local ID setting 1643
  - using DN to establish access 1642
  - viewing local DN 1643
- certificate, security 1427
- certificate, server 2200
- certificates
  - import 542
  - importing CRL 539
  - OSCP 534
  - root CA, installing 539
  - self-signed 533
  - signed server, installing 539
  - Single Sign On (SSO) 532
- Challenge-Handshake Authentication Protocol (CHAP)
  - 480, 1651
- changing unit's host name 1392
- channels
  - for 802.11a 883
  - for 802.11b 884
  - for 802.11n 5GHz 883
  - radio channels for wireless LANs 784
- CHAP 530, 1555
- Charging Data Function (CDF) 652
- Charging Data Record (CDR) 653
- checking windows version 2213
- Chunked Bypass 949
- CIDR 1411
- CIFS
  - protocol optimization 2598
- Cisco
  - router configuration 1534, 1549
  - switch configuration 1534, 1540, 1548
  - Cisco switch configuration 2397
  - Cisco VPN 1798
- Citrix
  - authentication 2678
- Classless Inter-Domain Routing (CIDR) 305
- CLI 1097, 1387
  - connecting 1405
  - connecting to from the web-based manager 1403
  - connecting to the 1405
  - Console widget 1406
  - session-pickup 1294, 1330
  - upgrading the firmware 1449
- CLI command
  - 1155, 1160, 1167, 1171, 1178, 1181, 1191, 1195, 1203, 1207, 1224, 1229, 1244, 1248, 1304
- CLI console 1399
- client 2680, 2697
  - using FortiWiFi unit as a WiFi client 879
  - WCCP 2703
- client certificates 2200
- client comforting 707
  - amount 708
- client IP
  - assigning with RADIUS 1713
- client mode 879
  - using FortiWiFi unit as a WiFi client 879
- client reputation
  - explicit web proxy 2680, 2697
- Client to FortiGate
  - SSL offloading 1903
- Client to FortiGate to Server
  - SSL offloading 1903
- cluster
  - adding a new FortiGate unit 1186
  - configuring in transparent mode 1165
  - connecting an HA cluster 1127
  - converting a standalone FortiGate unit 1184
  - definition 1145
  - distributed 1143, 1263
  - operating 1252
  - replacing a failed cluster unit 1187
  - virtual cluster 1217
- cluster configuration
  - backup 1276
  - restore 1276
  - troubleshoot
    - 1157, 1163, 1169, 1176, 1180, 1184, 1192, 1198, 1204, 1209, 1213, 1214, 1245, 1249
- cluster member 1076, 1268
  - cluster members list 1269
  - priority 1269
  - role 1269
- cluster name 1129
- cluster unit
  - connect to a cluster 1285
  - definition 1145
  - disconnect from a cluster 1285
  - getting information using SNMP 1266, 1267
  - getting cluster unit serial numbers 1267

- getting serial numbers using SNMP 1267
  - SNMP get 1266, 1267
- cluster units 1122
- CNAME 1581
- collector agent
  - Ignore User list 582
  - LDAP access 581
  - logs 597
  - settings 577
  - specifying 590
  - TCP ports 584
- column settings
  - configuring 1389
- Comfort Clients 948
- command 1409
  - abbreviation 1417
  - completion 1416
  - help 1416
  - multi-line 1417
- common name, LDAP servers 473
- concentrator 1684
  - IPSec tunnel mode 1635
  - IPSec VPN, policy-based 1635
- concepts
  - antivirus 2032
  - web filtering 2022
- concurrent sessions 1436
- configuration
  - backup 1276
  - locking 1504
  - restore 1276
  - revisions 1446
  - synchronization 1307
- configuration synchronization 1292
  - disabling 1308
- configure
  - FortiGuard 1428
  - restore 1445
- configuring
  - collector agent 577
  - dynamic DNS VPN 1698
  - firewall policy authentication 521
  - FortiClient dialup-client VPN 1714
  - FortiClient in dialup-client VPN 1718
  - FortiGate dialup-client VPN 1727
  - FortiGate firewall policies 557, 591
  - FortiGate in dialup-client IPsec VPN 1729
  - gateway-to-gateway IPsec VPN 1667
  - hub-and-spoke IPsec VPN 1679
  - IPsec VPN authentication 528
  - L2TP VPN authentication 531
  - LDAP server, FortiGate unit 555, 589
  - local users 489
  - peer user groups 507
  - peer users 493
  - PPTP VPN authentication 530
  - SSL VPN authentication 528
  - transparent mode IPsec VPN 1770
  - WAN optimization peer 2611
  - XAuth authentication 529
  - configuring a FortiGate unit for HA operation 1125
  - connected monitored interfaces
    - primary unit selection 1132
  - connecting
    - to the CLI using SSH 1407
    - to the CLI using Telnet 1408
    - to the console 1405
    - web-based manager 1426
  - connecting a FortiGate HA cluster 1127
  - connectionless 2272
    - session pickup 1334
  - conserve mode 2172
  - console 1405
  - console messages
    - synchronization fails 1310
  - contact-fixup 2546
  - content archive
    - metadata 718
    - MMS protection profile 718
    - summary 718
  - content blocking 636
  - content provider (CP) 671
  - content scanning
    - SSL 2173
  - Content-Length
    - SIP header 2563
  - control plane 373
  - control plane (GTP-C) 745
  - controlled upgrade 1454
  - conventions 1409
  - convergence 304, 374
  - cookie 668
    - persistence 1899
  - cookie persistence
    - HTTP host-based load balancing 1902
  - country code 669
  - coverage 789
  - cp1252 1421
  - CPU load 2351
  - CPU usage
    - HA statistics 1272
    - weight 1350
  - cpu usage
    - weighted load balancing 1350
  - cross site scripting (XSS) 525
  - Cross-Site Scripting
    - protection from 1389
  - cryptographic load 1816
  - custom login page 2212
  - custom login screen 2203

**D**

  - dampening 372
    - reachability half-life 372
  - dashboards, adding 1391
  - Data Leak Prevention 946
  - data leak prevention (DLP), see DLP
  - date 2292, 2317
    - quarantine files list 1941

- date and time 1427
- DB-9 1405
- DC
  - quarantine files list 1941
- DC Agent mode 572, 573, 580
  - dcagent.dll 600
- DCE-RPC 1606
- dcerps
  - session helper 1606
- Dead entry timeout
  - collector agent configuration 579
- dead gateway detection 332, 447, 1325, 1678
- Dead Peer Detection (DPD) 1650, 1678
  - Phase 1 1650
- debug
  - diagnose 1313
- decryption 1078, 1097, 1098, 1099, 1100
- dedicated monitoring
  - interface 1254
- dedicated to FortiAP 1484
- deep scan 2176
- deep SIP message inspection 2558
- default
  - password 2578
  - VoIP profile 2515
- default password 587
- default port
  - RADIUS servers 470
  - TACACS+ servers 480
- default route
  - NAT/Route example 2375
  - VDOM example 2378
  - VLAN 1533
- default TTL
  - web cache 2649
- defaultinfooriginate 442
- defaults 1446
- Define the problem 2321
- definitions 1409
- delaying session pickup 1331
- delete
  - local users from FortiGate configuration 493
  - user group from FortiGate configuration 508
- delete, shell command 1413
- Denial of Service (DoS) 2277
- dense mode 1031
- DENY 952
- deployment 789
- DES 1096
- designated router (DR), OSPF 1810
- Designated Routers (DRs) 1030
- destination NAT
  - SIP 2542
- device
  - failure 1291
- device authentication
  - explicit web proxy 2680, 2697
- device failover 1290, 1292
  - active-active 1344
  - active-passive 1128
  - configuration synchronization 1292
  - definition 1145
  - HA heartbeat 1292
  - IPsec SA synchronization 1292
  - route synchronization 1292
  - virtual MAC address 1292
- device priority 1274
  - primary unit selection 1131, 1136
  - subordinate unit 1273
- DH Group
  - IPsec interface mode 1656
  - Phase 1 1647
  - Phase 2 1653
- DH key size, FIPS-CC 1647
- DHCP 1142
  - configuring on an interface 1492
  - exclude range 1579
  - for WiFi clients 801
  - IPv6 1578, 1852
  - lease breaking 1580
  - lease time 1578
  - relay 1142
  - server 1142
  - servers and relays 1577
  - service 1578
- DHCP relay
  - in FortiGate dialup client configuration 1725
- DHCP server 1654
  - in FortiClient dialup-client configuration 1717
- DHCP-IPsec
  - IPSec VPN, phase 2 1631
  - phase 2 1654
- diagnose
  - firewall vip realserver 1891
  - firewall vip virtual-server 1891
  - sys ha reset-uptime 1135
  - sys ha showcsum 1313
- diagnose debug 1313
- diagnose hardware deviceinfo nic
  - 1155, 1160, 1167, 1171, 1178, 1181, 1191, 1195, 1203, 1207, 1224, 1229, 1244, 1248
  - CLI command 1305
- diagnose sys ha dump-by 1134
- diagnose sys sip 2511, 2517
- diagnose sys sip debug-mask 2511
- diagnose sys sip dialog 2511
- diagnose sys sip mapping list 2511
- diagnose sys sip status 2511
- diagnose sys sip-proxy calls 2516
- diagnose sys sip-proxy filter 2517
- diagnose sys sip-proxy log-filter 2517
- diagnose sys sip-proxy meters 2517
- diagnose sys sip-proxy stats 2517
- diagnostics
  - debug the packet flow 2443
  - packet sniffing 2442
  - traceroute 2379
  - tracert 2379

- diagnostics, tracer 1541
- dialog
  - SIP 2489, 2495
- dialup users
  - configuring authentication for 528
- dialup-client IPsec configuration
  - dialup server for FortiClient dialup clients 1714
  - dialup server for FortiGate dialup clients 1727
  - FortiGate client configuration 1729
  - FortiGate dialup client configuration 1727
  - requirements for FortiClient access 1713
  - requirements for FortiGate client access 1726
- dictionary
  - RADIUS attributes 469
- differentiated services 2251
  - mapping 2257
- Diffie-Hellman algorithm 1647, 1653
- Digital Encryption Standard 1617
- Dijkstra's algorithm 406
- directory
  - LDAP servers 472
- Directory Service
  - servers 481
  - user groups 507
- Directory service 481
- disabled mode
  - real server 1886
- disabling 1605
- disclaimer 1439
  - customized 522
  - default 522
- disconnecting a unit from a cluster
  - override 1142
- disk space
  - byte cache 2717
  - web cache 2717
- disk status, viewing 1599
- distance between cluster units 1143, 1263
- distance vector protocols 306
- distinguished names
  - elements 473
  - LDAP servers 473
  - list of 475
  - max size 474
- distributed cluster 1143, 1263
- Distributed Computing Environment Remote Procedure Call (DCE-RPC) 1606
- DLP 946, 2130
  - archiving 637, 2144
  - default rules 2143, 2144
  - explicit FTP proxy 2697
  - explicit web proxy 2680
  - MMS 726
- DLP archive
  - displaying on dashboard 2178
  - HTTPS, IMAPS, POP3S, SMTPS 2178
- DNAT 1095, 2404
- DNS 1580, 1606
  - CNAME 1581
  - external servers 1580
  - local domains 1580
  - master server 1581
  - public 1582
  - recursive 1583
  - server
    - server, DNS 1581
  - shadow 1582
  - slave 1581
  - split 1582
- DNS lookups 2361
- DNS server, dynamic DNS configuration 1698
- dns-tcp, session helper 1606
- dns-udp, session helper 1606
- domain component, LDAP servers 473
- Domain Name Identifier (DNI) 556
- domain name server 1580
- domain name, dynamic DNS configuration 1697, 1699
- DoS sensor
  - SCCP 2565
  - SIP 2565
- dotted decimal 1411
- downloading
  - quarantine logs 1942
- downloading firmware 1447
- dual internet connection 1568
- dual WAN
  - link redundancy 1568
  - load sharing 1571
- duplicate MAC 1552
- dynamic data chunking
  - byte caching 2599
- Dynamic DNS (DDNS) 1695
  - configuration steps 1698
  - domain name configuration 1699
  - overview 1695
  - remote VPN peer configuration 1704
- dynamic IP address
  - for remote host 1709
  - FortiGate DDNS peer 1697
  - FortiGate dialup client 1724
- dynamic IP pool
  - SIP 2543
- dynamic routing
  - failover 1315
- dynamic VPN address
  - mode-cfg 1735
- dynamic-gateway 1695

## E

- E Reload 2650
- earthing 1457
- ECMP 278
- eDirectory 165, 481, 564
  - servers 588
- eDirectory - see Directory Service
- edit, shell command 1413
- EI (Enhanced Extension Interface) 1094
- EICAR 2046

