



FortiADC - Handbook

Version 5.3.0

FORTINET DOCUMENT LIBRARY

<https://docs.fortinet.com>

FORTINET VIDEO GUIDE

<https://video.fortinet.com>

FORTINET BLOG

<https://blog.fortinet.com>

CUSTOMER SERVICE & SUPPORT

<https://support.fortinet.com>

FORTINET TRAINING & CERTIFICATION PROGRAM

<https://www.fortinet.com/support-and-training/training.html>

NSE INSTITUTE

<https://training.fortinet.com>

FORTIGUARD CENTER

<https://fortiguard.com/>

END USER LICENSE AGREEMENT

<https://www.fortinet.com/doc/legal/EULA.pdf>

FEEDBACK

Email: techdoc@fortinet.com



September 27, 2019

FortiADC 5.3.0 Handbook

01-530-600000-20190927

TABLE OF CONTENTS

Change Log	15
Introduction	16
Features	16
Basic network topology	16
Scope	18
Chapter 1: What's New	19
FortiADC 5.3.0	19
FortiADC 5.2.3	22
FortiADC 5.2.2	23
FortiADC 5.2.1	23
FortiADC 5.2.0	25
FortiADC 5.1.0	27
FortiADC 5.0.2	30
FortiADC 5.0.1	31
FortiADC 5.0.0	31
FortiADC 4.8.4	33
FortiADC 4.8.3	33
FortiADC 4.8.2	33
FortiADC 4.8.1	33
FortiADC 4.8.0	34
FortiADC 4.7.3	35
FortiADC 4.7.2	35
FortiADC 4.7.1	36
FortiADC 4.7.0	36
FortiADC 4.6.2	37
FortiADC 4.6.1	37
FortiADC 4.6.0	37
FortiADC 4.5.3	39
FortiADC 4.5.2	39
FortiADC 4.5.1	39
FortiADC 4.5.0	40
FortiADC 4.4.0	41
FortiADC 4.3.1	42
FortiADC 4.3.1	43
FortiADC 4.3.0	43
FortiADC 4.2.3	44
FortiADC 4.2.1	44
FortiADC 4.2.0	45
FortiADC 4.1	45
FortiADC 4.0 Patch 2	45
FortiADC 4.0 Patch 1	45

FortiADC 4.0	46
FortiADC 3.2.0	46
FortiADC 3.1.0	46
FortiADC 3.0.0	47
FortiADC 2.1.0	47
Chapter 2: Key Concepts and Features	48
Server load balancing	48
Feature summary	48
Authentication	49
Caching	49
Compression	50
Decompression	50
Content rewriting	50
Content routing	50
Scripting	50
SSL transactions	50
Link load balancing	51
Global load balancing	51
Security	51
High availability	52
Virtual domains	52
Chapter 3: Getting Started	53
Step 1: Install the appliance	53
Step 2: Configure the management interface	54
Step 3: Configure basic network settings	57
Step 4: Test connectivity to destination servers	60
Step 5: Complete product registration, licensing, and upgrades	60
Validating a VM license with no internet connection	62
Step 6: Configure a basic server load balancing policy	63
Step 7: Test the deployment	65
Step 8: Back up the configuration	67
Chapter 4: Server Load Balancing	68
Server load balancing basics	68
Server load balancing configuration overview	73
Configuring virtual servers	76
Two Options for virtual server configuration	76
Using content rewriting rules	82
Overview	83
Configuring content rewriting rules	83
Example: Redirecting HTTP to HTTPS	85
Example: Rewriting the HTTP response when using content routing	92
Example: Rewriting the HTTP request and response to mask application details	93
Example: Rewriting the HTTP request to harmonize port numbers	95
HSTS and HPKP support	96
HSTS	97
HPKP	97

Implementation of HSTS/HPKP	98
Configuring content routes	100
Using source pools	102
Configuring source pools	102
Example: DNAT	104
Example: full NAT	105
Example: NAT46 (Layer 4 virtual servers)	107
Example: NAT64 (Layer 4 virtual servers)	109
Example: NAT46 (Layer 7 virtual servers)	111
Example: NAT64 (Layer 7 virtual servers)	113
Using schedule pools	115
How to use the "schedule pool" feature	116
Configuring schedule pools	116
Using clone pools	116
Configuring a clone pool	117
To configure a clone pool:	118
Configuring Application profiles	119
WebSocket load-balancing	137
Configuring MySQL profiles	139
Single-master mode	139
Sharding mode	141
Creating a MySQL profile	143
Creating a MySQL configuration object	143
Specifying the MySQL user account	144
Configuring MySQL rules	145
Configuring sharding	145
Configuring client SSL profiles	147
Configuring HTTP2 profiles	151
Configuring load-balancing (LB) methods	152
Configuring persistence rules	153
Configuring error pages	158
Configuring decompression rules	159
Using decompression with script data body manipulation	161
Creating a PageSpeed configuration	163
Creating PageSpeed profiles	165
PageSpeed support and restrictions	167
Supported	167
Restrictions	167
Not Supported	167
Configuring compression rules	167
Compression and decompression	169
Using caching features	169
Static caching	169
Dynamic caching	171
Configuring caching rules	171
Using real server pools	172

Configuring real server pools	173
Example: Using port ranges and the port 0 configuration	176
Configuring real servers	178
Configuring real server SSL profiles	178
Using predefined scripts and commands	183
Create a script object	187
Import a script	187
Export a script	187
Delete a script	187
Linking multiple scripts to the same virtual server	188
Setting script priority	188
Compiling principles	190
Special notes	191
Predefined scripts and commands in v5.x.x	191
Configuring an L2 exception list	193
Creating a Web Filter Profile configuration	194
Using the Web Category tab	194
Configuring certificate caching	195
Configuring a certificate caching object	195
TCP multiplexing	195
Chapter 5: Link Load Balancing	197
Link load balancing basics	197
Using link groups	197
Using virtual tunnels	199
Link load balancing configuration overview	200
Configuring link policies	202
Configuring a link group	204
Configuring gateway links	206
Configuring persistence rules	207
Configuring proximity route settings	209
Configuring a virtual tunnel group	210
Chapter 6: Global Load Balancing	212
Global load balancing basics	212
Global load balancing configuration overview	214
Configuring servers	216
Configuring link	219
Configuring data centers	220
Configuring hosts	221
Configuring wizard	222
Configuring virtual server pools	225
Configuring location lists	226
Logical Topology	227
Configuring a Global DNS policy	228
Configuring DNS zones	229
Configuring general settings	233

Configuring the trust anchor key	234
Configuring DNS64	234
Configuring the DSSET list	235
Configuring an address group	236
Configuring remote DNS servers	237
Configuring the response rate limit	238
Chapter 7: Network Security	239
Security features basics	239
Managing IP Reputation policy settings	239
Configure IP reputation exception	241
Configure IP reputation black list	241
Using the Geo IP block list	242
Using the Geo IP whitelist	244
Special Geo codes	245
Enabling denial of service protection	245
Configuring an IPv4 firewall policy	246
Configuring an IPv6 firewall policy	247
Configuring an IPv4 connection limit policy	249
Configuring an IPv6 connection limit policy	250
Anti-virus	251
Important Notes	252
Creating an AV profile	252
Configure AV profiles from the GUI	252
Configure AV profiles from the Console	254
Setting AV quarantine policies	254
Configuring AV quarantine policies from the GUI	255
Configuring AV quarantine policies from the Console	256
Viewing the quarantine monitor	256
Setting AV service level	257
Configure AV service level from the GUI	257
Configure AV service level from the Console	257
Configuring IPS	258
Chapter 8: DoS Protection	263
Configuring DoS Protection Profile	263
Configuring HTTP access limit policy	265
Configuring HTTP connection flood policy	266
Configuring an HTTP request flood policy	266
Configuring an IP fragmentation policy	267
Configuring a TCP SYN flood protection policy	268
Configuring a TCP slow data flood protection policy	268
Chapter 9: Web Application Firewall	270
Web application firewall basics	270
Web application firewall configuration overview	271
Configuring a WAF Profile	272

Configuring a Web Attack Signature policy	274
Configuring a URL Protection policy	279
Configuring an Advanced Protection policy	280
Configuring an HTTP Protocol Constraint policy	282
Configuring CSRF protection	285
Configuring brute force attack detection	286
Configuring an SQL/XSS Injection Detection policy	288
Configuring WAF Exception objects	290
Configuring a Bot Detection policy	290
Configuring a Cookie Security policy	292
Configuring sensitive data protection	294
Configuring XML Detection	296
Configuring JSON detection	299
Importing XML schema	301
Uploading WSDL files	302
Importing JSON schema	302
Input Validation	303
Web Vulnerability Scanner	307
Chapter 10: User Authentication	314
Configuring AD FS Proxy	314
Configuring authentication policies	317
Configuring user groups	319
Using the local authentication server	320
Using an LDAP authentication server	321
LDAP bind messages	321
LDAP over SSL (LDAPS) and StartTLS	322
Configuring LDAP binding	322
Configuring a RADIUS authentication server	323
Configuring Duo authentication server support	324
Using Kerberos Authentication Relay	325
Authentication Workflow	326
FortiADC Kerberos authentication implementation	327
Configure Authentication Relay (Kerberos)	327
Two-factor authentication	329
Configuring FortiAuthenticator for two-factor authentication	329
Creating user accounts on FortiAuthenticator	329
Configuring FortiADC a user group	330
Set FortiADC as a RADIUS Service client	331
Configuring FortiADC for two-factor authentication	331
Creating a RADIUS server configuration using FortiAuthenticator	331
Adding admin user accounts with RADIUS authentication	332
Two-factor authentication in action	332
Using HTTP Basic SSO	333
Configure HTTP Basic SSO	334
SAML and SSO	334
Configure a SAML service provider	335
Import IDP Metadata	336

Chapter 11: Shared Resources	337
Configuring health checks	337
Monitoring health check status	345
Creating schedule groups	346
Creating IPv4 address objects	346
Configuring IPv4 address groups	347
Creating IPv6 address objects	348
Configuring IPv6 address groups	349
Managing ISP address books	350
Create an ISP address book object	352
Creating service objects	352
Creating service groups	354
Configuring WCCP	355
Chapter 12: Basic Networking	357
Configuring network interfaces	357
Physical interfaces	357
VLAN interface	358
Aggregate interface	359
Loopback interface	359
Softswitch	359
Configuring network interfaces	360
Configuring management interface	364
"Dedicated HA Management IP" vs. "Management Interface"	365
Configuring static routes	366
Configuring policy routes	368
Chapter 13: System Management	370
Configuring basic system settings	371
Configuring system time	372
Updating firmware	373
Upgrade considerations	373
Updating firmware using the web UI	374
Updating firmware using the CLI	375
Configuring an SMTP mail server	377
Configuring FortiGuard service settings	377
Pushing/pulling configurations	379
Configuring FortiSandbox service	381
FortiCloud Sandbox file upload limits	382
Backing up and restoring configuration	382
Run a manual backup	384
Restore a backup configuration	384
Schedule auto backups	384
Schedule auto backups onto FortiADC:	385
Schedule auto backups from the Console	385
SCP support for configuration backup	386
Rebooting, resetting, and shutting down the system	386
Create a traffic group	387

Create a traffic group via the command line interface	388
Create a traffic group from the Web GUI	388
Manage administrator users	389
Administrator user overview	389
Configure access profiles	391
Enable password policies	394
Configuring SNMP	394
Download SNMP MIBs	396
Configure SNMP threshold	396
Configure SNMP v1/v2	397
Configure SNMP v3	398
Configuring central management	399
Manage and validate certificates	402
Overview	402
Certificates and their domains	403
Prerequisite tasks	403
Manage certificates	404
Generating a certificate signing request	405
Importing local certificates	407
Creating a local certificate group	408
Importing intermediate CAs	410
Creating an intermediate CA group	411
OCSP stapling	412
Validating certificates	413
Configure a certificate verification object	413
Importing CRLs	415
Adding OCSPs	416
OCSP caching	419
Configure OCSP caching from the Console	420
Importing OCSP signing certificates	420
Importing CAs	421
Creating a CA group	422
System alerts	423
Configuring alert actions	423
Configuring alert policies	424
Creating alert configurations	425
Configuring SNMP trap servers	428
Configuring an email alert object	429
Configuring a syslog object	429
HSM Integration	430
Integrating FortiADC with SafeNet Network HSM	430
Preparing the HSM appliance	431
Generating a certificate-signing request on FortiADC	432
Downloading and uploading the certificate request (.csr) file	435
Uploading the server certificate to FortiADC	435
Chapter 14: Logging and Reporting	437
Downloading logs	437
Using the security log	438

Using the traffic log	445
Using the script log	452
Configuring local log settings	452
Configuring syslog settings	454
Configuring fast stats log settings	456
Enabling real-time statistics	456
Configuring report email	457
Configuring reports	457
Configuring Report Queries	459
Configuring fast reports	461
Display logs via CLI	463
Chapter 15: High Availability Deployments	465
HA feature overview	465
HA system requirements	470
HA configuration synchronization	470
Configuring HA settings	471
Monitoring an HA cluster	475
Updating firmware for an HA cluster	476
Deploying an active-passive cluster	478
Overview	478
Basic steps	480
Best practice tips	480
Deploying an active-active cluster	481
Configuration overview	481
Basic steps	483
Expected behavior	483
Best practice tips	492
Advantages of HA Active-Active-VRRP	493
Deploying an active-active-VRRP cluster	493
Configuration overview	494
Basic steps	495
Best practice tips	496
Chapter 16: Virtual Domains	497
Virtual domain basics	497
Enabling the virtual domain feature	497
Creating a virtual domain	498
Assigning network interfaces and admin users to VDOMs	498
Virtual domain policies	499
Disabling a virtual domain	501
Chapter 17: SSL Transactions	502
SSL offloading	502
SSL decryption by forward proxy	504
Layer-7 deployments	504
Layer 2 deployments	506
SSL profile configurations	508

Certificate guidelines	513
SSL/TLS versions and cipher suites	513
Exceptions list	517
SSL traffic mirroring	517
Chapter 18: Advanced Networking	520
NAT	520
Configure source NAT	521
Configure 1-to-1 NAT	524
QoS	526
Configuring a QoS queue	527
Configuring the QoS IPv6 filter	527
Configuring the QoS filter	528
OSPF	529
ISP routes	533
Reverse path route caching	533
BGP	536
How BGP works	537
IBGP vs. EBGP	537
Route health injection (RHI)	540
Access list vs. prefix list	540
Configuring an Access List	541
Configuring an Access IPv6 List	542
Configuring a Prefix List	542
Configuring a Prefix IPv6 List	543
Transparent mode	543
Chapter 19: Best Practices and Fine Tuning	544
Regular backups	544
Security	544
Topology	545
Administrator access	545
Performance tips	546
System performance	546
Reducing the impact of logging on performance	546
Reducing the impact of reports on system performance	546
Reducing the impact of packet capture on system performance	546
High availability	547
Chapter 20: Troubleshooting	548
Logs	548
Tools	548
execute commands	548
diagnose commands	549
System dump	550
Packet capture	551
Diff	552
Save debug file	553
Solutions by issue type	554

Login issues	554
Connectivity issues	555
Resource issues	560
Resetting the configuration	561
Restoring firmware (“clean install”)	561
Additional resources	564
Chapter 21: System Dashboard	565
Widgets	566
Dashboard management tools	566
Adding a dashboard	567
Editing a dashboard	567
Deleting a dashboard	568
Adding Features	568
Chapter 22: FortiView	569
Physical Topology	569
HA Status	570
Server Load Balance	570
Logical Topology	570
Virtual Servers	576
Data Analytics	581
Traffic Logs	585
Link Load Balance	587
Logical Topology	587
Link Group	588
Global Load Balance	589
Logical Topology	589
Host	590
Security	590
Threat Map	590
Data Analytics	591
Security Logs	593
All Segments	594
Event Logs	594
Alerts	595
All Sessions	596
Appendix A: Fortinet MIBs	597
Appendix B: Port Numbers	599
Appendix C: Scripts	601
Events and actions	601
Predefined scripts	602
Predefined commands	605
Control structures	624
Operators	625
String library	626
Special characters	627

Log and debug	627
HTTP data body commands	628
Examples	629
Select content routes based on URI string matches	629
Rewrite the HTTP request host header and path	630
Rewrite the HTTP response Location header	631
Redirect HTTP to HTTPS using Lua string substitution	631
Redirect mobile users to the mobile version of a website	631
Insert random message ID into a header	632
General HTTP redirect	632
Use request headers in other events	632
Compare IP address to address group	633
Redirect HTTP to HTTPS	633
Rewrite HTTP to HTTPS in location	633
Rewrite HTTP to HTTPS in referer	634
Rewrite HTTPS to HTTP in location	634
Rewrite HTTPS to HTTP in referer	634
Fetch data from HTTP events	634
Replace HTTP body data	635
Run multiple scripts	636
Prioritize scripts	637
Appendix D: Maximum Configuration Values	639

Change Log

Date	Change Description
09/06/2019	Updated CSRF Protection
8/22/2019	FortiADC 5.3.0 Initial release with Template 4.0.
06/26/2020	TLS 1.3 support for production environment

Introduction

Welcome, and thank you for selecting Fortinet products for your network.

The FortiADC D-series family of application delivery controllers (ADC) optimizes the availability, user experience, performance and scalability of enterprise application delivery.

An ADC is like an advanced server load balancer. An ADC routes traffic to available destination servers based on health checks and load-balancing algorithms; full-featured ADC like FortiADC also improve application performance by assuming some of the server task load. Server tasks that can be handled by the FortiADC appliance include SSL encryption/decryption, WAF protection, Gzip compression, and routing processes, such as NAT.

Features

FortiADC uses Layer 4 and Layer 7 session information to enable an ADC policy and management framework for:

- Server load balancing
- Link load balancing
- Global load balancing
- Security


The FortiADC D-series family includes physical appliances and virtual appliances.

Basic network topology

Your network routing infrastructure should ensure that all network traffic destined for the backend servers is directed to the FortiADC appliance. Usually, clients access backend servers from the Internet through a firewall such as a FortiGate, so the FortiADC appliance should be installed between your servers and the firewall.

[Basic network topology on page 16](#) shows a basic Router Mode deployment. Refer to the [Basic Deployment Topologies](#) guide for an overview of the packet flow in Router Mode, One-Arm Mode, and Direct Server Return Mode deployments.

Basic network topology



Note: The deployment topology might be different for global load balancing (GLB) or high availability (HA) clusters. Refer to those chapters for a description of features and illustrations.

Scope

This document describes how to use the web user interface to:

- Get started with your deployment.
- Configure feature options.
- Configure network and system settings.
- Monitor the system.
- Troubleshoot issues.

The following topics are covered elsewhere:

- Appliance installation—Refer to the [QuickStart Guide](#) for your appliance model.
- Virtual appliance installation—Refer to the [FortiADC-VM Install Guide](#).
- CLI commands—Refer to the [FortiADC CLI Reference](#). In parts of this manual, brief CLI command examples or CLI syntax are shown to help you understand how the web UI configuration pages are related to the CLI commands.

Chapter 1: What's New

This chapter lists features and enhancements introduced in each of the FortiADC releases.

FortiADC 5.3.0

FortiADC 5.3.0 offers the following new features:

Security

Intrusion Prevention System (IPS) protection (Powered by FortiGuard)

IPS service will allow you to protect your virtual servers from the latest network intrusions by actively detecting and blocking external threats before they can reach potentially vulnerable devices. The combination of real-time threat intelligence updates and thousands of existing intrusion prevention rules delivers the industry's best IPS protection.

Application and Networking DDoS Protection

A distributed denial-of-service (DDoS) attack is a malicious attempt to disrupt the normal traffic of a targeted server, service or network by overwhelming the target or its surrounding infrastructure with a flood of Internet traffic. FortiADC support 2 layers DDoS protection:

1. Networking DoS protection

- IP fragmentation
The attacker sends a huge volume of large or uncompleted IP fragmentation packets to the victim, to exhaust the victim's resources. The IP fragmentation protection here limits the total IP fragmentation memory size to avoid memory exhaustion.
- TCP SYN flood
By enabling SYN-Cookie to all the SYN packets that exceed the threshold, the system will drop all the fake SYN packets sent to the virtual server.
- TCP slow data flood
The attacker uses very slow traffic to consume all the target server's resources; it is difficult to distinguish it from normal traffic. This protection will detect this type of attack by dynamically probing client 0 windows; if it comes in "last" several times, the FortiADC will reset this connection on server.

2. Application DoS protection

- HTTP access limit
Limits the amount of HTTP requests-per-second from a certain IP.
- HTTP connection flood
Limits the number of TCP connections with the same session cookie.

- HTTP request flood

Limits the number of HTTP requests-per-second with the same session cookie.

Web Application Firewall

FortiADC web application firewalls provide advanced features that defend web applications from known and zero-day threats. FortiADC offers a complete security coverage for your web-based applications from the OWASP Top 10 and many other threats.

1. Signature DB enhancement

Enhances WAF engine to more efficiently scan for packets, also significantly increasing the detection rate.

2. New WAF signature wizard on GUI

Helps customer configure the WAF signature profile.

3. WAF Action enhancement

Besides deny and pass, supports 2 more actions for all WAF modules: Redirect and Block period.

4. CSRF protection

A cross-site request forgery (CSRF) is an attack that exploits the trust that a site has in a user's browser to transmit unauthorized commands.

To protect back-end servers from CSRF attacks, FortiADC has two lists:

- Web pages to protect against CSRF attacks – for insert JS
- URLs found in the requests that the pages generate – for Token/cookie validation

5. Input validation

FortiADC provides advanced validation of input fields, including parameter validation, hidden field validation and file security. This function will verify the user input from scan points like URL parameter, HTML form, hidden fields, upload file. If the format isn't correct or other attacks exist, the request will be blocked.

6. Brute force detection

FortiADC can prevent brute force login attacks. Brute force attackers attempt to penetrate systems by the sheer number of clients, attempts, or computational power, rather than by intelligent insight or advance knowledge of application logic or data.

7. Data loss protection

The data loss prevention (DLP) feature allows FortiADC to prevent information leaks, damages and loss.

It provides desensitization and warning measures for sensitive information leaks on websites (SSN numbers, and credit card information) and the leakage of sensitive keywords.

8. Cookie Security

HTTP cookie is a small piece of data sent from a website and stored in the client's computer. In some cases, it will store some sensitive data inside, e.g. password.

If the client sends out the request that Fortiadc doesn't recognize, it will take corresponding action (alert/ deny/ period-block/ remove-cookie).

9. Page anti-defacement

The anti-defacement features monitor your websites for defacement attacks. If it detects a change, it can automatically reverse the damage.

This feature monitors the modification of customer's specified page; once the modification is consider as abnormal, the specified action will be triggered, such as "restore changed page," "send email," "acknowledge changed page," or "just record log."

10. Web scraping detection

FortiADC provides an advanced access control for customers who want to have agility within web application (specific IP, files, connections).

FortiADC checks the http header content-type and the response code; if it matches the occurrence limit and is over the match percentage, it will detect it as web scraping.

11. Web vulnerability scanner enhancement

- Supports exception
Able to add URL into the exception list.
- Supports form-based login
Supports form-based login for web servers.

Firewall policy support address book

FortiADC firewall now supports address book in the policy.

Server Load Balancing

Two Factor Authentication (with FortiToken and Google Authenticator)

Two-factor authentication is a type of multi-factor authentication. It is a method of confirming users' claimed identities by using a combination of two different factors. FortiADC can use script to do 2-step verification with FortiToken and Google Authenticator.

Health Check Enhancement

Adds more detailed report for each health check failure log, so the customer can quickly grasp why the health check failed and what happened on the real server.

Supports CLI "diagnose debug slb_hc_status" to show the health check status for all the SLB pool.

Cloud and Automation

Cloud platform (AWS/Azure/OCI)

The BYOL FortiADC images are listed on the AWS/Azure/OCI cloud marketplace now, and the customer can deploy them through these cloud marketplaces.

Ansible support

Ansible is an automation platform that makes your applications and systems easy to deploy. FortiADC modules allow the customer to automatically initiate the configuration or manage the configuration on any kind of FortiADC devices,

including physical devices, VM in hypervisor or cloud.

System

Export local generated unencrypted certificate

Both encrypted and unencrypted private key are allowed to be exported; it is necessary for the customer to move FortiADC hosted HTTPS services.

Supports TLS1.3 in SSL profiles

Supports TCP/TCP-SSL syslog server

Besides UDP-based syslog server, FortiADC supports TCP/TCP-SSL based remote syslog servers in case the customer needs more confidential security for the logs.

Allows global syslog server to be shared by all vdoms

In some multiple vdom deployments, some non-root vdom administrators may need to send logs to global syslog server in case of networking issues in their vdom. This feature allows the global syslog server to be shared among all non-root vdoms.

Support logical topology for LLB and GSLB

Shows all the LLB group/member status, and GSLB host status, by a topology graph on FortiView.

SSL Updated to OpenSSL version 1.1.1

Hardware

FortiADC support 2 new hardware models:

- FortiADC 300F
- FortiADC 400F

For more info on new hardware, please review the FortiADC Datasheet.

FortiADC 5.2.3

FortiADC 5.2.3 offers the following new features:

Add a “response-half-closed-request” option to HTTP/HTTPS/TCP/SSL/RDP load-balance profile

This option will allow the FortiADC to serve the request and send back the response even if the client closes the output channel.

In some cases, the client may close the output channel even after sending out the request; but at the same time the client will be waiting for a response. If this option is disabled, the FortiADC will abort, and will not serve the request anymore once it receives notice that the client has closed the channel. This may cause clients to complain of failures.

Forward SNI to RS under ssl-forward-proxy mode

In SSL forward deployment, the second ADC (HTTP->HTTPS) may not forward any SNI to backend Real Server, causing failure for some servers. In this feature, if "SNI forward flag" in server SSL is enabled, it will forward host in HTTP header as SNI to Real Server by default. If there is no host in HTTP header, it will forward the ssl-sni settings as SNI to Real Server.

FortiADC 5.2.2

FortiADC 5.2.2 offers the following new features:

Remove Memory Restriction on Cloud platform

Memory Restriction has been removed for all BYOL VM on AWS/GCP/Azure/OCI/Aliyun cloud platforms.

PROXY protocol

Support PROXY protocol for HTTP/HTTPS virtual server, to pass original client information, such as the client IP address, to the backend proxies or servers.

See the PROXY protocol [reference](#).

FortiADC 5.2.1

FortiADC 5.2.1 offers the following new features:

Security

Fortinet Security Fabric support

The Fortinet Security Fabric delivers broad protection and visibility to every network segment, device, and appliance, whether virtual, in the cloud, or on-premises. After adding FortiADC to Security Fabric, it will show the real-time visibility of FortiADC, including Virtual Server status, and various statistics.

Web Cache Communication Protocol (WCCP) support

The Web Cache Communication Protocol (WCCP) allows the server to be enabled for transparent redirection to discover, verify, and advertise connectivity to one or more web-caches. You can configure FortiADC as a WCCP server to redirect HTTP/HTTPS VS traffic to 3rd party device for caching or more security inspection.

Global Load Balance

DNS notification and zone transfer

Allows FortiADC DNS service to send zone notification to slave servers, and also receive and process incoming zone transfer message from slave servers.

Public/private IP support for SLB server behind NAT

Customer can provide a public IP address for the GLB discovered virtual server address, which is necessary for the deployment which whose server is behind NAT.

Allow multiple PTR DNS Resource Records with the same IP address

Service Load Balance

Radius Change of Authorization (CoA) message support

The Radius Change of Authorization (CoA), defined in RFC5176, provides a mechanism to dynamically change the attributes of an AAA session after the user or device is authenticated. By this feature, FortiADC can process CoA messages from external Radius server and send the traffic to the right dynamic authorization server through persistence.

System

CRLDP authentication protocol (RFC5280) support

Certificate Revocation List Distribution Point (CRLDP) defines how to get a CRL file from a distribution point, which is LDAP URI or HTTP/HTTPS URL, to verify client certificate.

Download CRL file from LDAP server

Support multiple CRL files for a single certificate verification object

Log reporting enhancement for more virtual server statistics

Collect statistics like RPS, CPS, transaction latency, session duration, throughput per virtual server/real server, and generate reports including these metrics.

Traffic log browser GUI redesign

Usually if you enable traffic log, there will be a huge volume of traffic logs. In this situation, to browse or filter traffic log is much too slow; with this feature, we redesign the traffic log browser page to show and locate logs quickly.

FortiADC 5.2.0

FortiADC 5.2.0 offers the following new features:

Server Load Balance

L2 TCP/UDP/IP VS support content routing

Supports specific routing (schedule pool, persistence, method) by source address

L7 FTP VS with FULLNAT/DNAT/Transparent mode support

Oracle DB health check support on VM platforms

Dynamic Load method enhancement

Prior to 5.2.0, all connections are cleared if RS is detected to be exceeding the threshold; now, however, when RS exceeds the threshold, the old connection is kept while not dispatching new connections

Fully ADFS proxy replacement

The ADFS Proxy is a service that brokers a connection between external users and internal ADFS servers, also called a Web Application Proxy (WAP). More and more ADFS require the proxy to support MS-ADFSPiP (ADFS Proxy Integration Protocol) which involves client certificate authentication between proxy and ADFS, trust establishment, header injection, and more. FADC from 5.2.0 has support for MS-ADFSPiP.

SIP VS enhancement:

- support NAT of Media server address
- keep client address of UDP traffic for SIP server

Script new support function:

- Authentication event and operation
- Cookie encrypt/decrypt
- AES encrypt/decrypt
- crypto hash/sign/verify
- URL encode/decode/parse
- Base32
- File operation
- Random generation
- get_pid
- HTTP:respond

Global Load Balance

New dispatch method by server CPU/Memory usage

The "Server-Performance" method dynamically dispatches the DNS request to the server with the lowest CPU/Memory usage

Security

Web Vulnerable Scanner report enhancement

JSON schema validation support

JSON Schema provides a contract for what JSON data is required for a given application and how to interact with it. This feature supports the user uploading a JSON schema to validate JSON data, just like the XML validation that we had before.

IP Reputation black list support

Now possible to upload a list of IPs or CIDRs to the IP reputation black list, then blocking them by enabling "IP reputation" in Application Profile for VS.

Antivirus quarantine monitor page on GUI

New function to show/delete quarantined files on FortiADC by GUI (Network Security -> Quarantine Monitor)

All the certificate private key file on the ADC are encrypted now for more security

Dynamic TLS record sizing support to improve SSL latency and throughput

GEO support more accurate province

System

AWS/GCP/Azure/Aliyun BYOL VM support

Now supports uploading and deploying VM images on these public cloud platforms; you can easily extend existing FortiADC services to the cloud.

HA failover enhancement to avoid unnecessary switch after slave(former master) return back

In HA AP scenarios, the slave device will become master if the master device is down, but after the former master comes back, there will be a new switchover (the former master takes the master role, and the current master, the

former slave, switches back to slave). This switchover is unnecessary and may impact traffic, so the enhancement here is to avoid doing the switchover after the former master comes back.

Debug enhancement, support collect all debug information and download by GUI

Before, in order to submit information to Help Support, the customer needed to gather files from different places; now, this debug enhancement automatically collects all necessary debug information into one file, so it's easier to submit to Help Support.

Support to upload/download a file to/from FADC by GUI

Support FortiADCManager

FortiADCManager is a central management tool to manage all your FortiADC devices in your network, providing visibility and the ability to create/edit server load balance configurations for all FortiADC devices.

Upgrade kernel to latest version

Support “| grep <filter-string>” to filter the output on CLI

FortiADC 5.1.0

FortiADC 5.1.0 offers the following new features and enhancements:

Integration with Oracle Cloud Infrastructure (OCI)

Oracle Cloud Infrastructure Compute provides bare metal compute capacity that delivers performance, flexibility, and control without compromise. It is powered by Oracle's next generation, internet-scale infrastructure designed to help you develop and run your most demanding applications and workloads in the cloud.

This release comes with the FortiADC image (BYOL) on Oracle OCI, which provides FortiADC's complete feature set, including but are not limited to the following:

- L4/L7 SLB
- Global LB
- High Availability
- Web Application FW
- And more...

See the [deployment guide](#) for more information.

FortiADC Connector for Cisco ACI

FortiADC Connector for Cisco ACI (Application Centric Infrastructure) is the Fortinet solution to provide seamless integration between Fortinet Application Delivery Controllers (FortiADC) deployments and the Cisco APIC (Application Policy Infrastructure Controller). This integration allows customers to perform single point of FortiADC configuration and Management operation through Cisco APIC.

See the [release notes](#) for more information.

Amazon Elastic Compute Cloud

Amazon Elastic Compute Cloud (Amazon EC2) provides scalable computing capacity in the Amazon Web Services (AWS) cloud. Using Amazon EC2 eliminates your need to invest in hardware up front, so you can develop and deploy applications faster. You can use Amazon EC2 to launch as many or as few virtual servers as you need, configure security and networking, and manage storage.

See the [deployment guide](#) for more information.

Application Load Balancing

Health check script

- Supports health check script for testing special/legacy application servers.
- Supports all shell basic syntax variables - if, else, case, while, for, func, array, dictionary, awk, etc.
- Supports common applications - curl, nslookup, netcat/nc, ping, ps, ip, iplink, telnet, traceroute, wc, etc.

Oracle Health Check Support (hardware model only)

- Health Check can now validate the functionality of Oracle databases.

Clone pool

- Supports Clone Pool which can be used for copying traffic (inbound/outbound) to a dedicated IDS or a sniffer device.
- Available on both Layer-4 and Layer-7 virtual servers (TCP, UDP, HTTP/HTTPS, etc.)

UDP Stateless LB

FortiADC now provides a UDP stateless mode, allowing you to perform load balance without attempting to match the packet to a pre-existing connection in the connection table. This feature is especially useful when loadbalancing syslog servers (FortiAnalyzer).

LDAP/RADIUS connectivity check

Provides authentication validation option, to verify if the configured credentials are correct and authentication is successful.

LLB traffic log support

Global Load Balancing

Auto Sync GLB

Support for auto sync when new virtual servers are added.

New predefined objects to GLB Configuration

- New predefined DEFAULT_DNS_SERVER to GLB server
- New predefined DEFAULT_DATA_CENTER to GLB datacenter
- New predefined DEFAULT_DNS_POLICY to Global DNS Policy

GLB configuration Wizard

FortiADC now provides a wizard (three-step procedure) to create GLB configurations.

GLB Data Analytic

Networking

No-NAT configuration

Support for no-NAT option (usually when using LLB/FWLB feature)

GUI enhancements

FortiView enhancement

- Physical Topology
- GLB Data-Analytics
- New LLB Traffic Log
- HTTP Statistics Enhancements
- AV Reports and Statistics

Web UI enhancements

FortiADC introduces a new WebUI theme, enhancements to FortiView, including new logs.

New Web UI Theme

New Dashboard template

New design and improvements

- Virtual Server design
- High Availability

New VDOM page

Security

Web Vulnerability Scanner

The Web Application Vulnerability Scanner is a automated tool which performs black box test on web applications to look for security vulnerabilities, such as cross-site scripting, SQL injection, command injection, source code disclosure, and insecure server configuration.

FortiADC now supports a variety of web frameworks and mixed-technology sites, such as

- Automatic learning capabilities
- Including blind injection vectors
- Full Reporting on vulnerability risks

Antivirus enhancement

FortiADC AV now supports HTTP/HTTPS and SMTP scanning protection.

WAF HTTP/HTML Decoder

FortiADC now supports several basic decoders to parse HTTP body for Web Application Firewall. They include, but are not limited to the following:

- Chunked and Multipart Body Decoder
- Compress and decompress
- Base64 & unicode
- HTML and JavaScript parser

System

SSL Update to OpenSSL version 1.1.0

OCSP stapling tunneling to an HTTP proxy server

Support HA for BGP/OSPF route injection

Support add/delete interface inside VDOM directly

FortiADC 5.0.2

FortiADC 5.0.2 offers the following new features and enhancements:

- Support for DUO Radius proxy.
- New console commands for aggregate interface LACP negotiation
- Allows the use of user-selected listening port other than the default TCP Port 5858 for GLB server.

FortiADC 5.0.1

FortiADC 5.0.1 offers the following new features and enhancements:

- Clone Pool Traffic — Supports TCP and UDP traffic mirroring, allowing you to copy Layer-4 traffic to a dedicated IDS or a sniffer device. See [Using clone pools](#).
- SCP support for configuration backup — Allows you to back up your configuration files via the SCP protocol. See [SCP support for configuration backup](#).
- Password-protection for configuration backup — Enables you to protect your FortiADC configuration with a password. See [Backing up and restoring configuration](#).

FortiADC 5.0.0

FortiADC 5.0.0 offers the following new features and enhancements:

Security Fabric

- FortiSandbox integration—You can now use a file upload restriction policy to submit uploaded files to FortiSandbox for evaluation. If FortiSandbox identifies a file as a threat, FortiADC generates a corresponding attack log message and blocks further attempts to upload the file.
- Antivirus—FortiADC now supports the FortiSandbox's Malware Signature Database on all of its hardware platforms, except FortiADC 60F.

Management, GUI, and Logs

- Dynamic Dashboard—You can customize the Dashboard according to your preferences
- Create or edit a dashboard
- Add or remove Dashboard widgets
- FortiView enhancement—Adding new statistics for
- Server load balancing—Caching, Compression, and SSL
- Link load balancing
- Global load balancing
- Alert system enhancement—Allow to configure alert threshold based SLB (BW, Client RTT, or Connection) and Interface Avg. Bandwidth.

Server Load Balance (SLB)

- Layer-4 virtual server tunnel—In tunnel mode, FortiADC encapsulates the packet within an IP datagram and forwards it to the chosen server.
- Diameter Load balancing SSL enhancement—FortiADC supports Diameter traffic over SSL (client SSL).
- Source Pool NAT in Layer 7—Now it's possible to configure pool NAT when using Layer-7 virtual servers.

Global Load Balance (GLB)

- Global load balancing authentication—Provide TCP-MD5SIG or authentication verify between two or more FortiADC appliances working in global load balancing.

Predefined scripts

Scripts

- CLASS_SEARCH_n_MATCH
- OPTIONAL_CLIENT_AUTHENTICATION
- UTILITY_FUNCTIONS_DEMO (updated)
- COOKIE_COMMANDS
- IP_COMMANDS
- MANAGEMENT_COMMANDS
- SSL_EVENTS_n_COMMANDS
- TCP_EVENTS_n_COMMANDS

Web Application Firewall (WAF)

- SOAP validation—Enhances FortiADC's WAF B2B features with SOAP messages validation. It allows you to perform SOAP validation using a Web Services Description Language (WSDL) document.