- elements, distinguished names 473
- \_email 1411
- Email Filtering 946
- email filtering
  - IMAPS, POP3S, SMTPS 2177
- email filtering, *see also* FortiGuard, AntiSpam 2023
- email monitor 2186
- Email Signature 949
- email token 495
- Enable perfect forward secrecy (PFS)
  - IPsec interface mode 1656
  - Phase 2 1654
- Enable replay detection
  - IPsec interface mode 1656
- Enable replay detection, Phase 2 1653
- enable session pickup 1330
- Encapsulation 1614
- encryption 1078, 1097, 1098, 1099, 1100
  - heartbeat 1299
- Encryption Algorithm
  - IPsec VPN, manual key 1634
- encryption policy
  - allow outbound and inbound 1661
  - defining IP addresses 1659
  - defining IPsec 1662
  - defining multiple for same IPsec tunnel 1662
  - evaluating multiple 1663
  - outbound and inbound NAT 1661
  - traffic direction 1661
- encryption types 785
- End User Address 697
- end, shell command 1413
- endpoint
  - configuring a profile 2006
- EndPoint Control 946
- Endpoint Control
  - blocked users 2005
  - monitoring endpoints 2010
- Endpoint Mapper (EPM) 1606
- Endpoint Protection
  - modifying download portal 2010
- Endpoint Protection portal
  - modifying replacement pages 2010
- endpoints
  - monitoring 2010
- engine algorithm
  - IPS 2070
- engine count
  - IPS 2070
- enhanced packet-matching 316
- entering text strings (names) 1389
- environment variables 1418
- equal cost multipath (ECMP) 278, 281, 297
- equipment
  - FortiAP unit 788
  - FortiWiFi unit 787
  - wireless 787
- escape sequence 1418
- ESP 1096
- Establish a baseline 2321
- example
  - inter-VDOM 2413
  - NAT/Route VDOM 2368
  - VDOM 2368
- example IPSec configurations 1097
- examples
  - hub-and-spoke VPN 1689
- exclude range, DHCP 1579
- execute
  - ha synchronize all 1308
- execute formatlogdisk 1264
- execute shutdown 1458
- exempt
  - web cache 2651
- expired objects
  - cache 2650
- explicit 2349
- explicit FTP proxy 2689
  - antivirus 2697
  - application control 2697
  - DLP 2697
  - incoming IP address 2696
  - intrusion protection 2697
  - ips 2697
  - outgoing IP address 2696, 2697
  - replacement message 2691
  - reverse 2586
  - security profile 2695
- explicit HTTP proxy
  - incoming IP address 2672
  - incoming IPv6 address 204, 1833, 2674
  - outgoing IP address 204, 1833, 2672, 2674
  - outgoing IPv6 address 205, 1833, 2674
- explicit mode
  - WAN optimization 2596, 2599
- explicit proxy 494, 2349
- explicit web proxy 1333, 2665
  - antivirus 2680
  - application control 2680
  - authentication 2677, 2678
  - authentication realm 2672
  - DLP 2680
  - FortiGuard web filtering 2680
  - FTP 2665
  - HTTPS 2665, 2689
  - intrusion protection 2680
  - IPS 2680
  - IPv6 2665
  - load balancing 1343
  - PAC 2665
  - proxy auto-config 2665
  - proxy chaining 2674
  - realm 2672
  - security profile 2669, 2694
  - SOCKS 2665
  - unknown HTTP version 2671
  - web filtering 2680
- Extended Authentication (XAuth)
  - 486, 1639, 1650, 1718



extended authentication (XAuth) 1640  
Exterior Gateway Protocol (EGP) 358  
exterior gateway protocol (EGP) 306

## F

FA2 1346  
FA2 (NP1) processor 1073  
factory reset 1446  
failed authentication attempts 516  
failed cluster unit  
    replacing 1187  
fail-open  
    IPS 2070  
failover 1130  
    active-passive 1290  
    and attached network equipment 1335  
    attached network equipment 1362  
    definition 1145  
    delayed 1362  
    device 1290, 1292  
    dynamic routing 1315  
    enabling session failover 1330  
    GTP and HA session failover 1334  
    HA 1122  
    heartbeat 1146  
    issues with layer-3 switches 1361  
    link 1147, 1290, 1319  
    monitoring cluster units 1277  
    session 1123, 1291, 1330  
    subsecond 1324  
failover protection 1217  
    active-passive operating mode 1128  
    virtual clustering 1217  
failure  
    definition 1145  
    device 1291  
    link 1291, 1319  
    multiple link failures 1323  
fast path  
    required session characteristics 1094, 1816  
fast path requirements 1083, 1094  
fast roaming 791  
FB4 1346  
FB8 1346  
FDN 1260  
FGCP  
    definition 1146  
    description 1117  
FGSP 254, 1118, 1370  
    filters 1371  
FGT\_ha\_admin  
    HA administrator account 1285  
field 1410  
    SIP 2500  
file  
    quarantine 1268  
file block  
    default patterns 711  
file filtering  
    antivirus 2022

    general configuration steps 2141  
file name  
    quarantine files list 1941  
file pattern 2033  
    creating 2142  
file sharing 2409  
file size 2033  
File transfer protocol (FTP) 1606  
file type 2033  
    creating 2142  
filter  
    APN 687  
    filtering information on web-based manager lists  
        1388  
    GTP message type 687  
    IPS 2064  
    list, carrier end point 705  
    message type, GTP 687  
    quarantine files list 1941  
    web-based manager lists 1388  
final  
    SIP response message 2498  
Firefox 525  
firewall  
    configuring user groups 504  
    creating user groups 504  
    IPsec VPN dialup user access 503  
    load balancing 1333  
    per-IP traffic shaping 2247  
    policy authentication 517, 521  
    protection profile 2389  
    schedule 2388  
    service group 2367  
    shared traffic shapers 2245  
    user authentication timeout 514  
    user groups 503  
firewall address 2373, 2388, 2394  
    NAT/Route VDOM example 2373  
    simple VDOM NAT/Route example 2376  
    VDOM NAT/Route example 2376  
firewall IP addresses  
    defining 1659  
firewall IP addresses, defining L2TP 1565  
firewall load balancing 2601  
Firewall policies 952  
firewall policies 809  
    and Endpoint Control 2008  
firewall policy 2374  
    defining for policy-based VPN 1661  
    defining for route-based VPN 1663  
    guaranteed bandwidth 2246  
    hub to spoke 1686  
    inter-VDOM 2400  
    maximum bandwidth 2246, 2247  
    policy-based, for FortiGate dialup client 1731  
    policy-based, for gateway-to-gateway 1671  
    policy-based, for hub-and-spoke 1683  
    route-based, for FortiGate dialup client 1730  
    route-based, for gateway-to-gateway 1669  
    route-based, for hub-and-spoke 1682

- spoke to spoke 1687
- using as route-based "concentrator" 1685
- VDOM 2367, 2368
- VDOM example 2374, 2377, 2394
- VLAN Transparent 2384
- firewall policy and strong authentication 544
- Firewall schedules 942
- firewall session setup rate 2304
- firewall traffic shaper monitor 2258
- firewall vip realserver
  - diagnose 1891
- firewall vip virtual-server
  - diagnose 1891
- firmware
  - backup and restore from USB 1453
  - download 1447
  - from system reboot 1450
  - installing 1450
  - revert from CLI 1452
  - reverting with web-based manager 1452
  - testing before use 1447
  - upgrade with web-based manager 1449
  - upgrading using the CLI 1449
- firmware install 1073
- firmware upgrade
  - HA 1274
- firmware, updating
  - FortiAP unit 819
- first alive
  - load balancing 1885
- Fixed Port 918
- flood 686, 720
- flow control 1406
- flow inspection 2273
- flow, reverse shaping 2249
- flow-based
  - inspection 2273
- format
  - hard disk 2716
- formatlogdisk 1264
- formatting hard disks 1264
- formatting log disks 1264
- formatting USB disks 1396
- FortiAccel (NP1) processor 1073
- FortiAnalyzer 1261
  - antivirus 2022
- FortiAP 1484
- FortiAP unit 788
  - connecting to CLI 821
  - updating firmware 819
- FortiASIC 1647, 1648
  - NP2 1816
- FortiClient 2221, 2224
  - auto connect 1633, 2201
  - download location 2005
  - keep alive 1633, 2201
  - required version 2005
  - save password 1633, 2201
- FortiClient dialup client configuration
  - example 1719
  - FortiClient dialup-client configuration
    - configuration steps 1714
    - FortiClient configuration 1718
    - overview 1709
    - VIP address assignment 1711
  - FortiClient Endpoint Security peer 462
  - FortiClient peer 2599
  - Forticlient VPN 1632
  - FortiFone
    - SIP phone 2490
  - FortiGate
    - authenticating users 488
    - authenticating with XAuth 529
    - configuring to use LDAP 474
    - configuring to use RADIUS 470
    - IPsec VPN 528
  - FortiGate Cluster Protocol
    - description 1117
  - FortiGate dialup client IPsec configuration
    - FortiGate acting as client 1724
    - policy-based firewall policy 1731
    - route-based firewall policy 1730
    - using DHCP relay in 1725
  - FortiGate LDAP configuration 589
  - FortiGate Session Life Support Protocol
    - 254, 1118, 1370
  - FortiGate unit
    - adding to a cluster 1186
    - converting from standalone to a cluster 1184
    - replacing a failed 1187
  - FortiGate unit serial number
    - primary unit selection 1137
  - FortiGuard 1260, 1428
    - AntiSpam 2023
    - Antispam 2578
    - Antivirus 2036, 2578
    - as source of antivirus signatures 2005
    - as source of application signatures 2005
    - as source of FortiClient installer 2005
    - push update 1464, 1465, 1466
    - Web Filtering 2022, 2177
    - HTTPS 2177
  - FortiGuard Analysis and Management Service 718, 719
  - FortiGuard Antispam 1260
  - FortiGuard Antivirus 1260
  - FortiGuard Center 2036
  - FortiGuard Distribution Network 1260
  - FortiGuard Distribution System (FDS) 2319
    - servers 2320
  - FortiGuard Intrusion Protection 1260
  - FortiGuard quota, monitoring 2187
  - FortiGuard service 2361
  - FortiGuard Services
    - analysis service options 1463
    - licenses 1397
    - management and analysis service options 1463
    - support contract 1463
    - web filtering and antispam options 1468
  - FortiGuard Web Filter quota 2101

- FortiGuard Web Filtering 1260
- FortiGuard web filtering
  - explicit web proxy 2680
- FortiGuard, backup and restore configuration 1396
- FortiGuard, Distribution Network 2036
- FortiManager
  - remote backup and restore options 1396
- Fortinet
  - Technical Support, registering with 2578
  - Technical Support, web site 2578
- Fortinet MIB 1520
- Fortinet Server Authentication Extension (FSAE) 460, 521
- Fortinet Single Sign On (FSSO) 165, 554
- Fortinet Single Sign On Agent (FSSO) 482
- FortiOS 3.0 MR7 585
- FortiToken 497
- FortiWiFi unit 787
- forward delay
  - spanning tree parameter 1362
- forwarding
  - MAC forwarding table 1323, 1361
- \_fqdn 1411
- FQDN Addressing 911
- Fragmented Messages 949
- fragmented packets 1084, 1095, 1817
- frame size 1075
- frame size, maximum 1073
- fresh factor
  - web cache 2649
- FRUP 1119
- FSAE 4.0 MR1 570
- FSAE version 3.5.27 585
- FSM
  - hard disk 1143
- FSSO
  - DC Agent mode 565
  - DNI 556
  - DNI field 598
  - guest 555, 559, 589, 594
  - identity-based policy 557, 591
  - logoff detection 584
  - Mac OS 566
  - nested groups 567
  - Polling Mode 567
  - Polling mode 565
  - trust relation 570
- FSSO agent
  - FortiGate configuration 591
- FSSO\_Guest\_Users 562, 600
- FSSO\_Guest\_Users group 559, 594
- FTP 1095
  - explicit web proxy 2665
  - protocol optimization 2598
- FTP proxy 2689
  - antivirus 2697
  - change the prompt 2691
  - DLP 2697
  - security profile 2695

- full mesh 1124
  - HA 1239
  - redundant HA heartbeat interfaces 1240
- full mesh HA 1122, 1239
  - configuration example 1241
  - definition 1146
  - troubleshooting 1251
- full mode
  - SSL offloading 1903
- fully qualified domain name (FQDN) 1411
- fuzzing protection
  - SIP 2558

## G

- Garbagetimer 442
- GARP 1300
- Gateway Function (CGF) 652
- Gateway GPRS Support Node (GGSN) 650
- gateway-to-gateway IPsec configuration
  - configuration steps 1667
  - overview 1665
  - policy-based firewall policy 1671
  - route-based firewall policy 1669
- GB2312 1421
- general configuration steps
  - file filtering 2141
- General Packet Radio Service (GPRS) 649, 747
  - System Node (GSN) 655
- generating
  - IPsec phase 1 keys 1647
  - IPsec phase 2 keys 1653
- Generic Access Network (GAN) 753
- Generic Routing Encapsulation (GRE) 1555
- geographical
  - distribution 1143, 1263
- Geography Based Addressing 911
- get
  - shell command 1413
  - test vs 1891
- get hardware nic
  - 1155, 1160, 1167, 1171, 1178, 1181, 1191, 1195, 1203, 1207, 1214, 1224, 1229, 1244, 1248, 1304
- get system performance status 1263
- get test ipldb 1890
- gigabit interfaces, SNMP 1517
- global 2305
- Global System for Mobile Communications (GSM) 649, 747
- Gr interface 653
- grace period
  - age difference margin 1133
  - changing 1133
- graceful restart 373
  - BGP 1316
  - OSPF 1316
- graphical user interface. See web-based manager
- gratuitous ARP packets 1300
- gratuitous arps 1301

- gratuitous-arps
  - CLI command 1302
- grayware 2033, 2036
  - scanning 2041
- GRE-over-IPsec VPN 1798
- grounding 1457
- group
  - replacement message 1598
- group filters
  - FortiGate, on collector agent 583
- group ID
  - changing 1306
  - HA configuration option 1130
  - virtual MAC address 1306
- group name
  - HA cluster name 1129
  - HA configuration option 1129
- group-id
  - CLI command 1306
- GSM EDGE Radio Access Network (GERAN) 753
- GTP
  - Access Point Name (APN) 687
  - application identifier 650
  - control plane 745
  - create pdp request 687
  - HA session failover 1334
  - Information Element (IE) 687
  - International Mobile Station Identity (IMSI) 689
  - log messages 763
  - mobile country code (MCC) 689
  - mobile network code (MNC) 689
  - mobile subscriber identification number (MSIN) 689
  - path management 745
  - Routing Area Code (RAC) 693
  - runtime statistics 761
  - Sequence Number 685
  - Serial Number (SNR) 693
  - Software Version Number (SVN) 693
  - tunnel 646
  - tunnel management 745
  - Type Allocation Code (TAC) 693
- GTP billing 746
- GTP profile 683
  - advanced filtering 690, 693
  - anti-overbilling 697
  - APN filtering 687
  - encapsulated IP traffic filtering 694
  - encapsulated non-IP end user traffic filtering 695
  - general settings 685
  - information element removal policy 694
  - logging 697
  - message type filtering 687
  - specifying logging types 699
- GTP protocol anomaly prevention 696
  - missing mandatory information elements 696
  - out of state 697
- GTP UDP
  - session failover 1334
- GTP' (GTP prime) 746
- GTP-in-GTP 740, 761