SSL

- OCSP verification caching—Allows to speed up OCSP checking using OCSP caching. The first time a client accesses FortiADC or FortiADC accesses a real server, FortiADC will query the certificate's status using OCSP and cache the response.
- Dual certificates (RSA and ECDSA) support—Allows you to create certificate groups included in parallel RSA and ECDSA certificates for improve SSL performance
- Support SSL renegotiation—FortiADC now supports SSL renegotiation between client and server. It allows the use of the existing SSL connection when client authentication is required.

System

- Openstack integration—FortiADC provides load balancing services for OpenStack cloud applications. With Openstack integration, FortiADC is able to provide load balancing functionality and advanced application delivery services within OpenStack.
- NVGRE and VXLAN support—FortiADC allow to use overlay tunnel with virtual network NVGRE and VXLAN segments in either multicast (VXLAN) and unicast (NVGRE/VXLAN) modes.
- BGP Route Health Injection (RHI)—Allows to advertising route to virtual address based on the health status of the corresponding service

Note:

Below are the maximum number of files per minute that can be uploaded to [[[Undefined variable FortinetVariables.ProductName20]]] Cloud by FortiADCplatform:

- FortiADC 60F/VM01 = 5 files per minute
- FortiADC 100—400/VM02 = 10 files per minute
- FortiADC 700D/VM04 = 20 files per minute
- FortiADC 1000—2000/VM08 = 50 files per minute
- FortiADC 4000 = 100 files per minute

FortiADC 4.8.4

FortiADC 4.8.4 is mainly a patch release, with the following feature enhancements:

- Support wildcard domain in GLB zone configuration.
- Support custom port mapping between VM and vCenter.

FortiADC 4.8.3

FortiADC 4.8.3 is a patch release only; no new feature or enhancement has been implemented in this release.

FortiADC 4.8.2

FortiADC 4.8.2 is a patch release only; no new feature or enhancement has been implemented in this release.

FortiADC 4.8.1

Management

FortiView—provides a real-time and historical traffic data from log devices by source, domain, destination, threat map, RTT, and application health check. You can filter the data by a variety of attributes, as well as by device and time period.

- Server load balance:
 - Client and server RTT
 - Performance (throughput, CPS, and requests)
 - Health check
 - Sessions and persistence
 - Top locations, browsers, domains, and OSs
- Security (Web Application Firewall, GEO IP, IP Reputation, and DDoS):
 - Threat map
 - Top attacks, Geo IP sources, IP Reputation attacks
- System:
 - System logs
 - Traffic logs
 - System alerts

Server load-balancing (SLB)

- Diameter Load-Balancing—offers the following features:
 - Dispatch Diameter messages to multiple servers
 - Server health monitoring and failover

- Session ID persistence and source address persistence
- Schedule Pool—supports schedule pool that determines the times the system uses pool servers
- RADIUS persistence enchantment—supports AND/OR persistence relationship for multiple RADIUS attributes
- HTTP Content Rewrite enhancement:
 - Supports add/delete user-defined HTTP header
 - Supports capture groups and back reference regular expressions - Support in rewrite host, URL, referrer, location
- HTTP to HTTPS redirection in one VS:
 - Able to redirect users using only one virtual server

Global load-balancing (GLB)

- GLB protocol extends to work across all FortiADC versions.

System

- Two-factor authentication
 - Supports admin access
 - Two-factor authentication and validation using token by FortiAuthenticator
- RADIUS wildcard
 - Allows admin user authentication wildcard on remote RADIUS and LDAP servers

New hardware platform

- FortiADC 200F

FortiADC 4.8.0

Management

- New Alert System — Automatically generates email notification, SNMP traps, or Syslog entries on any critical event that occurs on FortiADC hardware or software modules
- Data Analytics — Supports security statistics (WAF, GEO-IP, IP-Reputation and DDoS) in real time
- Getting Started Wizard — Makes configuring FortiADC a breeze for first-time users
- Cisco ACI — Supports full Layer-4 service integration with Cisco Application Centric Infrastructure (ACI) via a RESTful API

Server Load Balance (SLB)

- LUA Script
 - Supports HTTP body manipulation in HTTP request and response
 - Allows multiple scripts in the same virtual server (VS)
- PageSpeed
 - Optimizes your website to ensure that your clients receive a faster browsing experience by minimizing RTT and payload size and optimizing browser rendering
 - Supports minifying CSS, JS, HTML and image optimizations
- HTTP/2.0 (Supports HTTP/2 Gateway)
 - Converts from HTTP/2 (client side) to HTTP/1 (server side)
 - HTTP multiplexing of transactions from client side to server
 - SSL security with TLS v1.2

- OSCP Stapling — Supports Online Certificate Status Protocol (OCSP) stapling, an alternative approach to OCSP in which the certificate holder has to periodically request the revocation status of certificates of servers from OCSP servers and attach the time-stamped response to the initial SSL/TLS handshake between client and server.

Web Application Firewall (WAF)

- XML & JSON Validation
 - Supports XML & JSON validation and format check
 - XML schema validation
 - Supports XML & JSON XSS, SQLi and limit check

Global Load Balance (GLB)

- GLB authentication — Supports authentication between multiple FortiADC appliances across data centers

System

- FortiADC-VM License — Allows license validation without Internet connection (via proxy)
- DHCP — Support DHCP mode on data or management interfaces

New Hardware Platform

- FortiADC 60F (Note: No HSM or PageSpeed support. Available on July 1, 2017.)

FortiADC 4.7.3

FortiADC 4.7.3 is a patch release only; no new feature or enhancement has been implemented in this release.

FortiADC 4.7.2

FortiADC 4.7.2 offers the following new features or enhancements:

HSM support

- Register HSM server in config file
- Save Client certificate and key to CMDB
- Upload HSM server certificate to FortiADC
- Add registered partition
- Generate CSR with HSM
- View certificate information on the GUI
- Feature configuration supported on both the CLI and the GUI

Support for new hardware models

- FortiADC 1000F
- FortiADC 2000F
- FortiADC 4000F

FortiADC 4.7.1

FortiADC 4.7.1 is a patch release which has fixed some known issues discovered in previous releases. No new features or enhancements have been implemented in this release.

For more information, refer to [FortiADC 4.7.1 Release Notes](#).

FortiADC 4.7.0

Management

- Network Map 2.0
 - Includes SiteMap on link load balance (LLB) and global server load balance (GSLB) modules
- Real server global object
 - Standalone real server objects
 - Allows a single real server to be shared across multiple real server pools and virtual servers
- Configuration templates for Applications
 - Supports SharePoint, Exchange, Windows Remote Desktop, IIS, and Apache

Server load balance (SLB)

- Supports Real-Time Messaging Protocol (RTMP) & Real-Time Streaming Protocol (RTSP)
 - Layer 7 load-balancing
 - Health check
- Supports MySQL
 - Layer 7 load-balancing, user authentication, and persistence
 - Health check
 - MySQL rules
 - Decompression
 - Allows decompressed traffic from servers for Layer 7 manipulation (content rewrite), caching, and security (Web Application Firewall)
- Client SSL profile
 - Provides advanced client SSL offloading parameters

User authentication

- Supports LDAP authentication for Regular/Anonymous/LDAPS method
- Supports HTTP basic SSO with HTML Form Authentication/HTML Basic Authentication

High availability (HA)

- Supports HA sync traffic over aggregate ports
- Allows configuration from every device regardless of their HA status (backup vs. master)
- Separated management interface for each node in an HA cluster
- Allows to retrieve license on HA active-passive slave

System

- Transparent mode
- Support transparent mode installation (Layer 2 forwarding)

- Health check validation
- Allow testing health check policy before bidding it to a real server pool.
- Provide a list of predefined services (TCP, UDP, HTTP, and more)
- Allows to match a admin user to a multiple VDOMs
- Adds Loopback interface in BGB/OSPF defined as router ID
- Attack logs aggregated by date and attack category
- Advanced filters in SLB logs

FortiADC 4.6.2

This is a patch release; no new features or enhancements are implemented. Refer to the [Release Notes](#) for detail.

FortiADC 4.6.1

OpenSSL Library Upgrade

The Software OpenSSL Library has been upgraded to OpenSSL-1.0.2 on FortiADC appliances shipped with the Cavium SSL card, which include the following hardware models:

- FortiADC 400D
- FortiADC 700D
- FortiADC 1500D
- FortiADC 2000D
- FortiADC 4000D

StartTLS

- Supports offloading TLS encryption from back-end SMTP servers

Script

- Supports HTTP:rand_id() function for HTTP

FortiADC 4.6.0

Monitoring and Logs

- Dashboard
- Statistics and information
- Search bar in VS and RS
- Backup server visibility
- Network map
- Three mode views
- Data analytics

DNS load-balancing, security, and caching

- Load-balance DNS traffic (queries and IP addresses) to DNS server
- Sanity check on DNS queries according to RFC 1034, 1035, and 2671
- DNS caching for answer records

Dynamic Load-balancing algorithm

- Dynamic LB based Server Performance such CPU, Memory and Disk

Client certificate forwarding

- Sends client certificates to back-end server for authentication, without affecting SSL offloading

Script validation

- Provides more information in case of syntax error
- Checks content routing for virtual servers
- Generates log message
- Import/export script files

Kerberos Authentication Relay

- Enables authentication between client and server
- Protects against eavesdropping and replay attacks
- Allows nodes communicating over a non-secure network to verify each other's identity in a secure manner

SSL/HTTP visibility (mirroring)

- FortiADC's transparent IP, TCP/S and HTTP/S mirroring capabilities decrypt secure traffic for inspection and reporting by FortiGate or other third-party solutions
- IPv4/IPv6 support

Virtual server port enchantment

- Supports non-consecutive ports in port-range
- Allows Port 0 on TCP or UDP (to catch traffic on all ports)

Security Assertion Markup Language (SAML) 2.0

- Provides Service Provider (SP) and Meta Data of Identity Provider (IdP).
- Can access all VS web resources with user log-in until session expired.

Enhanced Global Load Balancing (GLB) proximity methodology

- Static proximity (GEO, GEO-ISP) and dynamic proximity (RTT, Least Connections, Connection-Limit, Bytes-Per-Second)
- Static match first, dynamic match second

HTTP/S health check

- Adds Username-password Authentication into HTTP/S health check (basic, digest and NTLM)
- Allows to choose SSL Version/Ciphers in HTTPS Health Check

Password policy

- Allows the Admin to control password length and string

VDOM enhancement

- Supports VDOMs restrictions (performance and configuration)
- Able to limit performance (throughput, CPS, SSL, etc.) on each VDOM

SNMP MIBs

- Allows users to download SNMP MIBs from the Web GUI

FortiADC 4.5.3

OpenSSL Library Upgrade

Software OpenSSL library has been upgraded to OpenSSL-1.0.2 on FortiADC appliances shipped with the Cavium SSL card, which include the following hardware models:

- FortiADC 400D
- FortiADC 700D
- FortiADC 1500D
- FortiADC 2000D
- FortiADC 4000D

FortiADC 4.5.2

Software OpenSSL library upgrade

- Software OpenSSL library has been upgraded to openssl-1.0.1s (the latest version) on all FortiADC platforms.
- It's fully functional on FortiADC software.

Enhanced certificate validation

- Support for multiple Online Certificate Status Protocol (OCSP) configurations.
- Support for multiple Certificate Revocation List (CRL) files.

"Description" field for child records in Geo IP Whitelist

- Allows the user to add a brief notation for each child record added to a parent record.

US-Government (USG) mode

- Allows the user to change the appliance from the default regular (REG) mode to USG mode via a special license key.
- Locks the FortiADC D-Series appliance to servers located within the US only.

FortiADC 4.5.1

Acceleration

- Speeds up compression of .PNG, .JPG, and .BMP image files. See
- Caching time definition based on HTTP status code (200/301/302/304)

Server Load Balancing

- SSL Health Check Client certificate selection using SSL Certification
- Support for SIPv6 traffic includes a new health check and virtual server profile
- URL Redirection based on server HTTP status code

High Availability (HA)

- HA-VRRP mode that supports floating IP, traffic group, and fail-over

Global Load Balancing

- Supports DNS SRV record

Miscellaneous

- Full BGP routing support
- Adds a "Description" field in GeoIP White List

FortiADC 4.5.0

SSL offloading

- Support ECDSA SSL cipher suites. See [Chapter 17: SSL Transactions](#).
- SSL certificate validation for server-side SSL connections. See [Configuring real server SSL profiles](#).
- L2 exception list can specify FortiGuard web filter categories. See [Creating a Web Filter Profile configuration](#).

Server Load Balancing

- SIP—Support for SIP traffic includes a new health check, virtual server profile, and persistence method. See [Configuring health checks](#), [Configuring Application profiles](#), and [Configuring persistence rules](#).
- RDP—Support for RDP traffic includes a new virtual server profile and persistence method. See [Configuring Application profiles](#) and [Configuring persistence rules](#).
- HTTP/HTTPS profile—HTTP mode option can be set to HTTP keepalive to support Microsoft SharePoint and other apps that require the session to be kept alive. See [Configuring Application profiles](#).
- Caching—New dynamic caching rules. See [Using caching features](#).
- Real server pool—Member default cookie name is now the real server name. You can change this to whatever you want. See [Using real server pools](#).
- Scripting—Added predefined scripts that you can use as templates. See [Using predefined scripts and commands](#).

Global Load Balancing

- Persistence—Option to enable persistence for specified hosts based on source address affinity. See [Persistence](#).
- Dynamic proximity—Optional configuration for proximity based on least connections. See [Configuring virtual server pools](#).
- Support for @ in zone records. See [Configuring DNS zones](#).
- Zone records (including dynamic records) displayed on zone configuration page. See [Configuring DNS zones](#).

Security

- Bot Detection—Integrated with FortiGuard signatures to allow "good bots" and detect "bad bots." See [Configuring a WAF Profile](#).

Monitoring and Logs

- Fast reports—Real-time statistics and reports for SLB traffic. See [Configuring fast reports](#).
- Session tables and persistence tables—Dashboard tabs for SLB session tables and persistence tables. See [Chapter 21: System Dashboard](#).
- Network map search—Dashboard network map now has search. See [Chapter 21: System Dashboard](#).

System

- New health checks for SIP and custom SNMP. See [Configuring health checks](#)
- Config push/pull (not related to HA). See [Pushing/pulling configurations](#).
- HA sync can be auto/manual. See [Configuring HA settings](#).
- HA status includes details on synchronization. See [Monitoring an HA cluster](#).
- SNMP community host configuration supports subnet address and restriction of hosts to query or trap (or both). See [Configuring SNMP](#).
- Support STARTTLS in email alerts. See [Configuring an SMTP mail server](#).
- Coredump utilities. See .

Platform

- Virtual machine (VM) images for Hyper-V, KVM, Citrix Xen, and opensource Xen. See the [FortiADC-VM Install Guide](#) for details.

FortiADC 4.4.0

Server Load Balancing

- New SSL forward proxy feature can be used to decrypt SSL traffic in segments where you do not have the server certificate and private key. See [Chapter 17: SSL Transactions](#).
- New server-side SSL profiles, which have settings for the FortiADC-to-server connection. This enables you to specify different SSL version and cipher suites for the server-side connection than the ones specified for the client-side connection by the virtual server profile. See [Configuring real server SSL profiles](#).
- Support for ECDHE ciphers, null ciphers, and user-specified cipher lists. See [Chapter 17: SSL Transactions](#).
- You can now specify a list of SNAT IP address pools in the virtual server configuration. This enables you to use addresses associated with more than one outgoing interface. See [Configuring virtual servers](#).
- Added a health check for UDP, and added hostname to the general settings configuration. In HTTP/HTTPS checks, you can specify hostname instead of destination IP address. See [Configuring health checks](#).
- UDP profiles can now be used with Layer 2 virtual servers. See [Configuring Application profiles](#).
- Server name added to real server pool member configuration. The name can be useful in logs. When you upgrade, the names will be generated from the pool member IP address. You can change that string to whatever you like. See [Using real server pools](#).
- Added a comments setting to the virtual server configuration so you can note the purpose of a configuration. See [Configuring virtual servers](#).

Link Load Balancing

- You can now specify ISP addresses, address groups, and service groups in LLB policies. Using groups adds Boolean OR logic within the elements of LLB rules. See [Configuring link policies](#).

Global Load Balancing

- Added "dynamic proximity" to the server selection algorithm. Dynamic proximity is based on RTT. See .
- Added an option to send only a single record in responses instead of an ordered list of records. See [Configuring hosts](#).

- Support for health checks of third-party servers. See [Configuring servers](#).
- Support for TXT resource records. See [Configuring DNS zones](#).

Security

- You can now specify exceptions per WAF profile or per policy. Exceptions identify specific hosts or URL patterns that are not subject to processing by WAF rules. See [Configuring a WAF Profile](#).
- Additional WAF HTTP protocol constraint rules. See [Configuring a WAF Profile](#).

Monitoring and Logs

- Added a Network Map tab to the dashboard. In the Network Map, each virtual server is a tree. The status of the virtual server and real server pool members is displayed. See [Chapter 21: System Dashboard](#).
- Added on-demand and scheduled reports for many common queries. You can also configure custom queries. See [Configuring reports](#).
- Added event log categories and added a column in logs to support future integration with FortiAnalyzer. Removed the Download Logs page. Each log category page now has a **Download** button. See [Downloading logs](#).
- Enhanced SNMP MIBs and traps. See [Appendix A: Fortinet MIBs](#) for information on downloading the vendor-specific and product-specific MIB files.

System

- Shared Resources—Merged the address and service configuration for firewall and LLB. Added address groups and service groups, which can be used in LLB policy rules. See [Chapter 11: Shared Resources](#).
- Routing—Support for OSPF authentication. See [OSPF](#).
- HA—Added option to actively monitor remote beacon IP addresses to determine if the network path is available. See [Configuring HA settings](#).
- System—Updated the web UI to match CLI configuration options for global administrator and access profile. See [Manage administrator users](#).
- Web UI—Support for Simplified Chinese. See [Configuring basic system settings](#).
- Troubleshooting—New commands: `diagnose debug flow`, `diagnose debug report`, `diagnose debug timestamp`, `execute checklogdisk`, and `execute fixlogdisk`. See the [FortiADC CLI Reference](#).
- CLI—Added `execute ssh` and `execute telnet` for connections to remote hosts.

API

- REST API—Remote configuration management with a REST API. See the [FortiADC REST API Reference](#).

FortiADC 4.3.1

- Server Load Balancing Persistence—Added a Match Across Servers option to the Source Address affinity method. This option is useful when the client session for an application has connections over multiple ports (and thus multiple virtual servers). This option ensures the client continues to access the same backend server through different virtual servers for the duration of a session.
- Server Load Balancing TCP Multiplexing—Added support for HTTPS connections.
- Global Load Balancing DNS Server—The negative caching TTL in the SOA resource record is now configurable.

FortiADC 4.3.1

- Virtual domains—Increased the maximum number of VDOMs on the following platforms:
 - FortiADC 700D — 30
 - FortiADC 1500D — 45
 - FortiADC 2000D — 60
 - FortiADC 4000D — 90
- Health checks—Added an HTTP Connect health check that is useful for testing the availability of web cache proxies, such as FortiCache.
- ISP address book—Added a province location setting to the ISP address book. The province setting is used in GLB deployments in China to enable location awareness that is province-specific. For example, based on location, the DNS server can direct a user to a datacenter in Beijing or Guangdong rather than the broader location China. Only a predefined set of Chinese provinces is supported.
- Advanced routing—Exception list for reverse path route caching.

FortiADC 4.3.0

- Authentication—Framework to offload authentication from backend servers.
- Geo IP blocking—Policy that takes the action you specify when the virtual server receives requests from IP addresses in the blocked country's IP address space.
- Web application firewall—Protect against application layer attacks with policies such as signatures, HTTP protocol constraints, request URL and file extension patterns, and SQL/XSS injection detection.
- Scripts—Support for Lua scripts to perform actions that are not currently supported by the built-in feature set.
- SSL/TLS—Support for PFS ciphers.
- Health check improvements—The SLB and LLB health check configuration has been combined and moved to System > Shared Resources. You can configure destination IP addresses for health checks. This enables you to test both the destination server and any related services that must be up for the server to be deemed available. Also added support for Layer 2 and SSH health checks.
- Port range—Support for virtual IP address with a large number of virtual ports.
- NAT46/64—Support for NAT46/64 by the SLB module.
- ISP address book—Framework for an ISP address book that simplifies the ISP route and LLB proximity route configuration.
- Proximity routes—Support for using ISP address book entries in the LLB proximity route table.
- Backup pool member—Support for designating a link group or virtual tunnel group member as a “backup” that joins the pool when all of the main members are unavailable.
- Global load balancing—New framework that leverages the FortiGuard Geolocation database or the FortiADC predefined ISP address books to direct clients to the closest available FortiADC virtual servers.
- Stateful firewall—If client-to-server traffic is allowed, the session is maintained in a state table, and the response traffic is allowed.
- Virtual server traffic—Many of the firewall module features can be applied to virtual server traffic.
- ISP Routes—ISP routes are used for outbound traffic and link load balancing traffic.
- HA upgrade—Simpler one-to-many upgrade from the primary node.
- HA status—HA status tab on the system dashboard.

- HA remote login—You can use the `execute ha manage` command to connect to the command-line interface of a member node. See the CLI reference.
- SNMPv3 support
- Statistics and log database to better support dashboard and report queries.
- Improved dashboard—New time period options for the virtual server throughput graphs.
- Improved reports—New report queries for SLB HTTP virtual server reports, including client IP address, client browser type, client OS, and destination URL.
- Backup & restore—Option to back up the entire configuration, including error page files, script files, and ISP address books.

New CLI commands to facilitate troubleshooting:

- `diagnose debug config-error-log`—Use this command to see debug errors that might be generated after an upgrade or major configuration change.
- `diagnose debug crashlog`—Use this command to manage crashlog files. Typically, you use these commands to gather information for Fortinet Services & Support.
- `execute statistics-db`—Use this command to reset or restore traffic statistics.
- `config system setting`—Use this command to configure log database behavior (overwrite or stop writing) when disk utilization reaches its capacity.

For details, see the CLI reference.

FortiADC 4.2.3

- HTTPS and TCPS Profiles—Support for SHA-256 ciphers suites.

FortiADC 4.2.2

- Content rewriting—Support for PCRE capture and back reference to write the Location URL in redirect rules.
- Web UI—You can clone configuration objects to quickly create similar configuration objects. If a configuration object can be cloned, the copy icon appears in the tools column for its summary configuration page.
- Web UI—You can sort many of the configuration summary tables by column values. If a configuration summary table can be sorted, it includes sort arrows in the column headings. For example, the Server Load Balance > Virtual Server configuration summary page can be sorted by Availability, Status, Real Server pool, and so on. You can also sort the Dashboard > Virtual Server > Real Server list by column values—for example, by Availability, Status, Total Sessions, or throughput bytes.

FortiADC 4.2.1

Bug fixes only.

FortiADC 4.2.0

- New web UI
- New log subtypes
- New dashboard and report features
- Additional load balancing methods—Support for new methods based on a hash of a full URI, domain name, hostname, or destination IP address.
- Predefined health checks—Helps you get started with your deployment.
- Predefined persistence rules—Helps you get started with your deployment.
- HTTP Turbo profile—Improves the performance of HTTP applications that do not require our optional profile features.
- Layer 2 load balancing—Support for TCP profiles.
- Granular SSL configuration—Specify the SSL/TLS versions and encryption algorithms per profile.
- Connection rate limiting—Set a connection rate limit per real server or per virtual server.
- HTTP transaction rate limiting—Set a rate limit on HTTP transactions per virtual server.
- Additional link load balancing methods—Support for new methods in link groups, including spillover and hash of the source IP address.
- Global load balancing—A new implementation of our DNS-based solution that enables you to deploy redundant resources around the globe that you can leverage to keep your business online when a local area deployment experiences unexpected spikes or downtime.
- HA active-active clustering—Support for active-active clusters.
- Administrator authentication enhancements—Support for authenticating users against LDAP and RADIUS servers.
- Multinetting—You can configure a secondary IP address for a network interface when necessary to support deployments with backend servers that belong to different subnets.
- High speed logging—Supports deployments that require a high volume of logging activity.
- Packet Capture—Support for tcpdump.

FortiADC 4.1

No design changes. Bug fixes only.

FortiADC 4.0 Patch 2

No design changes. Bug fixes only.

FortiADC 4.0 Patch 1

No design changes. Bug fixes only.

FortiADC 4.0

- VDOMs—Virtual domains (VDOMs) allow you to divide a FortiADC into two or more virtual units that are configured and function independently. The administrator for each virtual domain can view and manage the configuration for his or her domain. The `admin` administrator has access to all virtual domain configurations.
- Caching – A RAM cache is a cache of HTTP objects stored in FortiADC's system RAM that are reused by subsequent HTTP transactions to reduce the amount of load on the backend servers.
- IP Reputation—You can now block source IP addresses that have a poor reputation using data from the FortiGuard IP Reputation Service.
- Layer 2 server load balancing—FortiADC can now load balance Layer 3 routers, gateways or firewalls. This feature is useful when the request's destination IP is unknown and you need to load balance connections between multiple next-hop gateways. Supports HTTP, HTTPS and TCPS client-side connection profiles only.
- Open Shortest Path First (OSPF) support—The new OSPF feature allows FortiADC to learn dynamic routes from or redistribute routes to neighboring routers.
- HTTPS profile type for virtual servers—The HTTPS profile type provides a standalone HTTPS client-side connection profile.
- Consistent Hash IP – The persistence policy type Hash IP has changed to Consistent Hash IP. Consistent hashing allows FortiADC to achieve session persistence more efficiently than traditional hashing.
- Enhanced logs

FortiADC 3.2.0

- Link routing policies—You can now specify how FortiADC routes traffic for each available ISP link, including by source or destination address and port.
- Virtual tunnels—You can now use tunneling between two FortiADC appliances to balance traffic across multiple links to each appliance. A typical scenario is a VPN between a branch office and headquarters for application-specific access.
- Persistent routing—You can now configure connections that persist regardless of the FortiADC link load balancing activity. You can configure persistence based on source IP, destination IP, and subnet.
- Proximity-based routing—Maximize WAN efficiency by using link proximity to determine latency between FortiADC and remote WAN sites so that FortiADC can choose the best route for traffic.
- Scheduled link load balancing—You can now apply a link load balancing policy during a specific time period.
- One-to-one (1-to-1) NAT—You can now fully define how each individual source and destination IP address will be translated. This feature is useful when you require a different NAT range for each ISP.
- PPPoE interface support—To support DSL connectivity, you can now configure interfaces to use PPPoE (Point-to-Point Protocol over Ethernet) to automatically retrieve its IP address configuration.

FortiADC 3.1.0

- Custom error page—You can now upload a custom error page to FortiADC that it can use to respond to clients when HTTP service is unavailable.

- Full NAT for Layer 3/4 load balancing—Layer 3/4 load balancing now supports full NAT (translation of both source and destination IP addresses). FortiADC can now round robin among a pool of source IP addresses for its connections to backend servers.
- Standby server—You can now configure FortiADC to forward traffic to a hot standby (called a Backup Server) when all other servers in the pool are unavailable.
- Log cache memory—To avoid hard disk wear and tear, FortiADC can cache logs in memory and then periodically write them to disk in bulk. Previously, FortiADC always wrote each log message to disk instantaneously.
- HA sync for health check status with IPv6—For high availability FortiADC clusters, the Layer 4 health check status of IPv6-enabled virtual servers is now synchronized.

FortiADC 3.0.0

- Link load balancing—FortiADC now supports load balancing among its links, in addition to distributing among local and globally distributed servers. Depending on if the traffic is inbound or outbound, different mechanisms are available: outbound can use weighted round robin; inbound can use DNS-based round robin or weighted round robin.
- HTTP response compression—FortiADC now can compress responses from your backend servers, allowing you to off load compression from your backend servers for performance tuning that delivers faster replies to clients.
- Quality of service (QoS)—FortiADC now can guarantee bandwidth and queue based upon source/destination address, direction, and network service.
- Source NAT (SNAT)—When applying NAT, FortiADC can now apply either static or dynamic source NAT, depending on your preference.
- Session persistence by source IP segment—FortiADC now can apply session persistence for entire segments of source IPs such as 10.0.2.0/24. Previously, session persistence applied to a single source IP.
- Health check enhancements—FortiADC now supports additional health check types for servers that respond to these protocols: email (SMTP, POP3, IMAP), TCPS, TCP `SYN` (half-open connection), SNMP, and UDP.
- HA enhancements—FortiADC HA now synchronizes Layer 3/4 and Layer 7 sessions and connections for session persistence and uninterrupted connections when the standby assumes control of traffic.

FortiADC 2.1.0

Support for FortiADC 200D and FortiADC VM—FortiADC software has been released to support these new platforms.

Chapter 2: Key Concepts and Features

This chapter includes the following topics:

- [Server load balancing on page 48](#)
- [Link load balancing on page 51](#)
- [Global load balancing on page 51](#)
- [Security on page 51](#)
- [High availability on page 52](#)
- [Virtual domains on page 52](#)

Server load balancing

Server load balancing (SLB) features are designed to give you flexible options for maximizing performance of your backend servers. The following topics give an overview of SLB features:

- [Feature summary](#)
- [Authentication](#)
- [Caching](#)
- [Compression](#)
- [Decompression](#)
- [Content rewriting](#)
- [Content routing](#)
- [Scripting](#)
- [SSL transactions](#)

Feature summary

[Server load balancing features on page 48](#) summarizes server load balancing features.

Server load balancing features

Features	Summary
Methods	• Round robin • Weighted round robin • Least connections • Fastest response • Hash of URI, domain, host, destination IP
Health check	Checks based on Layer 3, Layer 4, or Layer 7 data.

Features	Summary
Server management	• Warm up • Rate limiting • Maintenance mode with session ramp down
Persistence	Based on: • Cookies • TCP/IP header matches • A hash of TCP/IP header values • TLS/SSL session ID • RADIUS attribute • RDP Session Broker cookie • SIP caller ID
Layer 7	Profiles: HTTP, HTTPS, HTTP Turbo, RADIUS, RDP, SIP, TCPS Content routing: HTTP Host, HTTP Referer, HTTP Request URL, SNI hostname, Source IP address Content rewriting: URL redirect, 403 Forbidden, or HTTP request/response rewrite
Layer 4	Profiles: FTP, TCP, UDP Content routing: Source IP address
Layer 2	Profiles: HTTP, HTTPS, TCP, TCPS, UDP Note: Layer 2 load balancing is useful when the request's destination IP is unknown and you need to load balance connections between multiple next-hop gateways.

For detailed information, see [Chapter 4: Server Load Balancing](#).

Authentication

FortiADC SLB supports offloading authentication from backend servers. The auth policy framework supports authentication against local, LDAP, and RADIUS authentication servers, and it enables you to assign users to groups that are authorized to access protected sites.

For configuration details, see [Configuring authentication policies](#).

Caching

FortiADC SLB supports both static and dynamic caching. Caching reduces server overload, bandwidth saturation, high latency, and network performance issues.

When caching is enabled for a virtual server profile, the FortiADC appliance dynamically stores application content such as images, videos, HTML files and other file types to alleviate server resources and accelerate overall application performance.

For configuration details, see [Using caching features](#).

Compression

FortiADC SLB supports compression offloading. Compression offloading means the ADC handles compression processing instead of the backend servers, allowing them to dedicate resources to their own application processes.

When compression is enabled for a virtual server profile, the FortiADC system intelligently compresses HTTP and HTTPS traffic. Reducing server reply content size accelerates performance and improves response times. FortiADC supports both industry standard GZIP and DEFLATE algorithms.

For configuration details, see [Configuring compression rules](#).

Decompression

FortiADC SLB also supports decompression of HTTP request body before sending it to the Web Application Firewall (WAF) for scanning according to the content-encoding header. Upon receiving a compressed HTTP request body, FortiADC first uses the zlib library to extract the HTTP body to a temporary buffer and then sends the buffer to the WAF engine for scanning.

Content rewriting

FortiADC SLB supports content rewriting rules that enable you to rewrite HTTP requests and responses so that you can cloak the details of your internal network. You can also create rules to redirect requests.

For configuration details and examples, see [Using content rewriting rules](#).

Content routing

FortiADC SLB supports content routing rules that direct traffic to backend servers based on source IP address or HTTP request headers.

For configuration details, see [Configuring content routes](#).

Scripting

FortiADC SLB supports Lua scripts to perform actions that are not currently supported by the built-in feature set. Scripts enable you to use predefined script commands and variables to manipulate the HTTP request/response or select a content route. The multi-script support feature enables you to use multiple scripts by setting their sequence of execution.

For configuration details, see [Using predefined scripts and commands](#).

SSL transactions

FortiADC SLB supports SSL offloading. SSL offloading means the ADC handles SSL decryption and encryption processing instead of the backend servers, allowing the backend servers to dedicate resources to their own application processes.

SSL offloading results in improved SSL/TLS performance. On VM models, acceleration is due to offloading the cryptographic processes from the backend server. On hardware models with ASIC chips, cryptography is also hardware-accelerated: the system can encrypt and decrypt packets at better speeds than a backend server with a general-purpose CPU.

FortiADC SLB also supports SSL decryption by forward proxy in cases where you cannot copy the server certificate and private key to the FortiADC, either because it is impractical or impossible (in the case of outbound traffic to unknown Internet servers).

For detailed information, see [Chapter 17: SSL Transactions](#).

Link load balancing

Link load balancing (LLB) features are designed to manage traffic over multiple ISP or WAN links. This enables you to provision multiple links, resulting in reduced risk of outages and additional bandwidth to relieve traffic congestion.

For detailed information, see [Chapter 5: Link Load Balancing](#).

Global load balancing

Global load balancing (GLB) makes your network reliable and available by scaling applications across multiple data centers to improve application response times and be prepared for disaster recovery.

You can deploy DNS to direct traffic based on application availability and location.

For detailed information, see [Chapter 6: Global Load Balancing](#).

Security

In most deployment scenarios, we recommend you deploy FortiGate to secure your network. Fortinet includes security functionality in the FortiADC system to support those cases when deploying FortiGate is impractical. FortiADC includes the following security features:

- Firewall—Drop traffic that matches a source/destination/service tuple you specify.
- Security connection limit—Drop an abnormally high volume of traffic from a source/destination/service match.
- IP Reputation service—Drop or redirect traffic from source IPs that are on the FortiGuard IP Reputation list.
- Geo IP—Drop or redirect traffic from source IPs that correspond with countries in the FortiGuard Geo IP database.
- Web application firewall—Drop or alert when traffic matches web application firewall attack signatures and heuristics.
- Denial of service protection—Drop half-open connections to protect the system from a SYN flood attack.

For detailed information, see [Chapter 7: Network Security](#).

High availability

The FortiADC appliance supports high availability features like active-passive, active-active cluster, active-active-VRRP cluster, failure detection, and configuration synchronization. High availability deployments can support 99.999% service level agreement uptimes. For detailed information, see [Chapter 15: High Availability Deployments](#).

Virtual domains

A virtual domain (VDOM) is a complete FortiADC instance that runs on the FortiADC platform. The VDOM feature supports multitennant deployments. To do this, you create a virtual domain configuration object that contains all of the system and feature configuration options of a full FortiADC instance, and you provision an administrator account with privileges to access and manage only that VDOM. For detailed information, see [Chapter 16: Virtual Domains](#).

Chapter 3: Getting Started

This chapter provides the basic workflow for getting started with a new deployment.

Basic steps:

1. Install the appliance.
2. Configure the management interface.
3. Configure the following basic network settings:
 - Administrator password
 - System date and time
 - Network interfaces
 - DNS
4. Test connectivity.
5. Complete product registration, install your license, and update the firmware.
6. Configure a basic load balancing policy.
7. Test the deployment with load to verify expected behavior.
8. Back up this basic configuration so that you have a restore point.



Tips:

- Configuration changes are applied to the running configuration as soon as you save them.
 - Configuration objects are saved in a configuration management database. You cannot change the name of a configuration object after you have initially saved it.
 - You cannot delete a configuration object that is referenced in another configuration object (for example, you cannot delete an address if it is used in a policy).
-

Step 1: Install the appliance

This Handbook assumes you have already installed the appliance into a hardware rack or the virtual appliance into a VMware environment.

For information on hardware appliances, refer to the FortiADC hardware manuals.

For information on the virtual appliance, refer to the FortiADC-VM Install Guide.

To download these documents, go to:

<http://docs.fortinet.com/fortiadc-d-series/hardware>

Step 2: Configure the management interface

You use the management port for administrator access. It is also used for management traffic (such as SNMP or syslog). If your appliance has a dedicated management port, that is the port you configure as the management interface; otherwise, it is the convention to use port1 for the management interface.

You configure the following basic settings to get started so that you can access the web UI from a remote location (like your desk):

- **Static route**—Specify the gateway router for the management subnet so you can access the web UI from a host on your subnet.
- **IP address**—You typically assign a static IP address for the management interface. The IP address is the host portion of the web UI URL. For example, the default IP address for the management interface is 192.168.1.99 and the default URL for the web UI is <https://192.168.1.99>.
- **Access**—Services for administrative access. We recommend HTTPS, SSH, SNMP, PING.

Before you begin:

- You must know the IP address for the default gateway of the management subnet and the IP address that you plan to assign the management interface.
- You need access to the machine room in which a physical appliance has been installed. With physical appliances, you must connect a cable to the management port to get started.
- You need a laptop with an RJ-45 Ethernet network port, a crossover Ethernet cable, and a web browser (a recent version of Chrome or Firefox).
- Configure the laptop Ethernet port with the static IP address 192.168.1.2 and a netmask of 255.255.255.0. These settings enable you to access the FortiADC web UI as if from the same subnet as the FortiADC in its factory configuration state.

To connect to the web UI:

1. Use the crossover cable to connect the laptop Ethernet port to the FortiADC management port.
2. On your laptop, open the following URL in your web browser:

<https://192.168.1.99/>

The system presents a self-signed security certificate, which it presents to clients whenever they initiate an HTTPS connection to it.

3. Verify and accept the certificate, and acknowledge any warnings about self-signed certificates.

The system displays the administrator login page. See [Login page on page 54](#).


Login page



4. Enter the username **admin** and no password.

The system displays the dashboard. See [Dashboard after initial login on page 55](#).

Dashboard after initial login



To complete the procedures in this section using the CLI:

1. Use an SSH client such as PuTTY to make an SSH connection to 192.168.1.99 (port 22).
2. Acknowledge any warnings and verify and accept the FortiADC SSH key.
3. Enter the username **admin** and no password.
4. Use the following command sequence to configure the static route:

```
config router static
edit 1
set gateway <gateway_ipv4>
end
end
```

5. Use the following command sequence to configure the management interface:

```
config system interface
edit <interface_name>
set ip <ip&netmask>
set allowaccess {http https ping snmp ssh telnet}
end
end
```

The system processes the update and disconnects your SSH session because the interface has a new IP address. At this point, you should be able to connect to the CLI from a host on the management subnet you just configured. You can verify the configuration remotely.



Step 3: Configure basic network settings

The system supports network settings for various environments.

To get started, you configure the following basic settings:

- Administrator password—You must change the password for the **admin** account.
- System date and time—We recommend you use NTP to maintain the system time.
- Network interfaces—You must configure interfaces to receive and forward the network traffic to and from the destination servers.
- DNS—You must specify a primary and secondary server for system DNS lookups.

Before you begin:


- You must know the IP address for the NTP servers your network uses to maintain system time.
- You must know the IP addresses that have been provisioned for the traffic interfaces for your FortiADC deployment.
- You must know the IP address for the primary and secondary DNS servers your network uses for DNS resolution.

To change the admin password:

1. Go to System > Administrator to display the configuration page.
2. Double-click the key icon  in the row for the user **admin** to display the change password editor. See [System administrator change password editor on page 57](#).
3. Change the password and save the configuration.

For detailed information on configuring administrator accounts, refer to the online help or see [Manage administrator users](#).

System administrator change password editor




CLI commands:

```
FortiADC-VM # config system admin
FortiADC-VM (admin) # edit admin
FortiADC-VM (admin) # set password <string>
Current password for 'admin':
FortiADC-VM (admin) # end
```

To configure system time:

1. Go to System > Settings.
2. Click the **Maintenance** tab to display the configuration page. See [System time configuration page on page 58](#).
3. Enter NTP settings and save the configuration.

For detailed information, refer to the online help or see [Configuring system time](#).

System time configuration page

The screenshot shows the 'System Time' configuration page. At the top, there's a 'Manual' tab. Below it, the 'System Time' section displays the current date and time: 2018-09-18 07:13:17. The 'Daylight Saving Time' is set to OFF. The 'Time Zone' is set to GMT+9:00/Osaka,Sapporo,Tokyo,Seoul. The 'NTP' section is turned ON. The 'NTP Server' is set to pool.ntp.org. The 'Synchronizing Interval' is set to 60 minutes. At the bottom, there are 'Save' and 'Cancel' buttons.

CLI commands:

```
config system time ntp
set ntpsync enable
set ntpserver {<server_fqdn> | <server_ipv4>}
set syncinterval <minutes_int>
end
```



Or use a command syntax similar to the following to set the system time manually:


```
config system time manual
set zone <timezone_index>
set daylight-saving-time {enable | disable}
end
execute date <MM/DD/YY> <HH:MM:SS>
```

To configure network interfaces:

1. Go to Networking > Interface to display the configuration page.
2. Double-click the row for port2, for example, to display the configuration editor. See [Network interface configuration page on page 58](#).
3. Enter the IP address and other interface settings and save the configuration.

For detailed information, refer to the online help or see [Configuring network interfaces](#).

Network interface configuration page



The screenshot shows the 'Management Interface' configuration window. At the top, 'Management Status' is set to 'ON'. Below, 'Management Interface' is set to 'port1'. 'Management IP' is '10.106.154.178/23'. 'Management MAC Address' is 'c2:38:e9:82:b7:a7'. Under 'Management IP Allow Access', checkboxes for HTTPS, Ping, SSH, SNMP, HTTP, and Telnet are all checked. A 'Save' button is at the bottom right.



CLI commands:


```
config system interface
edit <interface_name>
set ip <ip&netmask>
set allowaccess {http https ping snmp ssh telnet}
end
end
```

To configure DNS:

1. Go to System > Settings to display the Basic configuration page. See [DNS configuration page on page 59](#).
2. Enter the IP address for a primary and secondary DNS server; then save the configuration.

For detailed information on configuring DNS, refer to the online help or see [Configuring basic system settings](#).

DNS configuration page



The screenshot shows the 'Basic' configuration page. Fields include: Hostname (FAD11HF3615000004), Language (English), HTTPS Port (443), HTTPS Server Cert (Factory), SSH Port (22), Primary DNS (208.91.112.53), Virtual Domain (OFF), Serial Number (FADV040000132844), Idle Timeout (480), HTTP Port (80), Default Intermediate CA Group (Click to select), Telnet Port (23), Secondary DNS (208.91.112.52), and Config Sync (OFF). A 'Save' button is at the bottom right.



CLI commands:

```
config system dns
set primary <address_ipv4>
set secondary <address_ipv4>
end
```

Step 4: Test connectivity to destination servers

Use ping and traceroute to test connectivity to destination servers.

To test connectivity from the FortiADC system to the destination server:

Run the following commands from the CLI:

```
execute ping <destination_ip4>
execute traceroute <destination_ipv4>
```

To test connectivity from the destination server to the FortiADC system:

1. Enable ping on the network interface.
2. Use the ping and traceroute utilities available on the destination server to test connectivity to the FortiADC network interface IP address.

For troubleshooting tips, see [Chapter 20: Troubleshooting](#).

Step 5: Complete product registration, licensing, and upgrades


Your new FortiADC appliance comes with a factory image of the operating system (firmware). However, if a new version has been released since factory imaging, you might want to install the newer firmware before continuing the system configuration.

Before you begin:


- Register—Registration is required to log into the Fortinet Customer Service & Support site and download firmware upgrade files. For details, go to <http://kb.fortinet.com/kb/documentLink.do?externalID=12071>.
- Check the installed firmware version—Go to the dashboard. See [License upgrade page on page 61](#).
- Check for upgrades—Major releases include new features, enhancements, and bug fixes. Patch releases can include enhancements and bug fixes.
- Download the release notes at <http://docs.fortinet.com/fortiadc-d-series/>.
- Download firmware upgrades at <https://support.fortinet.com/>.

To upgrade your license:

1. Go to the System > Settings.
2. Find your way to the **FortiGuard** tab.
3. There, under Status, click **Upgrade License**, to upload or locate the license file.

License upgrade page**To upgrade your firmware:**

1. Go to the System > Settings.
2. Find your way to the **Maintenance** tab.
3. Go down to **Firmware**.
4. Click **Upgrade Firmware**.



For detailed information, refer to the online help or see [Updating firmware](#).

Validating a VM license with no internet connection

If a FortiADC-VM is in a standalone environment with no Internet connection, it will not be able to connect to the FortiGuard Distribution Network (FDN) to validate its license. To validate the license of a standalone FortiADC-VM with no Internet connection, you must configure the FortiADC-VM to send the license request to a proxy server that is connected to the Internet. The proxy server will then send the license request to the FDN and return the license status to the FortiADC-VM.


Before you begin, you must:

- Have a proxy server connected to the Internet.
- Have Read-Write permission for System settings.

To configure a proxy server to validate a FortiADC-VM license

1. Go to System > Settings and select the **FortiGuard** tab.
2. Under the **Update Schedule** pane, find the edit function on the top right, the pencil.
3. Go into the window enable **Tunneling status**. See [FortiGuard configuration page on page 62](#).

FortiGuard configuration page



4. Complete the configuration as described in [Proxy server configuration on page 63](#).

Proxy server configuration

Settings	Guidelines
Tunneling address	Enter the IP address of the proxy server.
Tunneling port	Enter the port of the proxy server.
Tunneling username	If access control is enabled on the proxy server, enter the proxy server's username.
Tunneling password	If access control is enabled on the proxy server, enter the proxy server's password.

5. Click **Save**.

You can also configure the FortiADC-VM to communicate with the proxy server using the CLI. For more information, see the *CLI Reference*:

<http://docs.fortinet.com/fortiadc-d-series/reference>

Step 6: Configure a basic server load balancing policy

A FortiADC server load balancing policy has many custom configuration options. You can leverage the predefined health check, server profile, and load balancing method configurations to get started in two basic steps:

1. Configure the real server pool.
2. Configure the virtual server features and options.

For complete information on server load balancing features, start with [Server load balancing basics](#).

To configure the server pool:

1. Go to Server Load Balance > Real Server Pool to display the configuration page.
2. Click **Create New** to display the configuration editor. See [Real server pool basic configuration page on page 63](#).
3. Complete the basic configuration and click **Save**.
4. Go to your new server pool, and click on the edit function. It will open up the a dialogue where you can add members.
5. Under Member, click **Create New** to display the Edit Member configuration editor. See [Step 6: Configure a basic server load balancing policy on page 63](#).
6. Complete the member configuration and click **Save**.

For detailed information, refer to the online help or see [Configuring real server pools](#).

Real server pool basic configuration page

The screenshot shows the 'Real Server Pool' configuration window. It has a dark header bar with a menu icon and the title 'Real Server Pool'. The main area contains the following fields and controls:

- Name:** A text input field with a placeholder 'Required config name. No spaces.'
- Address Type:** Two buttons, 'IPv4' (highlighted in green) and 'IPv6'.
- Health Check:** A toggle switch currently set to 'OFF'.
- Real Server SSL Profile:** A dropdown menu currently showing 'NONE'.
- Member:** A section with a blue text message: 'Please save parent record first!'.

At the bottom right, there are two buttons: 'Save' (blue) and 'Cancel' (grey).

To configure the virtual server:

1. Go to Server Load Balance > Virtual Server to display the configuration page.
2. Click **Create New** to display the configuration editor. Choose between Advanced Mode and Basic Mode. See [Virtual server configuration page on page 64](#).
3. Complete the configuration and click **Save**.

For detailed information, refer to the online help or see [Configuring virtual servers](#).

Virtual server configuration page

The screenshot shows the 'Virtual Server' configuration page. It has a dark header bar with a menu icon and the title 'Virtual Server'. Below the header is a tab bar with three tabs: 'Basic' (active), 'General', and 'Monitoring'. The main area contains the following fields and controls:


- Name:** A text input field with a placeholder 'Required config name. No spaces.'
- Type:** Three buttons, 'Layer 7', 'Layer 4' (highlighted in green), and 'Layer 2'.
- Status:** Three buttons, 'Disable', 'Enable' (highlighted in green), and 'Maintain'.
- Address Type:** Two buttons, 'IPv4' (highlighted in green) and 'IPv6'.
- Traffic Group:** A dropdown menu currently showing 'default'.
- Specifics:** A section with a horizontal line separator.
 - Schedule Pool:** A toggle switch currently set to 'OFF'.
 - Content Routing:** A toggle switch currently set to 'OFF'.
 - Packet Forwarding Method:** A dropdown menu currently showing 'DNAT'.

At the bottom right, there are two buttons: 'Save' (blue) and 'Cancel' (grey).

Step 7: Test the deployment

You can test the load balancing deployment by emulating the traffic flow of your planned production deployment. [Basic network topology on page 65](#) shows a basic network topology.

Basic network topology



To test basic load balancing:

1. Send multiple client requests to the virtual server IP address.
2. Go to the dashboard to watch the dashboard session and throughput counters increment.

3. Go to Log & Report > Log Browsing > Event Log > Health Check to view health check results.
4. Go to Log & Report > Log Browsing > Traffic Log > SLB HTTP (for example) to view traffic log. It includes throughput per destination IP address.
5. Go to Log & Report > Report to view reports. It has graphs of top N policies and servers.

Step 8: Back up the configuration

Once you have tested your basic installation and verified that it functions correctly, create a backup. This “clean” backup is a reference point that has many benefits, including:

- Troubleshooting—You can use a diff tool to compare a problematic configuration with this baseline configuration.
- Restarting—You can rapidly restore your system to a simple yet working point.
- Rapid deployment—You can use the configuration file as a template for other FortiADC systems. You can edit use any text editor to edit the plain text configuration file and import it into another FortiADC system. You should change unique identifiers, such as IP address and sometimes other local network settings that differ from one deployment to another.

To backup the system configuration:

1. Go to System > Settings.
2. Click the **Backup & Restore** tab to display the backup and restore page.
3. Click **Back Up**.

For detailed information, refer to the online help or see [Backing up and restoring configuration](#).

Chapter 4: Server Load Balancing

This chapter includes the following topics:

- [Server load balancing basics on page 68](#)
- [Server load balancing configuration overview on page 73](#)
- [Configuring virtual servers on page 76](#)
- [Using content rewriting rules on page 82](#)
- [HSTS and HPKP support on page 96](#)
- [Configuring content routes on page 100](#)
- [Using source pools on page 102](#)
- [Using schedule pools on page 115](#)
- [Using clone pools on page 116](#)
- [Configuring Application profiles on page 119](#)
- [Configuring MySQL profiles on page 139](#)
- [Configuring client SSL profiles on page 147](#)
- [Configuring HTTP2 profiles on page 151](#)
- [Configuring load-balancing \(LB\) methods on page 152](#)
- [Configuring persistence rules on page 153](#)
- [Configuring error pages on page 158](#)
- [Configuring decompression rules on page 159](#)
- [Creating a PageSpeed configuration on page 163](#)
- [Creating PageSpeed profiles on page 165](#)
- [PageSpeed support and restrictions on page 167](#)
- [Configuring compression rules on page 167](#)
- [Using caching features on page 169](#)
- [Using real server pools on page 172](#)
- [Configuring real servers on page 178](#)
- [Configuring real server SSL profiles on page 178](#)
- [Using predefined scripts and commands on page 183](#)
- [Configuring an L2 exception list on page 193](#)
- [Creating a Web Filter Profile configuration on page 194](#)
- [Using the Web Category tab on page 194](#)
- [Configuring certificate caching on page 195](#)

Server load balancing basics

An application delivery controller (ADC) is like an advanced server load balancer. An ADC routes traffic to available destination servers based on health checks and load-balancing algorithms. ADCs improve application availability and performance, which directly improves user experience.

The physical distance between clients and the servers in your backend server farm has a significant impact on server response times. Besides physical distance, the most important factors contributing to server performance are:


- Number of simultaneous connections and requests that the servers can handle
- Load distribution among the servers

The purpose of an ADC is to give you multiple methods for optimizing server response times and server capacity.

After you have deployed an ADC, traffic is routed to the ADC *virtual server* instead of the destination *real servers*.

[Basic network topology on page 69](#) shows an example of a basic load balancing deployment. The FortiADC appliance is deployed in front of a server farm, and the network interfaces are connected to three subnets: a subnet for management traffic; a subnet that hosts real servers A, B, and C; and a different subnet that hosts real servers D, E, and F. The FortiADC system performs health checks on the real servers and distributes traffic to them based on system logic and user-defined settings.


Basic network topology



Optionally, you can further improve application security and performance by offloading system processes from the server and having them handled transparently by the ADC. Server tasks that can be handled by the FortiADC appliance include SSL encryption/decryption, WAF protection, Gzip compression, and routing processes, such as NAT.

[FortiADC processing on page 71](#) shows the order in which the FortiADC features process client-to-server and server-to-client traffic.

FortiADC processing



In the client-to-server direction:

- If SNI or SSL decryption is applicable, the system acts on those exchanges.
- Then, security module rules filter traffic, and traffic not dropped continues to the virtual server module.
- Virtual server security features are applied. Traffic not dropped continues for further processing.
- If a caching rule applies, the FortiADC cache serves the content and the request is not forwarded to a backend server.
- If the system selects a destination server based on a persistence rule, content route, or script, the load balancing rules are not applied.
- After selecting a server, the system performs any rewriting and re-encryption actions that are applicable, and then forwards the packets to the server.

In the server-to-client direction:

- WAF HTTP response, NAT, rewriting, persistence, and caching rules are applied.
- If applicable, the FortiADC compresses and encrypts the server response traffic.

Server load balancing configuration overview

The configuration object framework supports the granularity of FortiADC application delivery control rules. You can configure specific options and rules for one particular type of traffic, and different options and rules for another type.

[Server load balancing configuration steps on page 75](#) shows the configuration objects used in the server load balancing configuration and the order in which you create them.

Basic steps

1. Configure health check rules and real server SSL profiles.
This step is optional. In many cases, you can use predefined health check rules and predefined real server SSL profiles. If you want to use custom rules, configure them before you configure the pools of real servers.
2. Configure server pools.
This step is required. Server pools are the backend servers you want to load balance and specify the health checks used to determine server availability.
3. Configure persistence rules, optional features and policies, profile components, and load balancing methods.
You can skip this step if you want to select from predefined persistence rules, profiles, and methods.
4. Configure the virtual server.
When you configure a virtual server, you select from predefined and custom configuration objects.


Example workflow

For a members-only HTTPS web server farm, you might have a workflow similar to the following:

1. Configure security module firewall rules that allow only HTTPS traffic from untrusted subnets to the virtual server.
2. Import server SSL certificates, configure a local certificate group, and a certificate verification policy.
3. Configure HTTPS health checks to test the availability of the web servers.
4. Configure the server pools, referencing the health check configuration object.
5. Configure authentication:
 - Create a RADIUS or LDAP server configuration.
 - Create user groups.
 - Create an authentication policy.

6. Configure an HTTPS profile, referencing the certificate group and certificate verification policy and setting SSL version and cipher requirements.
7. Configure an application profile and client SSL profile if needed.
8. Configure the virtual server, using a combination of predefined and user-defined configuration objects:
 - Predefined: WAF policy, Persistence, Method
 - User-defined: Authentication Policy, Profile

Server load balancing configuration steps



Configuring virtual servers

The virtual server configuration supports three classes of application delivery control:

- Layer 7—Persistence, load balancing, and routing are based on Layer-7 objects, such as HTTP headers, cookies, and so on.
- Layer 4—Persistence, load balancing, and network address translation are based on Layer-4 objects, such as source and destination IP addresses.
- Layer 2—This feature is useful when the request's destination IP is unknown and you need to load-balance connections among multiple next-hop gateways.

Before you begin:

- You must have a deep understanding of the backend servers and your load-balancing objectives.
- You must have configured a real server pool and other configuration objects that you can incorporate into the virtual server configuration, such as persistence rules, user-defined profiles, content routes and rewriting rules, error messages, authentication policies, and source IP address pools if you are deploying NAT.
- You must have Read-Write permission for load-balance configurations.



Unlike virtual IPs on FortiGate or virtual servers on FortiWeb, virtual servers on FortiADC are activated as soon as you have configured them and set their status to **Enable**. You do not need to apply them by selecting them in a policy.

Two Options for virtual server configuration

FortiADC provides two options for configuring virtual servers—Basic Mode and Advanced Mode.

In Basic Mode, you are required to specify only the basic parameters needed to configure a virtual server. FortiADC automatically configures those advanced parameters using the default values when you click the Save button. The Basic Mode is for less experienced users who may not have the skills required to configure the advanced features on their own.

The Advanced Mode, on the other hand, is ideal for experienced or "power" users who are knowledgeable and comfortable enough to configure all the advanced features, in addition to the basic ones, on their own.

All virtual servers you have added, whether they are configured through Basic Mode or Advanced Mode, end up on the Load Balance > Virtual Server page. You can view the configuration details of a virtual server by clicking the entry.

Basic virtual server configuration

This option is used mostly for beginners who have less experience with FortiADC.

To configure a virtual server using Basic Mode:

1. Click Server Load Balance > Virtual Server.
2. Click Add > Basic Mode to open the Basic Mode configuration editor.
3. Complete the configuration as described in [Virtual server configuration Basic Mode on page 77](#).
4. Click Save.

Virtual server configuration Basic Mode

Settings	Guidelines
Name	Specify a unique name for the virtual server configuration object. Valid characters are A-Z, a-z, 0-9, _, and -. No space is allowed. This name appears in reports and in logs as the SLB "policy". Note: Once saved, the name of a virtual server configuration cannot be changed
Application	Select an application from the list menu: • Microsoft SharePoint Application • Microsoft Exchange Server Application • IIS • Apache • Windows Remote Desktop • HTTPS H2 • HTTPS H2C • HTTP(S) • TCPS • HTTP Turbo • RADIUS • DNS • SIP • TCP • UDP • FTP • IP • RTSP • RTMP • SMTP • DIAMETER
Address	Specify the IP address provisioned for the virtual server.
Port	Accept the default port number (80) or specify a port , ports, or a range of ports of your preference. Note: The virtual server will use the specified port or ports to listen for client requests. You can specify up to eight ports or port ranges separated by space. Valid values are from 0 to 65535. Port 0 applies to Layer-4 virtual servers only,
Interface	Select a network interface from the list menu, or specify a new one.
Real Server Pool	Select a real server pool (if you have one already configured) or create a new one.
SSL	Applicable to HTTP(S) applications only. Note: SSL is disabled by default, you must check the check box to enable it. Once SSL is enabled, you must select an profile from the Client SSL Profile drop-down menu below.
Client SSL Profile	Note: This setting applies to HTTPS, TCPS, HTTP2 H2, and SMTP applications only. In the case of HTTPS, it becomes available only when SSL is enabled. Select a client SSL profile from the drop-down menu.

Settings	Guidelines
Protocol	Note: This setting becomes available only when Application is set to IP. Enter up to eight numeric values or value ranges corresponding to the protocols you'd like to use, separated by space.
Domain Name	Note: This field becomes available only when Application is set to SMTP. Specify the FQDN.

Advanced virtual server configuration

This option is used mostly by advanced users of FortiADC.

To configure a virtual server using the Advanced Mode:

1. Go to Server Load Balance > Virtual Server.
2. Click **Add > Advanced Mode** to display the configuration editor.
3. Complete the configuration as described in [Virtual server configuration in Advanced Mode on page 78](#).
4. Save the configuration.

Virtual server configuration in Advanced Mode

Settings	Description
Basic	
Name	Enter a unique name for the virtual server. Valid characters are A-Z, a-z, 0-9, _, and -. No space is allowed. This name appears in reports and in logs as the SLB "policy". Note: Once you have saved the configuration, you cannot edit the virtual server name.
Status	• Enable—The virtual server can receive new sessions. • Disable—The server does not receive new sessions and closes any current sessions as soon as possible. • Maintain—The server does not receive new sessions, but maintains its current connections.
Type	• Layer 7—Persistence, load balancing, and routing are based on Layer-7 objects, such as HTTP headers, cookies, and so on. • Layer 4—Persistence, load balancing, and network address translation are based on Layer-4 objects, such as source and destination IP addresses. • Layer 2—This feature is useful when the request's destination IP is unknown and you need to load-balance connections among multiple next-hop gateways.
Address Type	• IPv4 • IPv6 Note: IPv6 is not supported for FTP, HTTP Turbo, RDP, or SIP profiles.
Traffic Group	Select the traffic group of your choice if you have one already configured, or create a new one by clicking Create New. Note: FortiADC will use the "default" if you do not choose or create a traffic group of your own.

Settings	Description
Specifics	Note: Some of the settings in this part of the GUI apply to both Layer-7 and Layer-4 virtual servers, and some apply to Layer-7 virtual servers only, but none of them applies to Layer-2 virtual servers.
Schedule Pool	OFF (disabled) by default. Click the button to enable it.
Schedule Pool List	Available only when Schedule Pool is enabled. (See above). Follow the instructions onscreen to: 1. Select the schedule pool(s). 2. Arrange them in a desired order.
Content Routing	OFF (disabled) by default. Click the button to enable it. Note: • When content routing is enabled, FortiADC will route packets to backend servers based on IP address (Layer-4 content) or HTTP header (Layer-7 content). • Content-routing rules override static or policy routes. • This option does NOT apply to SIP profiles.
Content Routing List	Available only when Content Routing is enabled. Follow the instructions onscreen to: 1. Select the content-routing rules. 2. Arrange them in a desired order. Note: You can select multiple content routing rules in virtual server configuration. Rules that you add are checked from top to bottom. The first rule to match is applied. If the traffic does not match any of the content-routing rule conditions specified in the virtual server configuration, the system will show some unexpected behaviors. Therefore, it is important that you create a “catch-all” rule that has no match conditions. In the virtual server configuration, this rule should be ordered last so it can be used to forward traffic to a default pool. See Configuring content routes.
Content Rewriting	OFF (disabled) by default. Click the button to enable it. Note: • This option applies to Layer-7 only. • This option does NOT apply to SIP profiles.
Content Rewriting List	Available only when Content Rewriting is enabled. Follow the instructions onscreen to: 1. Select the content rewriting rules. 2. Arrange them in a desired order. Note: You can select multiple content rewriting rules in the virtual server configuration. Rules that you add are consulted from top to bottom. The first rule to match is applied. If the traffic does not match any of the content rewriting rule conditions, the header is not rewritten. See Using content rewriting rules.
Transaction Rate Limit	Note: This setting applies to Layer-7 virtual servers only. It is not supported for HTTP Turbo profiles. Set a limit to the number of HTTP requests per second that the virtual server can process. Valid values are from 0 to 1,048,567. The default is 0 (disabled).

Settings	Description
	The system counts each client HTTP request against the limit. When the HTTP request rate exceeds the limit, the virtual server sends an HTTP 503 error response to the client.
Packet Forwarding Method	Note: This setting applies to Layer-4 virtual servers only. Select one of the following packet forwarding methods: • Direct Routing—Forwards the source and destination IP addresses with no changes. Note: For FTP profiles, when Direct Routing is selected, you must also configure a persistence method. • DNAT—Replaces the destination IP address with the IP address of the backend server selected by the load balancer. The destination IP address of the initial request is the IP address of the virtual server. Be sure to configure FortiADC as the default gateway on the backend server so that the reply goes through FortiADC and can also be translated. • Full NAT—Replaces both the destination and source IP addresses. IPv4 to IPv4 or IPv6 to IPv6 translation. • Tunneling—(For Layer-4 IPv4 virtual servers) Allows FortiADC to send client requests to real servers through Layer-4 IP tunnels. See Layer-4 Virtual server IP tunneling on page 1. • NAT46—(If Address Type is IPv4) Replaces both the destination and source IP addresses, translating IPv4 addresses to IPv6 addresses. • NAT64—(If Address Type is IPv6) Replaces both the destination and source IP addresses, translating IPv6 addresses to IPv4 addresses. For Full NAT, NAT46, and NAT64, the source IP address is replaced by an IP address from the pool you specify. The destination IP address is replaced with the IP address of the backend server selected by the load balancer
NAT Source Pool List	If you are configuring a Layer 4 virtual server and enable Full NAT or NAT46, select one or more source pool configuration objects. See Using source pools .
General	
Configuration	
Address	Enter the IP address provisioned of the virtual server. Note: You do not specify an IP address for a Layer 2 virtual server. A Layer 2 virtual server is not aware of IP addresses. Instead of routing data for a specific destination, this type of server simply forwards data from the specified network interface and port.
Port	Accept the default port or specify a port, ports, or port ranges of your preference. Note: The virtual server will use the specified port or ports to listen for client requests. You can specify up to eight ports or port ranges separated by space. Valid values are from 0 to 65535. Port 0 applies to Layer-4 virtual servers only, The port range option is useful in deployments where it is desirable to have a virtual IP address with a large number of virtual ports, such as data centers or web hosting companies that use port number to identify their specific customers. Statistics and configurations are applied to the virtual port range as a whole and not to the individual ports within the specified port range.

Settings	Description
	Note: If a Layer 2 virtual server is assigned a network interface that uses port 80 or 443, ensure that the HTTPS and HTTP administrative access options are not enabled for the interface. Setting a port range is not supported for FTP, HTTP Turbo, RADIUS, or Layer 2 TCP profiles.
Connection Limit	Set a limit to the number of concurrent connections. The default is 0 (disabled). Valid values are from 1 to 1,048,576. You can apply a connection limit per real server and per virtual server. Both limits are enforced. Attempted connections that are dropped by security rules are not counted. Note: This feature is NOT supported for FTP or SIP profiles.
Connection Rate Limit	With Layer 4 profiles, and with the Layer-2 TCP profile, you can limit the number of new connections per second. The default is 0 (disabled). Valid values are from 1 to 86,400. You can apply a connection rate limit per real server and per virtual server. Both limits are enforced. Attempted connections that are dropped by security rules are not counted. Note: Not supported for FTP profiles.
Interface	Network interface that receives client traffic for this virtual server.
Public IP type	IPv4 or IPv6 Set the Public IP type for the virtual server.
Public IP address	Virtual server public IP address.
Resources	
Profile	Select a predefined or user-defined profile configuration object. See Configuring Application profiles .
Persistence	Select a predefined or user-defined persistence configuration object. See Configuring persistence rules .
Method	Select a predefined or user-defined method configuration object. See .
Real Server Pool	Select a real server pool configuration object. See Configuring real server pools .
Auth Policy	Select an authentication policy configuration object. HTTP/HTTPS only. See Configuring authentication policies .
Scripting	Check the box to enable scripting. Add scripts to the Scripting List using the script menu. HTTP/HTTPS only. Note: FortiADC allows you to combine multiple individual scripts into one combined script so that you can execute them all at once. In that situation, you can set the order in which the scripts are executed by assigning the scripts with different priorities. For more information, see Support for multiple scripts.
L2 Exception List	Select an exception configuration object. Layer 2 HTTPS/TCPs only. See Configuring an L2 exception list .
HTTP Redirect to HTTPS	This option becomes available when an HTTPS server load-balancing profile is selected. It's disabled by default. Click the button to enable.

Settings	Description
	Note: If enabled, it opens HTTP service on an HTTPS virtual server which redirects traffic to an HTTP virtual server.
Redirect Service Port	This option becomes available when HTTP Redirect to HTTPS is enabled for an HTTPS type of server load-balancing profile, as described above. You can either accept the default port (80), or specify up to eight ports or ranges of ports of your preference.
Error Page	
Error Page	Select an error page configuration object. See Configuring error pages . Note: Not supported for SIP profiles.
Error Message	If you do not use an error page, you can enter an error message to be returned to clients in the event no server is available. 1023 bytes maximum. Note: Not supported for SIP profiles.
Security	
WAF Profile	AV profile can support HTTP/HTTPS/SMTP Select a WAF profile configuration object or create a new one. See Configuring a WAF Profile .
AV Profile	Select an existing AV profile from the drop-down menu or create a new one.
SSL Traffic Mirror	
SSL Traffic Mirror	This field applies to HTTPS and TCPS only. Select the check box to enable it. Then select the ports from the list of Available Items.
Application Optimization	
	Select a page speed optimization profile. See [PAGE SPEED TOPIC]. (Add in when Jack completes topic)
Traffic Log	
Log	Enable to record traffic logs for this virtual server. Note: Local logging is constrained by available disk space. We recommend that if you enable traffic logs, you monitor your disk space closely. We also recommend that you use local logging during evaluation and verification of your initial deployment, and then configure remote logging to send logs to a log management repository.
Comments	A string to describe the purpose of the configuration, to help you and other administrators more easily identify its use.

Using content rewriting rules

This section includes the following topics:

- [Overview](#)
- [Configuring content rewriting rules](#)
- [Example: Redirecting HTTP to HTTPS](#)
- [Example: Rewriting the HTTP response when using content routing](#)

- [Example: Rewriting the HTTP request and response to mask application details](#)
- [Example: Rewriting the HTTP request to harmonize port numbers](#)

Overview

You might rewrite the HTTP request/response and HTTP headers for various reasons, including the following:

- Redirect HTTP to HTTPS
- External-to-internal URL translation
- Other security reasons

[HTTP header rewriting on page 83](#) summarizes the HTTP header fields that can be rewritten.

HTTP header rewriting

Direction	HTTP Header
HTTP Request	• Host • Referer
HTTP Redirect	Location
HTTP Response	Location

The first line of an HTTP request includes the HTTP method, relative URL, and HTTP version. The next lines are headers that communicate additional information. The following example shows the HTTP request for the URL `http://www.example.com/index.html`:

```
GET /index.html HTTP/1.1
Host: www.example.com
Referer: http://www.google.com
```

The following is an example of an HTTP redirect including the HTTP Location header:

```
HTTP/1.1 302 Found
Location: http://www.iana.org/domains/example/
```

You can use literal strings or regular expressions to match traffic to rules. To match a request URL such as `http://www.example.com/index`, you create two match conditions: one for the Host header `www.example.com` and another for the relative URL that is in the GET line: `/index.html`.

For HTTP redirect rules, you can specify the rewritten location as a literal string or as a regular expression. For all other types or rules, you must specify the complete URL as a literal string.

Configuring content rewriting rules

Before you begin:

- You must have a good understanding of [HTTP header fields](#).
- You must have a good understanding of Perl-compatible regular expressions ([PCRE](#)) if you want to use them in rule matching or rewriting.
- You must have Read-Write permission for Load Balance settings.

After you have configured a content rewriting rule, you can select it in the virtual server configuration.

Note: You can select multiple content rewriting rules in the virtual server configuration. Rules you add to that configuration are consulted from top to bottom. The first to match is applied. If the traffic does not match any of the content rewriting rule conditions, the header is not rewritten.

To configure a content rewriting rule:

1. Go to Server Load Balance > Virtual Server.
2. Click the **Content Rewriting** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Content rewriting rule guidelines on page 84](#).
5. Save the configuration.

Content rewriting rule guidelines

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the virtual server configuration. Note: After you initially save the configuration, you cannot edit the name.
Comments	A string to describe the purpose of the configuration, to help you and other administrators more easily identify its use.
Action Type	Select whether to rewrite the HTTP request or HTTP response.
HTTP Request Rewrite Actions	
Rewrite HTTP Header	Host—Rewrites the Host header by replacing the hostname with the string you specify. For Host rules, specify a replacement domain and/or port. URL—Rewrites the request URL and Host header using the string you specify. For URL rules, specify a URL in one of the following formats: • Absolute URL — <code>https://example.com/content/index.html</code> • Relative URL — <code>content/index.html</code> If you specify a relative URL, the host header is not rewritten. Referer—Rewrites the Referer header with the URL you specify. For Referer rules, you must specify an absolute URL. Note: The rewrite string is a literal string. Regular expression syntax is not supported.
Redirect	Sends a redirect with the URL you specify in the HTTP Location header field. For Redirect rules, you must specify an absolute URL. For example: <code>https://example.com/content/index.html</code> Note: The rewrite string can be a literal string or a regular expression.
Send 403 Forbidden	Sends a 403 Forbidden response instead of forwarding the request.
Add HTTP Header	Adds user-defined HTTP header in content-rewriting rules in HTTP request. Header Name—Specify the HTTP header name Header Value—Specify the HTTP header value Note: • The HTTP header name and value must conform to RFC 2616.

Settings	Guidelines
	- The HTTP header and value must conform to PCRE regular expression. - This feature works with HTTP and HTTPS server load-balance profiles only.
Delete HTTP Header	Deletes user-defined HTTP header in content-rewriting rules in HTTP request. Header Name—See above. Header Value—See above Note: See above.
HTTP Response Rewrite Actions	
Rewrite HTTP Location	Rewrites the Location header field in the server response. For Location rules, you must specify an absolute URL. For example: <code>https://example.com/content/index.html</code> Note: The rewrite string is a literal string. Regular expression syntax is not supported.
Add HTTP Header	Adds user-defined HTTP header in content-rewriting rules in HTTP response. Note: Refer to HTTP Request Rewrite Actions > Add HTTP Header above.
Delete HTTP Header	Deletes user-defined HTTP header in content-rewriting rules in HTTP response. Note: Refer to HTTP Request Rewrite Actions > Delete HTTP Header above.
Match Condition	
Object	Select content matching conditions based on the following parameters: - HTTP Host Header - HTTP Location Header - HTTP Referer Header - HTTP Request URL - Source IP Address Note: When you add multiple conditions, FortiADC joins them with an AND operator. For example, if you specify both a HTTP Host Header and HTTP Request URL to match, the rule is a match only for traffic that meets both conditions.
Type	- String - Regular Expression
Content	Specify the string or PCRE syntax to match the header or IP address.
Reverse	Rule matches if traffic does not match the expression.

Example: Redirecting HTTP to HTTPS

You can use the content rewriting feature to send redirects. One common case to use redirects is when the requested resource requires a secure connection, but you accidentally type an HTTP URL instead of an HTTPS URL in the web browser.

For HTTP redirect rules, you can specify the rewritten location as a literal string or regular expression.

[Redirecting HTTP to HTTPS \(literal string\) on page 86](#) shows a redirect rule that matches a literal string and rewrites a literal string. In the match condition table, the rule is set to match traffic that has the Host header domain

example.com and the relative URL /resource/index.html in the HTTP request URL. The redirect action sends a secure URL in the Location header: https://example.com/resource/index.html.

Redirecting HTTP to HTTPS (literal string)

Content Rewriting

Name

1

Action Type

Request Response

Action

Rewrite HTTP Header

Specifics

Rewrite Host

OFF

Rewrite URL

OFF

Rewrite Referer

OFF

Comments

comments

Match Condition (Empty Match Condition will match anything)

Add Filter

Create New

ID	Object	Type	Content	Reverse	
1	HTTP Host Header	String	match	Disable	  
4	HTTP Host Header	String	match	Disable	  

Showing 1 to 2 of 2 entries

Show 10 entries

Previous 1 Next

Save

Cancel

Regular expressions are a powerful way of denoting all possible forms of a string. They are very useful when trying to match text that comes in many variations but follows a definite pattern, such as dynamic URLs or web page content.

[Redirecting HTTP to HTTPS \(regular expression\) on page 87](#) shows a redirect rule that uses PCRE capture and back reference syntax to create a more general rule than the previous example. This rule sends a redirect for all connections to the same URL but over HTTP. In the match condition table, the first regular expression is (.*). This expression

matches any HTTP Host header and stores it as capture 0. The second regular expression is `^/(.*)$`. This expression matches the path in the Request URL (the content after the /) and stores it as capture 1. The regular expression for the redirect action uses the back reference syntax `https://$0/$1`.

Redirecting HTTP to HTTPS (regular expression)

Content Rewriting

Name

1

Action Type

Request Response

Action

Rewrite HTTP Header

Specifics

Rewrite Host

OFF

Rewrite URL

OFF

Rewrite Referer

OFF

Comments

comments

Match Condition (Empty Match Condition will match anything)

Add Filter

Create New

ID	Object	Type	Content	Reverse	
3	HTTP Host Header	Regular Expression	match	Disable	  
4	HTTP Host Header	Regular Expression	match	Disable	  

Showing 1 to 2 of 2 entries

Show 10 entries

Previous 1 Next

Save

Cancel

[Common PCRE syntax elements on page 88](#) describes commonly used PCRE syntax elements.

[PCRE examples submitted to the FortiGate Cookbook on page 91](#) gives examples of useful and relevant expressions that were originally submitted to the FortiGate Cookbook. For a deeper dive, consult a PCRE reference.



Regular expressions can involve very computationally intensive evaluations. For best performance, you should only use regular expressions where necessary, and build them with care.