- GTP-U (GTP user data tunnelling) 761
- GTPv1 632
- guaranteed bandwidth 2238, 2244, 2246, 2247
  - firewall policy 2246
  - traffic shaping 2246
- guest account 555, 559, 589, 594
- guest network 786
- GUI. See web-based manager
- gui-voip-profile 2515
- Gx interface 653
- Gz interface 653

**H**

- H.245 1606, 1607
- h245l
  - session helper 1606
- H323, session helper 1607
- HA 371, 1269
  - alert email 1277
  - changing firmware upgrade 1275
  - cluster member 1269
  - cluster members list 1268
  - configure weighted-round-robin weights 1348
  - configuring virtual clustering 1221, 1223, 1228
  - connect a cluster unit 1285
  - definition 1117
  - disconnect a cluster unit 1285
  - event log message 1261
  - FGT\_ha\_admin administrator account 1285
  - firmware upgrade 1274
  - full mesh and 802.3ad aggregate interfaces 1240
  - full mesh and redundant heartbeat interfaces 1240
  - full mesh HA configuration example 1241
  - GTP session failover 1334
  - GTP tunnels 739
  - hello state 1262
  - host name 1269
  - IPS processing 2068
  - link failover scenarios 1323
  - log message 1261
  - manage individual cluster units 1284
  - manage logs for individual cluster units 1261
  - monitor cluster units for a failover 1277
  - router monitor 272
  - routes 272
  - SIP session failover 2569
  - SNMP and reserved management interface 1255
  - standby state 1262
  - states 1262
  - subordinate unit device priority 1273
  - subordinate unit host name 1273
  - VDOM partitioning 1151
  - viewing HA statistics 1271
  - virtual cluster 1217
  - virtual clustering 1149
  - virtual domains 1217
  - work state 1262
- HA group ID
  - changing 1306
- HA group name 1129

- HA heartbeat 1292
  - definition 1146
- HA session offloading 1075
- HA statistics
  - active sessions 1272
  - back to HA monitor 1272
  - CPU usage 1272
  - intrusion detected 1273
  - memory usage 1272
  - monitor 1272
  - network utilization 1272
  - refresh every 1272
  - serial no 1272
  - status 1272
  - total bytes 1273
  - total packets 1272
  - up time 1272
  - virus detected 1272
- HA virtual MAC address
  - definition 1146
- HA, virtual cluster 2411
- ha-eth-type
  - CLI command 1297, 1363
- half mode
  - SSL offloading 1903
- hard disk
  - AMC 1143
  - byte cache storage 2716
  - formatting 2716
  - FSM 1143
  - Wan optimization storage 2716
- hard disks
  - formatting 1264
- hardware
  - get hardware nic command
    - 1155, 1160, 1167, 1171, 1178, 1181, 1191, 1195, 1203, 1207, 1224, 1229, 1244, 1248, 1304
- hardware acceleration
  - RTP 2497
- hardware revisions
  - ignoring 1212
- hardware switch 1488
- hash map 1296
- Hash-based Message Authentication Code (HMAC)
  - 1647
- hb-interval 1298
- hb-lost-threshold 1298
- hc-eth-type
  - CLI command 1297, 1363
- header
  - SIP 2500
  - SIP messages 2490
- health check
  - ping 1894
- health check monitor
  - configuring 1888
  - matched content 1889
  - real server 1888
- health monitor
  - proxy forwarding 2675
  - real server 1888
- heartbeat 1292
  - authentication 1299
  - changing the heartbeat interval 1298
  - changing the hello state hold-down time 1299
  - changing the lost heartbeat threshold 1298
  - definition 1146
  - encryption 1299
  - modifying heartbeat timing 1298
- heartbeat device
  - definition 1146
- heartbeat failover
  - definition 1146
- heartbeat interface 1121, 1293
  - best practice 1144
  - configuring 1294
  - priority 1294
  - selection 1295
  - virtual clustering 1217
- heartbeat interfaces 1217
- hello state
  - changing the time to wait 1299
  - definition 1146
- hello state hold-down time
  - changing 1299
- helo-holddown 1299
- heuristics 2033, 2036
- hierarchy
  - LDAP servers 473
- high availability
  - definition 1117, 1146
- high availability (HA) 1075
  - active-active 1095
  - load balancing 1095
- High Speed Packet Access (HSPA) 753
  - 600
- HMAC check offloading 1076
- HMAC-MD5 1647
- HMAC-SHA-1 1647
- HNT 2552
- hnt-restrict-source-ip 2557
- Home Location Register (HLR) 652, 653
- Home Network Identity (HNI) 752
- host check 2212
  - custom software 2213
  - introduction 2193
  - MAC address 2196
  - OS patch 2215
- host ID
  - peer 2592
- host name 1274, 1392
  - best practice 1143
- host-based
  - load balancing 1901
- hosted NAT traversal
  - See HNT 2552
- hosted-nat-traversal 2555
- hostname

- cluster members list 1269
- HTTP 2650
  - authentication 2678
  - persistence 1898
  - protocol optimization 2598
  - unknown HTTP sessions 2597
  - WCCP service ID 2704
- HTTP 1.1 conditionals 2650
- HTTP cookie
  - persistence 1899
- HTTP Header Field 668
- HTTP host
  - cookie persistence load balancing 1902
  - load balancing 1886, 1901
- HTTP multiplexing 1333
  - load balancing 1343
- HTTP port
  - web cache 238, 239, 2644, 2647
- HTTP redirect 1438
- HTTP rule
  - non-HTTP sessions 2597
- HTTPS 1387, 1435
  - antivirus 2176
  - antivirus quarantine 2176
  - data leak prevention 2177
  - DLP archive 2178
  - explicit web proxy 2665, 2689
  - FortiGuard Web Filtering 2177
  - load balancing 1343
  - persistence 1898
  - protocol recognition 2176
  - web filtering 2176
- HTTPS redirect 1438
- HTTP-User-Agent 525
- hub
  - HA schedule 1344
- hub-and-spoke
  - spoke subnet addressing 1680
- hub-and-spoke IPsec configuration
  - concentrator, defining 1684
  - configuration example 1689
  - hub configuration 1681
  - infrastructure requirements 1680
  - overview 1679
  - policy-based concentrator 1684
  - policy-based firewall policy 1683
  - route-based firewall policy 1682
  - route-based inter-spoke communication 1684
  - spoke configuration 1685
- humidity 1456
- I
- ICAP 946, 2278
  - example of ICAP 257
- ICMP 919, 922
  - session pickup 1334
- ICMP Types 923
- ICMP6 919
- ICMPv6 925, 1843
- ID tag 1524, 1527
- identify-based policies 2278
- Identity based policies (IBP) 460
- Identity based policy (IBP) 524
- identity-based policy 557, 591
  - Local ID 462
  - NTLM guest 525
  - NTLM user agent strings 525
- identity-based security policies 2602
- Idle timeout
  - VPN connection 464
- idle timeout
  - changing for the web-based manager 1403
- idle timeout setting 2216
- IEEE 802.1 2381
- IEEE 802.11a, channels 883
- IEEE 802.11b, channels 884
- IEEE 802.1Q 1523, 1527
- IEEE 802.1q 1094
- IEEE 802.3ad 1094
- ifHighSpeed 1517
- IF-MIB.ifSpeed 1517
- if-modified-since 2649
- IGMP
  - RFC 1112 1031
  - RFC 2236 1031
  - RFC 3376 1031
- ignore
  - web cache setting 2649
- Ignore User List 582
- IKE Configuration Method 1735
- IKE encryption key 1649
- IKE negotiation
  - parameters 1646
- IKEv2 1639
- IM 2023
  - load balancing 1343
- IMAPS
  - antivirus 2176
  - antivirus quarantine 2176
  - data leak prevention 2177
  - DLP archive 2178
  - email filtering 2177
  - predefined firewall services 2176
  - protocol recognition 2176
- inactivity timeout
  - SIP session 2518
- Inbound NAT, encryption policy 1661
- incoming-ip
  - explicit FTP proxy 2696
  - explicit HTTP proxy 2672
- incoming-ip6
  - explicit HTTP proxy 204, 1833, 2674
- incremental
  - synchronization 1308
- indentation 1410
- independent VDOM configuration 2408
- \_index 1411
- index 1296
- index number 1411

- Information Elements (IE) 696, 752
- informational
  - SIP response message 2498
- initiator 1677
- inspection
  - flow 2273
  - flow-based 2273
  - proxy 2274
  - security layers 2275
  - SSL 2173
  - stateful 2271
- inspection without address translation
  - SIP 2493, 2514
- installation 2578
- installation on Vista 2192
- Instant Messaging (IM) 2409
- \_int 1411
- IntegratedISIS 440
- interface
  - 802.1Q trunk 1530, 1540
  - accelerated NP 2443
  - dedicated monitoring 1254
  - external, VLAN NAT example 1535
  - external, VLAN NAT/Route example 1535
  - failover 1319
  - Gr 653
  - Gx 653
  - Gz 653
  - HA heartbeat 1294
  - heartbeat 1121, 1293
  - load balance virtual server 1884
  - loopback 279
  - maximum number 1523, 1554, 2382
  - monitor 1319
  - NTP server 1394
  - one-armed sniffer 1490
  - physical 2406
  - point-to-point 2403
  - proxy ARP 1884
  - reserved management interface 1254
  - security mode 1485
  - software switch 1486
  - VDOM link 2403
  - virtual interface 2401
  - VLAN subinterface 1530, 1534, 1535, 1540
- interface index
  - hash map order 1296
- interface mode 1099
- interface mode IPSec 1097
- interface monitoring 1145
  - aggregate interfaces 1188
  - definition 1146
  - redundant interfaces 1200
- interfaces
  - aggregate 1491
  - AMC card 1481
  - DHCP 1492
  - loopback 1489
  - MTU packet size 1496
  - physical 1480
- PPPoE 1493
- redundant 1489
- secondary IP address 1497
- virtual domains 1497
- virtual LANs 1498
- wireless 1495
- zones 1499
- interior gateway protocol (IGP) 306
- International characters 1420
- International Mobile Equipment Identity (IMEI) 754
- International Mobile Subscriber Identity (IMSI) 752
- Internet Assigned Numbers Authority (IANA) 314
- internet gateway protocol (IGP) 2367
- Internet Service Provider (ISP) 520
- Internet-browsing
  - configuring FortiClient 1740
- Internet-browsing firewall policy
  - VPN server 1738
- Internet-browsing IPsec configuration
  - FortiClient dialup-client configuration 1739
  - gateway-to-gateway configuration 1738
  - infrastructure requirements 1737
  - overview 1737
- interval
  - changing the heartbeat interval 1298
- interval, comfort clients 707
  - protection profile 708
- inter-VDOM
  - benefits 2400
  - Ethernet 2405
  - firewall policy 2410
  - independent configuration 2408
  - management configuration 2401
  - management VDOM 2409
  - meshed configuration 2401, 2410
  - NAT to TP 2405
  - PPP 2405
  - stand alone configuration 2401, 2407
  - virtual interface 2401
- intrusion detected
  - HA statistics 1273
- intrusion monitor 2184
- Intrusion Prevention System (IPS) 1095
- intrusion prevention system, see IPS
- Intrusion Protection 945
- intrusion protection
  - explicit FTP proxy 2697
  - explicit web proxy 2680
- intrusion protection system, see IPS
- Invalid Certificate Log 948
- iOS updates
  - caching 2644, 2648
- IP 920, 1344
  - load balance virtual server 1884
  - load balancing 1910
- IP address
  - multicasting 1032
  - overlapping 1531
  - peer 2592
- IP address conservation