Common PCRE syntax elements

Pattern	Usage	Example
()	Creates a capture group or sub-pattern for back-reference or to denote order of operations.	Text: /url/app/app/mapp Regular expression: (/app)* Matches: /app/app Text: /url?paramA=valueA&paramB=valueB Regular expression: (param)A=(value)A&\0B\1B Matches: paramA=valueA&paramB=valueB
\$0, \$1, \$2, ...	Only \$0, \$1,..., \$9 are supported. A back-reference is a regular expression token such as \$0 or \$1 that refers to whatever part of the text was matched by the capture group in that position within the regular expression. Back-references are used whenever you want the output/interpretation to resemble the original match: they insert a substring of the original matching text. Like other regular expression features, back-references help to ensure that you do not have to maintain a large, cumbersome list of all possible URLs. To invoke a substring, use $\\$n$ ($0 \leq n \leq 9$), where n is the order of appearance of capture group in the regular expression, from left to right, from outside to inside, then from top to bottom.	Let's say the regular expressions in a condition table have the following capture groups: (a)(b)(c(d))(e) This syntax results in back-reference variables with the following values: \$0 — a \$1 — b \$2 — cd \$3 — d \$4 — e
\	Escape character. Except, if it is followed by an alphanumeric character, the alphanumeric character is <i>not</i> matched literally as usual. Instead, it is interpreted as a regular expression token. For example, \w matches a word, as defined by the locale. Except, if it is followed by regular expression special character: *. ^\$?+\\(){}[]\	Text: /url?parameter=value Regular expression: \?param Matches: ?param

Pattern	Usage	Example
	When this is the case, the \ escapes interpretation as a regular expression token, and instead treats the character as a normal letter. For example, \\ matches the \ character.	
.	Matches any single character <i>except</i> \r or \n. Note: If the character is written by combining two Unicode code points, such as à where the core letter is encoded separately from the accent mark, this will <i>not</i> match the entire character: it will only match one of the code points.	Text: My cat catches things. Regular expression: c.t Matches: cat cat
+	Repeatedly matches the previous character or capture group, 1 or more times, as many times as possible (also called “greedy” matching) <i>unless</i> followed by a question mark (?), which makes it optional. Does not match if there is not at least 1 instance.	Text: www.example.com Regular expression: w+ Matches: www Would also match “w”, “ww”, “www”, or any number of uninterrupted repetitions of the character “w”.
*	Repeatedly matches the previous character or capture group, 0 or more times. Depending on its combination with other special characters, this token could be either: * — Match as <i>many</i> times as possible (also called “greedy” matching). *? — Match as <i>few</i> times as possible (also called “lazy” matching).	Text: www.example.com Regular expression: .* Matches: www.example.com All of any text, except line endings (\r and \n). Text: www.example.com Regular expression: (w)*? Matches: www Would also match common typos where the “w” was repeated too few or too many times, such as “ww” in w.example.com or “www” in www.example.com. It would still match, however, if no amount of “w” existed.
?	Makes the preceding character or capture group optional (also called “lazy” matching). This character has a different significance when followed by =.	Text: www.example.com Regular expression: (www\.)?example.com Matches: www.example.com Would also match example.com.
?=	Looks ahead to see if the next character or capture group matches and evaluate the match based upon them, but does <i>not</i> include those next characters in the returned match string (if any).	Text: /url?parameter=value&pack Regular expression: p(?=arameter) Matches: p, but only in “parameter, <i>not</i> in “pack”, which does not end with “arameter”.

Pattern	Usage	Example
	This can be useful for back-references where you do not want to include permutations of the final few characters, such as matching “cat” when it is part of “cats” but <i>not</i> when it is part of “catch”.	
^	Matches either: - the <i>position</i> of the beginning of a line (or, in multiline mode, the first line), <i>not</i> the first character itself - the inverse of a character, but only if ^ is the first character in a character class, such as [^A] This is useful if you want to match a word, but only when it occurs at the start of the line, <i>or</i> when you want to match anything that is <i>not</i> a specific character.	Text: /url?parameter=value Regular expression: ^/url Matches: /url, but <i>only</i> if it is at the beginning of the path string. It will <i>not</i> match “/url” in subdirectories. Text: /url?parameter=value Regular expression: [^u] Matches: /rl?parameter=value
\$	Matches the <i>position</i> of the end of a line (or, in multiline mode, the entire string), <i>not</i> the last character itself.	
[]	Defines a set of characters or capture groups that are acceptable matches. To define a set via a whole range instead of listing every possible match, separate the first and last character in the range with a hyphen. Note: Character ranges are matched according to their numerical code point in the encoding. For example, [0-2] matches any UTF-8 code points from 40 to 42 inclusive: @AB	Text: /url?parameter=value1 Regular expression: [012] Matches: 1 Would also match 0 or 2. Text: /url?parameter=valueB Regular expression: [A-C] Matches: B Would also match “A” or “C”. It would <i>not</i> match “b”.
{}	Quantifies the number of times the previous character or capture group may be repeated continuously. To define a varying number repetitions, delimit it with a comma.	Text: 1234567890 Regular expression: \d{3} Matches: 123 Text: www.example.com Regular expression: w{1,4} Matches: www If the string were a typo such as “ww” or “www”, it would also match that.
(?i)	Turns on case-insensitive matching for subsequent evaluation, until it is turned off or the evaluation completes.	Text: /url?Parameter=value Regular expression: (?i)param Matches: Param Would also match pArAM etc.

Pattern	Usage	Example
	Matches <i>either</i> the character/capture group before <i>or</i> after the pipe ().	Text: Host: www.example.com Regular expression: (\r\n)\n\r Matches: The line ending, regardless of platform.

PCRE examples submitted to the FortiGate Cookbook

Regular Expression	Usage
[a-zA-Z0-9]	Any alphanumeric character. ASCII only; e.g. does not match é or É.
[#\?](.*)	All parameters that follow a question mark or hash mark in the URL. e.g. #pageView or ?param1=valueA&param2=valueB...; In this expression, the capture group does not include the question mark or hash mark itself.
\b10\.\1\.\1\.\1\b	A specific IPv4 address.
\b(25[0-5] 2[0-4][0-9] [01]?[0-9][0-9]?)\.(25[0-5] 2[0-4][0-9] [01]?[0-9][0-9]?)\.(25[0-5] 2[0-4][0-9] [01]?[0-9][0-9]?)\.(25[0-5] 2[0-4][0-9] [01]?[0-9][0-9]?)\b	Any IPv4 address.
(?i)\b.*\.(a(c(d e(ro)? f(g i m n o q r s(ia)? t y w x z) b(a b d e f g h i(z)? j m n o r s t v w y z) c(a(t)? c d f g h i k l m n o(m)?(op)? r s u v x y z) d(e j k m o z) e(c d u e g h r s t u) f(i j k m o r) g(a b d e f g h i l m n ov p q r s t u w y) h(k m n r t u) i(d e l m n(fo)?(t)? o q r s t) j(e m o(bs)? p) k(e g h i m n p r w y z) l(a b c i k r s t u v y) m(a c d e g h i l m n o(bi)? p q r s t u(seum)? v w x y z) n(a(me)? c e(t)? f g i l o p r u z) o(m rg) p(a e f g h k l m n r(o)? s t w y) qa r(e o s u w) s(a b c d e g h i j k l m n o r s t u v y z) t(c d e f g h i j k l m n o p r(avel)? t v w z)	Any domain name.

Regular Expression	Usage
u(a g k s y z)	
v(a c e g i n u)	
w(f s)	
xxx	
y(e t u)	
z(a m w))\b	
(?i)\bwww\.\example\.\com\b	A specific domain name.
(?i)b(.*)\.\example\.\com\b	Any sub-domain name of example.com.

Example: Rewriting the HTTP response when using content routing

It is standard for web servers to have external and internal domain names. You can use content-based routing to forward HTTP requests to example.com to a server pool that includes server1.example.com, server2.example.com, and server3.example.com. When you use content routing like this, you should also rewrite the Location header in the HTTP response so that the client receives HTTP with example.com in the header and not the internal domain server1.example.com.

[Rewriting the HTTP response when masking internal server names on page 92](#) shows an HTTP response rule that matches a regular expression and rewrites a literal string. In the match condition table, the rule is set to match the regular expression `server.*\.\example\.\com` in the HTTP Location header in the response. The rewrite action specifies the absolute URL `http://www.example.com`.

Rewriting the HTTP response when masking internal server names

Content Rewriting

Name

CWR1

Action Type

RequestResponse

Action

Rewrite HTTP Location

Specifics

Location

http://example.com

Comments

comments

Match Condition (Empty Match Condition will match anything)

Add Filter

Create New

ID	Object	Type	Content	Reverse	
1	HTTP Host Header	Regular Expression	match	Disable	  

Showing 1 to 1 of 1 entries

Show 10 entries

Previous 1 Next

Save

Cancel

Example: Rewriting the HTTP request and response to mask application details

Another use case for external-to-internal URL translation involves masking pathnames that give attackers information about your web applications. For example, the unmasked URL for a blog might be `http://www.example.com/wordpress/?feed=rss2`, which exposes that the blog is a wordpress application. In this case, you want to publish an external URL that does not have clues of the underlying technology. For example, in your web pages, you create links to `http://www.example.com/blog` instead of the backend URL.

On FortiADC, you create two rules: one to rewrite the HTTP request to the backend server and another to rewrite the HTTP response in the return traffic.

[Rewriting the HTTP request when you mask backend application details on page 93](#) shows an HTTP request rule. In the match condition table, the rule is set to match traffic that has the Host header domain `example.com` and the relative URL `/blog` in the HTTP request URL. The rule action rewrites the request URL to the internal URL `http://www.example.com/wordpress/?feed=rss2`.

Rewriting the HTTP request when you mask backend application details

Content Rewriting

Name

CWR1

Action Type

RequestResponse

Action

Rewrite HTTP Header

Specifics

Rewrite Host

OFF

Rewrite URL

ON

URL Content

www.example.com

Rewrite Referrer

OFF

Comments

comments

Match Condition (Empty Match Condition will match anything)

Add Filter

Create New

ID	Object	Type	Content	Reverse	
1	HTTP Host Header	String	match	Disable	  

Showing 1 to 1 of 1 entries

Show 10 entries

Previous

1

Next

Save

Cancel

Rewriting the HTTP response when you mask backend application details on page 94 shows the rule for the return traffic. In the match condition table, the rule is set to match traffic that has the string `http://www.example.com/wordpress/?feed=rss2` in the Location header of the HTTP response. The action replaces that URL with the public URL `http://www.example.org`.

Rewriting the HTTP response when you mask backend application details

Content Rewriting

Name

1

Action Type

RequestResponse

Action

Rewrite HTTP Location

Specifics

Location

http://www.example.org

Comments

comments

Match Condition (Empty Match Condition will match anything)

Add Filter

Create New

ID	Object	Type	Content	Reverse	
1	HTTP Location Header	String	match	Disable	  

Showing 1 to 1 of 1 entries

Show 10 entries

Previous

1

Next

Save

Cancel

Example: Rewriting the HTTP request to harmonize port numbers

The HTTP [Host](#) header contains the domain name and port. You might want to create a rule to rewrite the port so you can harmonize port numbers that are correlated with your application service. For example, suppose you want to avoid parsing reports on your backend servers that show requests to many HTTP service ports. When you review your aggregated reports, you have records for port 80, port 8080, and so on. You would rather have all HTTP requests served on port 80 and accounted for on port 80. To support this plan, you can rewrite the HTTP request headers so that all the Host header in all HTTP requests shows port 80.

[Rewriting the HTTP request port number on page 95](#) shows an HTTP request rule that uses a regular expression to match HTTP Host headers for `www.example.com` with any port number and change it to port 80.

Rewriting the HTTP request port number

Content Rewriting

Name

1

Action Type

RequestResponse

Action

Rewrite HTTP Header

Specifics

Rewrite Host

ON

Host Content

www.example.org

Rewrite URL

OFF

Rewrite Referer

OFF

Comments

comments

Match Condition (Empty Match Condition will match anything)

Add Filter

Create New

ID	Object	Type	Content	Reverse	
1	HTTP Host Header	Regular Expression	www.example.org	Disable	

Showing 1 to 1 of 1 entries

Show 10 entries

Previous 1 Next

Save

Cancel

HSTS and HPKP support

Starting from its 4.8.1 release, FortiADC supports HSTS and HPKP to offer enhanced web security to its users.

HSTS

HSTS, or HTTP Strict Transport Security, is a web security mechanism used to guard websites against malicious attacks, such as protocol downgrading and cookie hijacking. Once implemented, HSTS enables the web server to force web browsers to use secure HTTPS connections when interacting with it, and prohibit the use of insecure HTTP connections.

An HSTS-enabled web application server communicates its HSTS policy to web browsers via an HTTPS header field called "Strict-Transport-Security". The policy dictates that web browsers should only connect to the server via a secure connection during the period of time (i.e., max-age) specified in the policy. Based on the HSTS policy, compliant web browsers either automatically convert insecure (i.e., HTTP) connections to secure (i.e., HTTPS) ones or show an error message and bar the user from accessing the server if it cannot ensure the security of the connection.

HSTS is used to address SSL/TSL-stripping attacks and prevent hackers from stealing your cookie-based web login credentials.

HSTS syntax:

- Strict-Transport-Security: max-age=<expire-time> [; includeSubDomains] [; preload]
- Preload validation and registration:
- <https://www.chromium.org/hsts>
- <https://hstspreload.org>

HPKP

HPKP, or HTTP Public Key Pinning, is a web security mechanism used to prevent HTTPS websites from impersonation via mis-issued or fraudulent security certificates.

The first time a client browser accesses an HTTPS web application server, the server sends to the client a set of public keys, which are the only ones that should be trusted for connections to the domain. This list of "pinned" public key hashes are used for subsequent connections between the client and the server, and are valid only for the period of time that is specified in the HPKP policy.

HPKP syntax

- Public-Key-Pins: pin-sha256="<pin-value>"; pin-sha256="<backup-pin-value>"; max-age=expireTime [; includeSubDomains] [; report-uri="reportURI"]
- Public-Key-Pins-Report-Only: pin-sha256="<pin-value>"; pin-sha256="<backup-pin-value>"; max-age=<expire-time> [; includeSubDomains] [; report-uri="<uri>"]

HPKP note and validation

- Note a host as the known pinned host:
- Identified only by its domain name, but never IP
- Three conditions:
 - PKP received over an error-free TLS, including possible HPKP validation
 - At least one intersection

- The host must set a backup pin.
- Pin Validation:
- Ignore superfluous certificates
- Check intersection, at least one
- Can be disabled for some hosts according to local policy

Good HPKP practices

- If used incorrectly, HPKP could lock out users for a long period of time. Using backup certificates and/or pinning the CA certificate is recommended.
- Use small value for max-age.
- When a certificate expires, generate a new certificate using the old key if pinning is done on the server certificate.

HPKP calculation

Use the following OpenSSL commands to calculate HPKP fingerprints:

- `openssl rsa -in my-rsa-key-file.key -outform der -pubout | openssl dgst -sha256 -binary | openssl enc -base64`
- `openssl ec -in my-ecc-key-file.key -outform der -pubout | openssl dgst -sha256 -binary | openssl enc -base64`
- `openssl req -in my-signing-request.csr -pubkey -noout | openssl pkey -pubin -outform der | openssl dgst -sha256 -binary | openssl enc -base64`
- `openssl x509 -in my-certificate.crt -pubkey -noout | openssl pkey -pubin -outform der | openssl dgst -sha256 -binary | openssl enc -base64`
- `openssl s_client -servername www.example.com -connect www.example.com:443 | openssl x509 -pubkey -noout | openssl pkey -pubin -outform der | openssl dgst -sha256 -binary | openssl enc -base64`

Implementation of HSTS/HPKP

Support for HSTS and HPKP can be implemented for both SSL offloading and forward proxy.

SSL offloading

On the Server Load Balance>Virtual Server>Content Rewriting page, do the following:

1. (Optional) Add a content rewriting rule to delete the HSTS or HPKP header received from the real server. Skip this step if the real server did not send any HSTS or HPKP header. See [Delete HTTP header \(optional\) on page 98](#).
2. Add one content rewriting to add an HSTS or HPKP header, customize the max-age and other optional fields. See [Add HTTP header on page 99](#).

Delete HTTP header (optional)

Content Rewriting

Name

sa

Action Type

Request

Response

Action

Delete HTTP Header

Specifics

Header Name

header-name

Header Value

Specify the header value.

Comments

comments

Match Condition (Empty Match Condition will match anything)

Add Filter

Create New

ID	Object	Type	Content	Reverse	
1	HTTP Host Header	String	match	Disable	

Showing 1 to 1 of 1 entries Show 10 entries Previous 1 Next

Save

Cancel

Add HTTP header

Content Rewriting

Name: sa

Action Type: Request Response

Action: Add HTTP Header

Specifics

Header Name: header-name

Header Value: Specify the header value.

Comments: comments

Match Condition (Empty Match Condition will match anything)

Add Filter

Create New

ID	Object	Type	Content	Reverse	
1	HTTP Host Header	String	match	Disable	

Showing 1 to 1 of 1 entries Show 10 entries Previous 1 Next

Save Cancel

Forward proxy

On the Server Load Balance>Virtual Server>Content Rewriting page, do the following:

1. (Optional) Add a content rewriting rule to delete the HSTS or HPKP header received from the real server. Skip this step if the real server did not send any HSTS or HPKP header.
2. Do nothing to HSTS header (let it pass through).

Configuring content routes

You can use the content routes configuration to select the backend server pool based on matches to TCP/IP or HTTP header values.

Layer 7 content route rules are based on literal or regular expression matches to the following header values:

- [HTTP Host](#)
- [HTTP Referer](#)
- [HTTP Request URL](#)

- [SNI](#)
- Source IP address

You might want to use Layer 7 content routes to simplify front-end coding of your web pages or to obfuscate the precise server names from clients. For example, you can publish links to a simple URL named `example.com` and use content route rules to direct traffic for requests to `example.com` to a server pool that includes `server1.example.com`, `server2.example.com`, and `server3.example.com`.

Layer 4 content route rules are based on literal or regular expression matches to the following header values:

- Source IP address

Before you begin:

- You must have a good understanding of [HTTP header fields](#).
- You must have a good understanding of Perl-compatible regular expressions ([PCRE](#)) if you want to use them in rule matching.
- You must have Read-Write permission for Load Balance settings.

After you have configured a content routing rule, you can select it in the virtual server configuration.

Note: You can select multiple content routing rules in the virtual server configuration. Rules you add to that configuration are consulted from top to bottom. The first rule to match is applied. If the traffic does not match any of the content routing rule conditions specified in the virtual server configuration, the system behaves unexpectedly. Therefore, it is important that you create a “catch all” rule that has no match conditions. In the virtual server configuration, this rule should be ordered last so it can be used to forward traffic to a default pool.

To configure a content route rule:

1. Go to Server Load Balance > Virtual Server.
2. Click the **Content Routing** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Content routes configuration guidelines on page 101](#).
5. Save the configuration.

Content routes configuration guidelines

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, <code>_</code> , and <code>-</code> . No spaces. You reference this name in the virtual server configuration. Note: After you initially save the configuration, you cannot edit the name.
Type	• Layer 4 • Layer 7
Real Server	Select a real server pool.
Persistence Inherit	Enable to use the persistence object specified in the virtual server configuration.
Persistence	If not using inheritance, select a session persistence type.
Method Inherit	Enable to use the method specified in the virtual server configuration.
Method	If not using inheritance, select a load balancing method type.

Settings	Guidelines
Comments	A string to describe the purpose of the configuration, to help you and other administrators more easily identify its use.
Layer 4 Specifics	
IPv4/Mask	Address/mask notation to match the source IP address in the packet header.
IPv6/Mask	Address/mask notation to match the source IP address in the packet header.
Layer 7 Match Condition	
Object	Select content matching conditions based on the following parameters: • HTTP Host Header • HTTP Referer Header • HTTP Request URL • SNI • Source IP Address Note: When you add multiple conditions, FortiADC joins them with an AND operator. For example, if you specify both a HTTP Host Header and HTTP Request URL to match, the rule is a match only for traffic that meets both conditions.
Type	• String • Regular Expression
Content	Specify the string or PCRE syntax to match the header or IP address. Note: An empty match condition matches any HTTP request.
Reverse	Rule matches if traffic does not match the expression.
Ignore Case	Allows user to let the match be case sensitive. Default is ignore case, disabled.

Using source pools

This topic includes a procedure for configuring the source IP address pools used in NAT, and examples of NAT deployments. It includes the following sections:

- [Configuring source pools](#)
- [Example: DNAT](#)
- [Example: full NAT](#)
- [Example: NAT46 \(Layer 4 virtual servers\)](#)
- [Example: NAT64 \(Layer 4 virtual servers\)](#)
- [Example: NAT46 \(Layer 7 virtual servers\)](#)
- [Example: NAT64 \(Layer 7 virtual servers\)](#)

Configuring source pools

You use the Source Pool page to create configuration objects for source IP addresses used for NAT in Layer 4 virtual server configurations.

In a Layer 4 virtual server configuration, you select a “packet forwarding method” that includes the following network address translation (NAT) options:

- Direct Routing—Does not rewrite source or destination IP addresses.
- DNAT—Rewrites the destination IP address for packets before it forwards them.
- Full NAT—Rewrites both the source and destination IP addresses. Use for standard NAT, when client and server IP addresses are all IPv4 or all IPv6.
- NAT46—Rewrites both the source and destination IP addresses. Use for NAT 46, when client IP addresses are IPv4 and server IP addresses are IPv6.
- NAT64—Rewrites both the source and destination IP addresses. Use for NAT 64, when client IP addresses are IPv6 and server IP addresses are IPv4.

In a Layer 7 virtual server configuration, you do not select a packet forwarding option. Layer 7 virtual servers use NAT46 and NAT64 to support those traffic flows, but they do not use the Source Pool configuration.

See the examples that follow the procedure for illustrated usage.

Before you begin:

- You must have a good understanding of NAT. You must know the address ranges your network has provisioned for NAT.
- Be sure to configure the backend servers to use the FortiADC address as the default gateway so that server responses are also rewritten by the NAT module.
- You must have Read-Write permission for Load Balance settings.

After you have configured a source pool IP address range configuration object, you can select it in the virtual server configuration. You can assign a virtual server multiple source pools (with the same or different source pool interface associated with it).

To configure a source pool:

1. Go to Server Load Balance > Virtual Server.
2. Click the **NAT Source Pool** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Source pool configuration on page 103](#).
5. Save the configuration.

Source pool configuration


Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the virtual server configuration. Note: After you initially save the configuration, you cannot edit the name.
Interface	Interface to receive responses from the backend server. The interface used for the initial client traffic is determined by the virtual server configuration.
Address Type	• IPv4 • IPv6
Address Range	The first address in the address pool.
To	The last address in the address pool.

Settings	Guidelines
Node Member	
Name	Create a node member list to be used in an HA active-active deployment. In an active-active deployment, node interfaces are configured with a list of IP addresses for all nodes in the cluster. You use this configuration to provision SNAT addresses for each of the nodes. Name is a configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the virtual server configuration. Note: After you initially save the configuration, you cannot edit the name.
Pool Type	IPv4 or IPv6.
Minimum IP	The first address in the address pool.
Maximum IP	The last address in the address pool.
Interface	Interface to receive responses from the backend server. The interface used for the initial client traffic is determined by the virtual server configuration.
HA Node Number	Specify the HA cluster node ID.

Example: DNAT

[Destination NAT on page 104](#) illustrates destination NAT (DNAT). The NAT module rewrites only the destination IP address. Therefore, if you configure destination NAT, you do not need to configure a source pool. In this DNAT example, the destination IP address in the packets it receives from the client request is the IP address of the virtual server—192.168.1.101. The NAT module translates this address to the address of the real server selected by the load balancer—in this example, 192.168.2.1. The system maintains this NAT table and performs the inverse translation when it receives the server-to-client traffic.

Destination NAT




Example: full NAT

Full NAT on page 106 illustrates full NAT. The source IP / destination IP pair in the packets received is SRC 192.168.1.1 / DST 192.168.1.101. The NAT module translates the source IP address to the next available address in

the source pool—in this example, 192.168.2.101. It translates the destination IP address to the address of the real server selected by the load balancer—in this example, 192.168.2.1.

The system maintains this NAT table and performs the inverse translation when it receives the server-to-client traffic.

Full NAT




Example: NAT46 (Layer 4 virtual servers)

[NAT46 \(Layer 4 virtual servers\)](#) on page 107 illustrates full NAT with NAT46. The IPv6 client connects to the virtual server IPv4 address. The source IP / destination IP pair in the packets received is SRC 192.168.1.1 / DST 192.168.1.101. The NAT module translates the source IP address to the next available IPv6 address in the source pool—in this example, 2002::2:1001. It translates the destination IP address to the IPv6 address of the real server selected by the load balancer—in this example, 2002::2:1.

The system maintains this NAT table and performs the inverse translation when it receives the server-to-client traffic.

NAT46 (Layer 4 virtual servers)



Limitations: NAT46 (Layer 4 virtual servers)


Features	Notes
Profile	Not Supported: FTP
ICMP	ICMP traffic is dropped.

Example: NAT64 (Layer 4 virtual servers)

[NAT64 \(Layer 4 virtual servers\)](#) on [page 109](#) illustrates full NAT with NAT64. The IPv6 client connects to the virtual server IPv6 address. The source IP / destination IP pair in the packets received is SRC 2001::1:1 / DST 2001::1:101. The NAT module translates the source IP address to the next available IPv4 address in the source pool—in this example, 192.168.2.101. It translates the destination IP address to the IPv4 address of the real server selected by the load balancer—in this example, 192.168.2.1.

The system maintains this NAT table and performs the inverse translation when it receives the server-to-client traffic.

NAT64 (Layer 4 virtual servers)



	Ingress	Egress
NAT64		
SRC	2001::1:1	192.168.2.101
DST	2001::1:101	192.168.2.1

Limitations: NAT64 (Layer 4 virtual servers)


Features	Notes
Profiles	Not Supported: FTP
ICMP	ICMP traffic is dropped.
Security	Not Supported: IP Reputation, DoS protection, Security logs and reports

Example: NAT46 (Layer 7 virtual servers)

[NAT46 \(Layer 7 virtual servers\) on page 111](#) illustrates full NAT with NAT46. The IPv4 client connects to the virtual server IPv4 address. The source IP / destination IP pair in the packets received is SRC 192.168.1.1 / DST 192.168.1.101. The NAT module translates the source IP address to the IPv6 address of the egress interface that has IPv6 connectivity with the real server—in this example, 2002::2:1001. It translates the destination IP address to the IPv6 address of the real server selected by the load balancer—in this example, 2002::2:1.

The system maintains this NAT table and performs the inverse translation when it receives the server-to-client traffic.

NAT46 (Layer 7 virtual servers)



Limitations: NAT46 (Layer 7 virtual servers)


Feature	Note
Profiles	Not Supported: RADIUS, HTTP Turbo
Profile options	Not supported: Source Address (Using the original source IP address for the connection to the real server is contrary to the purpose of NAT.)
Virtual server options	Not supported: Connection Rate Limit
Real server pool options	Not supported: Connection Rate Limit

Example: NAT64 (Layer 7 virtual servers)

[NAT64 \(Layer 7 virtual servers\)](#) on page 113 illustrates full NAT with NAT64. The IPv6 client connects to the virtual server IPv6 address. The source IP / destination IP pair in the packets received is SRC 2001::1:1 / DST 2001::1:101. The NAT module translates the source IP address to the IPv4 address of the egress interface that has IPv4 connectivity with the real server—in this example, 192.168.2.101. It translates the destination IP address to the IPv4 address of the real server selected by the load balancer—in this example, 192.168.2.1.

The system maintains this NAT table and performs the inverse translation when it receives the server-to-client traffic.

NAT64 (Layer 7 virtual servers)



Limitations: NAT64 (Layer 7 virtual servers)

Feature	Note
Profiles	Not Supported: RADIUS, HTTP Turbo
Profile options	Not supported: Source Address (Using the original source IP address for the connection to the real server is contrary to the purpose of NAT.)
Virtual server options	Not supported: Connection Rate Limit
Real server pool options	Not supported: Connection Rate Limit
Security	Not Supported: IP Reputation, DoS protection, Security logs and reports

Using schedule pools

A schedule pool is a list of configuration objects, each of which is tied to a specific real-server pool and schedule group. Used together with real-server pools, schedule groups, and content routing rules, schedule pools make it much easier for you to streamline the operation and management of your real servers. You set or change the working schedules of your real servers with ease.

The schedule pool feature takes the following two factors are taken into consideration:

First, there can be multiple pools in a virtual server or a content routing configuration. This does not mean to introduce a traffic distributing hierarchy to load-balance across the pools because all the pools of different schedule pools in a virtual server obey the same rule of traffic distribution. So the basic schema is not changed. The way it works is the same as a single pool does. We have the following specific confines:

- The same real server pool is not allowed to be used in different schedule pools which are configured in the same virtual server.
- The same real server is not allowed to be used in different real-server pools that are used by schedule pools configured in the same virtual server.
- When multiple schedule pools are active, all the real-server pools within them (schedule pools) are active, and traffic can be transmitted to all the real servers in the real-server pools as scheduled. In that case, all the real servers are placed in different pools for scheduling.
- The backup real servers are backed up for all the current active real servers from multiple schedule pools of a virtual server.

Second, a schedule pool can be scheduled inactive. The schedule daemon tracks the states of all the schedules. When a schedule's state changes, the schedule daemon updates the new state to all the related daemons. As soon as the state of a schedule pool goes active, the system will start to transmit traffic to members of the corresponding pool unless there are some other mechanisms keeping the schedule pool or some members of the pool in "not work" state, as in the case of health check failure or backup members of the pool. Once a schedule turns inactive, the system will stop transmitting traffic to all the members of the corresponding pool. Some or all members of the pool may be in "not work" state for various reasons when a schedule's state changes to inactive. Anyway, when members of a pool turn inactive, the system will react in the same way as it does when they fail their health check — immediately removes the connections involved and cuts off traffic to those connections at the same time.

The schedule-based pool can be applied to all kinds of virtual servers and all kinds of content routing configurations. It should also work well with all packet-forwarding methods, and can handle all the protocols that FortiADC now supports.

How to use the "schedule pool" feature

The following are the basic steps you need to follow to take advantage of the schedule pool feature:

1. Configure schedule groups (Shared Resources > Schedule Group).
2. Configure real servers (Server Load Balance > Real Server).
3. Configure real-server pools (Server Load Balance > Real Server Pool).
4. Configure schedule pools (Server Load Balance > Virtual Server > Schedule Pool).
5. Configure content routing rules (Server Load Balance > Virtual Server > Content Routing). *(Optional)*
6. Configure virtual servers (Server Load Balance > Virtual Server)

Configuring schedule pools

The following instructions assume that you have properly configured schedule groups, real servers, and real server pools, as mentioned in the preceding paragraph.

To configure schedule pools:

1. From the main menu, click **Server Load Balance > Virtual Server**.
2. Select the Schedule Pool tab.
3. Click **Create New** to open the Schedule Pool dialog box.
4. Specify a unique name for the schedule pool.
5. Select a real server pool.
6. Select a schedule group.
7. Click **Save** when done.
8. Repeat Steps 2 through 7 to create as many schedule pools as needed.

Using clone pools

A clone pool is a set of destinations, of monitor servers.

The FortiADC is tasked with protecting the real-server pools. Before allowing traffic to reach the servers, it will duplicate the traffic, sending a copy towards the clone pool, which holds onto it.

As such, the clone pool is assigned to a virtual server. In the clone pool is a farm of monitor servers; some of these monitor servers can be IDS servers - intrusion detection system (IDS) - which will analyze traffic to identify suspicious patterns. The IDS server does not perform fire wall functions, like blocking the traffic. However, the IDS server will send out, say, an email, indicating that the server

Important: A clone pool receives all of the same traffic that the server pool receives.


To configure a clone pool, you first create a pool of IDS or sniffer devices and then assign the pool as a clone pool to a virtual server. The clone pool feature is the recommended method for copying production traffic to IDS systems or sniffer devices. Note that when you create the clone pool, the service port that you assign to each node is irrelevant; you can choose any service port. Also, when you add a clone pool to a virtual server, the system copies only new connections; existing connections are not copied.

You can configure a virtual server to copy client-side traffic, server-side traffic, or both:

- A client-side clone pool causes the virtual server to replicate client-side traffic (prior to address translation) to the specified clone pool.
- A server-side clone pool causes the virtual server to replicate server-side traffic (after address translation) to the specified clone pool.

[Clone pool topology on page 117](#) illustrates how clone pools work.

Clone pool topology



The following steps show the process in which FortiADC clones packets and sends them to the monitor servers:

1. Duplicates the packet data structure.
2. Looks up the route table by monitor server IP to find out the next-hop IP address and output device, if necessary.
3. Looks up the neighbors by the next-hop IP address, if necessary.
4. Updates packet headers with specified values or results of route and ARP look-up.
5. Sends the packets out to the monitor servers.

Configuring a clone pool

Before starting to create clone pools, keep the following in mind:

- Only one clone pool can be configured for the virtual server.
- The clone pool can have at most four members. The traffic will be duplicated and sent to each of the members.
- Only IPv4 addresses are supported.

- There are four modes by which you may update and send the packets.
- When the clone pool is added to the virtual server, the traffic (of old sessions and new) is duplicated and sent to the monitor servers in the clone pool.
- The following is true:
 - If the virtual server is of the type L7, then the profiles TURBOHTTP, HTTP, HTTPS, TCPS, RDP, are supported.
 - If the virtual server is of the type L2, then the profiles TCP, UDP, IP, HTTP, HTTPS, TCPS, are supported.
 - If the virtual server is of the type L4, then the profiles TCP, UDP, FTP, are supported.
- Traffic of both client and server sides may be cloned. For the client-side, traffic is replicated BEFORE the packet's address undergoes Network Address Translation (NAT) such that it may reach the clone members. For the server-side, however, NAT has already happened; the packet has already gone through the virtual server. Thus the traffic is replicated AFTER the packet address has been translated.

To configure a clone pool:

The following instructions assume that you have properly configured schedule groups, real servers, and real server pools.

1. Go to **Server Load Balance > Virtual Server > Clone Pool**.
2. Click **Create New**.
3. Return to Clone Pool tab and select your clone pool, and click **edit**.
4. Click **Create New** to create a member inside your clone pool. Create as many members as four.
5. Refer to the table below for entries and/or selections required for creating a clone pool.

Parameters for clone pool configuration

Entry/Selection	Description
Clone Pool	
Name	Specify a unique clone pool name
Pool Member	
Name	Specify a unique pool member name. Note: A pool member is a clone server. So this name is essentially the name you give to the clone server.
Interface	Select the interface (port) FortiADC uses to send out packets to the clone server.
Mode	The headers of duplicated packets need to be updated when sent to monitor servers. There are several modes in which this occurs. Select one of the following: • Mirror Interface—This mode does not change the packet header at all. It is most commonly used; with it, the monitor does not look at the content of the packet, neither does it receive the payload, it merely looks at how much data is being passed, and counts the bytes of the data. The original Layer 2 Destination Address (DA) or Source Address (SA) and Layer 3 IP Addresses are left intact. In this mode the FortiADC simply sends the packets "as is" out from the specified

Entry/Selection	Description
	interface. • Mirror Destination MAC Address Update—This mode uses Layer 2 forwarding. With the incoming packet, the ADC replaces the destination MAC address with the specified destination MAC address. It is preferred when connecting the ADC to end devices like the IDS. • Mirror Source MAC Update—This mode replaces the source MAC address in the incoming packet with the specified MAC address on the FortiADC device. This option is recommended where not changing the source MAC address could cause a loop. • Mirror Source Destination MAC Update—This mode replaces both the source and destination MAC addresses at Layer 2, but does not change the Layer-3 IP addressing information. • Mirror IP Update—This mode replaces the incoming packet's IP address with the specified IP address and then forwards the duplicated packet to those servers. This mode may also change the Layer 4 source and destination ports. If the virtual server port isn't set to wildcard port 0 while the port is specified, the Layer 4 destination port on the duplicated packets will be changed to the specified value. This option is recommended for scenarios in which monitor servers are not directly connected to the ACOS device.

Configuring Application profiles

An application profile is a configuration object that defines how you want the FortiADC virtual server to handle traffic for specific protocols.

[Application profile usage on page 119](#) describes usage for by application profile type, including compatible virtual server types, load balancing methods, persistence methods, and content route types.

Application profile usage

Profile	Usage	VS Type	LB Methods	Persistence
FTP	Use with FTP servers.	Layer 4	Round Robin, Least Connections, Fastest Response	Source Address, Source Address Hash

Profile	Usage	VS Type	LB Methods	Persistence
HTTP	Use for standard, unsecured web server traffic.	Layer 7, Layer 2	Layer 7: Round Robin, Least Connections, URI Hash, Full URI Hash, Host Hash, Host Domain Hash, Dynamic Load Layer 2: Same as Layer 7, plus Destination IP Hash	Source Address, Source Address Hash, Source Address-Port Hash, HTTP Header Hash, HTTP Request Hash, Cookie Hash, Persistent Cookie, Insert Cookie, Embedded Cookie, Rewrite Cookie, Passive Cookie
HTTPS	Use for secured web server traffic when offloading TLS/SSL from the backend servers. You must import the backend server certificates into FortiADC and select them in the HTTPS profile.	Layer 7, Layer 2	Same as HTTP	Same as HTTP, plus SSL Session ID
TURBO HTTP	Use for unsecured HTTP traffic that does not require advanced features like caching, compression, content rewriting, rate limiting, Geo IP blocking, or source NAT. The profile can be used with content routes and destination NAT, but the HTTP request must be in the first data packet. This profile enables packet-based forwarding that reduces network latency and system CPU usage. However, packet-based forwarding for HTTP is advisable only when you do not anticipate dropped packets or out-of-order packets.	Layer 7	Round Robin, Least Connections, Fastest Response	Source Address
RADIUS	Use with RADIUS servers.	Layer 7	Round Robin	RADIUS attribute
RDP	Use with Windows Terminal Service(remote desktop protocol).	Layer 7	Round Robin, Least Connections	Source Address, Source Address Hash, Source Address-Port Hash, RDP Cookie

Profile	Usage	VS Type	LB Methods	Persistence
SIP	Use with applications that use session initiation protocol (SIP), such as VoIP, instant messaging, and video.	Layer 7	Round Robin, URI Hash, Full URI Hash	Source Address, Source Address Hash, Source Address-Port Hash, SIP Call ID
TCP	Use for other TCP protocols.	Layer 4, Layer 2	Layer 4: Round Robin, Least Connections, Fastest Response Layer 2: Round Robin, Least Connections, Fastest Response, Destination IP Hash	Source Address, Source Address Hash
TCPS	Use for secured TCP when offloading TLS/SSL from the backend servers. Like the HTTPS profile, you must import the backend server certificates into FortiADC and select them in the TCPS profile.	Layer 7, Layer 2	Layer 7: Round Robin, Least Connections Layer 2: Round Robin, Least Connections, Destination IP Hash	Source Address, Source Address Hash, Source Address-Port Hash, SSL Session ID
UDP	Use with UDP servers.	Layer 4, Layer 2	Layer 4: Round Robin, Least Connections, Fastest Response, Dynamic Load Layer 2: Same as Layer 4, plus Destination IP Hash	Source Address, Source Address Hash
IP	Combines with Layer 2 TCP/UDP/HTTP virtual server to balance the rest of the IP packets passed through FortiADC. When running the IP protocol 0 VS, the traffic always tries to match none protocol 0 VS first.	Layer 2	Round Robin only.	Source Address, Source Address Hash
DNS	Use with DNS servers.	Layer 7	Round Robin, Least Connections	Not supported yet.
SMTP	Use with SMTP servers.	Layer 7	Round Robin, Least Connections	Source Address, Source Address Hash
RTMP	A TCP-based protocol used for streaming audio, video, and data over the Internet	Layer 7	Round Robin, Least Connection	Source Address, Source Address Hash

Profile	Usage	VS Type	LB Methods	Persistence
RTSP	A network control protocol used for establishing and controlling media sessions between end points	Layer 7	Round Robin, Least Connection	Source Address, Source Address Hash
MySQL	MySQL network protocol stack (i.e., MySQL-Proxy) which parses and builds MySQL protocol packets	Layer 7	Round Robin, Least Connection	N/A
DIAMETER	A successor to RADIUS, DIAMETER is the next-generation Authentication, Authorization and Accounting (AAA) protocol widely used in IMS and LTE.	Layer 7	Round Robin	Source Address. DIAMETER Session ID (default)

[Predefined profiles on page 122](#) shows the default values of the predefined profiles. All values in the predefined profiles are view-only, and cannot be modified. You can select predefined profiles in the virtual server configuration, or you can create user-defined profiles, especially to include configuration objects like certificates, caching settings, compression options, and IP reputation.