- NAT 2544
- IP address range
  - setting for L2TP VPN 531
  - setting for PPTP VPN 530
  - setting for SSL VPN 528
- IP addresses, tunnel mode 2196
- IP Based authentication 2349
- IP filter, carrier end point 706
- IP header, differentiated services 2252
- IP monitoring
  - remote 1325
- IP pool
  - proxy ARP 1884
  - SIP 2543
- IP port
  - HA schedule 1344
- IP protocol 108 1824
- IP stack validation 2276
- IP, protocol 89 403
- IPcomp 1824
- IPS 945, 2572
  - buffer size 2071
  - concepts 2062
  - custom signature keywords 2081
  - custom signature syntax 2080
  - disabling for pinholes 2572
  - engine algorithm 2070
  - engine count 2070
  - explicit FTP proxy 2697
  - explicit web proxy 2680
  - fail-open 2070
  - filter 2064
  - in an HA cluster 2068
  - overview 2022
  - packet logging 2072
  - protocol decoders 2071
  - scanning 2064
  - sensor 2064
  - session count accuracy 2070
- IPS for GTP-U 761
- IPS, one-armed 2400
- IPSec 1074, 1075, 1078, 1095, 1097, 1098, 1099
  - interface mode 1097
  - tunnel 1096
  - tunnel mode 1097
- IPsec 1816, 2245
  - SAs 1318
  - security associations 1318
  - server type 530
  - tunnel 1816
- IPSec Interface Mode 1097, 1100
- IPsec monitor 1635
- IPSec VPN
  - adding manual key 1633
  - Autokey list 1624
  - concentrator list 1635
  - configuring phase 1 1625
  - configuring phase 1 advanced options 1626
  - configuring phase 2 1629
  - configuring phase 2 advanced options 1629
  - Manual Key list 1633
- IPsec VPN
  - and PCI DSS 776
  - authentication methods 1642
  - authentication options 1642
  - backup 1765
  - certificates 1642
  - configuring authentication for 528
  - DDNS routing 1695
  - dialup users, access to 504
  - dialup users, configuring authentication for 528
  - extended authentication (XAuth) 1650
  - firewall IP addresses, defining 1659
  - firewall IPsec policy 1661
  - keeping tunnel open 1654
  - load balancing 1343
  - logging events 1825
  - monitoring, dialup connection 1822
  - monitoring, static or DDNS connection 1822
  - peer 534
  - peer identification 1645
  - phase 1 parameters 1637
  - phase 2 parameters 1653
  - role of encryption policy 1662
  - route-based firewall policy 1663
  - testing 1823
  - troubleshooting 1826
- IPsec VPN SA
  - synchronization 1292
- ips-rtp 2572
- IPv4 1094
  - \_ipv4 1411
  - \_ipv4/mask 1411
  - \_ipv4mask 1411
  - \_ipv4range 1411
- IPv6
  - DHCP 1578, 1852
  - explicit web proxy 2665, 2673
  - session failover 1333
  - SIP 2558
  - \_ipv6 1411
- IPv6 IPsec configurations
  - certificates 1773, 1859
  - configuration 1773, 1859
  - firewall policies 1774, 1860
  - IPv4-over-IPv6 example 1778, 1864
  - IPv6-over-IPv4 example 1782, 1868
  - IPv6-over-IPv6 example 1774, 1860
  - overview 1772, 1858
  - phase 1 1773, 1859
  - phase 2 1773, 1859
  - routing 1774, 1860
  - \_ipv6mask 1411
- IPX, layer-2 forwarding 1551
- ISAKMP 1096, 1650
- ISISProtocol 438
- ISO 8859-1 1421
- ITU-T E.164 753



## J

jumbo frames 1073

## K

K-12 1585

keep alive 1633, 2201

keepalive 1650

Keepalive Frequency, Phase 1 1649, 1650

key 1408

key size

certificate 2175

Keylife

IPsec interface mode 1657

keylife 1654

Keylife, Phase 1 1647

Keylife, Phase 2 1654

keyword 1138

keywords

IPS custom signatures 2081

## L

l2ep-eth-type

CLI command 1297, 1363

L2TP 462, 1334

port 1701 1786

VPN, configuring authentication for 531

L2TP Access Concentrator (LAC) 1786

L2TP VPN

authentication method 1565

configuration steps 1564

enabling 1565

firewall IP addresses, defining 1565

infrastructure requirements 1564

network configuration 1564

security policy, defining 1566

VIP address range 1565

L2TP-over-IPsec 1786

LACP 1188

active-passive HA mode 1189

LACPDU 1189

lacp-ha-slave

CLI keyword 1189, 1363

LAG 1188

language

changing the web-based manager language 1403

Layer 2 Tunneling Protocol (L2TP) 1786

Layer 3 1094

Layer 4 1095

layer 4 2276

layer-2 1524, 1527, 1530

example 1524

forwarding 1551

layer-2 loops 2381

layer-2 switch

troubleshooting 1361

layer-3 1527

layer-3 switch

failover issues 1361

LDAP 587, 1259, 1260, 2318

access, collector agent 581

fnbamd 479

wildcard admin 476

LDAP authorization 1432

LDAP server, external

for L2TP 1565

for PPTP 1557

LDAP servers 471, 588

authenticating users with 489

binding 472

common name 473

configuring FortiGate unit to use 474

directory 472

Distinguished Name Query list 475

distinguished names 473

domain component 473

hierarchy 473

protocols 472

RFC compliance 472

lease breaking

DHCP 1580

lease time 1578

least round trip time

load balancing 1885

least RTT

load balancing 1885

least session

load balancing 1886

Least-Connection

HA schedule 1344

license key 2352

licenses

viewing 1397

life of a packet 1621, 2271

UDP 2271

Lightweight Directory Access Protocol (LDAP) 165, 564

FortiGate configuring 555, 589

XAuth authentication with 529

limited bandwidth 2236

limiting

number of SIP dialogs 2566

line endings 1423

link

failure 1291

multiple link failures 1323

link aggregation 1094

Link Aggregation Control Protocol 1188

link failover 1290, 1319

active-active 1344

active-passive 1128

aggregate interfaces 1188

definition 1147

not detected by high-end switches 1323

redundant interfaces 1200

link failure

remote 1325

link redundancy 1568

link-failed-signal 1323

CLI 1323

link-state advertisement (LSA) 306

- Linux 546
- listen on interfaces 1394
- lists
  - using web-based manager 1388
- load balance
  - according to loading 1350
  - cpu usage 1350
  - explicit web proxy 1343
  - first alive 1885
  - health check monitor 1888
  - health check monitoring 1888
  - health monitoring 1888
  - how busy 1350
  - HTTP host 1886
  - HTTP multiplexing 1343
  - HTTPS 1343
  - IM 1343
  - IPsec VPN 1343
  - least RTT 1885
  - least session 1886
  - memory usage 1350
  - P2P 1343
  - proxy UTM sessions 1350
  - round robin 1885
  - schedule 1347
  - source IP hash 1885
  - SSL offloading 1343
  - SSL VPN 1343
  - static 1885
  - virtual server IP 1884
  - VoIP 1343
  - WAN optimization 1343
  - WCCP 1343
  - weighted 1885
- load balancing 1075, 1095, 1122, 1123, 1881, 2601
  - active-active 1129, 1343
  - basic example 1893
  - definition 1147
  - HTTP host-based 1901
  - IP 1910
  - load-balance-all 1343
  - monitoring 1890
  - real servers 1886
  - SSL 1902
  - SSL offloading 1903
  - TCP 1910
  - UDP 1910
- load sharing 1571
- load-balance-all 1346
  - best practice 1143
  - enabling 1343
- local
  - console access 1405
  - domain name 1580
- Local certificates
  - generating request 536, 545
  - installing signed 539
- Local Gateway IP 1096, 1097, 1098, 1099, 1100, 1816
- local host 1078, 1095, 1096
- Local ID
  - for certificates 1643
  - to identify FortiGate dialup clients 1725
- Local Interface
  - IPsec VPN, manual key 1634
- Local SPI
  - IPsec VPN, manual key 1634
- local users
  - configuring 489
  - creating 489
  - deleting from FortiGate configuration 493
  - removing from FortiGate configuration 493
- Location Area Identity (LAI) 653
- location server
  - SIP 2492
- locking configuration 1504
- lockout
  - administrator 1437
- log disk
  - formatting 1264
- log file 587
- log message
  - HA 1261
- log message, FortiGate 1931
- log messages
  - HA 1262
- logging 1261, 2361
  - downloading quarantine logs 1942
  - enabling SSL VPN events 2216
  - GTP settings 697
  - HA log messages 1262
  - management practices 1953
  - MMS profile 673
  - setting event-logging parameters 2216
  - specifying GTP packets 699
  - viewing quarantine logs 1940
- logging in, security messages 533
- logging out
  - web-based manager 1404
- logging VPN events 1825
- login 1439
  - restricting unwanted 1436
- login page, custom 2212
- login screen 2203
- login, one time 2198, 2202
- logon blackout period 516
- logon events
  - logging to memory 582
- logs 1935
  - antivirus 1934
  - application control 1934
  - archives, DLP 1936
  - attack, IPS 1935
  - data leak prevention 1933
  - email filter 1935
  - event 1932
  - managing for individual cluster units 1261
  - nac quarantine 1933
  - network scan 1936
  - other traffic 1932
  - packet 1935
  - traffic 1931

- loopback interface 279
- loopback interfaces 1489
- lost heartbeat threshold
  - changing 1298
- M**
- MAC
  - MAC forwarding table 1323, 1361
- MAC address 1552, 2196
  - aggregate interfaces 1188
  - redundant interfaces 1201
  - virtual 1300
  - VRP virtual 1365
- MAC filter, wireless 804
- MAC forwarding tables 1323, 1361
- Mac OS 546, 566
- MAC table 2383
- Main Interface IP 1097
- main mode 1637
- maintenance
  - configuration revision 1453
  - disk 1599
- malformed-request-line 2561
- manage cluster units
  - HA 1284
- management access 1435
- management configuration 2409
- Management Information Base (MIB) 1514
- management interface
  - reserved 1254
- management IP address
  - changing 1393
- management services 2353
- management VDOM
  - 603, 2333, 2337, 2353, 2355, 2356, 2401
- Manual Key
  - IPSec VPN 1633
- MAPI 2580
  - protocol optimization 2598
- Martian addresses 279
- master DNS server 1581
- master unit 1075
  - See Also primary unit 1122
- matched content 1890
  - HTTP health check monitor 1889
- max cache object size
  - web cache 2648
- Max HTTP message length
  - web cache 2649
- Max HTTP request length
  - web cache 2649
- max TTL
  - web cache 2649
- max-body-length 2563
- max-dialogs 2566
- maximum age
  - spanning tree parameter 1362
- maximum bandwidth 2238, 2244, 2246, 2247
  - firewall policy 2246, 2247

- maximum connections
  - real server 1886
- maximum file size
  - antivirus 2039
- maximum frame size 1073
- MD5 1096, 1816
- memory 1554, 2351
- memory constraints 718
- memory usage
  - HA statistics 1272
  - WAN Optimization 2602, 2648
  - web caching 2602, 2648
  - weight 1350
  - weighted load balancing 1350
- meshed configuration 2401, 2410
- meshed VPN 1665
- message
  - SIP 2495
- message fingerprint 731
- message flood 735
  - alert notifications 722
  - threshold 722, 724
- message length
  - SIP 2563
- message request-line
  - SIP 2500
- message start line
  - SIP 2500
- message status-line
  - SIP 2500
- message, warning 1439
- MGCP 1607
  - session helper 1607
- MIB 1264
  - FortiGate 1520
  - HA 1264
  - RFC 1213 1520
  - RFC 2665 1520
- MicroSoft Challenge-Handshake Authentication
  - Protocol v1 (MSCHAP) 480
- Microsoft Point-to-Point Encryption (MPPE) 1556
- Microsoft Windows 1623
- Microsoft Windows VPN 1786
- middle-man 2278
- min TTL
  - web cache 2649
- missing MED 366
- MMS
  - address translation 637
  - carrier end point filtering 664
  - content checksum 674
  - DLP archive 726
  - file filtering 710
  - flood prevention 720
  - notifications 637
  - profile 661
  - virus scanning 702
- MMS message flood 678
- MMS profile

- DLP archive 672
- logging 673
- MMS address translation 668
- MMS bulk email filtering 665
- MMS notification 669
- scanning 663
- MMS protection profile 644
  - adding to a security policy 645
  - content archive 718
- MMS Service Provider Network (MSPN) 633
- Mobile Country Code (MCC) 752
- mobile device
  - Symbian Series 60 701
- Mobile Network Code (MNC) 752
- Mobile Station (MS) 646
- Mobile Station Identification Number (MSIN) 752
- Mobile Subscriber Integrated Services Digital Network (MSISDN) 703, 722, 753
- Mobile Subscriber Integrated Services Digital Network Number (MSISDN) 667
- mode
  - operation 2578
  - real server 1886
- Mode, Phase 1 1641
- modem 1574
  - modes 1574
  - routing 1576
- modes of operation
  - overview 2191
  - port forwarding 2192
  - tunnel mode 2191
  - web-only mode 2191
- monitor
  - application control 2155
  - HA statistics 1272
  - interface 1319
  - load balancing 1890
  - port 1319
- monitored interface
  - definition 1147
  - primary unit selection 1132
- monitoring
  - administrators 1397
  - antivirus 2184
  - applications 2187
  - archives and dlp 2186
  - attacks 2184
  - DHCP 1579
  - email activity 2186
  - endpoints 2188
  - FortiGuard quota 2187
  - ips 2184
  - IPsec sessions 1635
  - proxy forwarding 2675
  - RAID 1401
  - rogue APs 847
  - traffic shapers 2258
  - WAN Optimization 2615
  - WAN optimization 2603
  - web activity 2185
  - web caching 2652
  - wireless clients 846
- more 1422
- Mozilla 525
- MS RPC 1606
- MS Windows Active Directory (AD) 460, 524
- MS-CHAP 480
- MSIE 525
- MSISDN 737
- MTU (Maximum Transmission Unit)
  - 1073, 1084, 1095, 1817
- MTU packet size, interface 1496
- multicast
  - dense mode 1031
  - IGMP 1031
  - RFC 3973 1030
  - RFC 4601 1030
- multicast enhancement 805
- multicast-enable command 1036
- multicasting
  - debugging example 1047
  - enabling 1036
  - IP addresses 1032
  - RIPv2 1034
  - security policies 1036
- Multi-Exit Discriminator (MED) 366
- multi-line command 1417
- Multimedia Broadcast and Multicast Services (MBMS) 746
- Multimedia Message Service Center (MMSC) 633
- Multipath routing 279
- multiple account sessions, limits 526
- multiple group enforcement 505
- multiple pages 1422