Predefined profiles

Profile	Defaults
LB_PROF_DIAMETER	Identity—Blank Realm—Blank Vendor ID—Blank Product Name—Blank Idle Timeout—300 (seconds) (Note: This refers to the built-in session ID persistence timeout.) Server Close Propagation—OFF (Note: This means that the connection on the client side stays open when the server closes any connection on its side.)
LB_PROF_TCP	Timeout TCP Session—100 Timeout TCP Session after FIN—100 IP Reputation—Disabled Customized SSL Ciphers Flag—Disabled Geo IP block list—None Geo IP Whitelist—None
LB_PROF_UDP	Timeout UDP Session—100 IP Reputation—Disabled Stateless—Disabled Customized SSL Ciphers Flag—Disabled

Profile	Defaults
	Geo IP block list—None Geo IP Whitelist—None
LB_PROF_HTTP	Client Timeout—50 Server Timeout—50 Connect Timeout—5 Queue Timeout—5 HTTP Request Timeout—50 HTTP Keepalive Timeout—50 Buffer Pool—Enabled Source Address—Disabled X-Forwarded-For—Disabled X-Forwarded-For Header—Blank IP Reputation—Disabled HTTP Mode—Keep Alive Customized SSL Ciphers Flag—Disabled Compression—None. Decompression—None Caching—None Geo IP Block List—None Geo IP Whitelist—None Geo IP Redirect URL—http://
LB_PROF_HTTP_SERVERCLOSE	Client Timeout—50 Server Timeout—50 Connect Timeout—5 Queue Timeout—5 HTTP Request Timeout—50 HTTP Keepalive Timeout—50 Buffer Pool—Enabled Source Address—Disabled X-Forwarded-For—Disabled X-Forwarded-For Header—None IP Reputation—Disabled HTTP Mode—Server Close Customized SSL Ciphers Flag—Disabled Compression—None Decompression—None Caching—None Geo IP Block List—None Geo IP Whitelist—None

Profile	Defaults
	Geo IP Redirect URL—http://
LB_PROF_TURBOHTTP	Timeout TCP Session—100 Timeout TCP Session after FIN—100 IP Reputation—Disabled Customized SSL Ciphers Flag—Disabled Geo IP Block List—None Geo IP Whitelist—None
LB_PROF_FTP	Timeout TCP Session—100 Timeout TCP Session after FIN—100 IP Reputation—Disabled Customized SSL Ciphers Flag—Disabled Geo IP Block List—None Geo IP Whitelist—None Source Address—Off
LB_PROF_RADIUS	Dynamic Auth—Disable Session Timeout—300
LB_PROF_SIP	SIP Max Size—65535 Server Keepalive Timeout—30 Server Keepalive—Enabled Client Keepalive—Disabled Client Protocol—UDP Server Protocol—None Failed Client Type—Drop Failed Server Type—Drop Insert Client IP—Disabled Customized SSL Ciphers Flag—Disabled Geo IP Block List—None Geo IP Whitelist—None Source Address—Off Media Address—0.0.0.0
LB_PROF_RDP	Client Timeout—50 Server Timeout—50 Connect Timeout—5 Queue Timeout—5 Buffer Pool—Enabled Source Address—Disabled IP Reputation—Disabled Customized SSL Ciphers Flag—Disabled

Profile	Defaults
	Geo IP Block List—None Geo IP Whitelist—None
LB_PROF_IP	IP Reputation—Disabled Customized SSL Ciphers Flag—Disabled Geo IP Block List—None Geo IP Whitelist—None Timeout IP Session—100
LB_PROF_DNS	DNS Cache Flag—Enabled DNS Cache Ageout Time—3600 DNS Cache Size—10 DNS Cache Entry Size—512 DNS Cache Response Type—All Records DNS Malform Query Action—Drop DNA Max Query Length—512 DNS Authentication Flag—Disabled
LB_PROF_TCPS	Client Timeout—50 Server Timeout—50 Connect Timeout—5 Queue Timeout—5 Buffer Pool—Enabled Source Address—Disabled IP Reputation—Disabled Dynamic Auth—Disabled Customized SSL Ciphers Flag—Disabled Client SNI Required—Disabled Geo IP block list—None Client SNI Required—disabled Certificate Group—LOCAL_CERT_GROUP Certificate Verify—None
LB_PROF_HTTPS	Client Timeout—50 Server Timeout—50 Connect Timeout—5 Queue Timeout—5 HTTP Request Timeout—50 HTTP Keepalive Timeout—50 Buffer Pool—Enabled Source Address—Disabled X-Forwarded-For—Disabled X-Forwarded-For Header—None

Profile	Defaults
	IP Reputation—Disabled HTTP Mode—Keep Alive SSL Proxy Mode—Disabled Customized SSL Ciphers Flag—Disabled Client SNI Required—Disabled Compression—None Decompression—None Caching—None Geo IP Block List—None Geo IP Whitelist—None Geo IP Redirect URL—http:// Certificate Group—LOCAL_CERT_GROUP Certificate Verify—None
LB_PROF_HTTPS_SERVERCLOSE	Client Timeout—50 Server Timeout—50 Connect Timeout—5 Queue Timeout—5 HTTP Request Timeout—50 HTTP Keepalive Timeout—50 Buffer Pool—Enabled Source Address—Disabled X-Forwarded-For—Disabled X-Forwarded-For Header—None IP Reputation—Disabled HTTP Mode—Server Close SSL Proxy Mode—Disabled Customized SSL Ciphers Flag—Disabled SSL Cipher—Shows all available SSL ciphers, with the default ones selected Allow SSL Versions—SSLv3, TLSv1.0, TLS1.1, TLSv1.2 Client SNI Required—Disabled Compression—None Decompression—None Caching—None Geo IP Block List—None Geo IP Whitelist—None Geo IP Redirect URL—http:// Certificate Group—LOCAL_CERT_GROUP Certificate Verify—None

Profile	Defaults
LB_PROF_SMTP	Starttls Active Mode—require Customized SSL Ciphers Flag—Disabled SSL Ciphers—Shows all available SSL Ciphers, with the defaults ones selected Allow SSL Versions —SSLv3, TLSv1.0, TLSv1.1, TLSv1.2 Forbidden Command—expn, turn, vrfy Local Certificate Group—LOCAL_CERT_GROUP
LB_PROF_RTSP	Max Header Size—Default is 4096. Valid values range from 2048 to 65536. Source Address—Disabled by default. When enabled, FortiADC will use the client address to connect to the server pool.
LB_PROF_RTMP	Source Address—Disabled by default. When enabled, FortiADC will use the client address to connect to the server pool.

Before you begin:

- You must have already created configuration objects for certificates, caching, and compression if you want the profile to use them.
- You must have Read-Write permission for Load Balance settings.

To configure custom profiles:

1. Go to Server Load Balance > Application Resources. Click the Application Profile tab.
2. Click **Create New** to display the configuration editor.
3. Give the profile a name, select a protocol type; then complete the configuration as described in [Profile configuration guidelines on page 127](#).
4. Save the configuration.



You can clone a predefined configuration object to help you get started with a user-defined configuration.

To clone a configuration object, click the clone icon  that appears in the tools column on the configuration summary page.

Profile configuration guidelines

Type	Profile Configuration Guidelines
TCP	

Type	Profile Configuration Guidelines
Timeout TCP Session	Client-side timeout for connections where the client has not sent a FIN signal, but the connection has been idle. The default is 100 seconds. The valid range is 1 to 86,400.
Timeout TCP Session after FIN	Client-side connection timeout. The default is 100 seconds. The valid range is 1 to 86,400.
IP Reputation	Enable to apply the FortiGuard IP reputation service. See Managing IP Reputation policy settings .
Geo IP Block List	Select a Geo IP block list configuration object. See Using the Geo IP block list .
Geo IP Whitelist	Select a whitelist configuration object. See Using the Geo IP whitelist .
IP	
IP Reputation	Enable to apply FortiGuard IP reputation service. IP reputation. See Managing IP Reputation policy settings .
Geo IP Block List	Select a Geo IP block list configuration object. See Using the Geo IP block list .
Geo IP Whitelist	Select a whitelist configuration object. See Using the Geo IP whitelist .
Timeout IP Session	Client-side session timeout. The default is 100 seconds. The valid range is 1 to 86,400.
DNS	
Customized SSL Ciphers Flag	Enable or disable the Customized SSL Ciphers Flag.
DNS Cache Flag	Enable/Disable DNS cache flag.
Geo IP Block List	Select a Geo IP block list configuration object. See Using the Geo IP block list .
Geo IP Whitelist	Select a whitelist configuration object. See Using the Geo IP whitelist .
DNS Cache Flag	Enable or disable the DNS Cache Flag.
DNS Cache Ageout Time	Enter a value from 0 to 65,535. The default is 3,600.
DNS Cache Size	Enter a value from 1 to 100. The default is 10.
DNS Cache Entry Size	Enter a value from 256 to 4,096. The default is 512.
DNS Malform Query Action	Choose either of the following: • Drop • Forward
DNS Max Query Length	Enter a value from 256 to 4,096. The default is 512.
DNS Authentication Flag	Enable or disable DNS authentication flag.
<i>Special Note</i>	With the 4.8.1 release, FortiADC supports DNS zone transfer, i.e., DNS traffic over TCP from servers and server-oriented requests from inside the server cluster.
UDP	

Type	Profile Configuration Guidelines
Stateless	Enable to apply UDP stateless function.
Timeout UDP Session	Client-side session timeout. The default is 100 seconds. The valid range is 1 to 86,400.
IP Reputation	Enable to apply the FortiGuard IP reputation service. See Managing IP Reputation policy settings .
Geo IP Block List	Select a Geo IP block list configuration object. See Using the Geo IP block list .
Geo IP Whitelist	Select a whitelist configuration object. See Using the Geo IP whitelist .
HTTP	
Client Timeout	This timeout is counted as the amount of time when the client did not send a complete request HTTP header to the FortiADC after the client connected to the FortiADC. If this timeout expires, FortiADC will send a 408 message to client and close the connection to the client.
Server Timeout	This timeout is counted as the amount of time when the server did not send a complete response HTTP header to the FortiADC after the FortiADC sent a request to server. If this timeout expires, FortiADC will close the server side connection and send a 503 message to the client and close the connection to the client.
Connect Timeout	This timeout is counted as the amount of time during which FortiADC tried to connect to the server with TCP SYN. After this timeout, if TCP connection is not established, FortiADC will drop this current connection to server and respond with a 503 message to client side and close the connection to the client.
Queue Timeout	This timeout is counted as the amount of time during which the request is queued in the dispatched queue. When the request cannot be dispatched to a server by a load balance method (for example, the server's connection limited is reached), it will be put into a queue. If this timeout expires, the request in the queue will be dropped and FortiADC will respond with a 503 message to client side and close the connection to the client.
HTTP Send Timeout	This timeout is counted as the amount of time it took FortiADC to send a response body data (not including the header); the time is counted starting from when the body is transferred. If this timeout expires, FortiADC will close the connection of both side.
HTTP Request Timeout	This timeout is counted as the amount of time the client did not send a complete request (including both HTTP header and request body) to FortiADC after the client connected to FortiADC. If this timeout expires, FortiADC will send a 408 message to client and close the connection to the client.
HTTP Keepalive Timeout	This timeout is counted as the time FortiADC can wait for a new request after the previous transaction is completed. This is an idle timeout if the client does not send anything in this period. If this timeout expires, FortiADC will close the connection to the client.

Type	Profile Configuration Guidelines
Source Address	Use the original client IP address as the source address when connecting to the real server.
X-Forwarded-For	Append the client IP address found in IP layer packets to the HTTP header that you have specified in the X-Forwarded-For Header setting. If there is no existing X-Forwarded-For header, the system creates it. If you only enable http-x-forwarded-for and do not configure http-x-forwarded-for-header, the default is to add such a header: X-Forwarded-For: <client's ip>
X-Forwarded-For Header	Specify the HTTP header to which to write the client IP address. Typically, this is the X-Forwarded-For header, but it is customizable because you might support traffic that uses different headers for this. Examples: Forwarded-For, Real-IP, or True-IP. If http-x-forwarded-for-header <string> is configured, the added header is: <string>: <client's ip>.
IP Reputation	Enable to apply the FortiGuard IP reputation service. See Managing IP Reputation policy settings .
HTTP Mode	• Server Close—Close the connection to the real server after each HTTP transaction. • Once Only—An HTTP transaction can consist of multiple HTTP requests (separate requests for an HTML page and the images contained therein, for example). To improve performance, the "once only" flag instructs the FortiADC to evaluate only the first set of headers in a connection. Subsequent requests belonging to the connection are not load balanced, but sent to the same server as the first request. • Keep Alive—Do not close the connection to the real server after each HTTP transaction. Instead, keep the connection between FortiADC and the real server open until the client-side connection is closed. This option is required for applications like Microsoft SharePoint.
Compression	Select a compression configuration object. See Configuring compression rules .
Caching	Select a caching configuration object. See Using caching features .
Geo IP Block List	Select a Geo IP block list configuration object. See Using the Geo IP block list .
Geo IP Whitelist	Select a whitelist configuration object. See Using the Geo IP whitelist .
Geo IP Redirect URL	For HTTP, if you have configured a Geo IP redirect action, specify a redirect URL.
Tune Buffer Size	Adjust the value of the HTTP/HTTPS VS's connection buffer size. • For every session, there are two connection buffers. • The default size is 8030, it is not recommended that you edit it. It's hidden in the Advance tab, and when you edit it you will get a warning message. • Tuning this option is dangerous because it may lead to concurrent session number reduction or other unpredictable problems.
Max HTTP Headers	Adjust the max header number that HTTP/HTTPS VS can process for every request or response. If a request or response has a header over this limit, it will be dropped, and return error message 400.

Type	Profile Configuration Guidelines
	- The default value is 100, it's not recommended that you edit it. It is hidden in the Advance tab, and when you edit it you will get a warning message. - Tuning this option is dangerous and may lead to concurrent session number reduction or other unpredictable problems.
FTP	
Timeout TCP Session	Client-side timeout for connections where the client has not sent a FIN signal, but the connection has been idle. The default is 100 seconds. The valid range is 1 to 86,400.
Timeout TCP Session after FIN	Client-side connection timeout. The default is 100 seconds. The valid range is 1 to 86,400.
IP Reputation	Enable to apply the FortiGuard IP reputation service. See Managing IP Reputation policy settings .
Customized SSL Ciphers Flag	Enable or disable the Customized SSL Ciphers Flag.
Geo IP Block List	Select a Geo IP block list configuration object. See Using the Geo IP block list .
Geo IP Whitelist	Select a whitelist configuration object. See Using the Geo IP whitelist .
RADIUS	
Timeout RADIUS Session	The default is 300 seconds. The valid range is 1 to 3,600.
Dynamic Auth	Enable or disable Dynamic Authorization for RADIUS Change of Authorization(CoA)
Dynamic Auth Port	Configures the UDP port for CoA requests. The default is 3799.
RDP	
Client Timeout	Client-side TCP connection timeout. The default is 50 seconds. The valid range is 1 to 3,600.
Server Timeout	Server-side IP session timeout. The default is 50 seconds. The valid range is 1 to 3,600.
Connect Timeout	Multiplexed server-side TCP connection timeout. Usually less than the client-side timeout. The default is 5 seconds. The valid range is 1 to 3,600.
Queue Timeout	Specifies how long connection requests to a backend server remain in a queue if the server has reached its maximum number of connections. If the timeout period expires before the client can connect, FortiADC drops the connection and sends a 503 error to the client. The default is 5 seconds. The valid range is 1 to 3,600.
Buffer Pool	Enable or disable buffering.
Source Address	Use the original client IP address as the source address in the connection to the real server.
IP Reputation	Enable to apply the FortiGuard IP reputation service. See Managing IP Reputation policy settings .
Customized SSL Ciphers Flag	Enable or disable the Customized SSL Ciphers Flag.

Type	Profile Configuration Guidelines
Geo IP Block List	Select a Geo IP block list configuration object. See Using the Geo IP block list .
Geo IP Whitelist	Select a whitelist configuration object. See Using the Geo IP whitelist .
TCPs	
Client Timeout	Client-side TCP connection timeout. The default is 50 seconds. The valid range is 1 to 3,600.
Server Timeout	Server-side IP session timeout. The default is 50 seconds. The valid range is 1 to 3,600.
Connect Timeout	Multiplexed server-side TCP connection timeout. Usually less than the client-side timeout. The default is 5 seconds. The valid range is 1 to 3,600.
Queue Timeout	Specifies how long connection requests to a backend server remain in a queue if the server has reached its maximum number of connections. If the timeout period expires before the client can connect, the system drops the connection and sends a 503 error to the client. The default is 5 seconds. The valid range is 1 to 3,600.
Buffer Pool	Enable or disable buffering.
Source Address	Use the original client IP address as the source address in the connection to the real server.
IP Reputation	Enable to apply the FortiGuard IP reputation service. See Managing IP Reputation policy settings .
Customized SSL Ciphers Flag	Enable or disable the use of user-specified cipher suites.
Customized SSL Ciphers	If the customize cipher flag is enabled, specify a colon-separated, ordered list of cipher suites. An empty string is allowed. If empty, the default cipher suite list is used.
SSL Ciphers	Ciphers are listed from strongest to weakest: • ECDHE-ECDSA-AES256-GCM-SHA384 • ECDHE-ECDSA-AES256-SHA384 • ECDHE-ECDSA-AES256-SHA • ECDHE-ECDSA-AES128-GCM-SHA256 • ECDHE-ECDSA-AES128-SHA256 • ECDHE-ECDSA-AES128-SHA • ECDHE-ECDSA-DES-CBC3-SHA • ECDHE-ECDSA-RC4-SHA • ECDHE-RSA-AES256-GCM-SHA384 • ECDHE-RSA-AES256-SHA384 • ECDHE-RSA-AES256-SHA • DHE-RSA-AES256-GCM-SHA384 • DHE-RSA-AES256-SHA256 • DHE-RSA-AES256-SHA • AES256-GCM-SHA384 • AES256-SHA256

Type	Profile Configuration Guidelines
	• AES256-SHA • ECDHE-RSA-AES128-GCM-SHA256 • ECDHE-RSA-AES128-SHA256 • ECDHE-RSA-AES128-SHA • DHE-RSA-AES128-GCM-SHA256 • DHE-RSA-AES128-SHA256 • DHE-RSA-AES128-SHA • AES128-GCM-SHA256 • AES128-SHA256 • AES128-SHA • ECDHE-RSA-RC4-SHA • RC4-SHA • RC4-MD5 • ECDHE-RSA-DES-CBC3-SHA • EDH-RSA-DES-CBC3-SHA • DES-CBC3-SHA • eNULL We recommend retaining the default list. If necessary, you can deselect the SSL ciphers that you do not want to support.
Allow SSL Versions	You have the following options: • SSLv2 • SSLv3 • TLSv1.0 • TLSv1.1 • TLSv1.2 We recommend retaining the default list. If necessary, you can deselect SSL versions you do not want to support. Note: FortiADC does not support session reuse for SSLv2 at the client side. Instead, a new SSL session is started.
Client SNI Required	Require clients to use the TLS server name indication (SNI) extension to include the server hostname in the TLS client hello message. Then, the FortiADC system can select the appropriate local server certificate to present to the client.
Geo IP Block List	Select a Geo IP block list configuration object. See Using the Geo IP block list .
Geo IP Whitelist	Select a whitelist configuration object. See Using the Geo IP whitelist .
Local Certificate Group	A configuration group that includes the certificates this virtual server presents to SSL/TLS clients. This should be the backend servers' certificate, NOT the appliance's GUI web server certificate. See Manage certificates .
Certificate Verify	Select a certificate validation policy. See Manage and validate certificates .
HTTPS	
HTTPS	Same as HTTP, plus the certificate settings listed next. See Chapter 17: SSL Transactions for an overview of HTTPS features.

Type	Profile Configuration Guidelines
SSL Proxy Mode	Enable or disable SSL forward proxy.
Customized SSL Ciphers Flag	Enable or disable use of user-specified cipher suites.
Customized SSL Ciphers	If the customize cipher flag is enabled, specify a colon-separated, ordered list of cipher suites. An empty string is allowed. If empty, the default cipher suite list is used.
SSL Ciphers	We recommend retaining the default list. If necessary, you can deselect ciphers you do not want to support.
Allow SSL Versions	We recommend retaining the default list. If necessary, you can deselect SSL versions you do not want to support. Note: FortiADC does not support session reuse for SSLv2 at the client side. Instead, a new SSL session is started.
Client SNI Required	Require clients to use the TLS server name indication (SNI) extension to include the server hostname in the TLS client hello message. Then, the FortiADC system can select the appropriate local server certificate to present to the client.
Local Certificate Group	A configuration group that includes the certificates this virtual server presents to SSL/TLS clients. This should be the backend servers' certificate, NOT the appliance's GUI web server certificate. See Manage certificates .
Certificate Verify	Select a certificate validation policy. See Manage and validate certificates .
TURBO HTTP	
Timeout TCP Session	Client-side timeout for connections where the client has not sent a FIN signal, but the connection has been idle. The default is 100 seconds. The valid range is 1 to 86,400.
Timeout TCP Session after FIN	Client-side connection timeout. The default is 100 seconds. The valid range is from 1 to 86,400.
IP Reputation	Enable to apply the FortiGuard IP reputation service.
Customized SSL Ciphers Flag	Enable or disable the Customized SSL Ciphers Flag.
Geo IP Block List	Select a Geo IP block list configuration object. See Using the Geo IP block list .
Geo IP Whitelist	Select a whitelist configuration object. See Using the Geo IP whitelist .
SIP	
SIP Max Size	Maximum message size. The default is 65535 bytes. The valid range is from 1 to 65,535.
Server Keepalive Timeout	Maximum wait for a new server-side request to appear. The default is 30 seconds. The valid range is 5-300.
Server Keepalive	Enable/disable a keepalive period for new server-side requests. Supports CRLF ping-pong for TCP connections. Enabled by default.

Type	Profile Configuration Guidelines
Client Keepalive	Enable/disable a keepalive period for new client-side requests. Supports CRLF ping-pong for TCP connections. Disabled by default.
Client Protocol	Client-side transport protocol: - TCP - UDP (default)
Server Protocol	Server-side transport protocol. - TCP - UDP Default is "unset", so the client-side protocol determines the server-side protocol.
Failed Client Type	Action when the SIP client cannot be reached: - Drop—Drop the connection. - Send—Drop the connection and send a message, for example, a status code and error message.
Failed Server Type	Action when the SIP server cannot be reached: - Drop—Drop the connection. - Send—Drop the connection and send a message, for example, a status code and error message.
Insert Client IP	Enable/disable option to insert the client source IP address into the X-Forwarded-For header of the SIP request.
Client-Request-Header-Insert (maximum 4 members)	
Type	- Insert If Not Exist—Insert before the first header only if the header is not already present. - Insert Always—Insert before the first header even if the header is already present. - Append If Not Exist—Append only if the header is not present. - Append Always—Append after the last header.
HeaderName:Value	The header:value pair to be inserted.
Client-Request-Header-Erase (maximum 4 members)	
Type	- All—Parse all headers for a match. - First—Parse the first header for a match.
HeaderName	Header to be erased.
Client-Response-Header-Insert (maximum 4 members)	
Type	- Insert If Not Exist—Insert before the first header only if the header is not already present. - Insert Always—Insert before the first header even if the header is already present. - Append If Not Exist—Append only if the header is not present. - Append Always—Append after the last header.


Type	Profile Configuration Guidelines
HeaderName:Value	The header:value pair to be inserted.
Client-Response-Header-Erase (maximum 4 members)	
Type	- All—Parse all headers for a match. - First—Parse the first header for a match.
HeaderName	Header to be erased.
Server-Request-Header-Insert (maximum 4 members)	
Type	- Insert If Not Exist—Insert before the first header only if the header is not already present. - Insert Always—Insert before the first header even if the header is already present. - Append If Not Exist—Append only if the header is not present. - Append Always—Append after the last header.
HeaderName:Value	The header:value pair to be inserted.
Server-Request-Header-Erase (maximum 4 members)	
Type	- All—Parse all headers for a match. - First—Parse the first header for a match.
HeaderName	Header to be erased.
Server-Response-Header-Insert (maximum 4 members)	
Type	- Insert If Not Exist—Insert before the first header only if the header is not already present. - Insert Always—Insert before the first header even if the header is already present. - Append If Not Exist—Append only if the header is not present. - Append Always—Append after the last header.
HeaderName:Value	The header:value pair to be inserted.
Server-Response-Header-Erase (maximum 4 members)	
Type	- All—Parse all headers for a match. - First—Parse the first header for a match.
HeaderName	Header to be erased.
SMTP	
Starttls Active Mode	Select one of the following: - Allow—The client can either use or not use the STARTTLS command. - Require—The STARTTLS command must be used to encrypt the connection first. - None—The STARTTLS command is NOT supported.
Forbidden Command	Select any, all, or none of the commands (i.e., expn, turn, vrfy).

Type	Profile Configuration Guidelines
	If selected, the command or commands will be rejected by FortiADC; otherwise, the command or commands will be accepted and forwarded to the back end.
Domain Name	Specify the domain name.
Local Certificate Group	LOCAL_CERT_GROUP.
Certificate Verify	Specify the certificate verify configuration object.
RTMP	
Source Address	When enabled, specify the client address to be used to connect to the server pool.
RTSP	
Max Header Size	Specify the maximum size of the RTSP header.
Source Address	When enabled, specify the client address to be used to connect to the server pool.
MySQL	Note: The system does not provide default MySQL profiles as it does with the other protocols.
Single Master	If selected, the profile will use the single-master mode. You will then need to specify and configure the master server and slave servers.
Sharding	If selected, the profile will use the sharding mode to load-balance MySQL traffic.
DIAMETER	FortiADC comes with a default load-balancing profile titled "LB_PROF_DIAMETER". If it is selected, FortiADC will not change Diameter packets except the host IP address AVP, which means that FortiADC functions as a relay agent.
Identity	Leave blank. If defined, FortiADC will change the Origin-Host AVP of the Diameter packet.
Realm	Leave blank. If defined, FortiADC will change the Origin-Realm AVP of the Diameter packet.
Vendor ID	Leave blank. If defined, FortiADC will change the Vendor-ID AVP of the Diameter packet.
Product Name	Leave blank. If defined, FortiADC will change the Product-Name AVP of the Diameter packet.
Idle Timeout	300 (seconds) by default. Valid values range from 1 to 86,400.
Server Close Propagation	OFF by default, which means that the connection on the client side stays open when the server closes the connection on its side.

WebSocket load-balancing


The WebSocket protocol provides full duplex communication between client and server over a single TCP connection. The initial handshake occurs over the HTTP protocol, while subsequent WebSocket message frames layer over the TCP protocol, as illustrated in [WebSocket load-balancing on page 137](#).

WebSocket load-balancing



You can configure FortiADC in such a way that it is able to load-balance Layer-7 virtual servers with HTTP or HTTPS profiles to the WebSocket protocol without any change to the default configuration. During the setup phase, the virtual server works in HTTP mode, processing Layer-7 information. It automatically detects the connection and upgrade exchange, and is able to switch to tunnel mode when the upgrade negotiation succeeds. When the WebSocket is established, and the virtual server fails over to tunnel mode in which no data is analyzed anymore (and anyway, WebSocket does not communicate in HTTP). See [WebSocket with FortiADC on page 138](#).

WebSocket with FortiADC



If you want to configure your FortiADC appliance to perform HTTP inspection and WebSocket traffic load-balancing, you must use a Layer-7 virtual server with an HTTP profile. If WebSocket traffic is over the transport layer security protocol,

you must use a Layer-7 virtual server with an HTTPS profile and choose an appropriate server SSL profile in the real-server pool.

If you only want WebSocket load-balancing, use a Layer-4 or Layer-7 virtual server with a TCP profile.

For more information, see <https://en.wikipedia.org/wiki/WebSocket> and <http://tools.ietf.org/html/rfc6455>.

Configuring MySQL profiles

FortiADC (Version 4.7.0 and later) supports MySQL server load-balancing.

MySQL application profiles are user-specific and must be configured only by the user on a case by case basis. For this reason, FortiADC does not provide any default predefined MySQL application profiles that you can use out of the box. So you must configure your own MySQL load-balancing application profiles to take advantage of this feature.


FortiADC supports two MySQL database load-balancing modes: single master and data sharding.

Single-master mode

The single-master mode is a database server configuration in which a single master MySQL server is responsible for all write operations (i.e., create, update, or delete requests), and one or more slave servers handle all read-only operations. The master server replicates data to the slave servers in a close to real-time fashion. This mode can improve database performance to a certain extent by offloading read-intensive operations to slave servers. It is ideal for load-balancing database traffic that involves more read operations.

[Single-master mode on page 139](#) illustrates the network topology of database server load-balancing in single-master mode.

Single-master mode




By default, FortiADC passes all write requests to the master server and all read requests (such as select) to the slave servers. So once you have created a MySQL server load-balancing profile, FortiADC will automatically apply this default mode when load-balancing MySQL traffic on the network. However, if you do not like the default behavior, you can change it by setting up your own MySQL server load-balancing rules when configuring your MySQL application profile. For more information, see [Configuring MySQL rules on page 145](#).

Sharding mode

Database sharding is a "shared-nothing" database partitioning technique that breaks down a large database involving a number of database servers into small database chunks and spread them across a number of distributed servers. It's a highly scalable approach to improving the throughput and performance of large enterprise business applications that are transaction-extensive and database-centric because it provides scalability across independent servers, each having its own CPU, memory, and disks.

[Sharding mode on page 141](#) illustrates MySQL server load-balancing in data-sharding mode.

Sharding mode



In sharding mode, FortiADC stores global data on the Master Global—it sends all requests that do not belong to any group to global servers. Using the keys that you have specified, it sends part of the requests to Group 0 and some to Group 1. It supports split read/write in every group.

It must be noted that Data Manipulation Language (DDL) is not supported in sharding mode.

Creating a MySQL profile

Creating a MySQL profile involves the following steps:

1. Create a MySQL configuration object.
2. Specify the existing user name and password of the MySQL database to be used by the MySQL profile configuration object.
3. Configure MySQL Rule (for single-master mode, optional) or MySQL Sharding (for database sharding mode).

Note: You can create MySQL profiles from either the GUI or the CLI. The following paragraphs discuss how to configure a MySQL profile using the GUI. For instructions on how to create MySQL profiles from the CLI, refer to the [FortiADC 4.7.0 CLI Reference](#).

Before you begin:

- You must have already created MySQL database objects to be used the MySQL profile.
- You must have read-write permission for load-balance settings.

Creating a MySQL configuration object

1. Go to Server Load Balance > Application Resources.
2. Select the Application Profile tab if it is not already selected.
3. Click **Create New** to open the Application Profile configuration editor.
4. In the Name field, enter a unique profile name.
5. In the Type field, click the down arrow and select MySQL from the drop-down menu.
6. For MySQL Mode, select Single Master or Sharding. Refer to [MySQL profile configuration guidelines on page 146](#).
7. Click Save. Your newly created MySQL profile configuration object is automatically appended to the bottom of the Server Load Balancing > Application Resources > Application Profile page.
8. Click the newly created MySQL profile to open it. See [MySQL application profile configuration on page 143](#).

MySQL application profile configuration

The screenshot shows the 'Application Profile' configuration page in FortiADC. The profile is named '1' and is of type 'MySQL'. Under the 'Specifics' section, 'MySQL Mode' is set to 'Single Master'. The 'MySQL Account' section is empty, with a 'Create New' button. Below this is a table for MySQL accounts, which is currently empty. The 'MySQL Rule' section is also empty, with a 'Create New' button. At the bottom are 'Save' and 'Cancel' buttons.

Application Profile

Name: 1

Type: MySQL

Specifics

MySQL Mode: ☒ Single Master ☐ Sharding

MySQL Account

+ Add Filter Create New

ID	Username
No data available in table	

Showing 0 to 0 of 0 entries Show 10 entries Previous Next

MySQL Rule

+ Add Filter Create New

ID	Type	Database List	User List	Table List
No data available in table				

Showing 0 to 0 of 0 entries Show 10 entries Previous Next

Save Cancel

Note: The image above shows a sample MySQL profile configuration object named "1". Once a MySQL profile is created, you need to specify the MySQL database user account, and create MySQL Rule or Sharding depending on which MySQL mode you choose to use. The following paragraphs discuss the procedures for each of those tasks.

Specifying the MySQL user account

Once a MySQL profile is created, you must specify a MySQL user account to be used with the profile by entering the user name and password of that account.

It's important to note that you are asked to provide the user name and password of an existing MySQL account. *So do not try to create a new user account here.*

To specify a MySQL user account:

1. In the MySQL User Password pane (see the illustration above), click Create New. The Edit MySQL User Password dialog opens.
2. Enter the user name and password of the MySQL database account,
3. Click Save.

Configuring MySQL rules

When configuring a MySQL rule, you first need to decide whether you want FortiADC to send requests to the Master database server or the Slave database server(s). Then you can set a few conditions (rules) to tell FortiADC how to send the requests. It must be noted that all the conditions are of an "OR" relationship.

To configure a MySQL rule:

1. In the MySQL Rule pane, click Create New. The Application Profile > Edit MySQL Rule dialog opens.
2. Make the desired entries or selections as described in [MySQL profile configuration guidelines on page 146](#).
3. Click Save.

Configuring sharding

FortiADC supports two types of database-sharding: by range or by hash. In the former case, FortiADC distributes the data to different groups according to the key range. In the latter case, it first hashes the keys and then automatically distributes the data to different groups.

To configure MySQL sharding:

1. In the MySQL Sharding pane, click Create New. The Application Profile > Edit MySQL Sharding dialog opens.
2. Make the desired entries or selections as described in [MySQL profile configuration guidelines on page 146](#).
3. Click Save.

Note: When configuring pool members in the CLI to match the real server pool members on the GUI, you can use the `set mysql-group-id` command to set the groups that match the pool members:

```
config load-balance pool
edit "sharding"
set real-server-ssl-profile NONE
config pool_member
edit 1
set pool_member_service_port 3306
set pool_member_cookie rs
set real-server master
next
edit 2
set pool_member_service_port 3306
set pool_member_cookie rs2
set real-server master2
set mysql-group-id 1
next
edit 3
set pool_member_service_port 3306
set pool_member_cookie rs3
set real-server slave
```

```

set mysql-read-only enable
next
edit 4
set pool_member_service_port 3306
set pool_member_cookie rs4
set real-server slave2
set mysql-read-only enable
set mysql-group-id 1
next
end
next
end

```



You can clone a predefined configuration object to help you get started with a user-defined configuration.

To clone a configuration object, click the clone icon  that appears in the tools column on the configuration summary page.

MySQL profile configuration guidelines

Parameter	Description
Application Profile	
Name	A unique name for the MySQL profile you are creating.
Type	MySQL
MySQL Mode	Select either of the following: Single Master—If selected, FortiADC will configure the MySQL profile in single-master mode. See Single-master mode . Sharding—If selected, FortiADC will configure the MySQL profile in database-sharding mode. See Sharding mode .
MySQL User Password	
User Name	The user name of the MySQL database.
Password	The password for the MySQL user name you've entered above.
MySQL Rule	
Type	Select either of the following: - Master—If selected, FortiADC will send all data specified in the MySQL rule to the master MySQL database server. - Slave—If selected, FortiADC will send all data specified in the MySQL rule to the slave MySQL database server.
Database List	A list of up to eight MySQL database names separated by space

Parameter	Description
User List	A list of up to eight user names separated by space
Table List	A list of up to eight MySQL Database tables separated by space
Client IP List	A list of up to eight FortiADC client IP addresses separated by space
SQL List	A list of up to eight MySQL statements separated by space
Sharding	
Type	Select either of the following: - Range—If selected, FortiADC will send data in the data tables to different groups based on the specified range of the keys. - Hash—If selected, FortiADC will perform hash calculations and then automatically send data to different groups.
Database	The database name
Table	The table name
Key	The column name
Group List	A list of up to eight group IDs Note: The group IDs must match the real server pool members.

Configuring client SSL profiles

A client SSL profile is used to manage the SSL session between the client and the proxy. It allows FortiADC to accept and terminate client requests sent via the SSL protocol. The Client SSL Profile page provides the settings for configuring client-side SSL connections, and displays all the client SSL profiles that have been configured on the system.

Before you begin creating a client SSL profile:

- You must have already created configuration objects for certificates, certificate caching, and certificate verify if you want to include them in the profile.
- You must have Read-Write permission for Load Balance settings.

To configure custom profiles:

- Go to Server Load Balance > Application Resources. Click the Client SSL Profile tab.
- Click **Create New** to display the configuration editor.
- Complete the configuration as described in [Client SSL profile configuration guidelines on page 148](#).
- Save the configuration.



You can clone a predefined client SSL profile to help you get started with a user-defined configuration.

To clone a configuration object, click the clone icon  that appears in the tools column on the configuration summary page.

Client SSL profile configuration guidelines

Type	Profile Configuration Guidelines
Name	Specify a unique name for the client SSL profile.
Customized SSL Ciphers Flag	Enable or disable the use of user-specified cipher suites. If enabled, you must specify a colon-separated, ordered list of a customized SSL cipher suites. See below.
Customized SSL Ciphers	Available only when the Customized SSL Cipher Flag is enabled (see above). Specify a colon-separated, ordered list of a customized SSL cipher suites. Note: FortiADC will use the default SSL cipher suite if the field is left empty.
SSL Ciphers	Ciphers are listed from strongest to weakest: • ECDHE-ECDSA-AES256-GCM-SHA384 • ECDHE-ECDSA-AES256-SHA384 • ECDHE-ECDSA-AES256-SHA • ECDHE-ECDSA-AES128-GCM-SHA256 • ECDHE-ECDSA-AES128-SHA256 • ECDHE-ECDSA-AES128-SHA • ECDHE-ECDSA-DES-CBC3-SHA • ECDHE-ECDSA-RC4-SHA • ECDHE-RSA-AES256-GCM-SHA384 • ECDHE-RSA-AES256-SHA384 • ECDHE-RSA-AES256-SHA • DHE-RSA-AES256-GCM-SHA384 • DHE-RSA-AES256-SHA256 • DHE-RSA-AES256-SHA • AES256-GCM-SHA384 • AES256-SHA256 • AES256-SHA • ECDHE-RSA-AES128-GCM-SHA256 • ECDHE-RSA-AES128-SHA256 • ECDHE-RSA-AES128-SHA • DHE-RSA-AES128-GCM-SHA256 • DHE-RSA-AES128-SHA256 • DHE-RSA-AES128-SHA • AES128-GCM-SHA256 • AES128-SHA256 • AES128-SHA

Type	Profile Configuration Guidelines
	• ECDHE-RSA-RC4-SHA • RC4-SHA • RC4-MD5 • ECDHE-RSA-DES-CBC3-SHA • EDH-RSA-DES-CBC3-SHA • DES-CBC3-SHA • eNULL Note: We recommend retaining the default list. If necessary, you can deselect the SSL ciphers that you do not want to support.
Allowed SSL Versions	You have the following options: • SSLv3 • TLSv1.0 • TLSv1.1 • TLSv1.2 • TLSv1.3 We recommend retaining the default list. If necessary, you can deselect SSL versions you do not want to support. Note: FortiADC does not support session reuse for SSLv2 at the client side. Instead, a new SSL session is started.
Client Certificate Verify	Select the client certificate verify configuration object.
Client Certificate Forward	Disabled by default. When enabled, you must specify the client certificate forward header. See below.
Client Certificate Forward Header	When Client Certificate Forward is enabled (see above), specify the client certificate forward header.
Forward Proxy	By default, (SSL) Forward Proxy is disabled. When enabled, you'll have to configure additional settings noted below.
Client SNI Required	Require clients to use the TLS server name indication (SNI) extension to include the server hostname in the TLS client hello message. Then, the FortiADC system can select the appropriate local server certificate to present to the client.
Local Certificate Group	Select a local certificate group that includes the certificates this virtual server presents to SSL/TLS clients. This should be the backend servers' certificate, NOT the appliance's GUI web server certificate. See Manage certificates .
Reject OCSP Stapling with Missing Nextupdate	This flag is meaningful only when you have configured OCSP stapling in Local Certificate Group. By default, this option is disabled (unselected). In that case, FortiADC accepts all OCSP responses, including those in which the next update field is not set. If enabled, and the next update field is not set in an OCSP stapling response, FortiADC will not load this OCSP stapling response or present it to clients during the SSL/TLS handshake.
Renegotiation	Enable or disable SSL renegotiation from the client side. Note:

Type	Profile Configuration Guidelines
Renegotiation Interval	- The feature is disabled by default. - When enabled, you must configure the options below. Specify the minimum interval between two successive client-initiated SSL renegotiation requests. The unit of measurement can be second, minute, or hour, e.g., 100s, 20m, or 1h. Note: - The default is -1, which disables the function. 0 means 'Indefinite'. - FortiADC will terminate the connection once the threshold is exceeded.
SSL DH Parameter Size	Specify the pubkey length in Diffie Hellman. Default is 1024.
SSL Renegotiate Period	Specify the period in second (default), minute, or hour at which FortiADC will initiate SSL renegotiation. Note: The default is 0, which disables the function.
SSL Renegotiate Size	Specify the amount (MB) of application data that must have been transmitted over the SSL connection when FortiADC initiates SSL renegotiation. Note: The default is 0, which disables the function.
Secure Renegotiation	Select one of the following: - Request—FortiADC requests secure renegotiation of SSL connections. - Require—(Default) Specifies that FortiADC requires secure renegotiation of SSL connections. In this mode, FortiADC permits initial SSL handshakes from clients, but terminates renegotiation requests from clients that do not support secure renegotiation. - Require Strict—FortiADC requires strict secure renegotiation of SSL connections. In this mode, FortiADC denies initial SSL handshakes from clients that do not support secure renegotiation.
Dynamic record sizing	Allows ADC to dynamically adjust the size of TLS records based on the state of the connection, in order to prevent bottlenecks caused by the buffering of TLS record fragments. Note: The feature is disabled by default.
Note: The following fields become available only when Forward Proxy is enabled.	
Forward Proxy Certificate Caching	Select a Forward Proxy Certificate Caching rule.
Forward Proxy Local Signing CA	Select a Forward Proxy Local Signing CA.
Forward Proxy Intermediate CA Group	Select a Forward Proxy Intermediate CA Group.
Backend SSL SNI Forward	Disabled by default. Enable it to let FortiADC forward Server Name Indication (SNI) from the client to the back end.
Backend Customized SSL Ciphers Flag	Enabled by default. In this case, you must specify the backend customized SS ciphers. See below.

Type	Profile Configuration Guidelines
Backend Customized SSL Ciphers	Specify the customized SSL ciphers to be supported at the back end.
Backend Allowed SSL Versions	We recommend retaining the default list. If necessary, you can deselect SSL versions you do not want to support. Note: FortiADC does not support session reuse for SSLv2 at the client side. Instead, a new SSL session is started.

Configuring HTTP2 profiles

You can now create application profiles that support HTTP2. To do so, you must first create an HTTP2 Profile, then use that profile when creating a new application profile.

To configure HTTP2 profiles:

1. Go to Server Load Balance > Application Resources. Click the HTTP2 Profile tab.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in [HTTP2 profile configuration guidelines on page 151](#).
4. Save the configuration.

HTTP2 profile configuration guidelines

Type	Profile Configuration Guidelines
Name	Specify a unique name for the HTTP2 profile.
Priority Mode	Set to Best Effort. Not configurable.
Upgrade Mode	Set to Upgradeable. Not configurable.
Max Concurrent Stream	Specify the maximum number of concurrent streams available at one time. The default number is 5.
Max Receive Window	Specify the maximum number of bytes that can be received without sending an acknowledgment response. The default is 65535 bytes.
Max Frame Size	Specify the max size of the data frames, in bytes that the HTTP2 protocol sends to the client. Setting a large frame size improves network utilization, but it can also affect concurrency. The default is 16384 bytes.
Header Table Size	Specify the size of the header table, in KB. A larger table size allows for better HTTP header compression, but it requires more memory. The default is 4096.
Header List Limitation	Specify the size of the name value length , in bytes, that the HTTP2 protocol sends in a single header frame. The default is 65536.
SSL Constraint	Enable or disable SSL constraint. If enabled, the following conditions must be met: • The TLS implementation supports Server Name Indication.

Type	Profile Configuration Guidelines
	• The TLS implementation disables compression. • The TLS implementation disables renegotiation. • Renegotiation takes place before the connection preface is sent. • HTTP/2 uses cipher suites with ephemeral key exchange. • Ephemeral key exchange has a size of at least 2048 bits (for DHE) or a security level of at least 128 bits (for ECDHE). • Clients accept DHE no smaller than 4096 bits. • Stream or block ciphers are not used with HTTP.

Configuring load-balancing (LB) methods

The system includes predefined configuration objects for all supported load balancing methods, and there is no need to create additional configuration objects. You may choose to do so, however, for various reasons, for example, to use a naming convention that makes the purpose of the configuration clear to other administrators.

[Predefined LB methods on page 152](#) describes the predefined methods.

Predefined LB methods

Predefined	Description
LB_METHOD_ROUND_ROBIN	Selects the next server in the series: server 1, then server 2, then server 3, and so on.
LB_METHOD_LEAST_CONNECTION	Selects the server with the least connections.
LB_METHOD_FASTEST_RESPONSE	Selects the server with the fastest response to health check tests.
LB_METHOD_URI	Selects the server based on a hash of the URI found in the HTTP header, excluding hostname.
LB_METHOD_FULL_URI	Selects the server based on a hash of the full URI string found in the HTTP header. The full URI string includes the hostname and path.
LB_METHOD_HOST	Selects the server based on a hash of the hostname in the HTTP Request header Host field.
LB_METHOD_HOST_DOMAIN	Selects the server based on a hash of the domain name in the HTTP Request header Host field.
LB_METHOD_DEST_IP_HASH	Selects the next hop based on a hash of the destination IP address. This method can be used with the Layer 2 virtual server.
LB_METHOD_DYNAMIC_LOAD	Selects the server with the highest weight assigned to it based on its SNMP health check.

Predefined	Description
	Note: Dynamic load-balancing is a load-balancing method in which FortiADC (the load-balancer) actively polls server pool members, and then assigns a weighted value to each member based on a set of default or user-defined thresholds. The value ranges from 1 to 256, and determines the amount of traffic FortiADC directs to a member. The greater the value that FortiADC assigns to a member, the more client requests it (the member) receives. Dynamic load-balancing relies on the status of SNMP health check to calculate the load on each real server. The health check covers a real server's CPU, memory, and disk usage. When a real server has exceeded its health check thresholds, it will be marked as "down". If that happens, FortiADC will stop sending client requests to that server.

Before you begin:

- You must have Read-Write permission for Load Balance settings.

To configure a load-balancing method configuration object:

1. Go to Server Load Balance > Virtual Server > Application Resources.
2. Click the LB **Method** tab.
3. Click **Create New** to display the configuration editor.
4. Give configuration object a name and select the load-balancing type.
5. Save the configuration.

Configuring persistence rules

Persistence rules identify traffic that should not be load balanced, but instead forwarded to the same backend server that has seen requests from that source before. Typically, you configure persistence rules to support server transactions that depend on an established client-server session, like e-commerce transactions or SIP voice calls.

The system maintains persistence session tables to map client traffic to backend servers based on the session attribute specified by the persistence rule.

The persistence table is evaluated before load balancing rules. If the packets received by the ADC match an entry in the persistence session table, the packets are forwarded to the server that established the connection, and load balancing rules are not applicable.

Most persistence rule types have a timeout. When the time that has elapsed since the system last received a request from the client IP address is greater than the timeout, the system does not use the mapping table to forward the request. Instead, it again selects the server using the method specified in the virtual server configuration. Hash-based rule types have a timeout built into the hash algorithm. For other types, you can specify the timeout.

[Predefined persistence rules on page 154](#) describes the predefined persistence rules. You can get started with these commonly used persistence methods or create custom objects.

Predefined persistence rules

Predefined	Description
LB_PERSIS_SIP	Persistence based on source IP address or subnet.
LB_PERSIS_CONSISTENT_SIP	Persistence based on a hash of source IP address.
LB_PERSIS_HASH_SRC_ADDR_PORT	Persistence based on a hash that includes source IP address and port.
LB_PERSIS_HASH_COOKIE	Persistence based on a hash of a session cookie provided by the backend server.
LB_PERSIS_RDP_COOKIE	Persistence based on RDP cookie sent by RDP clients in the initial connection request.
LB_PERSIS_SSL_SESS_ID	Persistence based on the SSL session ID.
LB_PERSIS_SIP_CALL_ID	Persistence based on the SIP call ID.
LB_PERSIS_PASSIVE_COOKIE	Persistence based on a passive cookie generated by the server. FortiADC does not generate or manage the cookie, but only observes it in the HTTP stream, thus the name "passive cookie". Also known as "server cookie".

Before you begin:

- You must have a good understanding and knowledge of the applications that require persistent sessions and the methods that can be used to identify application sessions.
- You must have Read-Write permission for Load Balance settings.

After you have configured a persistence rule, you can select it in the virtual server configuration.

To configure a persistence rule:

1. Go to Server Load Balance > Application Resources.
2. Click the **Persistence** tab.
3. Click **Create New** to display the configuration editor.
4. Give the rule a name, select the type, and specify rule settings as described in [Persistence rule guidelines on page 155](#).
5. Save the configuration.



You can clone a predefined configuration object to help you get started with a user-defined configuration.

To clone a configuration object, click the clone icon  that appears in the tools column on the configuration summary page.

Persistence rule guidelines

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the virtual server configuration. Note: After you initially save the configuration, you cannot edit the name.
Type	Select a persistence type.
Source Address	
Source Address	Persistence is based on source IP address.
Timeout	Timeout for an inactive persistence session table entry. The default is 300 seconds. The valid range is 1-86,400.
Subnet Mask Bits (IPv4)	Number of bits in a subnet mask to specify a network segment that should follow the persistence rule. For example, if IPv4 maskbits is set to 24, and the backend server A responds to a client with the source IP 192.168.1.100, server A also responds to all clients from subnet 192.168.1.0/24.
Subnet Mask Bits (IPv6)	Number of bits in a subnet mask to specify a network segment that should follow the persistence rule.
Match across servers	Enable so clients continue to access the same backend server through different virtual servers for the duration of a session. For example, a client session with a vSphere 6.0 Platform Services Controller (PSC) has connections on the following ports: 443, 389, 636, 2012, 2014, 2020. A FortiADC deployment to load balance a cluster of vSphere PSCs includes Layer 4 virtual server configurations for each of these ports. To ensure a client's connections for a session go to the same backend real server: 1. Create a persistence object based on Source Address affinity and select the Match Across Servers option. 2. Select this persistence object in each of the Layer 4 virtual servers configured to load balance the vSphere PSC pool. 3. Select the same real server pool object in each of the Layer 4 virtual servers configured to load balance the vSphere PSC pool. When these options are enabled, FortiADC dispatches the initial connection to a real server destination (for example, RS1) based on the virtual server's load balancing method, and the persistence object is noted in the connection table. Subsequent connection attempts with the same source IP address to any FortiADC virtual server that has this persistence object and real server pool are dispatched to RS1, as long as the session is active. Note: In the Layer 4 virtual server configuration, you specify a packet forwarding method. You can use Source Address persistence with Match Across Servers with any combination of Direct Routing, DNAT, and Full NAT packet forwarding methods. However, with NAT46 and NAT64 packet forwarding methods, the source address type is different from the real server address type. To use Match Across Servers with NAT46 or NAT64, all virtual servers for the application must be configured with the same packet forwarding method: all NAT46 or all NAT64.

Settings	Guidelines
Source Address Hash	
Source Address Hash	Persistence is based on a hash of the IP address of the client making an initial request.
Source Address-Port Hash	
Source Address-Port Hash	Persistence is based on a hash of the IP address and port of an initial client request.
HTTP Header Hash	
HTTP Header Hash	Persistence is based on a hash of the specified header value found in an initial client request.
Keyword	A value found in an HTTP header.
HTTP Request Hash	
HTTP Request Hash	Persistence is based on a hash of the specified URL parameter in an initial client request.
Keyword	A URL parameter.
Cookie Hash	
Cookie Hash	Persistence is based on a hash of the cookie provided by the backend server.
Persistent Cookie	
Persistent Cookie	Persistence is based on the cookie provided in the backend server response. It forwards subsequent requests with this cookie to the original backend server.
Keyword	Backend server cookie name.
Timeout	Timeout for an inactive persistence session table entry. The default is 300 seconds. The valid range is 1-86,400.
Insert Cookie	
Insert Cookie	Persistence is based on a cookie inserted by the FortiADC system. The system inserts a cookie whose name is the value specified by Keyword and whose value is the real server pool member Cookie value and expiration date (if the client does not already have a cookie). For example, if the value of Keyword is <code>sessid</code> and the real server pool member Cookie value is <code>rs1</code>, FortiADC sends the cookie <code>sessid=rs1 U6iFN</code> to the client, where <code>U6iFN</code> is the expiration date as a base64 encoded string.
Keyword	Specifies the cookie name.
Timeout	Timeout for an inactive persistence session table entry. The default is 300 seconds. The valid range is 1-86,400.
Rewrite cookie	
Rewrite Cookie	Persistence is based on the cookie provided in the backend server response, but the system rewrites the cookie.

Settings	Guidelines
	The system checks the HTTP response for a <code>Set-Cookie:</code> value that matches the value specified by Keyword. It replaces the keyword value with the real server pool member Cookie value. For example, the value of Keyword in the persistence configuration is <code>sessid</code>. The real server pool member Cookie value is <code>rs1</code>. After an initial client request, the response from the server contains <code>Set-Cookie: sessid=666</code>, which FortiADC changes to <code>Set-Cookie: sessid=rs1</code>. FortiADC uses this rewritten value to forward subsequent requests to the same backend server as the original request.
Keyword	Specifies a <code>Set-Cookie:</code> value to match.
Embedded Cookie	
Embedded Cookie	Persistence is based on the cookie provided in the backend server response. Like Rewrite Cookie, the system checks the HTTP response for a <code>Set-Cookie:</code> value that matches the value specified by Keyword in the persistence configuration. However, it preserves the original value and adds the real server pool member Cookie value and a ~ (tilde) as a prefix. For example, the value of Keyword is <code>sessid</code>. The real server pool member Cookie value is <code>rs1</code>. After an initial client request, the response from the server contains <code>Set-Cookie: sessid=666</code>, which the system changes to <code>Set-Cookie: sessid=rs1~666</code>. It uses this rewritten value to forward subsequent requests to the same backend server as the original request.
Keyword	Specifies a <code>Set-Cookie:</code> value to match.
RADIUS Attribute	
Type	Select RADIUS Attribute.
Timeout	Specify the timeout for an inactive persistence session table entry. The default is 300 seconds, and valid values range from 1 to 86,400.
Match Across Virtual Servers	OFF (disabled) by default. Click the button to enable it. If enabled, clients will continue to access the same backend server through different virtual servers for the duration of a session.
Override Connection Limit	OFF (disabled) by default, which means that when the connection limit is reached, new connections will still be persistently forwarded to the real server. If enabled, new connections will be forwarded to another node (load-balancing) until all nodes are full.
RADIUS Attribute Relation	RADIUS persistence rule supports multiple RADIUS settings, which can be either of the following relations: • AND (Default) — The persistence condition is true if all RADIUS attributes are found. • OR—The persistence condition is true if any of the attributes is found.

Settings	Guidelines
RADIUS Attribute	After you have saved the RADIUS-type persistence configuration object, you can open the Persistence configuration editor and add up to four (4) RADIUS attributes to it. Note: If you choose to use the 26-Vendor-Specific attribute, you need to specify the Vendor ID and Vendor Type.
RDP Cookie	
RDP Cookie	Persistence based on RDP cookie sent by RDP clients in the initial connection request.
SSL Session ID	
SSL Session ID	Persistence is based on SSL session ID.
Timeout	Timeout for an inactive persistence session table entry. The default is 300 seconds. The valid range is 1-86,400.
SIP Call ID	
SIP Call ID	Persistence is based on SIP Call ID. For SIP services, you can establish persistence using Source Address, Source Address Hash, or SIP caller ID.
Timeout	Timeout for an inactive persistence session table entry. The default is 300 seconds. The valid range is 1-86,400.

Configuring error pages

When backend real servers are unavailable, FortiADC can respond to clients attempting HTTP/HTTPS connections with an HTML error page. Once you create an HTML error page, you can select it in virtual server configurations.

You do not have to create an HTML error page if you want to simply send a basic text error message when backend servers are unavailable. Instead, you can enter an error message in a text box from within the virtual server configuration. See [Error Message on page 82](#).

Before you begin:

- You must have Read-Write permission for Server Load Balance settings.
- Copy the error message file to a location you can reach from your browser; the error message file must be named `index.html` and contained in a tar, tar.gz, or zip file with a maximum size of 1MB.

To upload an error message file:

1. Go to Server Load Balance > Application Resources.
2. Click the **Error Page** tab.
3. Click **Create New** to display the configuration editor.
4. Enter the name of the error page. You will use this name to select the error page in virtual server configurations. No spaces.
5. Click **Choose File** and browse and select the error message tar, tar.gz, or zip file of up to 1MB in size.
6. Save the configuration.



It is possible to modify error pages that you have already created. To do so, double-click the error page or select the  (edit) icon in the row of the error page that you want to modify, and upload a new error message tar, tar.gz, or zip file as above.

Note: While it is possible to modify the error message file, once an error page is created, you cannot modify its name.

Configuring decompression rules

If the HTTP/HTTPS request body is compressed, FortiADC cannot pass it to the other functional modules which perform inspection or modification.

To allow FortiADC to pass compressed HTTP/HTTPS client requests to other modules for inspection or modification before forwarding it to the back-end server, you must create a FortiADC decompression policy.

You can configure FortiADC to temporarily decompress the body of a request based on its file type, which can be specified by the HTTP/HTTPS Content-Type: header. The appliance can then inspect or modify the traffic. If no inspection or modification is needed, it will allow the compressed version of the request to pass to the back-end server.

FortiADC supports HTTP/HTTPS request decompression in either gzip or deflate format. Upon receiving a compressed HTTP/HTTPS request body, FortiADC first extracts the HTTP/HTTPS request body to a temporary buffer and then sends the buffer to the other modules.

Note that, for the current release, decompression only works for Web Application Firewall (WAF) and Scripting functions.

FortiADC supports decompression of the following content-type files:

- application/javascript
- application/soap+xml
- application/x-javascript
- application/xml
- text/css
- text/html
- text/javascript
- text/plain
- text/xml
- custom

Before you begin:

- You must have a good understanding of HTTP decompression and knowledge of the content types served from the backend real servers.
- You must have Read-Write permission for Load Balance settings.

Decompression is not enabled by default. After you have configured a decompression rule, you can select it in the profile configuration. To enable decompression, select the profile when you configure the virtual server.

To configure a decompression rule:

1. Click Server Load Balance > Application Resources.
2. Click the **Decompression** tab.

3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Decompression configuration on page 160](#).
5. Save the configuration.

Decompression configuration

Settings	Guidelines
Name	Specify a unique name for the decompression rule. Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the profile configuration. Note: After you initially save the configuration, you cannot edit the name.
URI List Type	• Include— Select this option to create a decompression inclusion rule. HTTP/HTTPS responses that match the URIs and content types specified in this rule will be decompressed by FortiADC before being passed to the client. • Exclude—Select this option to create a decompression exclusion rule. HTTP/HTTPS responses that match the URIs and content types specified in this rule will not be decompressed by FortiADC before being passed to the client.
URI List	Click Add and specify URIs to build the list.
Content Types	Click Add and select from the following content types to build the list: • application/javascript • application/soap+xml • application/x-javascript • application/xml • text/css • text/html • text/javascript • text/plain • text/xml • custom Note: The "custom" option allows you to specify almost any content/media type, including image files in .JPG, .PNG, and .BMP formats. The default is */*, which means any content/media type.

You can use the CLI to configure decompression rules:

```
config load-balance decompression
edit <name>
set cpu-limit {enable | disable}
set max-cpu-usage [1-100]
set uri-list-type {include | exclude}
config uri_list
edit <ID>
set uri <refex_pattern>
next
end
config content-types
edit <ID>
set content-type <types>
{
  application/javascript
  application/soap+xml
  application/x-javascript
  application/xml
  custom <plain-string>
  text/css
  text/html
  text/javascript
  text/plain
  text/xml
}
next
end
```



You can use the CLI to select a decompression rule in a server load balance profile (HTTP):

```
config load-balance profile
edit <name>
...
set decompression <decompression name>
...
next
end
```

Using decompression with script data body manipulation

Script data body manipulation can work in tandem with compression or decompression rules in a rather transparent way. When a decompression rule is configured and used with scripting, FortiADC will decompress HTTP data first, then apply script data body manipulation, and then re-compress the data before sending it to clients.

So, if HTTP data is compressed before being sent out from the real server, you must create a decompression rule if you want to access the original data and use it in a script. This can be done either via the GUI or the Console. The following paragraphs show you the basic steps for configuring decompression rules to work with script data body manipulation.

From the GUI

Step 1: Creating a decompression rule

1. Click Server Load Balance > Application Resources > Decompression.
2. Click Create New to open the Decompression configuration dialog.
3. For Name, specify a unique name for the decompression rule.
4. For URI Rule Type, select Include or Exclude.
5. Click Save. The dialog closes and the decompression rule appears in the Decompression table.
6. Double-click the decompression rule (or click the corresponding Edit button) to open it.
7. In the URI Rule section, make the desired configuration. (Optional)
8. In the Content Types sections, make the desired configuration. (Optional)
9. Click Save.
10. Repeat the above steps to create as many decompression rules as needed.

Step 2: Configuring a load balance profile

1. Click Server Load Balance > Application Resources > Application Profile.
2. Click **Create New** to open the Application Profile configuration dialog.
3. For Type, click the down arrow and select HTTP or HTTPS from the list menu.
4. For Decompression, click the down arrow and select a decompression rule from the list menu.
5. Complete all the other fields required for load-balancing profile configuration.
6. Click Save.

Step 3: Enabling scripting in virtual server configuration

1. Click Server Load Balance > Virtual Server > Virtual Server.
2. Click Add > Advanced Mode.
3. For Type (under the Basic section), be sure to select Layer 7.
4. For Profile (under the General section), be sure to select an HTTP or HTTPS profile associated with the decompression rules that you have configured.
5. For Scripting, be sure to turn it on (enable it), and then select the desired script or scripts.
6. Complete all the other fields required for virtual server configuration.
7. Click Save.

From the Console

Use the following example commands as a reference when configuring decompression and script data body manipulation from the Console.

Step 1: Creating a decompression rule

```
config load-balance decompression
edit "decompress"
set uri-list-type include
config uri_list
edit 1
```

```
set uri /
next
end
config content_types
edit 1
set content-type text/html
next
end
next
end
```

Step 2: Configuring a load balance profile

```
config load-balance profile
edit "http"
set type http
set decompression decompress
next
end
```

Step 3: Enabling scripting in virtual server configuration

```
config load-balance virtual-server
edit "vs"
set load-balance-profile http
set scripting-flag enable
set scripting-list data
next
end
```

Creating a PageSpeed configuration

A PageSpeed configuration sets the rule(s) that FortiADC follows when rendering web pages. Creating a PageSpeed configuration object involves the following:

- Specify the inode/file cache limits
- Choose a PageSpeed profile (Must be configured in advance)
- Set page control
- Set resource control

To create a PageSpeed configuration object:

1. Click Server Load balance > Application Optimization.
2. Select the Page Speed tab.
3. Make the entries or selections as described in [PageSpeed configuration on page 163](#).
4. Click Save when done.

PageSpeed configuration

Parameter	Description
PageSpeed	

Parameter	Description
Name	Enter a name for the PageSpeed configuration object that you are creating.
File Cache Inode Limit	Specify the maximum number of inodes that can be cached on FortiADC for this virtual server. The default is 10,000. Valid values range from 1 to 100,000. Note: An <i>inode</i> is a data structure with information about files or directories on a filesystem on Linux or other Unix-type operating systems. It's generated when a filesystem is created. Within a filesystem, every file and directory has a corresponding inode identified by an inode number. Each inode contains the attributes and disk block location(s) of the file's or directory's data, which may include metadata (e.g., access mode, times of last change, modification) and user, ownership, and permission data. A filesystem has a set number of inodes, which indicates the maximum number of files or directories it can hold. A FortiADC appliance can support up to 100,000 inodes. Every time you open a file, the kernel of the server reads the file's inode. The more files and directories you have, the more inodes the server uses. And the more inodes the server uses, the more system resources it consumes. So it is always a good practice to try to limit the number of inodes a host has on a shared server. This will prevent it from using all system resources. To ensure efficient use of its resources, FortiADC cleans its cache every 10 minutes. It cleans the cache only when either of the following conditions is met: • The virtual server has reached its set inode cache limit. • The virtual server has reached its file size cache limit. When performing cache clean-up, FortiADC will use the "first-in first-out" (FIFO) principle to remove the oldest cached inodes or files until the cached data is reduced to less than 75% of its set inode- or file-cache limit(s).
File Cache Size Limit	Specify the maximum file size that can be cached on FortiADC for this virtual server. The default is 128. Value values range from 1 to 512 (MB).
PageSpeed Profile	Select a PageSpeed profile from the list menu. Note: You must have PageSpeed profiles created before you start to create a PageSpeed rule. For instructions on how to create a PageSpeed profile, refer to Creating PageSpeed profiles on page 165
Page Control	
Type	Select either of the following page control types: • Include — If selected, FortiADC will process Web pages associated with the URI specified below. • Exclude — If selected, FortiADC will skip Web pages associated with the URI specified below.
URI Pattern	Specify the full URI in regular expression. For example, <pre>(http(s)://)*example.com/*/htmls/*.html</pre> Note: In the HTTP response body, HTML sometimes is linked to a certain resource URL. If the resource contains a domain name, then FortiADC will do the fetch according to the fetch-domain setting or the rewrite-domain setting.

Parameter	Description
	Wildcards include * (asterisk) which matches any 0 (zero) or more characters, and ? (question mark) which matches exactly one character. Unlike Unix shells, the / directory separator is not special, and can be matched by either * or ?. The resources are always expanded into their absolute form before expanding. A wildcard will be matched against the full URL, including any query parameters. For example, you can use <code>"*.jsp"</code> to match <code>http://example.com/index.jsp?test=xyz</code>.
Resource Control	
Origin Domain Patten	Specify the original domain pattern in regular expression in alphanumeric characters. For example, <code>(http(s)://)*.example.com</code> Note: Valid characters are 0–9, a–z, A–Z, . (period), : (colon), hyphen (-) and / (forward slash). The FortiADC 4.8.0 release only supports HTTP or HTTPS. To improve web page performance, PageSpeed will examine and modify the content of the resources referenced on web pages. It does that by fetching those resources using HTTP, according to the URL reference specified on an HTML page.
Rewrite Domain	Specify the fetch domain string. For example, <code>http://www.example.com</code> Valid characters are 0–9, a–z, A–Z, . (period), : (colon), hyphen (-) and / (forward slash). The FortiADC 4.8.0 release only supports HTTP or HTTPS.
Fetch Domain	Specify the rewrite domain string. For example, <code>http://www.example.com</code> Valid characters are 0–9, a–z, A–Z, . (period), : (colon), hyphen (-) and / (forward slash). The FortiADC 4.8.0 release only supports HTTP or HTTPS.

Creating PageSpeed profiles

PageSpeed provides a technology solution to speed up web application response and optimize web pages and resources in real time.

As a module on FortiADC device, PageSpeed is simple to deploy and does not require any integration into Web application servers or any client installation on end-user devices. With the PageSpeed feature, you can select the approach(es) to make your web site faster and more user-friendly.

A PageSpeed profile specifies the option(s) for optimizing the delivery of web applications. To take full advantage of the benefits that PageSpeed offers, you must first create your own PageSpeed profiles and then select the application optimization option(s) to add to them. Once you have your own PageSpeed profiles created, you can simply select them to include in any PageSpeed configurations you create.

FortiADC offers options for optimizing the delivery of the following web content:

- HTML
- CSS
- Image

For more information and instructions on how to use these options, see Table 1.

To create a PageSpeed profile:

1. Click Server Load balance > Application Optimization.
2. Select the Page Speed Profile tab.
3. Make the entries or selections as described in [Application optimization parameters on page 166](#).
4. Click Save when done.

Application optimization parameters

Parameter	Description
HTML	Disable (default) or enable HTML optimization. If enabled, you must also select a specific option(s) below. Note: FortiADC supports optimization of compressed HTML files.
Move CSS to Head	If selected, FortiADC will move CSS elements above script tags. Note: This ensures that the CSS styles are parsed in the head of the HTML page before any body elements are introduced. In so doing, it can effectively reduce the number of times web browsers have to re-flow HTML documents.
CSS	Disable (default)/enable CSS optimization. Note: If enabled, you must also select the specific option(s) below.
Combine CSS	If selected, FortiADC will combine multiple CSS elements into one. Note: This option replaces multiple CSS files with a combined CSS file that contains the contents of all individual CSS files. As a result, it can reduce the number of HTTP/HTTPS requests web browsers make during page refresh. This is particularly beneficial to older browsers that can handle only up to two connections per domain. Not only can this reduce the overhead for HTTP/HTTPS headers and communications warm-up, but also work well with TCP/IP slow-start because it increases the effective payload bit rate through a browser's network connection.
Maxi Combine CSS Byte	Specify the maximum number of CSS bytes that can be combined. The default is 4,096. Note: Valid values range from 1 to 10,240.
Image	Disable (default)/enable image optimization. Note: If enabled, you must also select the specific option(s) below.
Resize Image	Disabled by default. If enabled, this will reduce the dimension of an image to the "width=" and "height=" attributes defined in the tag or in the inline "style=attribute". Note: • The option will remove color profile and metadata. • The re-sized image may also be re-compressed or converted to a new format and quality based on user configuration. • This option applies to .jpg, .png, and .webp images only.
JPEG Sampling	Disabled by default. When enabled, it will apply 4:2:0 chroma subsampling to .jpg images, in which hue and saturation have only 25% as many samples as brightness. Because the human eye is less sensitive to hue and saturation than to brightness, this subsampling technique can greatly reduce image size with no noticeable effect on perception

PageSpeed support and restrictions

Implementation of PageSpeed is subject to the following conditions or restrictions.

Supported

PageSpeed is supported in the following use scenarios:

- Layer-7 server load balancing HTTP
- Layer-7 server load balancing HTTPS

Restrictions

Support for Layer-7 server load balancing HTTP/HTTPS is subject to the following conditions:

- Content-type must be text or html
- Data without compression

Not Supported

The following are not supported:

- Too many virtual servers using PageSpeed at the same time
- HTTP/2
- File cache sync for high availability (HA)

Note: Although it is possible to create more than 16 virtual machines with PageSpeed, you must do it with careful consideration. This is because virtual machines with PageSpeed consume more system memory, and your FortiADC appliance could quickly run out of memory as a result.

Configuring compression rules

To offload compression from your back-end servers, you can configure FortiADC to perform HTTP/HTTPS compression on behalf of the server.

The following content types can be compressed:

- application/javascript
- application/soap+xml
- application/x-javascript
- application/xml
- text/css
- text/html
- text/javascript
- text/plain

- text/xml
- custom

Not all HTTP/HTTPS responses should be compressed. Compression offers the greatest performance improvements when applied to URIs whose media types include repetitive text such as tagged HTML and JavaScript. Files that already contain efficient compression such as GIF images usually should not be compressed, as the CPU usage and time spent compressing them will result in an increased delay rather than network throughput improvement. Plain text files where no words are repeated, such as configurations with unique URLs or IPs, also may not be appropriate for compression.

FortiADC supports HTTP/HTTPS response compression in either gzip or deflate format.

Before you begin:

- You must have a good understanding of HTTP/HTTPS compression and knowledge of the content types served from the back-end real servers.
- You must have Read-Write permission for Load Balance settings.

Compression is not enabled by default. After you have configured a compression inclusion rule, you can select it in the profile configuration. To enable compression, select the profile when you configure the virtual server.

To configure a compression rule:

1. Click Server Load Balance > Application Optimization.
2. Click the **Compression** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Compression configuration on page 168](#).
5. Save the configuration.

Compression configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, __, and -. No spaces. You reference this name in the profile configuration. Note: After you initially save the configuration, you cannot edit the name.
URI List Type	• Include— Select this option to create a compression inclusion rule. HTTP/HTTPS responses that match the URIs and content types specified in this rule will be compressed by FortiADC before being passed to the client. • Exclude—Select this option to create a compression exclusion rule. HTTP/HTTPS responses that match the URIs and content types specified in this rule will not be compressed by FortiADC before being passed to the client.
URI Rule	Click Add and specify the URI to create the rule. Note: You must use a regular expression, e.g., <code>https://example.com/tmp/test.txt</code> .
Content Types	Click Add and select from the following content types to build the list: • application/javascript • application/soap+xml • application/x-javascript • application/xml • text/css • text/html

Settings	Guidelines
	• text/javascript • text/plain • text/xml • custom Note: The "custom" option allows you to specify almost any content/media type, including image files in .JPG, .PNG, and .BMP formats. The default is */*, which means any content/media type.



You can use the CLI to configure advanced options:

```
config load-balance compression
```

```
edit 1
```

```
set cpu-limit {enable | disable}
```

```
set max-cpu-usage <percent> -- max cpu usage for compression
```

```
set min-content-length <bytes> -- min bytes for compression
```

```
end
```

Compression and decompression

FortiADC supports HTTP/HTTPS response compression and request decompression with either gzip or deflate format.

You can offload HTTP/HTTPS response compression to FortiADC to save resources on your back-end servers, and let FortiADC to decompress compressed HTTP/HTTPS client requests for WAF inspection before passing them to your back-end servers.

Using caching features


The system RAM cache can store HTTP content and serve subsequent HTTP requests for that content without forwarding the requests to the backend servers, thereby reducing the load on the backend servers.

You can configure basic static caching or dynamic caching rules.

Static caching

[Static caching feature on page 169](#) illustrates the static caching feature.

Static caching feature



Before content is cached	After content has been cached
1. FortiADC receives the request from Client A and checks to see if it has a cached copy of the content. 2. If it does not, it forwards the request to a backend server. 3. The server sends content in response, and FortiADC caches the content. 4. FortiADC sends it to the client.	1. FortiADC receives the request from Client B and checks to see if it has a cached copy of the content. 2. It does, so it responds by sending the content to the client. The backend server is not contacted.

In general, the RAM cache conforms with the cache requirements described in sections 13 and 14 in [RFC 2616](#).

If caching is enabled for the profile that is applied to traffic processing, the system evaluates HTTP responses to determine whether or not to cache the content. HTTP responses with status codes 200, 203, 300, 301, 400 can be cached.

The following content is not cached:

- A response for a request that uses any method other than GET.
- A response for a request of which URI is contained in URI Exclude List or Dynamic Request URI Invalid list.
- A response for a request that contains any of the following headers: If-Match, If-Unmodified-Since, Authorization, Proxy-Authorization.
- A response that contains any of the following headers: Pragma, Vary, Set-Cookie, and Set-Cookie2.
- A response that does not include the Content-Length header. The Content-Length header must be 0.
- A response that does not contain the following headers: Cache-Control, Expires.
- A response with a Cache-Control header that does not have any of the following values: public, max-age, s-maxage.
- A response with a Cache-Control header that has one of the following values: no-cache, no-store, private.

In addition, content is not cached if the user-configured RAM cache thresholds described below are exceeded.

Dynamic caching

Dynamic caching is subject to rules you configure. In the Dynamic Caching Rules List, content that matches "caching invalid" URIs is never cached; otherwise, content that matches the Dynamic Cache Rule List of URIs is cached for the period you specify.