**N**

- \_name 1411
- naming rules 2354
- NAT 2404
  - keepalive frequency 1650
  - port translation (NAT-PT) 1608
  - SDP 2546
  - SIP ALG IP address conservation 2545
  - SIP ALG NAT tracing 2545
  - SIP contact headers 2545
  - SIP session helper NAT tracing 2545
  - traversal 1649, 1827
  - VLAN example 1535
  - with IP address conservation 2544
- NAT 64 903
- NAT 66 903
- NAT device
  - authentication 2678
- NAT mode 1392
- NAT/Route Mode 906
- NAT/Route mode 2600
  - general configuration steps 1153, 1189, 1201, 1222
  - HA network topology 1153
  - reserved management interface 1255

- web-based manager configuration steps
  - 1154, 1158, 1166, 1170, 1177, 1180
- NAT64
  - session failover 1333
- NAT66
  - session failover 1333
- nat-trace 2545
- Nat-traversal, Phase 1 1649
- negative response duration
  - web cache 2648
- negotiating
  - IPsec phase 1 parameters 1647
  - IPsec phase 2 parameters 1653
- nested tunnels 740
- NetBIOS, for Windows networks 1553
- network
  - train 1300
- Network Access Server (NAS) 468, 470
- Network Address Translation (NAT) 1649
- network equipment
  - failover time 1362
- Network Identifier 753
- network instability 1551
- Network Layer or Packet Filter Firewalls 891
- Network layer Service Access Point Identifier (NSAPI) 650
- network processing unit (NPU) 1078
- network processor accelerated interfaces
  - accelerate active-active HA 1346
- network processors
  - FA2 (NP1) 1073
  - FortiAccel (NP1) 1073
  - NP1 1073
  - NP2 1073
  - NP4 1073
  - NP6 1073
- Network Time Protocol (NTP) 2293, 2318
- Network Time Protocol server (NTP) 1394
- network topologies 814
- network topology
  - dynamic DNS 1695
  - FortiClient dialup-client 1709
  - FortiGate dialup client 1724
  - fully meshed network 1665
  - gateway-to-gateway 1665
  - hub-and-spoke 1679
  - Internet-browsing 1737
  - NAT/Route mode HA 1153
  - partially meshed network 1665
  - redundant-tunnel 1741
  - supported IPsec VPNs 1622
  - transparent mode VPN 1766
- network utilization
  - HA statistics 1272
- next 1415
- nic
  - get hardware nic
    - 1155, 1160, 1167, 1171, 1178, 1181, 1191, 1195, 1203, 1207, 1224, 1229, 1244, 1248, 1304
- none
  - HA schedule 1344
- non-HTTP sessions
  - HTTP rule 2597
- no-sdp-fixup 2546
- notification alerts 714, 728
- not-so-stubby area (NSSA) 273
- Novell eDirectory 165, 564
- Novell eDirectory - see Directory Service
- NP interface 2443
- NP interfaces 764
- NP1 1073, 1346
- NP1 processor 1073
- NP2 1073, 1346
- NP2 interface 2247
- NP2 network processor 1816
- NP2 network processors 1816
- NP2 processor 1073
- NP4 1346
- NP4 processor 1073
- NP6 processor 1073
- NT LAN Manager (NTLM) 165, 460, 524, 564
- NTLM 521
  - Guest profile access 525
  - HTTP-User-Agent 525
  - NTLM mode 569
  - user agent strings 525
  - XSS vulnerability 525
- NTLM enabled browsers 525
- NTLM statistics 585
- NTP server 1428
  - listen interfaces 1394
- NT-style domain mode implementation 566
- null modem 1405, 1407

**O**

- object 1410
- OID 1264, 1266, 1267
- ONC-RPC 1606, 1608
- one time passcode (OTP) 459, 483, 497
- one-armed IPS 2400
- one-armed sniffer 1490
- one-time login 2198, 2202
- Online Certificate Status Protocol (OCSP) 534, 535
- open shortest path first (OSPF). See routing, OSPF
- Open Systems Interconnect (OSI) 1524
- OpenLDAP 471
- OpenSSL 535
- Opera 525
- operating a cluster 1252
- operating mode
  - active-passive 1128
- operating temperature 1456
- operation mode 1393, 2578
  - active-active 1129
- option 1410

- order of operations for shapers 2243
- OS patch check 2213, 2215
- OSFP
  - graceful restart 1316
- OSPF
  - protecting with IPsec 1808
  - with redundant IPsec tunnels 1814
- ospf
  - adjacent routers 404, 409
  - area 403
  - area border router (ABR) 403
  - Dijkstra's algorithm 406
  - e1 273
  - e2 273
  - Hello packets 404
  - Hello protocol 404
  - IP datagrams 403
  - link-state 403
  - neighbor 404
  - NSSA 273
  - path cost 406
  - state of neighbor 409
- ospf AS 399
- out of band management 1254
- out of path
  - topology 2581
- Outbound NAT, encryption policy 1661
- outgoing-ip
  - explicit FTP proxy 2696, 2697
  - explicit HTTP proxy 204, 1833, 2672, 2674
- outgoing-ip6
  - explicit HTTP proxy 205, 1833, 2674
- overlap
  - resolving IP address 1725
  - resolving through FortiGate DHCP relay 1725
- overlapping VPN subnets 1672
- override 1138
  - and primary unit selection 1138
  - configuration changes lost 1141
  - disconnecting a unit from a cluster 1142
  - primary unit selection 1136, 1140
- oversize threshold 665, 708
- Oversized File Log 947
- Oversized File/Email Threshold 948
- P**
- P1 Proposal, Phase 1 1646
- P2 Proposal 1097
  - IPSec VPN, phase 2 1630
- P2 Proposal, Phase 2 1653
- P2P 2023
  - load balancing 1343
- PAC
  - explicit web proxy 2665
- packet
  - flow 2275
  - forwarding rate 1075, 1097
  - gratuitous ARP 1300
  - life of 2271
  - processing flow 1074
- packet capture 1512
- Packet Data Protocol (PDP) 745
- packet data protocol (PDP) 645
- packet flow 1074
- packet header 1507
- packet logging
  - IPS 2072
  - settings 2180
- packet rates 2238
- packet sniffer 764, 2441
  - verbosity level 765, 2442
- Packet verification 2276
- packets
  - layer-3 routing 1527
  - VLAN-tagged 1530
- Packettypes 441
- page controls
  - web-based manager 1388
- paging 1422
- PAP 480, 530, 1555
- parity 1406
- partially meshed VPN 1665
- password
  - administrator 2578
  - authentication group 2613
  - changing 1429
  - changing, administrator 1397
  - configuring authentication 1435
  - FortiClient 1633, 2201
  - HA configuration option 1130
- Password Authentication Protocol (PAP) 480, 1651
- password policy
  - PCI DSS requirements 779
- patch check, host OS 2215
- \_pattern 1411
- pattern 1411
  - carrier end point 681, 705
  - creating 2142
  - default file block list 711
- PCI DSS 775
  - defined 768
  - example network topology 772
  - firewall policy considerations 773
  - logging wireless network activity 776
  - objectives and requirements 768
  - scanning for rogue wireless APs 774
  - wireless guidelines 771
- peer 1274
  - accept any peer 2610
  - host ID 2592
  - IP address 2592
  - monitoring WAN optimization 2615
  - WAN optimization 2610
- peer authentication 2611
  - WAN optimization 2610
- peer ID
  - assigning to FortiGate unit 1644
  - enabling 1645
- Peer Options 1639

- peer user groups
  - configuring 507
  - creating 507
- peer users 488, 493
  - configuring 493
  - creating 493
- per policy shaper 2244, 2245
- perfect forward secrecy (PFS) 1675
- perfect forward secrecy, enabling 1654
- performance
  - improving session pickup performance 1331
- periodic
  - synchronization 1309
- per-IP 2247
  - NP2 interface 2247
- per-IP traffic shaping 2247
- Perl regular expressions
  - carrier end point pattern 681
- Perl regular expressions, using 1423
- permissions 1415
- persistence 1885
  - HTTP cookie 1899
  - HTTP/HTTPS 1898
- personal bookmarks 2203
- Phase 1629
- Phase 1 1096, 1097, 1098, 1099, 1100
- phase 1
  - IPSec VPN 1625, 1629
- phase 1 advanced options
  - IPSec VPN 1626
- phase 1 parameters
  - authenticating with certificates 1639
  - authenticating with preshared keys 1640
  - authentication method 1642
  - authentication options 1642
  - defining 1637
  - defining the tunnel ends 1638
  - IKE proposals 1647
  - main or aggressive mode 1638
  - negotiating 1647
  - overview 1637
  - peer identifiers 1644
  - user accounts 1645
- Phase 2 1078, 1097, 1098, 1099, 1100
- phase 2
  - Autokey keep alive 1654
  - IPSec VPN 1629
  - key expires 1654
  - PFS 1675
- phase 2 advanced options
  - IPSec VPN 1629
- phase 2 parameters
  - autokey keep alive 1654
  - auto-negotiate 1654
  - configuring 1655, 1656
  - defining 1653
  - DHCP-IPsec 1654
  - keylife 1654
  - negotiating 1653
  - perfect forward secrecy (PFS) 1654
  - quick mode selectors 1655
  - replay detection 1653
- Phase I 1096, 1816
- Phase II 1096, 1816
- physical interface 2406
- ping
  - health check monitor 1894
- ping server 1569, 1576
- pinhole
  - disabling IPS 2572
  - more secure 2549
  - RTP 2490, 2494
  - SIP 2490, 2494
  - smaller 2549
  - strict-register 2549
- PKI authentication - see peer users
- PKI certificate
  - PKCS 538
- PKI user 494
- planning VPN configuration 1621
- pmap
  - session helper 1608
- PMK caching 791
- point-to-point interface 2403
- Point-to-Point Tunneling Protocol (PPTP) 1555
- policies
  - multicast 1036
- policy 1095
  - guaranteed bandwidth 2246
  - maximum bandwidth 2246, 2247
- Policy order 954
- policy route
  - adding 270, 285
  - moving in list 287
- policy server, VPN
  - configuring FortiGate unit as 1717
- policy6\_list 894
- policy-based VPN
  - vs. route-based 1621
- Polling mode 567, 580
- polling mode
  - event log polling 566
- POP3S
  - antivirus 2176
  - antivirus quarantine 2176
  - data leak prevention 2177
  - DLP archive 2178
  - email filtering 2177
  - predefined firewall services 2176
  - protocol recognition 2176
- port
  - forwarding 2192
  - number, web-portal connections 2211
  - virtual server 1885
- port 10443 504
- port 1701 1786
- port 179 368
- port 21123 697
- port 212 745

- port 2152 746, 761
- port 445 566
- port 4500 1649
- port 47 1608
- port 500 1649
- port 8000 587
- port monitor 1319
  - virtual clustering 1219
- port monitoring 1145
  - aggregate interfaces 1188
  - redundant interfaces 1200
- port number
  - changing the port numbers that the SIP ALG listens on 2516
  - changing the port numbers that the SIP session helper listens on 2508
- port, RADIUS servers 470
- port, session helper 1603
- ports
  - port 1024 2320
  - port 1025 2320
  - port 53 2320
  - port 8888 2320
- power
  - security consideration 786
  - WLAN power level 784
- power off 1458
- PPP 1142
- PPPoE 1142
- PPPoE interface 1493
- PPTP 462, 1334
  - external server 1560
  - layer-2 forwarding 1551
  - session helper 1608
- PPTP PPTP VPN
  - configuring authentication for 530
- PPTP VPN
  - authentication method 1557
  - configuring pass through 1560
  - enabling 1558
  - FortiGate implementation 1555
  - IP address range 530
  - security policy, defining 1559
  - VIP address range 1558
- PRACK
  - SIP message 2490, 2498
- pragma-no-cache 2650
- pre-authentication 791
- predefined firewall services
  - IMAPS, POP3S, SMTPS 2176
- preempt mode
  - VRRP 1369
- preserve-override 2545
- pre-shared key
  - authenticating FortiGate unit with 1641
  - authentication group 2613
- preshared key 461, 1617
- primary cluster unit
  - definition 1147
- primary unit 1075, 1122
  - connected monitored interfaces 1132
  - definition 1147
  - getting information using SNMP 1264
  - override keyword 1138
  - recovery after a failover 1291
  - selection 1131
  - SNMP get 1264
- primary unit selection
  - age 1132, 1133
  - basic 1131
  - device priority 1131, 1136
  - FortiGate unit serial number 1137
  - interface monitoring 1132
  - monitored interfaces 1132
  - override 1136, 1138, 1140
  - serial number 1137
- priority
  - cluster members 1269
  - heartbeat interface 1294
- priority traffic 2244
- problem scope 2322
- profile
  - VoIP 2515
- proposal
  - IPSec VPN, phase 2 1630
- protection profile
  - amount, comfort clients 708
  - interval, comfort clients 707, 708
- protocol
  - ospf Hello 404
- protocol anomaly attacks 739
- protocol decoders 2071
- Protocol Independent Multicast (PIM) 1030
- protocol optimization 2580
  - CIFS 2598
  - FTP 2598
  - HTTP 2598
  - MAPI 2598
  - TCP 2598
- protocol recognition
  - HTTPS, IMAPS, POP3S, SMTPS 2176
- Protocol Types 919
- protocol, session helper 1603
- protocols
  - authentication 517
  - LDAP servers 472
- provisional
  - SIP response message 2498
- provisional response acknowledgement
  - SIP message 2490, 2498
- provisional-invite-expiry-time 2520
- proxy
  - antivirus 2680, 2697
  - DLP 2680, 2697
  - explicit web 1333
  - explicit web proxy authentication 2678
  - FortiGuard web filtering 2680
  - web filtering 2680
- proxy ARP 1884