Dynamic caching is useful for dynamic web app experiences, such as online stores. For example, suppose a site uses a shopping cart. The URL to list items in the shopping cart is as follows:

`http://customshop.com/cart/list`

The URLs to add or delete items in the cart is as follows:

`http://customshop.com/cart/add`

`http://customshop.com/cart/delete`

In this case, you never want to cache the added or deleted pages because the old content will be "invalidated" by the changes you make. You may want, however, to cache the list page, but only for the period of time that you specify. The dynamic "invalid" rules makes it possible for you to never cache added and deleted pages, whereas the Dynamic Cache Rule List allows you to cache the list page for a specified period of time.

Another case where dynamic caching is useful is when content on a page is dynamic. For example, suppose an online ticket vendor has the following URL that shows how many tickets remain available for an event: `http://customshop.com/tickets/get_remains`. The number of tickets available is updated by a backend database. In this case, you might want to invalidate caching the URL or give it a small age out time.

Configuring caching rules

Before you begin:

- You must have a good understanding of caching and knowledge about the size of content objects clients access on the backend servers.
- You must have deep and detailed knowledge of your website URIs if you want to create dynamic caching rules.
- You must have Read-Write permission for Load Balance settings.

Caching is not enabled by default. After you have configured caching, you can select it in the profile configuration. To enable caching, select the profile when you configure the virtual server.

To configure caching:

1. Click Server Load Balance > Application Optimization.
2. Click the **Caching** tab.
3. Click **Create New** to display the Caching configuration editor.
4. Complete the configuration as described in [Caching configuration on page 172](#).
5. Save the configuration.

Caching configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the profile configuration. Note: After you initially save the configuration, you cannot edit the name.
Maximum Object Size	The default is 1 MB. The valid range is 1 byte to 10 MB.
Maximum Cache Size	The default is 100 MB. The valid range is 1 byte to 500 MB.
Maximum Entries	The default is 10,000. The valid range is 1 to 262,144.
Maximum Age	The default is 43,200 seconds. The valid range is 60 to 86,400. The backend real server response header also includes a maximum age value. The FortiADC system enforces whichever value is smaller.
URI Exclude List	
URI	Specify URIs to build the list. You can use regular expressions. This list has precedence over the Dynamic Cache Rule List. In other words, if a URI matches this list, it is ineligible for caching, even if it also matches the Dynamic Cache Rule list.
Dynamic Cache Rule List	
ID	Enter a unique ID. Valid values range from 1 to 1023.
Age	Timeout for the dynamic cache entry. The default is 60 seconds. The valid range is 1-86,400. This age applies instead of any age value in the backend server response header.
URI	Pattern to match the URIs that have content you want cached and served by FortiADC. Be careful with matching patterns and the order rules in the list. Rules are consulted from lowest rule ID to highest. The first rule that matches is applied.
Invalid URI	Pattern to match URIs that trigger cache invalidation. Be careful with matching patterns and the order rules in the list. Rules are consulted from lowest rule ID to highest. The first rule that matches is applied. This list has precedence over the Dynamic Cache URI list. In other words, if a URI matches this list, it is ineligible for caching, even if it also matches the Dynamic Cache URI list.

Using real server pools

This section includes the following topics:

- [Configuring real server pools](#)
- [Example: Using port ranges and the port 0 configuration](#)

Configuring real server pools

Server pools are groups of real servers that host the applications that you load balance.

To configure a server pool:

1. Create a server pool object.
2. Add members.

Before you begin:

- You must have a good understanding and knowledge of the backend server boot behavior, for example, how many seconds it takes to “warm up” after a restart before it can process traffic.
- You must know the IP address and port of the applications.
- If you want to select user-defined health checks, you must create them before creating the pool configuration. See [Configuring health checks](#).
- If you want to select user-defined real server SSL profiles, you must create them before creating the pool configuration. See [Configuring real server SSL profiles](#).
- You must have Read-Write permission for Load Balance settings.

After you have configured a real server pool, you can select it in the virtual server configuration.

To configure a pool:

1. Go to Server Load Balance > Real Server Pool.
The configuration page displays the Real Server tab.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration and add members as described in [Real Server Pool configuration guidelines on page 173](#).
4. Save the configuration.

Real Server Pool configuration guidelines

Settings	Guidelines
Real Server Pool	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the virtual server configuration. Note: After you initially save the configuration, you cannot edit the name.
Address Type	• IPv4 • IPv6
Health Check	Enable health checking for the pool. You can override this for individual servers in the pool.
Health Check Relationship	• AND—All of the selected health checks must pass for the server to be considered available. • OR—One of the selected health checks must pass for the server to be considered available.
Health Check List	Select one or more health check configuration objects.

Settings	Guidelines
Real Server SSL Profile	Select a real server SSL profile. Real server profiles determine settings for communication between FortiADC and the backend real servers. The default is NONE, which is applicable for non-SSL traffic.
Member	
Status	• Enable—The server can receive new sessions. • Disable—The server does not receive new sessions and closes any current sessions as soon as possible. • Maintain—The server does not receive new sessions but maintains any current connections.
Real Server	Click the down arrow and select a real server configuration object from the list menu. Note: The name of the selected real server pool member will appear in logs and reports.
Port	Enter the backend server's listening port (number), as described below: • HTTP—80, • HTTPS—443 • FTP—21 • SMTP—25 • DNS—53 • POP3—110 • IMAP4—143 • RADIUS—1812 • SNMP—161 Tip: The system uses Port 0 as a “wildcard” port. When configured to use Port 0, the system uses the destination port from the client request. For example, if you specify 0, and the destination port in the client request is 50000, the traffic will be forwarded to Port 50000.
Weight	Assigns relative preference among members—higher values are more preferred and are assigned connections more frequently. The default is 1. The valid range is 1 to 256. All load balancing methods consider weight. Servers are dispatched requests proportional to their weight, relative to the sum of all weights. The following example shows the effect of weight on Round Robin: • Server A, Weight 2; Server B, Weight 1: Requests are sent AABAAB. • Server A, Weight 3; Server B, Weight 2: Requests are sent AABAB. For other methods, weight functions as a tie-breaker. For example, with the Least Connection algorithm, requests are sent to the server with the least connections. If the number of connections is equal, the request is sent to the server with the greater weight. For example: • Server A, Weight 1, 1 connection • Server B, Weight 2, 1 connection • The next request is sent to Server B.
Recover	Seconds to postpone forwarding traffic after downtime, when a health check indicates that this server has become available again. The default is 0 (disabled). The valid range is 1 to 86,400 seconds. After the recovery period elapses, the FortiADC assigns connections at the warm rate.

Settings	Guidelines
	Examples of when the server experiences a recovery and warm-up period: • A server is coming back online after the health check monitor detected it was down. • A network service is brought up before other daemons have finished initializing and therefore the server is using more CPU and memory resources than when startup is complete. To avoid connection problems, specify the separate warm-up rate, recovery rate, or both. Tip: During scheduled maintenance, you can also manually apply these limits by setting Status to Maintenance instead of Enable. Note: Not applicable for SIP servers.
Warm Up	If the server cannot initially handle full connection load when it begins to respond to health checks (for example, if it begins to respond when startup is not fully complete), indicate how long to forward traffic at a lesser rate. The default is 0 (disabled). The valid range is 1 to 86,400 seconds. Note: Not applicable for SIP servers.
Warm Rate	Maximum connection rate while the server is starting up. The default is 10 connections per second. The valid range is 1 to 86,400 connections per second. The warm up calibration is useful with servers that have the network service brought up before other daemons have finished initializing. As the servers are brought online, CPU and memory are more utilized than they are during normal operation. For these servers, you define separate rates based on warm-up and recovery behavior. For example, if Warm Up is 5 and Warm Rate is 2, the number of allowed new connections increases at the following rate: • 1st second—Total of 2 new connections allowed (0+2). • 2nd second—2 new connections added for a total of 4 new connections allowed (2+2). • 3rd second—2 new connections added for a total of 6 new connections allowed (4+2). • 4th second—2 new connections added for a total of 8 new connections allowed (6+2). • 5th second—2 new connections added for a total of 10 new connections allowed (8+2). Note: Not applicable for SIP servers.
Connection Limit	Maximum number of concurrent connections to the backend server. The default is 0 (disabled). The valid range is 1 to 1,048,576 concurrent connections. Note: Connection Limit is not supported for FTP or SIP servers.
Connection Rate Limit	Limit the number of new connections per second to this server. The default is 0 (disabled). The valid range is 1 to 86,400 connections per second. In Layer 4 deployments, you can apply a connection rate limit per real server and per virtual server. Both limits are enforced. Note: The connection rate limit applies only when the real servers belong to a Layer 4 virtual server. If you add a real server pool with this setting configured to a Layer 7 virtual server, for example, the setting is ignored. Note: Connection Rate Limit is not supported for FTP or SIP servers.

Settings	Guidelines
Cookie	Specify the cookie name to be used when cookie-based Layer 7 session persistence is enabled. The cookie is used to create a FortiADC session ID, which enables the system to forward subsequent related requests to the same backend server. If you do not specify a cookie name, it is set to the pool member server name string. Note: This option is NOT applicable for SIP servers.
MySQL Group ID	Specify the MySQL group ID.
MySQL Read Only	Disabled by default. Select the button to enable it.
Backup	Designate this as a backup server to which FortiADC will direct traffic when the other servers in the pool are down. The backup server receives connections when all the other pool members fail the health check or you have manually disabled them. Note: Not applicable for SIP servers.
Health Check Inherit	When enabled, FortiADC will use the pool's health check settings. If disabled, you must select a health check to use with this individual backend server. See below.
Health Check	Select this option to specify a health check configuration object for this server. Note: This option becomes available only when
Health Check Relationship	• AND—All of the selected health checks must pass for the server to be considered available. • OR—One of the selected health checks must pass for the server to be considered available.
Health Check List	Select one or more health check configuration objects. Shift-click to select multiple objects at the same time.
RS Profile Inherit	Enable to inherit the real server SSL profile from the pool configuration. Disable to specify the real server profile in this member configuration. See below.
RS Profile	If RS Profile Inherit (above) is disabled, you must specify a real server SSL profile. A real server SSL profile determines the settings for communication between FortiADC and backend real server. Note: This option becomes available only when RS Profile Inherit is disabled.

Example: Using port ranges and the port 0 configuration

In some deployments, it is advantageous to support listening port ranges for client requests. For example, data centers or web hosting companies sometimes use port numbers to identify their customers. Client A sends requests to port 50000, client B to port 50001, client C to port 50002, and so on.

To support this scenario:

1. On the real servers, configure the listening ports and port ranges according to your requirements.
2. On the FortiADC, when you configure the real server pool member, specify port 0 for the port. The system handles port 0 as a “wildcard” port. When configured to use port 0, the system uses the destination port from the client request. For example, if you specify 0, and the destination port in the client request is 50000, the traffic is forwarded to port 50000.

3. When you configure the virtual server, specify a listening port and port range. The port range is like an offset. If the specified port is 50000 and the port range is 10, the virtual server listens on ports 50000-50009.

[Real server port 0 configuration on page 177](#) and [Virtual server port range configuration on page 177](#) highlight the key FortiADC configuration elements.

Real server port 0 configuration

|

Virtual server port range configuration

The screenshot shows the 'Virtual Server' configuration page with the 'Basic' tab selected. The 'Name' field contains 'DNS'. The 'Type' is set to 'Layer 7'. The 'Status' has buttons for 'Disable', 'Enable', and 'Maintain'. The 'Address Type' is set to 'IPv4'. The 'Traffic Group' is set to 'default'. Under the 'Specifics' section, the 'Schedule Pool' is set to 'OFF' and the 'Transaction Rate Limit' is set to '0'.

The screenshot shows the 'Virtual Server' configuration page with the 'General' tab selected. The 'Configuration' section includes the 'Address' field (10.125.2.16), 'Interface' (port2), and 'Port' (53). The 'Resources' section includes 'Profile' (DNS), 'Method' (LB_METHOD_ROUND_ROBIN), and 'Real Server Pool' (dns). A red box highlights the 'Port' field and its description: 'Default: 80 Range: 0 or 1-65535. You can specify up to eight ports or port ranges separated by space, e.g., 80-90 100. Valid values are from 0 to 65535, with 0 for Layer-4 virtual servers only.'

Note: Ports shown on the Dashboard > Virtual Server > Real Server page are for the configured port, so in this case, port 0. The ports shown in traffic logs are the actual destination port, so in this case, port 50000.

Note: The real-server port must be 0 or the same as the virtual server port for Layer-4 virtual servers in tunnel mode.

Configuring real servers

Real servers are physical servers that are used to form real server pools. These dedicated servers provide clients with services such as HTTP or XML content, streaming audio or video, TFTP/FTP uploads and downloads, etc. You can start configuring a real server by giving it a unique configuration name, setting its status, and specifying its IP address.

After you have created your real server configuration objects, you can select them as members to form real server pools. At that stage, further configurations are needed as discussed in [Configuring real server pools on page 1](#).

To configure a real server configuration object:

1. Go to Server Load Balance > Real Server Pool > Real Server.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration and add members as described in [Real Server configuration on page 178](#).
4. Click Save.
5. Repeat the same steps to add as many real server configuration objects as needed.

Real Server configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the virtual server configuration. Note: After you initially save the configuration, you cannot edit the name.
Status	Select one of the options: • Enable—The server can receive new sessions. • Disable—The server does not receive new sessions and closes any current sessions as soon as possible. • Maintain—The server does not receive new sessions but maintains any current connections.
Address	For IPv4 real server, enter the real server's IP address in IPv4 address format.
Address6	For IPv6 real server, enter the real server's IP address in IPv6 address format.


Note: The instructions above only covers the basic configuration of real servers. More configuration tasks are needed when you use them to form real server pools.

Configuring real server SSL profiles

A real server SSL profile determines settings used in network communication on the FortiADC-server segment, in contrast to a virtual server profile, which determines the settings used in network communication on the client-FortiADC segment.

[SSL profiles on page 178](#) illustrates the basic idea of client-side and server-side profiles.

SSL profiles



[Predefined real server profiles on page 180](#) provides a summary of the predefined profiles. You can select predefined profiles in the real server pool configuration, or you can create user-defined profiles.

Predefined real server profiles

Profile	Defaults
LB_RS_SSL_PROF_DEFAULT	- Allow version: SSLv3, TLSv1.0, TLSv1.1, TLSv1.2, TLSv1.3 - Cipher suite list,
LB_RS_SSL_PROF_ECDSA	- Allow version: SSLv3, TLSv1.0, TLSv1.1, and TLSv1.2 - Cipher suite list: ECDHE-ECDSA-AES256-GCM-SHA384, ECDHE-ECDSA-AES256-SHA384, ECDHE-ECDSA-AES256-SHA, ECDHE-ECDSA-AES128-GCM-SHA256, ECDHE-ECDSA-AES128-SHA256, ECDHE-ECDSA-AES128-SHA, ECDHE-ECDSA-RC4-SHA, ECDHE-ECDSA-DES-CBC3-SHA,
LB_RS_SSL_PROF_ECDSA_SSLV3	- Allow version: SSLv3, - Cipher suite list: ECDHE-ECDSA-AES256-SHA, ECDHE-ECDSA-AES128-SHA, ECDHE-ECDSA-RC4-SHA, ECDHE-ECDSA-DES-CBC3-SHA,
LB_RS_SSL_PROF_ECDSA_TLS12	- Allow version: TLSv1.2, - Cipher suite list: ECDHE-ECDSA-AES256-GCM-SHA384, ECDHE-ECDSA-AES256-SHA384, ECDHE-ECDSA-AES128-GCM-SHA256, ECDHE-ECDSA-AES128-SHA256,
LB_RS_SSL_PROF_ENULL	- Allow version: SSLv3, TLSv1.0, TLSv1.1, and TLSv1.2, - Cipher suite list: eNull, Recommended for Microsoft Direct Access servers where the application data is already encrypted and no more encryption is needed.
LB_RS_SSL_PROF_HIGH	- Allow version TLSv1.2, - Cipher suite list: ECDHE-RSA-AES256-GCM-SHA384 ECDHE-RSA-AES256-SHA384 ECDHE-RSA-AES256-SHA DHE-RSA-AES256-GCM-SHA384 DHE-RSA-AES256-SHA256 AES256-GCM-SHA384 AES256-SHA256,
LB_RS_SSL_PROF_LOW_SSLV2	- Allow version: SSLv2, - Cipher suite list: RC4-MD5,
LB_RS_SSL_PROF_LOW_SSLV3	- Allow version SSLv3, - Cipher suite list,
LB_RS_SSL_PROF_MEDIUM	- Allow version: TLSv1.0, TLSv1.1, and TLSv1.2, - Cipher suite list,
NONE	SSL is disabled.

Before you begin:

- You must have Read-Write permission for Load Balance settings.

To configure custom real server profiles:

1. Go to Server Load Balance > Real Server Pool.
2. Click the **Server SSL** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Real Server SSL Profile configuration guidelines on page 181](#).
5. Save the configuration.



You can clone a predefined configuration object to help you get started with a user-defined configuration.

To clone a configuration object, click the clone icon  that appears in the tools column on the configuration summary page.

Real Server SSL Profile configuration guidelines

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the real server pool configuration. Note: After you initially save the configuration, you cannot edit the name.
SSL	Enable/disable SSL for the connection between the FortiADC and the real server.
Note: The following fields become available only when SSL is enabled. See above.	
Customized SSL Ciphers Flag	Enable/disable use of user-specified cipher suites. When enabled, you must select a Customized SSL Cipher. See below.
Customized SSL Ciphers	If the customize cipher flag is enabled, specify a colon-separated, ordered list of cipher suites. An empty string is allowed. If empty, the default cipher suite list is used. The names you enter are validated against the form of the cipher suite short names published on the OpenSSL website: https://www.openssl.org/docs/manmaster/apps/ciphers.html
SSL Cipher Suite List	Ciphers are listed from strongest to weakest: • ECDHE-ECDSA-AES256-GCM-SHA384 • ECDHE-ECDSA-AES256-SHA384 • ECDHE-ECDSA-AES256-SHA • ECDHE-ECDSA-AES128-GCM-SHA256 • ECDHE-ECDSA-AES128-SHA256 • ECDHE-ECDSA-AES128-SHA • ECDHE-ECDSA-DES-CBC3-SHA • ECDHE-ECDSA-RC4-SHA • ECDHE-RSA-AES256-GCM-SHA384

Settings	Guidelines
	• ECDHE-RSA-AES256-SHA384 • ECDHE-RSA-AES256-SHA • DHE-RSA-AES256-GCM-SHA384 • DHE-RSA-AES256-SHA256 • DHE-RSA-AES256-SHA • AES256-GCM-SHA384 • AES256-SHA256 • AES256-SHA • ECDHE-RSA-AES128-GCM-SHA256 • ECDHE-RSA-AES128-SHA256 • ECDHE-RSA-AES128-SHA • DHE-RSA-AES128-GCM-SHA256 • DHE-RSA-AES128-SHA256 • DHE-RSA-AES128-SHA • AES128-GCM-SHA256 • AES128-SHA256 • AES128-SHA • ECDHE-RSA-RC4-SHA • RC4-SHA • RC4-MD5 • ECDHE-RSA-DES-CBC3-SHA • EDH-RSA-DES-CBC3-SHA • DES-CBC3-SHA • EDH-RSA-DES-CBC-SHA • DES-CBC-SHA • eNULL We recommend retaining the default list. If necessary, you can deselect ciphers you do not want to support.
Allowed SSL Versions	Select SSL versions that are allowed for the connection.
Certificate Verify	Specify a Certificate Verify configuration object to validate server certificates. This Certificate Verify object must include a CA group and may include OCSP and CRL checks.
SNI Forward Flag	Enable/disable forwarding the client SNI value to the server. The SNI value will be forwarded to the real server only when the client-side ClientHello message contains a valid SNI value; otherwise, nothing is forwarded.
Session Reuse Flag	Enable/disable SSL session reuse.
Session Reuse Limit	The default is 0 (disabled). The valid range is 0-1048576.
TLS Ticket Flag	Enable/disable TLS ticket-based session reuse .
Renegotiation	This option controls how FortiADC responds to mid-stream SSL reconnection requests either initiated by real servers or forced by FortiADC. Note:

Settings	Guidelines
Renegotiation Period	- This option is enabled by default. - When disabled, you must select an option for Renegotiation-Deny-Action. Specify the interval from the initial connect time that FortiADC renegotiates an SSL session. The unit of measurement can be second (default), minute, or hour, e.g., 100s, 20m, or 1h. Note: - The default is 0, which disables the function. - If a custom value is set, FortiADC will renegotiate the SSL session accordingly. For example, if you set the renegotiate period to 3600s (or 3600, 60m, or 1h), FortiADC will renegotiate the SSL session at least once an hour.
Renegotiate Size	Specify the amount (in MB) of application data that must have been transmitted over the secure connection before FortiADC initiates the renegotiation of an SSL session. Note: The default is 0, which disables the function.
Secure Renegotiation	Select one of the following options: - Request—FortiADC requests secure renegotiation of SSL connections. - Require—FortiADC requires secure renegotiation of SSL connections. In this mode, FortiADC allows initial SSL handshakes from real servers, but terminates renegotiation from real servers that do not support secure renegotiation. - Require Strict—FortiADC requires strict secure renegotiation of SSL connections. In this mode, FortiADC denies initial SSL handshakes from real servers that do not support secure renegotiation.
Renegotiation-Deny-Action	This option becomes available when Renegotiation is disabled on the server side. In that case, you must select an action that FortiADC will take when denying an SSL renegotiation request: - Ignore (default)—Ignores SSL renegotiation requests. - Terminate— Terminates SSL connections.

Using predefined scripts and commands

You can use scripts to perform actions that are not supported by the current built-in feature set. Scripts enable you to use predefined script commands and variables to manipulate HTTP requests and responses, redirection, or select a content route.

[Predefined scripts and commands on page 183](#) describes FortiADC's predefined scripts and commands that you can copy and customize.

Predefined scripts and commands

Predefined script/command	Usage
Scripts	
CLASS_SEARCH_n_MATCH	Demonstrates how to use the <code>class_match</code> and <code>class_search</code> utility function.

Predefined script/command	Usage
COMPARE_IP_ADDR_2_ADDR_GROUP_DEMO	Compares an IP address to an address group to determine if the IP address is included in the specified IP group. For example ,192.168.1.2 is included in 192.168.1.0/24. Note: Do NOT use this script "as is". Instead, copy it and customize the IP address and the IP address group.
CONTENT_ROUTING_by_URI	Routes to a pool member based on URI string matches. You should not use this script as is. Instead, copy it and customize the URI string matches and pool member names.
CONTENT_ROUTING_by_X_FORWARDED_FOR	Routes to a pool member based on IP address in the X-Forwarded-For header. You should not use this script as is. Instead, copy it and customize the X-Forwarded-For header values and pool member names.
GENERAL_REDIRECT_DEMO	Redirects requests to a URL with user-defined code and cookie. Note: Do NOT use this script "as is". Instead, copy and customize the code, URL, and cookie.
HTTP_2_HTTPS_REDIRECTION	Redirects requests to the HTTPS site. You can use this script without changes.
HTTP_2_HTTPS_REDIRECTION_FULL_URL	Redirects requests to the specified HTTPS URL. Note: This script can be used directly, without making any change.
HTTP_DATA_FETCH_SET_DEMO	Collects data in HTTP request body or HTTP response body. In <code>HTTP_REQUEST</code> or <code>HTTP_RESPONSE</code> , you could collect specified size data with "size" in <code>collect()</code> . In <code>HTTP_DATA_REQUEST</code> or <code>HTTP_DATA_RESPONSE</code> . You could print the data use "content", calculate data length with "size", and rewrite the data with "set". Note: Do NOT use this script "as is". Instead, copy it and manipulate the collected data.
HTTP_DATA_FIND_REMOVE_REPLACE_DEMO	Finds a specified string, removes a specified string, or replaces a specified string to new content in HTTP data. Note: Do NOT use this script "as is". Instead, copy it and manipulate the collected data.
INSERT_RANDOM_MESSAGE_ID_DEMO	Inserts a 32-bit hex string into the HTTP header with a parameter "Message-ID". Note: You can use the script directly, without making any change.
MULTIPLE_SCRIPT_CONTROL_DEMO_1	Uses <code>demo_1</code> and <code>demo_2</code> script to show how multiple scripts work. <code>Demo_1</code> with priority 12 has a higher priority.

Predefined script/command	Usage
	Note: You could enable or disable other events. Do NOT use this script "as is". Instead, copy it and customize the operation.
MULTIPLE_SCRIPT_CONTROL_DEMO_2	Uses demo_1 and demo_2 script to show how multiple scripts work. Demo_2 with priority 24 has a lower priority. Note: You could enable or disable other events. Do NOT use this script "as is". Instead, copy it and customize the operation.
OPTIONAL_CLIENT_AUTHENTICATION	Performs optional client authentication. Note: Before using this script, you must have the following four parameters configured in the client-ssl-profile: - client-certificate-verify—Set to the verify you'd like to use to verify the client certificate. - client-certificate-verify-option—Set to optional - ssl-session-cache-flag—Disable. - use-tls-tickets—Disable.
REDIRECTION_by_STATUS_CODE	Redirects requests based on the status code of server HTTP response (for example, a redirect to the mobile version of a site). Do NOT use this script "as is". Instead, copy it and customize the condition in the server HTTP response status code and the URL values.
REDIRECTION_by_USER_AGENT	Redirects requests based on User Agent (for example, a redirect to the mobile version of a site). You should not use this script as is. Instead, copy it and customize the User Agent and URL values.
REWRITE_HOST_n_PATH	Rewrites the host and path in the HTTP request, for example, if the site is reorganized. You should not use this script as is. Instead, copy it and customize the "old" and "new" hostnames and paths.
REWRITE_HTTP_2_HTTPS_in_LOCATION	Rewrites HTTP location to HTTPS, for example, rewrite "Location:http://www.example.com" to "Location:https://www.example.com". Note: You can use the script directly, without making any change.
REWRITE_HTTP_2_HTTPS_in_REFERER	Rewrites HTTP referer to HTTPS, for example, rewrite "Referer: http://www.example.com" to "Referer: https://www.example.com". Note: You can use the script directly, without making any change.
REWRITE_HTTPS_2_HTTP_in_LOCATION	Rewrites HTTPS location to HTTP, for example, rewrite "Location:https://www.example.com" to "Location:http://www.example.com".

Predefined script/command	Usage
	Note: You can use the script directly, without making any change.
REWRITE_HTTPS_2_HTTP_in_REFERER	Rewrites HTTPS referer to HTTP, for example, rewrite "Referer: https://www.example.com" to "Referer: http://www.example.com". Note: You can use the script directly, without making any change.
SPECIAL_CHARACTERS_HANDLING_DEMO	Shows how to use those "magic characters" which have special meanings when used in a certain pattern. The magic characters are () . % + - * ? [] ^ \$
USE_REQUEST_HEADERS_in_OTHER_EVENTS	Stores a request header value in an event and uses it in other events. For example, you can store a URL in a request event, and use it in a response event. Note: Do NOT use this script "as is". Instead, copy it and customize the content you want to store, use <code>collect()</code> in <code>HTTP_REQUEST</code> to trigger <code>HTTP_DATA_REQUEST</code> , or use <code>collect()</code> in <code>HTTP_RESPONSE</code> to trigger <code>HTTP_DATA_RESPONSE</code> .
UTILITY_FUNCTIONS_DEMO	Shows the utility functions that can be applied to all events when <code>HTTP_REQUEST{</code>
Commands	
COOKIE_COMMANDS	Lists the two cookie commands and shows how to use them.
IP_COMMANDS	Lists the IP commands and shows how to use them.
MANAGEMENT_COMMANDS	Lists the management commands and shows how to use them.
SSL_EVENTS_n_COMMANDS	Lists the SSL events and commands.
TCP_EVENTS_n_COMMANDS	Lists the TCP events and commands.

You can type or paste the script content into the configuration page. After you have created a script configuration object, you can specify it in a virtual server configuration.

Before you begin:

- Create a script. See [Appendix C: Scripts](#).
- You must have Read-Write permission for System settings.

The following paragraphs show how to:

- [Multi-script support](#)
- [Create a script object](#)
- [Import a script](#)
- [Export a script](#)
- [Delete a script](#)
- [Predefined scripts and commands in 5.x.x](#)

Create a script object

To create a script configuration object:

1. Go to Server Load Balance > Scripting.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in [Script configuration on page 187](#).
4. Save the configuration.

Script configuration

Settings	Guidelines
Name	Unique group name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Input	Type or paste the script. Note: If you want the script to be part of a big multiple script and have it executed in a certain order, be sure to set its priority. For more information, see Support for multiple scripts .

Import a script

To import a script:

1. Click Import
2. Click Choose File to browse for the script file.
3. Click Save.

Export a script

To export a script:

1. Select the script of interest.
2. Click Export.

Delete a script

To delete a script:

1. Select the script of interest.
2. Click Delete.

Linking multiple scripts to the same virtual server

FortiADC supports the use of a single script file containing multiple scripts and applies them to a single virtual server in one execution. Different scripts can contain the same event. You can specify the priority for each event in each script file to control the sequence in which multiple scripts are executed or let the system to execute the individual scripts in the order they are presented in the multi-script file.


For the current release, you can add up to 16 individual scripts to create a big multi-script file.

If you'd like to, you can disable the processing of the rest of the scripts (e.g., you can disable the processing of the remaining scripts in the list in a script), and even totally disable the processing of a certain event (e.g., you can disable processing the HTTP RESPONSE event in a HTTP REQUEST script). FortiADC also supports multiple calls of HTTP:redirect(), HTTP:redirect_with_cookie(), LB:routing(), and HTTP:close() functions such that the final one prevails.

In practice, rather than building one big complicated script containing all the required logic, it might be more useful to break it down into smaller functional pieces in the form of individual scripts. This is because executing multiple scripts at the same time is more efficient than running them separately, one at a time. Also, breaking down a giant script into multiple small individual scripts makes it more flexible to apply scripts to different virtual servers because some virtual servers may use some of the scripts while others may use them all. With the small individual scripts at hand, you can simply pick and choose only the scripts you need to assemble a big multi-script file with a set priority for each script and apply them all at once to a virtual server.

[Apply multiple scripts on page 188](#) shows how to link multiple scripts to a single virtual server from the GUI.

Apply multiple scripts



Setting script priority

Priority in a multi-script is *optional*, but is highly recommended. When executing a big multiple-script file, care must be taken to avoid conflicting commands among the scripts. You can set the priority for each script using the script editor on FortiADC's GUI. Valid values range from 1 to 1,000, with 500 being the default. The smaller the value, the higher the priority. Below is an example script with a set priority:

```
when HTTP_REQUEST priority 100 {
  LB:routing("cr1")
}
```

To display the priority info in the GUI, you can define one and only one event in each script file, as shown below:

```
Script 1:
when HTTP_REQUEST priority 500 {
  LB:routing("cr1")
}
```



```
}  
Script 2:  
when HTTP_RESPONSE priority 500 {  
  HTTP:close()  
}  
Script 3:  
when HTTP_REQUEST priority 400 {  
  LB:routing("cr2")  
}  
Script 4:  
when HTTP_RESPONSE priority 600 {  
  HTTP:close()  
}
```

Individual script files are loaded separately into the Lua stack. A numeric value (starting from 1) is appended to each event (e.g., for HTTP_REQUEST event, there are functions HTTP_REQUEST1, HTTP_REQUEST2, and so on so forth).

To support multiple scripts, FortiADC:

- Supports multiple calls of redirect/routing/close function, making them re-entrant so that the final one prevails. For that purpose, the system checks the behavior of multiple calls across `redirect()`, `close()`, and `routing()`. If `redirect()` comes first, followed by `close()`, then `close()` prevails. If `close()` comes first, followed by `redirect()`, then `redirect()` prevails. If you want to `close()`, you must disable the event after `close()`.
- Allows enabling or disabling events. There are times when you may want to disable the processing of the remaining scripts while a multi-script file is being executed, or want to disable processing the response completely. The mechanism serves that purpose.
- Allows enabling or disabling automatic event-enabling behavior. In the HTTP keep-alive mode, the system by default re-enables HTTP REQUEST and HTTP RESPONSE processing for the next transaction (even if they are disabled in the current transaction using the above enable or disable event mechanism). Now you can disable or enable this automatic enabling behavior.

[Script priority on page 189](#) shows a sample multi-script with priority information.

Script priority

Name: MULTIPLE_SCRIPT_CONTROL_DEMO_1

```

1  when RULE_INIT priority 14 {
2      --This is one of the script to demo the control of multiple scripts
3      --please change the priority of each event according to your need
4      debug("INIT in script 1\n");
5  }
6
7  when HTTP_REQUEST priority 12 {
8      debug("HTTP_REQUEST in script 1\n");
9      --add your own manipulation here
10
11     --you can disable rest of the HTTP_REQUEST events from executing
12     --by disabling this event
13     t={};
14     t["event"]="req"; -- can be "req", "res", "data_req", and "data_res"
15     t["operation"]="disable"; -- can be "enable", and "disable"
16     HTTP:set_event(t);
17     debug("disable rest of the HTTP_REQUEST events in script 1\n");
18     --you can also disable other events, say HTTP_RESPONSE, data events
19
20     --in the case of keep-alive, all events will be re-enabled automatically
21     --even though they are disabled in previous TRANSACTION using the
22     --HTTP:set_event(t) command. To disable this automatic re-enabling
23     --behavior, you can call HTTP:set_auto(t) as below
24     t={};
25     t["event"]="req"; -- can be "req", "res", "data_req", and "data_res"

```

Compiling principles

- All individual scripts should be pre-compiled when they are linked to a virtual server, where they can be combined into one big multi-script.
- For the same event, combine the commands in different scripts according to their priorities and orders.
- For commands of different priorities, FortiADC processes the high-priority commands first, and then the low-priority ones; for commands of the same priority, it processes them in the order they appear in the combined script.
- And if you are using multiple scripts with overlapping events for bidirectional traffic, you must ensure that the response traffic traverses the overlapping events in the expected order. By default, the scripts applied to the same virtual server will run in the order in which they are applied, regardless of the direction of traffic flow.
- For a specified event, you must make sure to avoid the conflict commands in different scripts. For example, if you have multiple scripts applied to the same virtual server and the scripts contain both request and response logic, the default execution order is like this:



but NOT like this:



As shown above, FortiADC cannot control the order in which events in the scripts are executed. The only way to enforce the execution order for response traffic is to use the event priority command, as we have discussed above. When setting the priorities, pay special attention to both request and response flows.

Special notes

When using the multi-script feature, keep the following in mind:


- The multi-script feature is supported on all FortiADC hardware platforms.
- Currently, the feature can be applied to layer-2 and Layer-7 virtual servers on HTTP/HTTPS protocol only.
- Scripts are VDOM-specific, and cannot be shared among different VDOMs.
- Session tables set up using scripts must be synced through high-availability (HA) configuration.
- Each multi-script script can contain up to 256 individual scripts, each being no more than 32 kilobytes.

Predefined scripts and commands in v5.x.x

The 5.x.x release comes with more predefined scripts and commands. You can view and use these scripts and commands by clicking Server Load Balance>Scripting.

[v5.x.x Scripts and predefined commands on page 191](#) highlights the functions of these scripts and commands and shows how to use them.

v5.x.x Scripts and predefined commands



Note:

- `UTILITY_FUNCTIONS_DEMO` and `CLASS_SEARCH_n_MATCH` provide various utility commands.
- `MULTIPLE_SCRIPT_CONTROL_DEMO_1` and `MULTIPLE_SCRIPT_CONTROL_DEMO_2` show how to use multiple-script support.
- `HTTP_DATA_FIND_REMOVE_REPLACE_DEMO` and `HTTP_DATA_FETCH_SET_DEMO` show how to manipulate HTTP data.
- `SPECIAL_CHARACTERS_HANDLING_DEMO` shows how to handle certain special characters.
- `INSERT_RANDOM_MESSAGE_ID_DEMO` shows how to generate random message IDs.
- `OPTIONAL_CLIENT_AUTHENTICATION` shows how to perform optional client authentication based on a request URL.
- `COMPARE_IP_ADDR_2_ADDR_GROUP_DEMO` shows how to perform IP address match.
- `USE_REQUEST_HEADERS_in_OTHER_EVENTS` shows how to share information across events.
- Many more predefined scripts are provided for load balance content routing, HTTP redirection, and HTTP content rewriting.

Configuring an L2 exception list

In some jurisdictions, SSL interception and decryption is disfavored for some types of websites or disallowed entirely. You use the L2 Exception List configuration to define such destinations. You can leverage FortiGuard web filter categories, and you can configure a list of additional destinations.

Before you begin:

- You must have created a Web Filter Profile configuration that includes the web categories to exclude from SSL decryption.
- You must have hostname or IP address details on additional destinations you want to exclude from SSL decryption.
- You must have Read-Write permission for Load Balance settings.

After you have created an L2 exception list configuration object, you can select it in a Layer 2 virtual server configuration.

To configure an exception list:

1. Go to Server Load Balance > SSL-FP Resources.
2. Click the **L2 Exception List** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [L2 exception list configuration on page 193](#).
5. Save the configuration.

L2 exception list configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the profile configuration. Note: After you initially save the configuration, you cannot edit the name.
Description	A string to describe the purpose of the configuration, to help you and other administrators more easily identify its use.
Web Filter Profile	Select a Web Filter Profile configuration.
Member	
Type	How you want to define the exception: • Host • IP
Host Pattern	Specify a wildcard pattern, such as *.example.com.
IP/Netmask	Specify the IP address and CIDR-formatted subnet mask, separated by a forward slash, such as 192.0.2.0/24. Note: • Dotted quad formatted subnet masks are not accepted. • IPv6 addresses are not supported.

Creating a Web Filter Profile configuration

You use the web filter profile configuration to create groups of FortiGuard categories that you want to include in the SSL forward proxy "L2 Exception List" configuration. The web filter profile should include categories that should not be processed by the outbound L2 SSL forward proxy feature. To address privacy concerns, you can include categories such as "Personal Privacy", "Finance and Banking", "Health and Wellness", and Medicine.

Before you begin:

- Learn about FortiGuard web filter categories. Go to <http://fortiguard.com/webfilter>.
- You must have Read-Write permission for Load Balance settings.

After you have created a web filter profile configuration object, you can select it in a L2 exception list configuration.

To create a web filter profile configuration:

1. Go to Server Load Balance > SSL-FP Resources.
2. Click the **Web Filter Profile** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Web Filter Profile configuration on page 194](#).
5. Save the configuration.

Web Filter Profile configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the profile configuration. Note: After you initially save the configuration, you cannot edit the name.
Description	A string to describe the purpose of the configuration, to help you and other administrators more easily identify its use.
Category-Members	
Category	Select a category or subcategory from the predefined list.

Using the Web Category tab

The Web Category tab displays the web filter categories imported from FortiGuard. You specify web categories when you create web filter groups.

For information on FortiGuard web categories, go to the FortiGuard website:

<http://fortiguard.com/webfilter>

Before you begin:

- You must have read permission for load balancing settings.