- FortiGate interface 1884
- IP pool 1884
- virtual IP 1884
- proxy auto-config
  - explicit web proxy 2665
- proxy chaining
  - explicit web proxy 2674
  - health monitoring 2675
- proxy forwarding server
  - explicit web proxy 2674
  - health checking 2675
- proxy FQDN
  - web cache 2649
- proxy inspection 2274
- proxy server
  - SIP 2491
- Proxy Servers 892
- proxy UTM sessions
  - weighted load balancing 1350
- proxy, web
  - IPv6 2673
- Pseudonodes 440
- public key cryptography standard (PKCS) 536
- public key infrastructure (PKI) 532
- Public Land Mobile Network (PLMN) 652
- publis DNS server 1582
- purge, shell command 1414
- push update 1464, 1465
  - override 1466

## Q

- Quality of Service 906
- quality of service 2234
- quarantine 733
  - file 1268
- quarantine files list
  - apply 1941
  - date 1941
  - DC 1941
  - file name 1941
  - filter 1941
  - service 1941
  - sorting 1940
  - status 1941
  - status description 1941
  - TTL 1941
  - upload status 1941
- quarantine logs 1940
- Query list
  - LDAP Distinguished Name 475
- Queuing 907
- queuing 2235
- Quick mode selectors, Phase 2 1655
- quota
  - FortiGuard Web Filter 2101

## R

- RADIUS 1259, 1260, 1735, 2318
  - assigning client IPs with 1713

- attributes 468
- authentication servers 466
  - long key 467
  - vendor ID 469
  - XAuth authentication with 529
- radius 470
- RADIUS server, external
  - for L2TP 1565
  - for PPTP 1557
- RADIUS servers
  - attribute dictionary 469
  - authenticating users with 489
  - changing default port 470
  - configuring FortiGate unit to use 470
  - default port 470
  - port 470
  - VSA 469
- random
  - HA schedule 1344
- RAS, session helper 1607
- rate limit
  - number of SIP dialogs 2566
- rate limiting
  - SCCP 2565
  - SIMPLE 2565
  - SIP 2565
- read & write access level
  - administrator account 1395
- read only access level
  - administrator account 1395
- real server
  - active mode 1886
  - adding 1887
  - disabled mode 1886
  - health check monitoring 1888
  - health monitoring 1888
  - load balancing 1886
  - maximum connections 1886
  - mode 1886
  - standby mode 1887
  - weight 1886
- Real Time Control Protocol 2520
- Real Time Protocol 2532
- realm
  - explicit web proxy 2672
- reboot, upgrade 1454
- recording log messages 1931
- recursive DNS 1583
- redirect 1438
- redirect server
  - SIP 2491
- redistributed routes
  - ospf e1/e2 273
- redundant
  - interfaces 1489, 1568
  - modem mode 1574
- redundant interface
  - active-active mode 1201
  - active-passive mode 1201
  - best practice 1144

- HA 1122, 1239
- HA MAC addresses 1201
- port monitoring 1200
- redundant UTM protocol 1119
- redundant VPNs
  - configuration 1742
  - example, fully redundant configuration 1745
  - example, partially-redundant configuration 1758
  - overview 1741
- refresh every
  - HA statistics 1272
- regex 2181
- register-contact-trace 152, 2549
- registering
  - with Fortinet Technical Support 2578
- registrar
  - SIP 2492
- Registration, Admission, and Status (RAS) 1607
- registry
  - key 585
  - remote service 597
- regular expression 1411
- regular expressions 2181
- relay
  - DHCP 1142, 1577, 1578
- Release Notes 564
- remote
  - administration 1435
  - FortiManager options 1396
  - L2TP VPN client 1566
  - shell 1609
- remote administrator 476
- remote client
  - authenticating with certificates 1639
  - FortiGate dialup-client 1724
  - in Internet-browsing IPsec configuration 1737
- Remote Gateway
  - IPSec manual key setting 1634
- remote gateway
  - dialup user 1654
- Remote Gateway, Phase 1 1639
- remote Internet access 2221
- remote IP monitoring 1325
- remote link failover
  - best practice 1145
  - virtual clustering 1219
- remote link failure 1325
- remote peer
  - authenticating with certificates 1639
  - dynamic DNS configuration 1704
  - gateway-to-gateway IPsec configuration 1667
  - manual key configuration 1633
  - transparent IPsec VPN configuration 1767
- Remote SPI
  - IPSec VPN, manual key 1634
- removing
  - local users from FortiGate configuration 493
  - user group from FortiGate configuration 508
- rename, shell command 1414
- Rendezvous Points (RPs) 1030
- replacement FortiGate unit
  - adding to a cluster 1187
- replacement message
  - explicit FTP proxy 2691
- replacement message groups 1598
- replacement messages
  - administration 1594
  - alert mail 1595
  - authentication 518
  - captive portal default 1596
  - Device Detection Portal 1596
  - Endpoint Control 1596
  - FortiGuard web filtering 1596
  - FTP 1596
  - HTTP 1596
  - IM, P2P 1597
  - images 1591
  - mail 1596
  - MM1 1598
  - MM3 1598
  - MM4 1598
  - MM7 1598
  - modifying 1592
  - NAC quarantine 1597
  - NNTP 1597
  - RSA
    - SecurID 519
  - SMS 636
  - spam 1597
  - SSL VPN 1597
  - tags 1592
  - traffic quota control 1598
  - user authentication 1595
  - viewing 1591
  - web proxy 1597
- replacing a broken cluster unit 1187
- replacing a failed cluster unit 1187
- replay detection
  - 1078, 1096, 1097, 1098, 1099, 1100, 1816
- replay detection, enabling 1653
- request
  - SIP 2489
- request messages 2497
- Request-line
  - deep SIP message checking 2561
- request-line
  - SIP 2500
- reserved characters 1418
- reserved management interface 1254
  - NAT/Route mode 1255
  - transparent mode 1255
- reset age
  - command 1135
- reset uptime
  - command 1135
- reset-uptime
  - diagnose command 1135
- response
  - SIP 2489

- restore
  - cluster configuration 1276
- restore defaults 1446
- restoring configuration See widgets
- restricting login attempts 1436
- Return Material Authorization (RMA) 2330
- revalidated pragma-no-cache 2650
- reverse explicit FTP proxy 2586
- reverse path lookup 279
- reverse proxy
  - web cache 2588, 2655
  - with web caching 2644
- reverse shaping 2250
- reverting firmware 1452
- revisions 1446
- revocation list, importing 539
- RFC
  - 1213 1514, 1520
  - 2516 1493
  - 2665 1514, 1520
  - RFC 1349 286
  - RFC 1519 305
  - RFC 1771 358
  - RFC 1965 363
  - RFC 1966 362
  - RFC 1997 366
  - RFC 2385 360
  - RFC 2453 319
  - RFC 3065 363
  - RFC 3509 406
  - RFC 4271 358
  - RFC 4632 305
  - RFC 5237 285
  - RFC 791 286
  - SIP 2490
- RFC 1112 1031
- RFC 1700 696
- RFC 2236 1031
- RFC 2474 2251
- RFC 2475 2251
- RFC 2543 2571
- RFC 2986 536
- RFC 317 1824
- RFC 3286 921
- RFC 3376 1031
- RFC 3973 1030
- RFC 4601 1030
- RFC 4960 921
- RFC 5280 532
- RFC 791 2250
- RFC compliance
  - LDAP servers 472
- RFC4960 658
- RIP
  - hop count 326
  - RFC 1058 319, 438
  - RFC 2453 319
  - RIP Next Generation (RIPng) 320, 444
  - version 1 319, 438
  - version 2 319
- RIPv2 1034
- RJ-45 1405
- RJ-45-to-DB-9 1405, 1407
- roaming subscribers 687
- role
  - cluster members 1269
- Role Based Access Control (RBAC) 469, 503, 523
- root certificate, installing 539
- round robin
  - load balancing 1885
- Round Trip Time (RTT) 2320
- Round-Robin
  - HA schedule 1344
- route 1099, 1100, 1101
- route flap 373
  - HA 371
- route hold 1317
- route reflectors (RR) 362
- route synchronization 1292
- route-based VPN
  - firewall policy 1663
  - vs. policy-based 1621
- route-hold 1316
- Routeleaking 442
- router
  - WCCP 2703
- router monitor
  - HA 272
- route-ttl 1316
- route-wait 1316
- routing 2211
  - administrative distance 279
  - asymmetric 1554
  - BGP 1533, 2410
  - blackhole 278
  - configuring 2665, 2689
  - default for VLAN 1533
  - domain 314
  - ECMP 278
  - enhanced packet-matching 316
  - hop count 2366
  - loopback interface 279
  - modem 1576
  - multicast 2411
  - OSPF 1533, 2410
  - RIP 1533, 2411
  - routing table, searching 276
  - STP 1554
  - viewing information 271
- Routing Area Identifier (RAI) 754
- routing information protocol (RIP). See routing, RIP
- routing policy
  - protocol number 285
- routing table 406, 2277
  - removing routes 360
- routing table updates
  - synchronizing 1315
- routing, asymmetric 764

- routing, default 2375
- routing, default route
  - VDOM example 2375, 2378
- routing, transparent VPN IPsec configuration 1769
- RPF (Reverse Path Forwarding) 764, 2442
- RSA 483
  - X.509 461
- RSA SecurID 519
- rsh, session helper 1609
- RTCP 2490, 2520
- RTP 2490, 2532, 2542
  - hardware acceleration 2497
  - pinhole 2490, 2494
- RTSP, session helper 1610
- rule
  - non-HTTP sessions 2597
  - unknown HTTP sessions 2597

## S

- SA
  - IPsec 1318
- save password, FortiClient 1633, 2201
- scan buffer size
  - antivirus 2038
- scanning order
  - antivirus 2033
- SCCP
  - DoS sensor 2565
  - protection profile 2565
  - rate limiting 2565
  - VoIP profile 2514
- schedule
  - antivirus and attack definition updates 1465
  - load balance 1347
- Schedule Expiration 944
- Schedule groups 943
- school administration 1585
- SCP
  - authentication 1445
  - backup configuration 1443
  - client application 1444
  - restore configuration 1445
  - SSH access 1444
- screen resolution
  - minimum recommended 1387
- scripts
  - uploading 1600
- SCTP 655, 921
- SDP 2490, 2496, 2502
  - NAT 2546
  - session profile 2502
- Secure Certificate Enrollment Protocol (SCEP) 534
- secure HTTP (HTTPS) 533
- Secure Shell (SSH)
  - key 1408
- secure tunnelling 2580
- SecurID 483
  - firewall policy 485
- security 785

- security association
  - IPsec 1318
- Security Association (SA) 1653
- security association (SA) 1074, 1078, 1096, 1816
- security certificate 1427
- security IP addresses
  - defining L2TP 1565
- security layer
  - stateful inspection 1621
- security layers 2275
- security mode 1485
- security policies
  - multicast 1036
- security policy 2249
  - defining L2TP 1565, 1566
  - defining PPTP 1559
  - identity-based 2602
  - MMS protection profile 645
  - VLAN 1532
  - VLAN example 1537
  - VLAN transparent mode 1543, 1546
- security policy, web-only mode access 2204
- security processing modules 1071
  - configuring 2071
  - displaying information 1072
  - example configuration 2077
- security profile
  - explicit FTP proxy 2695
  - explicit web proxy 2669, 2694
  - ftp proxy 2695
  - web proxy 2669
- security vulnerabilities
  - monitoring for (PCI DSS related) 778
- selecting the primary unit 1131
- self-signed certificate 533
- self-validate 535
- sensor
  - IPS 2064
- serial communications (COM) port 1405
- serial no
  - HA statistics 1272
- serial number
  - getting using SNMP 1267
  - primary unit selection 1137
- Series 60 701
- server
  - DHCP 1142, 1577
  - WCCP 2703
- server certificate 2200
  - installing signed 539
  - obtaining 538
- server comfoting 708
- server load balance port forwarding virtual IP
  - adding 1917
- server load balance virtual IP
  - adding 1911
- servers
  - configuring XAuth authentication using 529
- service

- quarantine files list 1941
- Service Categories 919
- service group
  - VDOM Transparent example 2393
  - WCCP 2704
- service ID
  - WCCP 2704
- service number
  - WCCP 2704
- service, DHCP 1578
- Serving GPRS Support Node (SGSN) 650
- session
  - failover 1291
    - key 1074
  - session count accuracy 2070
  - Session creation 2276
  - session description protocol
    - See SDP 2502
  - session failover 1123, 1128, 1330
    - active-active 1344
    - definition 1147
    - enabling 1330
    - failover
      - session 1290
    - GTP and HA 1334
    - IPv6 1333
    - NAT64 1333
    - NAT66 1333
    - SIP 1333, 2569
  - session helper
    - 1095, 1605, 1606, 1607, 1608, 1609, 1611, 2278
    - changing the configuration 1603
    - changing the port numbers that the SIP session helper listens on 2508
    - dcerpc 1606
    - disabling the SIP session helper 2506
    - DNS 1606
    - enabling the SIP session helper 2506
    - H.245 1606
    - h245O 1606
    - h323 1607
    - mgcp 1607
    - pmap 1608
    - port 1603
    - PPTP 1608
    - protocol 1603
    - ras 1607
    - rsh 1609
    - rtsp 1610
    - sip 1610
    - TFTP 1610
    - tns 1611
    - viewing 1602
  - session hijacking 544
  - Session Initiation Protocol for Instant Messaging and Presence Leveraging Extensions
    - See SIMPLE 2514
  - Session Initiation Protocol. See SIP
  - session pick-up
    - definition 1147
    - session pickup 1128
      - best practice 1143
      - connectionless 1334
      - delay 1331
      - enable 1330
      - enhancing performance 1331
      - ICMP 1334
      - improving performance 1331
      - selecting FortiGate interfaces to use 1331
      - UDP 1334
    - session profile
      - SDP 2502
    - session synchronization
      - between two standalone FortiGate units 1370
      - improving performance 1331
      - using multiple FortiGate interfaces 1331
    - session tables 2278
    - session timeout 507
    - session-based authenticated user 2349
    - session-helper 1602
    - session-pickup 1263
      - CLI command 1294, 1330
    - session-pickup-delay 1263, 1331
    - session-sync-dev 1331
    - set 1415
    - setting
      - authentication protocols 517
      - firewall policy authentication 517
      - firewall user authentication timeout 514
      - SSL VPN authentication timeout 514, 528
    - setting administrative access for SSH or Telnet 1406
    - settings 1441
      - administrators 1441
    - SHA1 1096, 1816
    - SHA-256 1647
    - SHA-256, 384, 512 1648
    - shadow DNS server 1582
    - shaper
      - all policies 2244, 2245
      - application control 2248
      - per policy 2244, 2245
      - per-IP 2247
      - processing order 2243
      - security policy 2249
      - shared 2244
    - shared secret 461
    - shared shaper 2244
    - shared traffic shapers 2245
    - sharing
      - WAN optimization tunnels 2601
    - shell command
      - delete 1413
      - edit 1413
      - end 1413
      - get 1413
      - purge 1414
      - rename 1414
      - show 1414
    - shielded twisted pair 1457

- Shift-JIS 1420, 1421
- Shortest Path First (SPF) 406
- show 1415
  - shell command 1414
- shut down 1458
- signature-based IPS 2277
- signatures, update 1429
- SIMPLE
  - protection profile 2565
  - rate limiting 2565
  - VoIP profile 2514
- single login 2198, 2202
- Single Sign On (SSO) 532, 2197
- SIP
  - accepting register response 2527
  - blocking requests 2563
  - changing the port numbers that the SIP ALG listens on 2516
  - changing the port numbers that the SIP session helper listens on 2508
  - contact headers and NAT 2545
  - deep header inspection 2558
  - deep message inspection 2558
  - destination NAT 2542
  - dialog 2489, 2495
  - different source and destination NAT for SIP and RTP 2544
  - disabling the SIP session helper 2506
  - DoS sensor 2565
  - enabling the SIP session helper 2506
  - fields 2500
  - fuzzing protection 2558
  - HA session failover 2569
  - headers 2500
  - inspection without address translation 2493, 2514
  - IP address conservation 2545
  - IPv6 2558
  - location server 2492
  - message request-line 2500
  - message sequence 2495
  - message start line 2500
  - message status-line 2500
  - NAT tracing 2545
  - NAT with dynamic IP pool 2543
  - NAT with IP address conservation 2544
  - pinhole 2490, 2494
  - protection profile 2565
  - proxy server 2491
  - rate limiting 2565
  - redirect server 2491
  - registrar 2492
  - request 2489
  - request-line
    - SIP 2500
  - response 2489
  - RFCs 2490
  - session failover 1333
  - source NAT 2542
  - start line
    - SIP 2500
  - status-line
    - SIP 2500
  - Transparent mode 2493, 2514
  - user element 2489
  - VoIP profile 2515
- SIP ALG
  - changing the port numbers that the SIP ALG listens on 2516
  - NAT tracing 2545
- SIP dialogs
  - limiting the number 2566
- SIP message
  - body 2490
  - final 2498
  - headers 2490
  - informational 2498
  - PRACK 2490, 2498
  - provisional 2498
- SIP phone
  - FortiFone 2490
- SIP requests 2563
- SIP session
  - inactivity timeout 2518
- SIP session helper
  - changing the port numbers that the SIP session helper listens on 2508
  - disabling 2506
  - enabling 2506
  - NAT tracing 2545
- SIP, session helper 1610
- sip-nat-trace 2545
- sip-ssl-port 2516
- sip-tcp-port 136, 2516
- sip-udp-port 136, 2516
- Skinny Call Control Protocol
  - See SCCP 2514
- Skinny Call Control Protocol. See SCCP
- slave DNS server 1581
- slave unit 1076, 1122
  - See Also subordinate unit 1122
- SMS text message 636
- SMS token 496
- SMTSPS
  - antivirus 2176
  - antivirus quarantine 2176
  - data leak prevention 2177
  - DLP archive 2178
  - email filtering 2177
  - predefined firewall services 2176
  - protocol recognition 2176
- SNAT 1095, 2404
- sniffer 1490, 1512
- SNMP 1264, 2361
  - Agent 1515
  - configuring community 1518
  - get command 1521
  - gigabit interfaces 1517
  - HA reserved management interface 1255
  - manager 1514, 1518
  - MIB 1264

- MIBs 1520
- queries 1516, 1518, 1519
- RFC 12123 1520
- RFC 2665 1520
- trap 1264
- v3 1514, 1515
- SNMP get
  - any cluster unit 1266, 1267
  - primary unit 1264
  - subordinate unit 1266, 1267
- snmpget 1258, 1259, 1266
- SOCKS
  - explicit web proxy 2665
- soft-switch 1486
- sorting
  - quarantine files list 1940
- source IP hash
  - load balancing 1885
- source NAT
  - SIP 2542
- spam 731
- spanning tree
  - forward delay 1362
  - maximum age 1362
- spanning tree protocol 1362
  - settings and HA 1362
- Spanning Tree Protocol (STP) 1551, 1553, 2381
- special characters 1418
- Spill-over 300
- split brain 1294
  - heartbeat 1144
- split DNS 1582
- split tunnel 2223
- split tunneling 2224
- SQL 467
- SQLNET
  - session helper 1611
- SSH 1406, 1407, 1435
  - key 1408
- SSID
  - described 793
  - whether to broadcast 785
- SSL 533
  - antivirus 2176
  - antivirus quarantine 2176
  - certificate 2174
  - content inspection 2173
  - content scanning 2173
  - data leak prevention 2177
  - DLP archive 2178
  - email filtering 2177
  - example 2046
  - FortiGuard Web Filtering 2177
  - HTTPS 2177
  - inspection 2173
  - load balancing 1902
  - predefined firewall services 2176
  - protocol recognition 2176
  - settings, all 2175
  - supported FortiGate models 2173
  - web filtering 2176
- SSL Client Certificate Restrictive option 545
- SSL offloading 239, 1333, 2580, 2596, 2645
  - certificates.certificate
    - SSL offloading 1904
  - Client to FortiGate 1903
  - Client to FortiGate to Server 1903
  - full mode 1903
  - half mode 1903
  - load balancing 1343, 1903
- SSL VPN
  - allow/deny client renegotiation 2211
  - authentication timeout 514, 528
  - checking client certificates 2200
  - event logging 2216
  - FortiClient 2224
  - host check 2212
  - host OS check 2215
  - load balancing 1343
  - personal bookmarks 2203
  - specifying server certificate 2200
  - specifying timeout values 2200
  - split tunneling 2224
  - Subsession 609, 2223
  - user authentication 528
  - user groups, configuring 504
  - user groups, creating 504
  - user groups, IPsec VPN dialup users 504
  - Virtual Desktop 2219
  - web portal 2199
- SSL VPN user groups 504
- ssl.root 2278
- SSO (Single Sign On) 2197
- standalone FortiGate unit
  - adding to a cluster 1186
  - converting to a cluster 1184
- standalone mode 1574
- standalone session synchronization 1370
- standby mode
  - real server 1887
- standby state
  - definition 1148
  - HA 1262
- start line
  - SIP 2500
- state
  - hello 1146
  - standby 1148
  - work 1148
- state synchronization
  - definition 1148
- Stateful Firewalls 891
- stateful inspection 739, 765, 1621, 2271, 2442
- stateful SIP tracking 2518
- stateless 2271
- Stateless Firewalls 891
- static
  - load balancing 1885
- static route 1099, 1100, 1101
  - adding policy 270, 285

- administrative distance 279
  - moving in list 287
  - table priority 280
  - table sequence 280
- static weight 1348
- statistics
  - viewing HA statistics 1271
- status
  - HA statistics 1272
  - quarantine files list 1941
- status description
  - quarantine files list 1941
- status-line
  - SIP 2500
- STP 1362
- STP, forwarding 1554
- \_str 1411
- strict
  - VoIP profile 2515
- strict-register 2549
- string 1411
- strong authentication 543
  - for administrators 544
  - for SSL VPN users 544
- sub second failover 1324
- sub-command 1409, 1412
- subinterface
  - VLAN NAT/Route 1530
- subordinate cluster unit
  - definition 1148
- subordinate unit 1122
  - definition 1148
  - getting information using SNMP 1266, 1267
  - getting serial numbers using SNMP 1267
  - SNMP get 1266, 1267
- sub-second failover 1324
- subsecond failover 1324
- supernetting 364
- switch
  - hardware 1488
  - link failover 1323
  - troubleshooting layer-2 switches 1361
- switching vdoms 1403
- Symbian OS version 6 701
- synchronization
  - configuration 1307
  - failure console messages 1310
  - incremental 1308
  - IPsec VPN SA 1292
  - periodic 1309
  - route 1292
  - sessions between standalone FortiGate units 1370
  - TCP sessions between standalone FortiGate units 1370
- synchronize all
  - CLI command 1308
- synchronizing routing table updates 1315
- synchronizing the configuration
  - disabling 1308

- syntax 1409
  - IPS custom signatures 2080
- sys ha showcsum
  - diagnose 1313
- SYSLOG 2318
- system
  - idle timeout 1435
  - reboot, installing 1450
  - session-helper 1602
  - time 1394
  - viewing resources 1399
- system resources
  - memory constraints 718

## T

- table 1410
  - arp 1323, 1361
  - MAC forwarding table 1323, 1361
- TACACS+ 1259
- TACACS+ server
  - authentication 1440
- TACACS+ servers 480
  - ASCII 480
  - authenticating users with 489
  - authentication protocols 480
  - changing default port 480
  - CHAP 480
  - default port 480
  - max number 480
  - MS-CHAP 480
  - PAP 480
  - port 480
- tags, replacement messages 1592
- TCP 920
  - load balancing 1910
  - port 111 1603
  - port 135 1606
  - port 1720 1603
  - port 1723 1603, 1608
  - port 21 1606
  - port 512 1603
  - port 514 1603
  - protocol optimization 2598
- TCP header flags 2271
- TCP port 587
  - WAN optimization tunnels 2600
  - web cache 238, 239, 2644, 2647
- TCP port 49 480
- TCP ports
  - for collector agent 584
- TCP session synchronization 254, 1118
  - between two standalone FortiGate units 1370
- TCP sessions
  - load-balance-all 1343
- TCP SYN packets 2277
- TCP WinNuke 1077
- TCP/IP stack 2278
- TCP/UDP/SCTP 919
- Technology Assistance Center (TAC) 2324
- Telnet 1406, 1408



- Terminal Access Controller Access-Control System (TACACS+) 480
- test vs
  - get 1891
- testing
  - VDOM 2379
  - VDOM transparent mode 1550
  - VLAN 1541
- testing VPN connections 1823
- text strings (names) 1389
- TFTP 1073
  - server 1450
  - session helper 1610
- third-party products 1360
- threshold
  - oversize 665, 708
- time 2292, 2317
  - and date 1427
  - configuring 1394
  - NTP 1428
  - protocol 1428
  - zone 1427
- time server
  - NTP 1394
- time to live for routes 1317
- timeout
  - session 507
  - user group 506
- timeout values 2200
- Timeouttimer 443
- timer
  - provisional invite 2520
- timing
  - modifying heartbeat timing 1298
- TKIP 802
- TNS 1611
- tns
  - session helper 1611
- top sessions
  - viewing 1400
- topology
  - out of path 2581
- ToS 2250
  - byte value 2237
  - mapping 2257
- total bytes
  - HA statistics 1273
- total packets
  - HA statistics 1272
- trace
  - SIP ALG NAT tracing 2545
  - SIP session helper NAT tracing 2545
- traceroute 2379
- tracert 1541, 2379
- traffic
  - policing 2235
  - priority 2244
  - reverse shaping 2249
  - reverse shaping only 2250
  - shaping 2235
- traffic offloading 1094
- Traffic policing 906
- Traffic shaper, FSAE
  - minimum bandwidth 566, 572
- Traffic Shaping 907
- traffic shaping 2602
- traffic shaping offloading 1096
- traffic statistics 1075
- train the network 1300
- Transparent
  - advanced example 2384
  - firewall address 2388, 2394
  - firewall policy 2384
  - firewall schedule 2388
  - VDOM example 2387, 2397
- Transparent Mode 906
- Transparent mode 2600
  - configuring an active-active HA cluster 1165
  - general configuration steps 1165
  - SIP 2493, 2514
  - VLAN subinterface 2384
- transparent mode 1393, 1541
  - management IP address 1393
  - reserved management interface 1255
  - security policy 1543, 1546
  - VDOM example 1545, 1548, 1549
  - VLAN example 1544
  - VLAN subinterface 1542
  - WAN optimization 2596, 2599
- transparent mode VPN configuration
  - configuration steps 1770
  - infrastructure requirements 1769
  - overview 1766
  - prerequisites to configuration 1769
- transport mode
  - setting 1800
- trap
  - SNMP 1264
- troubleshoot
  - cluster configuration
    - 1157, 1163, 1169, 1176, 1180, 1184, 1192, 1198, 1204, 1209, 1213, 1214, 1245, 1249
  - VPN 1825
- troubleshooting 1512
  - BFD 374
  - bgp 370
  - communication sessions lost after a failover 1317
  - dampening 372
  - debug packet flow 2443
  - full mesh HA 1251
  - graceful restart 373
  - holddown timer 371
  - layer-2 loops 2381
  - layer-2 switch 1361
  - packet sniffing 764, 2441
  - route flap 370
  - routing table 282
- troubleshooting VPNs 1826
- trunk

- interface 1530, 1540
- links 1524
- TTL
  - quarantine files list 1941
  - web cache default 2649
  - web cache maximum 2649
  - web cache minimum 2649
- TTL reduction 1095
- tunnel
  - bi-directional initiation 1661
  - sharing WAN optimization tunnels 2601
  - TCP port 2600
  - WAN optimization 2600
- Tunnel Endpoint Identifier (TEID) 650, 743
- tunnel mode 2191
  - configuring FortiGate server 2206
  - IP address range 2196
  - routing 2211
  - SSL VPN IP range 528
- tunnel mode IPsec 1097
- tunnel request 2611
- tunnel-non-http 2597
- two-factor authentication 459, 494
  - email 495
  - FortiToken 497
  - SMS 496
- type of service 2250
- Type of service (TOS) 286
- types of user groups 503
- types of users 488

**U**

- UA 2489
- UAC 2489
- UAS 2489
- UDP 920, 2271
  - GTP session failover 1334
  - load balancing 1910
  - port 111 1603
  - port 135 1606
  - port 1719 1607
  - port 2427 1607
  - port 2727 1607
  - session pickup 1334
- UE
  - See UA 2489
- UMTS Terrestrial Radio Access Network (UTRAN) 753
- unicast reverse path forwarding (uRPF) 271
- Unicode 1420
- Unified Threat Management 893
- Unified Threat Management, *see* UTM
- unit operation
  - viewing 1399
- Universal Mobile Telecommunications System (UMTS) 649, 747
- universal unique identifier (UUID) 1606
- unknown action 1409
- unknown HTTP sessions 2597
- unknown HTTP version
  - explicit web proxy 2671
- unset 1415
- unwanted login attempts 1436
- up time
  - HA statistics 1272
- update signatures 1429
- Updatetimers 443
- updating
  - antivirus and IPS, web-based manager 1429
- updating switch arp tables 1323
- upgrade
  - after reboot 1454
- upgrading
  - firmware using the CLI 1449
- upload status
  - quarantine files list 1941
- uploading scripts 1600
- URL block
  - web filter 2123
- URL filtering 2022
- URL formats 2110
- usage-based ECMP 300
- USB
  - backup 1453
- USB disks, formatting 1396
- USB Modem widget 1400
- user accounts 2195
- User Agent 2489
- User Agent Client 2489
- User Agent Server 2489
- user authentication
  - IPsec VPN dialup users 528
  - L2TP VPN 531
  - logon blackout period 516
  - PPTP VPN 530
  - protocols 517
  - SSL VPN 528
  - timeout 514
  - XAuth 529
- user data tunnelling (GTP-U) 745
- user element
  - See UA 2489
- user group for wireless users 808
- user groups 503, 2195
  - creating 504
  - different access permissions 2226
  - Directory Service 507
  - firewall 503
    - on authentication servers 504
    - on FortiGate unit 557, 591
  - peer, configuring 507
  - peer, creating 507
  - types of 503
  - Windows AD 577
- user IP address 585
- User Location Information (ULI) 754
- user logoff
  - ports 139 and 445 584
- usergroup timeouts 506

- username overlap 503
- Users 518
- users 488
  - banned 502
  - local, creating 489
  - local, deleting from FortiGate configuration 493
  - local, removing from FortiGate configuration 493
  - peer, configuring 493
  - peer, creating 493
  - types of 488
- users, number of concurrent 2349
- using the CLI 1405
- UTF-8 1420
- UTM
  - overview 2019
  - sessions continue after active-active HA failover 1335
  - VDOM 2172
- UTM profiles 944, 2023
- UTM proxy
  - weight 1350
- UTM Proxy Option Components 947
- UTM scanning process 905

## V

- \_v4mask 1411
- \_v6mask 1411
- value 1410
- value-added-service (VAS) ID 671
- value-added-service-provider (VASP) ID 671
- vcluster 2411
- VDOM 2245, 2304, 2305, 2324
  - configuration 2387
  - firewall policy 2367, 2368
  - independent configuration 2408
  - limited resources 1554, 2351
  - management 603
  - management configuration 2401, 2409
  - management services 2353
  - management VDOM 2333, 2337, 2355, 2356
  - maximum interface 2382
  - maximum interfaces 1523, 1554
  - maximum number 2351
  - meshed configuration 2401, 2410
  - partitioning (HA) 1149
  - simple VDOM NAT/Route example 2373
  - stand alone configuration 2401, 2407
  - status 2355
  - Transparent mode 2380
  - transparent mode 1541
  - UTM 2172
  - VDOM example 2370, 2376
  - VLAN subinterface 2362
  - VPN settings 2368
- VDOM partitioning
  - HA 1151
- VDOMs 2600
- vdoms, switching 1403
- Vendor specific attributes (VSA) 469
- Verifications of IP options 2276

- viewing
  - Alert Message Console 1399
  - carrier end point IP filter list 649, 747, 752
  - configuration revisions 1453
  - disk status 1599
  - FortiGuard support contract 1463
  - IPSec VPN auto key list 1624
  - IPSec VPN concentrator list 1635
  - IPSec VPN manual key list 1633
  - licenses 1397
  - quarantine logs 1940
  - session history, widget 1400
  - system information 1391
  - system resources 1399
  - top sessions 1400
  - unit operation 1399
- VIP address
  - L2TP clients 1565
  - PPTP clients 1558
- VIP address, FortiClient dialup clients 1711
- VIP addresses 1654
- virtual
  - domains 1497
  - LANs 1498
- virtual AP
  - creating 798
- virtual cluster 1217
  - and virtual domains 1217
  - configuring 1221, 1223, 1228
- virtual clustering 1122, 1123
  - definition 1148
  - port monitoring 1219
  - remote link failover 1219
- Virtual Desktop 2219
- virtual domain, transparent VPN IPsec configuration 1769
- virtual domains 2600
- virtual domains (VDOMs) 603
- virtual interface 2401
- virtual interfaces 2245
- virtual IP 1333, 1884
  - assigning with RADIUS 1713
  - WAN optimization 2601
- virtual IP address (VIP) 1672
- virtual MAC address 1292, 1300
  - definition 1146
  - group ID 1306
  - how its determined 1302
  - VRRP 1365
- Virtual Private Network (VPN) 1614
- Virtual Private Network, see VPN.
- virtual router MAC address
  - VRRP 1365
- Virtual Router Redundancy Protocol 1364
- Virtual Router Redundancy Protocol (VRRP) 2384
- virtual server 1333
  - arp-reply 1885
  - interface 1884
  - IP 1884
  - port 1885

- virus
  - explicit web proxy 2680, 2697
  - MMS scanning 702
  - viral marketing 731
- virus detected
  - HA statistics 1272
- virus scan 2033, 2036
- Visitor Location Register (VLR) 653
- VLAN 1094, 2245
  - adding to VDOM 2362
  - application 1523
  - jumbo traffic frames 1496
  - maximum number 1523, 1554, 2382
  - on WiFi-Ethernet bridge 838
  - security policy 1532
  - subinterface 1530, 1534, 1535, 1540
  - tagged packets 1530
  - Transparent mode 2380
  - transparent mode 1541
- VLAN ID 1527
  - range 1524
  - tag 1524
- VLAN subinterface
  - Transparent mode 2384
  - transparent mode 1542
  - VDOM example 2387, 2392
  - VDOM NAT/Route 2362
  - VDOM transparent mode example 1545
  - VLAN NAT example 1535
  - VLAN NAT/Route example 1535
- VoIP 946, 1608
  - load balancing 1343
  - profile 2515
- VoIP Profile
  - SCCP 2514
  - SIMPLE 2514
- VoIP profile
  - default 2515
  - strict 2515
- VoIP support
  - enabling on the web-based manager 2515
- VPN 1078
  - authentication 528
  - auto connect 1633, 2201
  - backup 1765
  - client-based authentication 463
  - gateway 1099
  - idle timeout 464
  - IPsec 528
  - keep alive 1633, 2201
  - L2TP 531
  - logging events 1825
  - monitoring, dialup connection 1822
  - monitoring, static or DDNS connection 1822
  - planning configurations 1621
  - policy-based vs. route-based 1621
  - PPTP 530
  - preparation steps 1623
  - SSL 528
  - testing 1823
  - troubleshooting 1826
  - VDOM 2368
- vpn
  - initiator 1677
- VPN encryption/decryption offloading 1078
- VPN policy server
  - configuring FortiGate unit as 1717
- VPN, configuring L2TP 1564
- VPNs
  - Forticlient 1632
- VRRP 1119, 1364
  - adjusting the advertisement message interval 1369
  - advertisement messages 1364
  - Configuring 1366
  - destination IP address 1369
  - example 1366, 1367
  - preempt mode 1369
  - startup time 1369
  - virtual MAC address 1365
- VSA
  - dictionary 469
  - RADIUS servers 469
- vulnerability
  - Cross-Site Scripting 1389
  - XSS 1389
- vulnerability scan
  - configuring scans 171, 2013
  - viewing results 171, 2013

## W

- WAN Optimization
  - web caching 238, 2644
- WAN optimization 1333, 1335
  - and virtual IPs 2601
  - explicit mode 2596
  - FortiGate models supported 2578
  - load balancing 1343
  - memory usage 2602, 2648
  - monitoring 2603
  - peer authentication 2610
  - peers 2610
  - storage 2716
  - transparent mode 2596
- WAN optimization peer
  - configuring 2611
  - monitoring 2615
- WAP 636
- warning message 1439
- warning to install FortiClient 2004
- WCCP 1333, 2703
  - cache engine 2703
  - client 2703
  - load balancing 1343
  - router 2703
  - server 2703
  - service group 2704
  - service ID 2704
  - service number 2704
  - topology 2580, 2588
  - well known service 2704
- WCCP service ID

- HTTP 2704
- wdigets
  - unit operation 1399
- web cache 2580, 2648
  - always revalidate 2648
  - changing the relative amount of disk space 2717
  - default TTL 2649
  - exempt 2651
  - fresh factor 2649
  - HTTP port 238, 239, 2644, 2647
  - max cache object size 2648
  - max HTTP message length 2649
  - max HTTP request length 2649
  - maximum TTL 2649
  - minimum TTL 2649
  - monitoring 2652
  - negative response duration 2648
  - proxy FQDN 2649
  - reverse proxy 2588, 2644, 2655
  - storage 2716
  - TCP port 238, 239, 2644, 2647
  - WAN Optimization 238, 2644
- Web Cache Communication Protocol
  - See WCCP 2703
- web caching 2643
  - memory usage 2602, 2648
- web content filtering 2023
- web filter 1935
  - how URL formats are detected, HTTP 2110
  - how URL formats are detected, HTTPS 2110
  - quota 2101
  - URL block 2123
  - URL filter 2124
- web filter profile 2119
- Web Filtering 945
- web filtering 2022
  - explicit web proxy 2680
  - HTTPS 2176
- web filtering service 1594
- web monitor 2185
- web portal
  - customize login 2199
  - customizing login page 2212
  - setting login page port number 2211
- web proxy 1333, 2665
  - antivirus 2680
  - authentication 2677, 2678
  - DLP 2680
  - FortiGuard web filtering 2680
  - IPv6 2665, 2673
  - security profile 2669
  - security profilesecurity profile
    - web proxy 2694
  - web filtering 2680
- web site, content category 1592
- Web UI. See web-based manager
- web-based manager 1387, 1426
  - changing the language 1403
  - connecting to the CLI 1403
  - idle timeout 1403
  - logging out 1404
  - pages 1387
  - screen resolution 1387
  - using web-based manager lists 1388
- web-based manager configuration steps
  - NAT/Route mode
    - 1154, 1158, 1166, 1170, 1177, 1180
- web-based manager, lock 1504
- web-based manager, switching vdoms 1403
- web-based user authentication 463
- webcache-storage-percentage 2717
- web-only mode 2191
  - security policy for 2204
- weight
  - real server 1886
  - static 1348
- weighted
  - load balancing 1885
- weighted round-robin
  - HA schedule 1344
- weighted-round-robin
  - configuring weights 1348
- well known service
  - WCCP 2704
- widget
  - USB modem 1400
- widgets 1397
  - alert message console 1399
  - CLI console 1399
  - licence information 1397
  - RAID monitor 1401
  - session history 1400
  - system information 1391
  - system resources 1399
  - top sessions 1400
- WiFi controller
  - discovery methods 820
- wild cards 1411
- wildcard
  - carrier end point pattern 681
- wildcard admin configuration 476
- wildcard pattern matching 1424
- wildcards 2181
- Windows 2008 467
- Windows Active Directory (AD)
  - forest 570
  - trust relation 570
- Windows networks
  - enabling NetBIOS 1553
- Windows Terminal Server
  - authentication 2678
- Windows updates
  - caching 2644, 2648
- windows version check 2213
- Windows VPN 1786
- WINS 1553
- wire speed 1075
- wireless 1495
  - client mode 879

- Wireless LAN (WLAN) 753
- wireless security considerations 775
- WLAN
  - firewall policies 809
- word boundary, Perl regular expressions 1424
- work state
  - definition 1148
  - HA 1262
- Workstation verify interval
  - collector agent configuration 579
- worm-generated messages 731

## X

- X.509 532, 544
  - managing security certificates 535
  - server certificate 533
- X.509 security certificates 2198
- XAUTH 528

- configuring authentication with 529
- XAuth (extended authentication)
  - authenticating users with 1650
  - FortiClient application as client 1718
  - FortiGate unit as server 1651
- XD4 1346
- x-mms-response-status 703
- x-mms-response-text 703
- XSS vulnerability
  - protection from 1389
- XSS vulnerability characters 525

## Z

- zero bandwidth 2244
- zone
  - using as route-based "concentrator" 1684
- zones 1499