To display web categories:

1. Go to Server Load Balance > SSL-FP Resources.
2. Click the **Web Category** tab.

To manage how long the URL lists from FortiGuard are cached:

1. Go to System > FortiGuard.
2. Under Web Filter Configure, adjust caching settings as desired.

Configuring certificate caching

Certificate caching allows the system to cache the certificates presented to it for later use. Once cached, the certificates can be readily retrievable from the cache so that the system does not have to reload them when clients requesting service. In so doing, system performance can be greatly improved.

Configuring a certificate caching object

1. Click Server Load Balance > SSL-FP Resources.
2. Click the Certificate Caching tab.
3. Click Create New to open the certificate caching editor.
4. Make the desired entries as described in [Certificate caching configuration guidelines on page 195](#).
5. Click Save.


Certificate caching configuration guidelines

Settings	Guidelines
Name	Enter a unique name for the certificate caching rule.
Maximum Certificate Cache Size	Specify the maximum size of the certificate caching object. The default is 100 M.
Maximum entries	Specify the maximum number of real servers whose certificates (RSA + ECDSA) are to be cached. The default is 100,000.

TCP multiplexing

The TCP multiplexing option enables Layer 7 load balancing virtual servers to “reuse” existing TCP connections between FortiADC and backend real servers. Using this connection pool can reduce TCP overhead and improve web server and application performance. See [Client requests handled using connections from the connection pool on page 195](#).

Client requests handled using connections from the connection pool



Note: The feature is not supported for profiles with the Source Address option enabled.

You can enable and configure this option using the CLI only.

To configure a connection pool and assign it to a virtual server:

Use the following command to configure the connection pool:

```
config load-balance connection-pool
edit <name>
set age <integer>
set reuse <integer>
set size <integer>
set timeout <integer>
next
end
```

Settings	Guidelines
age	Maximum duration of a connection in seconds. The recommended value is 3000.
reuse	Maximum number of times that the virtual server can reuse the connection. The recommended value is 2000.
size	Maximum number of connections in the connection pool. The recommended value is 0, which specifies that there is no limit on the connection size.
timeout	Maximum number of seconds a connection can be idle before the system deletes it. The recommended value is 30.

To assign the connection pool configuration to a virtual server, enter the following command:

```
config load-balance virtual-server
edit <virtual-server_name>
set type l7-load-balance
set connection-pool <pool_name>
end
```

where:

<pool_name> is the name of the connection pool.

Chapter 5: Link Load Balancing

This chapter includes the following topics:

- [Link load balancing basics on page 197](#)
- [Link load balancing configuration overview on page 200](#)
- [Configuring gateway links on page 206](#)
- [Configuring persistence rules on page 207](#)
- [Configuring proximity route settings on page 209](#)
- [Configuring a link group on page 204](#)
- [Configuring a virtual tunnel group on page 210](#)
- [Configuring link policies on page 202](#)

Link load balancing basics

The link load balancing (LLB) features are designed to manage traffic over multiple internet service provider (ISP) or wide area network (WAN) links. This enables you to subscribe to or provision multiple links, resulting in reduced risk of outages, additional bandwidth for peak events, and potential cost savings if your ISP uses billing tiers based on bandwidth rate or peak/off-peak hours.

In most cases, you configure link load balancing for outgoing traffic. Outbound traffic might be user or server traffic that is routed from your local network through your ISP transit links, leased lines, or other WAN links to destinations on the Internet or WAN. You configure link policies that select the gateway for outbound traffic.

When the FortiADC system receives outbound traffic that matches a source/destination/service tuple that you configure, it forwards it to an outbound gateway link according to system logic and policy rules that you specify.


The LLB feature supports load balancing among link groups or among virtual tunnel groups.

Using link groups

The link group option is useful for ISP links. It enables you to configure multiple ISP links that are possible routes for the traffic. The LLB picks the best route based on health checks, LLB algorithms, bandwidth rate thresholds, and other factors you specify, including a schedule.

[LLB link groups on page 197](#) shows an example topology when FortiADC is deployed to support link groups.

LLB link groups




Using virtual tunnels

A virtual tunnel is a good choice when you want to load balance traffic from applications that embed the source address in the packet payload, like VPN and VoIP traffic. Such traffic can be difficult to load balance using traditional LLB methods. Virtual tunnels enable reliable, site-to-site connectivity using Generic Routing Encapsulation (GRE). The local FortiADC appliance encapsulates traffic so that it can be routed according to your link policy rules. The link policy rules use LLB techniques to identify the best available route among a group of links. If one of the links breaks down, the traffic can be rerouted through another link in the tunnel group. When traffic egresses the remote FortiADC appliance, it is decapsulated and the original source and destination IP addresses are restored.


[WAN connectivity over single leased lines on page 199](#) shows an example of a deployment that does not use LLB. It uses dedicated leased lines for its WAN links, which are reliable, but expensive.

WAN connectivity over single leased lines



[LLB virtual tunnels on page 199](#) shows the same network deployed with FortiADC appliances. The LLB link policy load balances traffic among more affordable ADSL links.

LLB virtual tunnels



Depending on your business, you might use the link group option, the virtual tunnel option, or both.



The FortiADC system evaluates traffic to determine the routing rules to apply. With regard to link load balancing, the system evaluates rules in the following order and applies the first match:


1. LLB link policy
2. Policy route
3. Static/Dynamic route
4. LLB default link group

Link load balancing configuration overview

The system has a configuration framework that enables granular link load balancing rules.

[LLB configuration summary on page 200](#) shows the configuration objects used in the LLB configuration and the order in which you create them. A *link policy* specifies the source/destination/service matches to which the policy applies. You apply a link policy to a *link group* or a *virtual tunnel*.

LLB configuration summary



The granular configuration of the gateway configuration includes health checks and bandwidth thresholds. The granular configuration of link groups includes load balancing methods, persistence rules, and proximity routes.

The granular configuration of virtual tunnels includes load balancing methods. In the virtual tunnel configuration, you can enable health check tests, but you do not use health check configuration objects.

Basic steps

1. Add address, address group, service, service group, and schedule group configuration objects that can be used to match traffic to link policy rules. This step is recommended. If your policy does not use match criteria, it will not have granularity.

2. Configure optional features. If you want to use health check rules, configure them before you configure the gateway links. If you want to use persistence rules or proximity routes, configure them before you configure a link group.
3. Configure gateway links.
4. Configure link groups or virtual tunnels.
5. Configure the link policy. When you configure a link policy, you set the source/destination/service matching tuple for your link groups or virtual tunnels.

Configuring link policies

A link policy matches traffic to rules that select a link group or virtual tunnel.

The policy uses a matching tuple: source, destination, service, and schedule. The policy match is a Boolean AND—All must match for the rule to be applied.

The elements of the tuple support specification by group objects. This is a Boolean OR—If source IP address belongs to member 1 OR member 2, then source matches.

The logical combinations enable you to subscribe multiple address spaces or services to a group of links, and create load balancing rules on that group basis.

The policy table is consulted from top to bottom. The first rule to match is applied.



The FortiADC system evaluates traffic to determine the routing rules to apply. With regard to link load balancing, the system evaluates rules in the following order and applies the first match:

1. LLB link policy
2. Policy route
3. Static/Dynamic route
4. LLB default link group

Before you begin:

- You must have configured any address, service, and schedule objects that you want to use as match criteria for your policy.
- You must have configured a link group or virtual tunnel group.
- You must have Read-Write permission for Link Load Balance settings.

To configure a link policy:

1. Go to Link Load Balance > Link Policy.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in [Link policy configuration on page 203](#).
4. Save the configuration.
5. Reorder rules, as necessary.

Link policy configuration

Option	Guidelines
Default Link Group	Select a link group configuration object that is used as the default when traffic does not match policy rules.
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Ingress Interface	Select the network interface to which the policy applies.
Source Type	Whether to use address, address group, or ISP address objects for this rule.
Source, Source ISP, or Source Group	Select an address object to match source addresses. If you do not specify a source address, the rule matches any source address. See Configuring IPv4 address groups .
Destination Type	Whether to use address, address group, or ISP address objects for this rule.
Destination, Destination ISP, or Destination Group	Select an address object to match destination addresses. If you do not specify a destination address, the rule matches any destination. See Configuring IPv4 address groups .
Service Type	Whether to use service or service group objects for this rule.
Service or Service Group	Select a service object to match destination services. If you do not specify a service, the rule matches any service. See Creating service groups .
Schedule	Select the schedule object that determines the times the system uses the logic of this configuration. The link policy is active when the current time falls in a time period specified by one or more schedules in the schedule group. If you do not specify a schedule, the rule applies at all times. See Creating schedule groups .
Group Type	- Link Group—Policy applies to a link group. Select the option, then the link group. See Configuring a link group. - Virtual Tunnel—Policy applies to a virtual tunnel. Select the option, then the virtual tunnel. See Configuring a virtual tunnel group.
Link Group	Select a link group.
Reordering	
	After you have saved a rule, reorder rules as necessary. The rules table is consulted from top to bottom. The first rule that matches is applied and subsequent rules are not evaluated.
Hit Counts	Hit Counts: For monitor only. The value indicates the link policy hit times.

Configuring a link group

Link groups include ISP gateways your company uses for outbound traffic. Grouping links reduces the risk of outages and provisions additional bandwidth to relieve potential traffic congestion. See [Using link groups](#).

The link group configuration specifies the load balancing algorithm and the gateway routers in the load balancing pool. You can enable LLB options, such as persistence rules and proximity routes.

Before you begin:

- You must have configured gateway links and persistence rules and before you can select them in the link group configuration.
- You must have Read-Write permission for Link Load Balance settings.

After you have configured a link group configuration object, you can select it in the link policy configuration.

To configure a link group:

1. Go to Link Load Balance > Link Group.
The configuration page displays the Link Group tab.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration and add members as described in [Link group configuration on page 204](#).
4. Save the configuration.

Link group configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the LLB policy configuration. Note: After you initially save the configuration, you cannot edit the name.
Address Type	IPv4 Note: IPv4 is selected by default, and cannot be changed.
Route Method	• Weighted Round Robin—Dispatches new connections to link members using a weighted round-robin method. • Least Connections—Dispatches new connections to the link member with the lowest number of connections. • Least New Connections per Second—Dispatches new connections to the link member that has the lowest rate of new connections per second. • Least Throughput Outbound—Dispatches new connections to the link member with the least outbound traffic. • Least Throughput Inbound—Dispatches new connections to the link member with the least inbound traffic. • Least Throughput Total—Dispatches new connections to the link member with the least total traffic (that is, inbound plus outbound). • Spillover Throughput Outbound—Dispatches new connections according to the spillover list based on outbound traffic. • Spillover Throughput Inbound—Spillover list based on inbound traffic. • Spillover Throughput Total—Spillover list based on total traffic (that is, inbound plus

Settings	Guidelines
	outbound). Source Address Hash—Selects the gateway link based on a hash of the source IP address.
Persistence	Select a persistence configuration. Optional.
Proximity Route	- Enable—The system uses the proximity route logic and configuration when determining routes. - Disable—The system does not use the proximity route configuration.
Add member	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Gateway	Select a gateway configuration object. See Configuring gateway links .
Weight	Assigns relative preference among members—higher values are more preferred and are assigned connections more frequently. The default is 1. The valid range is 1 to 255. All load balancing methods consider weight, except spillover, which uses its own priority configuration. Servers are dispatched requests proportional to their weight, relative to the sum of all weights. The following example shows the effect of weight on WRR: - Sever A, Weight 2; Server B, Weight 1: Requests are sent AABAAB. - Sever A, Weight 3; Server B, Weight 2: Requests are sent AABAB. For other methods, weight functions as a tie-breaker. For example, with the Least Connection algorithm, requests are sent to the server with the least connections. If the number of connections is equal, the request is sent to the server with the greater weight. For example: - Server A, Weight 1, 1 connection - Server B, Weight 2, 1 connection The next request is sent to Server B.
Spillover Priority	Assigns a priority to the link when using a spillover load balancing method. Higher values have greater priority. When a spillover method is enabled, the system dispatches new connections to the link that has the greatest spillover priority until its threshold is exceeded; then it dispatches new connections to the link with the next greatest priority until its threshold is exceeded, and so on. If multiple links in a link group have the same spillover priority, the system dispatches new connections among those links according to round robin. The default is 0. The valid range is 0-9.
Status	- Enable—The member is considered available for new traffic. - Disable—The member is considered unavailable for new traffic.
Backup	Enable to designate the link as a backup member of the group. All backup members are inactive until all main members are down.

Configuring gateway links

The gateway link configuration enables you to specify health checks, bandwidth rate thresholds, and spillover threshold behavior for the gateway links you add to link groups.

Before you begin:

- You must know the IP addresses of the ISP gateway links used in the network segment where the FortiADC appliance is deployed.
- You must have added health check configuration objects that you want to use to check the gateway links.
- You must have Read-Write permission for Link Load Balance settings.

After you have configured a gateway link configuration object, you can select it in the link group configuration.

To configure a gateway link:

1. Go to Link Load Balance > Link Group.
2. Click the **Gateway** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [LLB gateway configuration on page 206](#).
5. Save the configuration.

LLB gateway configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the link group configuration. Note: After you initially save the configuration, you cannot edit the name.
Address	IP address of the gateway link.
Health Check	Enable health checks.
Health Check Relationship	• AND—All of the selected health checks must pass for the link to be considered available. • OR—One of the selected health checks must pass for the link to be considered available.
Health Check List	Select one or more health check configuration objects.
Inbound Bandwidth	Maximum bandwidth rate for inbound traffic through this gateway link.
Outbound Bandwidth	Maximum bandwidth rate for outbound traffic to this gateway link. If traffic exceeds this threshold, the FortiADC system considers the gateway to be full and does not dispatch new connections to it. The default is 2,000,000 Kbps. The valid range is 1 to 2,147,483,647. We recommend you tune bandwidth thresholds strategically, using the bandwidth rate and price structure agreement you have with your ISP to your advantage.
Inbound Spillover Threshold	Maximum inbound bandwidth rate for a link in a spillover load balancing pool.

Settings	Guidelines
Outbound Spillover Threshold	Maximum outbound bandwidth rate for a link in a spillover load balancing pool. If you enable spillover load balancing in the link group configuration, the system maintains a spillover list. It dispatches new connections to the link with the greatest priority until its spillover threshold is exceeded; then dispatches new connections to the link with the next greatest priority until its threshold is exceeded, and so on. The default is 2,000,000 Kbps. The valid range is 1 to 2,147,483,647.
Total Spillover Threshold	Maximum total bandwidth rate (inbound plus outbound) for a link in a spillover load balancing pool.

Configuring persistence rules

Persistence rules identify traffic that should be ignored by load balancing rules and instead be forwarded to the same gateway each time the traffic traverses the FortiADC appliance.

You should use persistence rules with applications that use a secure connection. Such applications drop connections when the server detects a change in a client's source IP address.

[Persistence rules used in link load balancing on page 207](#) describes the types of persistence rules you can configure.

Persistence rules used in link load balancing

Persistence	Description
Source-Destination Pair	Packets with the same source IP address and destination IP address take same outgoing gateway.
Source-Destination Address	Packets with a source IP address and destination IP address that belong to the same subnet take the same outgoing gateway.
Source Address	Packets with a source IP address that belongs to the same subnet take the same outgoing gateway.
Destination Address	Packets with a destination IP address that belongs to the same subnet take same outgoing gateway.

Before you begin:

- You must have an awareness of the types of outbound traffic from your network. Persistence rules are useful for traffic that requires an established session, such as secure connections (HTTPS and SSH, for example).
- You must have knowledge of the source and/or destination subnets to which the persistence rules should apply.
- You must have Read-Write permission for Link Load Balance settings.



You can use persistence rules in link groups but not virtual tunnels.

To configure a persistence rule:

1. Go to Link Load Balance > Link Group.
2. Click the **Persistence** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Persistence rule configuration on page 208](#).
5. Save the configuration.

Persistence rule configuration

Type	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the link group configuration. Note: After you initially save the configuration, you cannot edit the name.
Type	Select one of the persistence types, as described below.
Source-Destination Pair	
Timeout	The default is 300 seconds.
Source-Destination Address	
Timeout	The default is 300 seconds.
Source IPv4 Netmask Bits	Number of bits in a subnet mask to specify a network segment that should follow the persistence rule.
Destination IPv4 Netmask Bits	Number of bits in a subnet mask to specify a network segment that should follow the persistence rule. For example, if you set this to 24, and the system chooses a particular gateway router for destination IP 192.168.1.100, the system will select that same gateway for traffic to all destination IPs in subnet 192.168.1.0/24.
Source Address	
Timeout	The default is 300 seconds.
Source IPv4 Netmask Bits	Number of bits in a subnet mask to specify a network segment that should follow the persistence rule. The default is 32, but you can set it to any value between 1 and 32. For example, if you set this to 24, and the system chooses a particular gateway router for client IP 192.168.1.100, the system will select that same gateway for subsequent client requests when the subsequent client belongs to subnet 192.168.1.0/24.
Destination Address	
Timeout	The default is 300 seconds.
Destination IPv4 Netmask Bits	Number of bits in a subnet mask to specify a network segment that should follow the persistence rule.

Configuring proximity route settings

The proximity route feature enables you to associate link groups with efficient routes. Proximity routes can improve user experience over the WAN because traffic is routed over fast routes.

You can use either or both of these methods:

- **Static Table**—You specify the gateways to use for traffic on destination networks.
- **Dynamic Detection**—The system polls the network for efficient routes. The algorithm selects a gateway based on latency.

If you configure both, the system checks the static table first for a matching route and, if any, uses it. If there is no matching static route, the system uses dynamic detection.

Before you begin:

- You must have knowledge of IP addresses used in outbound network routes to configure a static route.
- You must have Read-Write permission for Link Load Balance settings.

To configure a proximity route:

1. Go to Link Load Balance > Link Group.
2. Click the **Proximity Route** tab.
3. Complete the configuration as described in [Proximity route rule configuration on page 209](#).
4. Save the configuration.

Proximity route rule configuration

Type	Guidelines
Mode	• Static Table First—Consult the static table first. If no match, use dynamic detection. • Static Table Only—Use the static table; do not use dynamic detection. • Dynamic Detect Only—Use dynamic detection; do not use the static table. • Disable—Do not use the proximity route configuration.
Static Table	
Type	• ISP—Use an ISP address object. • Subnet—Specify an IP netmask manually. Routes that are specified manually have priority over ISP address object entries.
ISP Name	If you use the ISP configuration type, select an ISP address book configuration object. If an address exists in multiple ISP address books, the route entries have priority as follows: 1. User-defined entries. 2. Entries from an address book that has been imported. 3. Entries from the predefined address book (default for the firmware image).
IP Subnet	If you use the Subnet configuration type, specify a destination IP address and netmask.
Gateway	Select a gateway configuration object. The gateway must be able to route packets to the destination IP address that you have specified.
Dynamic Detect	

Type	Guidelines
Protocol	- ICMP—Use ICMP to detect routes. Calculate proximity by the smaller RTT. - ICMP and TCP—Some hosts do not respond to ICMP requests. Specify this option to use both ICMP and TCP to detect routes and RTT. For TCP detection, port 7 (TCP echo) is used. A connection refused or connection reset by the destination is treated as successful detection.
Aging Period	The default is 86,400 seconds (24 hours).
Retry Number	The default is 3.
Retry Interval	The default is 3.

Configuring a virtual tunnel group

Virtual tunnels enable reliable, site-to-site connectivity using Generic Routing Encapsulation (GRE) to tunnel traffic between pairs of FortiADC appliances. See [Using virtual tunnels](#).

The virtual tunnel group configuration sets the list of tunnel members, as well as load balancing options like algorithm and weight.

When you add members to a virtual tunnel configuration, you specify a local and remote IP address. These addresses are IP addresses assigned to a network interface on the local and remote FortiADC appliance.

Before you begin:

- You must have Read-Write permission for Link Load Balance settings.

After you have configured a virtual tunnel configuration object, you can select it in the link policy configuration.

To configure a virtual tunnel:

- Go to Link Load Balance > Virtual Tunnel.
- Click **Create New** to display the configuration editor.
- Complete the configuration and add members as described in [Virtual tunnel configuration on page 210](#).
- Save the configuration.

Virtual tunnel configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the LLB policy configuration. Note: After you initially save the configuration, you cannot edit the name.
Method	- Weighted Round Robin—Dispatches packets to VT members using a weighted round-robin method. - Source-Destination Hash—Dispatches packets by source-destination IP address tuple.
Add member	

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Tunnel Local Address	IP address for the network interface this system uses to form a VPN tunnel with the remote system.
Tunnel Remote Address	IP address that the remote FortiADC system uses to form a VPN tunnel with this system.
Health Check	• Enable—Send probes to test whether the link is available. • Disable—Do not send probes to test the health of the link.
Weight	Assigns relative preference among members—higher values are more preferred and are assigned connections more frequently.
Status	• Enable—The member is considered available for new traffic. • Disable—The member is considered unavailable for new traffic.
Backup	Enable to designate the tunnel as a backup member of the group. All backup members are inactive until all main members are down.

Chapter 6: Global Load Balancing

This chapter includes the following topics:

- [Global load balancing basics on page 212](#)
- [Global load balancing configuration overview on page 214](#)
- [Configuring servers on page 216](#)
- [Configuring link on page 219](#)
- [Configuring data centers on page 220](#)
- [Configuring hosts on page 221](#)
- [Configuring wizard on page 222](#)
- [Configuring virtual server pools on page 225](#)
- [Configuring location lists on page 226](#)
- [Logical Topology on page 227](#)
- [Configuring an address group on page 236](#)
- [Configuring remote DNS servers on page 237](#)
- [Configuring the DSSET list on page 235](#)
- [Configuring DNS zones on page 229](#)
- [Configuring DNS64 on page 234](#)
- [Configuring the response rate limit on page 238](#)
- [Configuring a Global DNS policy on page 228](#)
- [Configuring general settings on page 233](#)
- [Configuring the trust anchor key on page 234](#)

Global load balancing basics


The global load balancing (GLB) feature is a DNS-based solution that enables you to deploy redundant resources around the globe that you can leverage to keep your business online when a local area deployment experiences unexpected spikes or downtime. The FortiADC system implements a hardened BIND 9 DNS server that can be deployed as the authoritative name server for the DNS zones that you configure. Zone resource records are generated dynamically based on the global load balancing framework. The DNS response to a client request is an ordered lists of answers that includes all available virtual servers. A client that receives DNS response with a list of answers tries the first and only proceeds to the next answers if the first answer is unreachable. The response list is based on the following priorities:

1. **Virtual server health**—Availability is determined by real-time connectivity checking. When the DNS server receives a client request, it checks connectivity for all possible matches and excludes unavailable servers from the response list.
2. **Persistence**—You can enable persistence for applications that have transactions across multiple hosts. A match to the persistence table has priority over proximity algorithms.
3. **Geographic proximity**—Proximity is determined by matching the source IP address to either the FortiGuard Geo IP database or the FortiADC predefined ISP address book.

4. Dynamic proximity—Proximity is determined by application response time (RTT probes), least connections, or byte-per-second.
5. Weighted round robin—If proximity algorithms are not configured or not applicable, available virtual servers are listed in order based on a simple load balancing algorithm.

Global load balancing deployment on page 213 shows an example global load balancing deployment with redundant resources at data centers in China and the United States.

Global load balancing deployment



FortiADC-1 is the local SLB for the data center in China. FortiADC-2 is the local SLB for the data center in the United States. FortiADC-3 is a global SLB. It hosts the DNS server that is authoritative for `www.example.com`. When a client clicks a link to `www.example.com`, the local host DNS resolver commences a DNS query that is ultimately resolved by the authoritative DNS server on FortiADC-3. The set of possible answers includes the virtual servers on FortiADC-1 or FortiADC-2. The global load balancing framework uses health status and proximity algorithms to determine the set of answers that are returned, and the order of the answer list. For example, you can use the global SLB framework geoproximity feature to direct clients located in China to the virtual server in China; or if the virtual server in China is unavailable, then to the redundant resources in the United States.

You configure the global load balancing framework and DNS settings only on the global FortiADC (FortiADC-3 in the example above). The virtual server IP addresses and ports can be discovered by the FortiADC global SLB from the FortiADC local SLBs. The GLB DNS server uses the discovered IP addresses in the DNS response. The framework also supports third-party IP addresses and health checks for them.

The DNS server supports the following security features:

- DNSSEC—Domain Name System Security Extensions. DNSSEC provides authentication by associating cryptographically generated digital signatures with DNS resource record (RR) sets. The FortiADC system makes it easy to manage the keys that must be provided to DNS parent domains and the keys that must be imported from DNS child domains.
- Response rate limit—Helps mitigate DNS denial-of-service attacks by reducing the rate at which the authoritative name servers respond to high volumes of malicious queries.
- DNS forwarding—In a typical enterprise local area network, the client configuration has the IP address of an internal authoritative DNS server so that requests for internal resources can be answered directly from its zone data. Requests for remote resources are sent to another DNS server known as a forwarder. The internal server caches the results it learns from the forwarder, which optimizes subsequent lookups. Using forwarders reduces the number of DNS servers that must be able to communicate with Internet DNS servers.



Further reading:

BIND 9 reference manuals: <http://www.bind9.net/manuals>

RFC 1035 (DNS): <http://tools.ietf.org/html/rfc1035>

RFC 4033 (DNSSEC): <http://tools.ietf.org/html/rfc4033>

Global load balancing configuration overview

In a global load balancing deployment, you configure DNS server and global load balancing details only on the global FortiADC instance. The configuration framework enables granular administration and fine tuning of both the DNS server and the global load balancing framework.

[Global load balancing configuration summary on page 214](#) shows the basic configuration elements for global load balancing and the recommended order for creating the configuration objects. The order is important for initial configurations because complex configuration elements like policies often include references to simple configuration objects like the remote DNS servers (forwarders) or DNS64 rules, but the simple elements must be created first.

Global load balancing configuration summary



Basic steps (DNS server)

1. Configure address groups to use in your DNS policy matching rules. The system includes the predefined address groups **any** and **none**.
2. Configure remote DNS servers (forwarders) and the DSSET list that you might reference in the zone configuration.
3. Complete the zone configuration. The global load balancing framework generates the zone configuration for zones that include the FortiADC virtual servers.
4. Configure DNS64 or response rate limit configurations that you might reference in the DNS policy.
5. Configure the DNS policy that matches a source/destination tuple to a zone. You can also enable and configure DNSSEC in the DNS policy.
6. Configure general DNS settings to be applied when DNS requests do not match the DNS policy.

Basic steps (Global load balancing)


1. Create the data center, servers, virtual server pool, and host configurations that are the framework for associating locations with virtual servers and generating the DNS zone configuration and resource records. You can adjust the dynamic proximity and persistence settings at any time.
2. Review the generated DNS zone configuration.
3. Create a policy that matches traffic to the generated zone configuration.

Configuring servers

In the context of the global server load balance configuration, servers are the local SLB (FortiADC instances or third-party servers) to be load balanced. For FortiADC instances, the GLB checks status and synchronizes configuration from the local SLB so that it can learn the set of virtual servers that can be included in the GLB virtual server pool.

[Virtual server discovery on page 216](#) illustrates configuration discovery. Placement in this list does not include them in the pool. You also must name them explicitly in the virtual server pool configuration.

Virtual server discovery



Before you begin:

- You must have created the data center configuration objects that are associated with the local SLB.
- You must have created virtual server configurations on the local FortiADC SLB. In this procedure, the global SLB discovers them.
- You must have created gateway configuration objects on the local FortiADC SLB if you want to configure a gateway health check. In this procedure, the global SLB discovers them.
- You must have Read-Write permission for Global Load Balance settings.

After you have created a server configuration object, you can specify it the global load balancing virtual server pool configuration.

To configure servers:

1. Go to Global Load Balance > Global Object.
2. Click the **Server** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Server configuration on page 217](#).
5. Use the **Discover** utility to populate the Member list configuration with virtual server configuration details from the local FortiADC SLB.
6. Optional. Edit the populated list to select a discovered gateway configuration object if you want the GSLB to perform gateway health checks.
7. Save the configuration.

Server configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the virtual server pool configuration.

Settings	Guidelines
Note: After you initially save the configuration, you cannot edit the name.	
Type	- FortiADC SLB: A FortiADC instance. - Generic Host: A third party ADC or server.
Synchronization	Enable/disable synchronization of the virtual server status from the local FortiADC SLB. Disabled by default. If enabled, synchronization occurs whenever there is a change in virtual server status.
Auth Type	- None—No password. - TCP MD5SIG—With password, but can not be used if NAT is in between the client and server. This is because, when using the TCP MD5SIG authentication in a network with NAT in between, the IP layer is encrypted. So is every packet. Because the IP address will be changed, the encryption check will always fail. - Auth Verify—The authentication key is sent to the server after a three-way handshake. The key is encrypted and NAT in between will not affect the authentication.
Password	Enter the password to authenticate key. Note: This field appears only when TCP MD5SIG or Auth Verify is selected as the authentication type. The password you enter here must match the password configured on the FortiADC appliance in a global sever load-balancing configuration.
Address Type	IPv4 or IPv6.
IP Address	Specify the IP address for the FortiADC management interface. This IP address is used for synchronization and also status checks. If the management interface is unreachable, the virtual servers for that FortiADC are excluded from DNS answers.
Port	5858 by default.
Data Center	Select a data center configuration object. The data center configuration object properties are used to establish the proximity of the servers and the client requests.
Auto-sync	Enable/disable auto-sync from the server. Global load balancing will auto-sync the server member when enabled.
Health Check Control	If type is Generic Host, enable/disable health checks for the virtual server list. The health check settings at this configuration level are the parent configuration. When you configure the list, you can specify whether to inherit or override the parent configuration. Note: This option is available only when Generic Host is selected. See Type above. Health checking is built-in, and you can optionally configure a gateway health check.
Health Check Relationship	- AND—All of the specified health checks must pass for the server to be considered available. - OR—One of the specified health checks must pass for the server to be considered available.
Health Check List	Select one or more health check configuration objects.
Member	
Discover	Populate the member list with virtual servers from the local FortiADC configuration. After the list had been populated, you can edit the configuration to add a gateway health check.

Settings	Guidelines
Override	Select this option if you want to update the discovered virtual server configuration with the latest configuration information whenever you use the Discover utility (for example, additions or changes to previously discovered configurations). Unselect this option if you want to preserve the previously discovered configuration and not have it overwritten by the Discover operation.
Name	Must match the virtual server configuration name on the local FortiADC.
Address Type	IPv4 or IPv6.
IP Address	Virtual server IP address.
Gateway	Enable an additional health check: is the gateway beyond the FortiADC reachable? The list of gateway configuration objects is populated by discovery, but you must select the appropriate one from the list.
Health Check Inherit	If type is Generic Host, enable to inherit the health check settings from the parent configuration. Disable to specify health check settings in this member configuration.
Health Check Control	Enable health checking for the virtual server. Note: This option is available only when Health Check Inherit is disabled. In that case, you can enable this option and configure the Health Check Relationship and Health Check List fields below.
Health Check Relationship	• AND—All of the specified health checks must pass for the server to be considered available. • OR—One of the specified health checks must pass for the server to be considered available.
Health Check List	Specify one or more health check configuration objects.

Configuring link

To configure a global load balance link:

1. Go to Global Load Balance > Global Object.
2. Click the **Link** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Global load balance link configuration on page 219](#).
5. Save the configuration.

Global load balance link configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the global load balance servers configuration. Note: After you initially save the configuration, you cannot edit the name.
Data Center	Select a data center from the list.

Settings	Guidelines
	Note: You must the data center(s) configured ahead of time.
ISP	Select an ISP from the list.
Gateway	
Server	Select a server.
Gateway Name	Specify the name of a gateway.
or Select Here	Click the down arrow to select a gateway from the drop-down list. Note: Use this option only when you already have a list of gateways configured on the server.

Configuring data centers

The data center configuration sets key properties: Location and/or ISP and ISP province. These properties are used in the global load balancing algorithm that selects the FortiADC in closest proximity to the client.

Before you begin:

- If you want to select a user-defined ISP address book, you must create it before creating the data center configuration.
- You must have Read-Write permission for Global Load Balance settings.

After you have created a data center configuration object, you can specify it in the global load balance servers configuration.

To configure a data center:

1. Go to Global Load Balance > Global Object.
2. Click the **Data Center** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Data center configuration on page 220](#).
5. Save the configuration.

Data center configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the global load balance servers configuration. Note: After you initially save the configuration, you cannot edit the name.
Location	Select a location from the drop-down list menu. See the note below.
Description	Optional description to help administrators know the purpose or usage of the configuration.

Note: Starting from FortiADC 5.x.x, the GUI shows the full country or region names listed in alphabetical order for location list and data center configuration. The Console uses country or region name abbreviations instead. The abbreviations are done in accordance with the ISO standards. So if you configure a location list or data center from the

Console, be sure to consult [ISO-3166-1](#) and/or [ISO 3166-2:CN](#) for the correct abbreviations to use. See the following example commands:

```
config global-load-balance topology
edit "tpl"
set member ZZ US CN65
next
end
```


Where ZZ stands for Reserved, US for United States, and CN65 for China, Xingjiang

Configuring hosts

Host settings are used to form the zone configuration and resource records in the generated DNS zone used for global load balancing.

[Host configuration and the generated DNS zone on page 221](#) shows how the host settings are mapped to zone settings and resource records. Domain and hostname are used in both the configuration and the generated configuration name. The IP address and weight are derived from the virtual server pool.

Host configuration and the generated DNS zone



Before you begin:

- You must have created the global virtual server pools you want to use.
- You must have Read-Write permission for Global Load Balance settings.

After you have created a host configuration object, it can be used to form the zone and resource records in the generated DNS zone configuration.

To configure a host:

1. Go to Global Load Balance > FQDN Settings.
2. Click the **Host** tab.
3. Click **Create New** to display the configuration editor.

4. Complete the configuration as described in [Host configuration on page 222](#).
5. Save the configuration.

Host configuration


Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Note: After you initially save the configuration, you cannot edit the name.
Host Name	The hostname part of the FQDN, such as <code>www</code> . Note: You can specify the @ symbol to denote the zone root. The value substituted for @ is the preceding \$ORIGIN directive.
Domain Name	The domain name must end with a period. For example: <code>example.com.</code>
DNS Policy	Select the DNS policy you want the host to use.
Respond Single Record	Enable/disable an option to send a single record in response to a query. Disabled by default. By default, the response is an ordered list of records.
Persistence	Enable/disable the persistence table. Disabled by default. If you enable persistence, the client source address is recorded in the persistence table, and subsequent requests from the same network or the same host or domain are sent an answer with the virtual servers listed in the same order (unless a server becomes unavailable and is therefore omitted from the answer).
Virtual Server Pool Selection Method	- Weight—If selected, virtual server pool will be responded by weight. - DNS Query Origin—If selected, virtual server pool with the same topology information as the local DNS address will be responded. - Global Availability—If selected, virtual servers will be responded by their global availability: the first virtual server in queue will always be responded if it is globally available, and the next virtual server in queue will be responded if the preceding virtual server is unavailable.
Default Feedback IPv4	Specify an IP address to return in the DNS answer if no virtual servers are available.
Default Feedback IPv6	Specify an IPv6 address to return in the DNS answer if no virtual servers are available.
Virtual Server Pool	
Name	Enter the mkey.
Virtual Server Pool	Select a virtual server pool from the list, or create a new one.
Weight	Assign a weight. Valid values range from 1 to 255.
Topology	Select a topology from the list, or create a new one.
ISP	Select an ISP from the list or create a new one.

Configuring wizard

The GLB wizard is a step-by-step tool to help you configure a GLB object in GUI.

To use the GLB Wizard:

Go to **Global Load Balance**. Look for a green tab all the way in the top right corner, **Global Load Balance Wizard**.




Click the wizard, and open up the following dialogue.

To configure a GLB wizard:

1. In **Server**, configure the Name, Address, and Data Center Location.
2. In **Virtual Server Pool**, configure the Name, Preferred and Alternate. Discover the server members that were previously configured in **Server**, and select from the given list.
3. It is required to specify the Name, Host Name, and Domain Name. You can also, if you want, specify the Default Feedback IPv4 or Default Feedback IPv6.
4. After clicking **Finished**, the Global DNS Configuration will be complete. The **Global DNS Configuration** radio button in **GLB > Zone Tools > General Settings** will be enabled automatically.

In **Zone Tools**, three things will happen:

1. Under **Zone**, your new zone will appear.
2. Under **Global DNS Policy**, select the first policy. If more than one policy exists, choose the first one.



3. A dialogue opens. Under **Zone List**, you will see that your zone has moved automatically from **Available Items** to **Selected Items**.

 Global DNS Policy

Name
DEFAULT_DNS_POLICY

Source
any

Destination
any

Zone List
Selected Items
fqdn_generate_example.com.

<

>

Available Items

Double-click to deselect. Drag to reorder.

Double-click to select.

DNS64 List
Selected Items

<

>

Available Items

Double-click to deselect. Drag to reorder.

Double-click to select.

Recursion
☒ ON

DNSSEC
☐ OFF

DNSSEC Validation
☐ OFF

Forward
☒ First ☐ Only

Forwarders
Click to select

Response Rate Limit
Click to select

Save

Cancel

Configuring virtual server pools

The virtual server pool configuration defines the set of virtual servers that can be matched in DNS resource records, so it should include, for example, all the virtual servers that can be answers for DNS requests to resolve `www.example.com`.

You also specify the key parameters of the global load balancing algorithm, including proximity options, status checking options, load balancing method, and weight.

The DNS response is an ordered list of answers. Virtual servers that are unavailable are excluded. Available virtual servers are ordered based on the following priorities:

1. Geographic proximity
2. Dynamic proximity
3. Weighted round robin

A client that receives DNS response with a list of answers tries the first and only proceeds to the next answers if the first answer is unreachable.

Before you begin:

- You must have created GLB Servers configuration objects.
- You must have Read-Write permission for Global Load Balance settings.

After you have created a virtual server pool configuration object, you can specify it in the global load balancing host configuration.

To configure a virtual server pool:

1. Go to Global Load Balance > FQDN Settings.
2. Click the **Virtual Server Pool** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Virtual server pool configuration on page 225](#).
5. Save the configuration.

Virtual server pool configuration

Settings	Guidelines
Name	Specify a unique name for the virtual server pool configuration. Valid characters are A-Z, a-z, 0-9, <code>_</code> , and <code>-</code> . No spaces. You reference this name in the host configuration. Note: After you initially save the configuration, you cannot edit the name.
Preferred	• None—No preference. • Geo—If selected, virtual servers with the same GEO information as the local DNS address will respond. • Geo-ISP—If selected, virtual servers with the same ISP information as the local DNS address will respond first, and virtual servers with the same GEO information as the local DNS address will respond second. • RTT—Virtual servers with the shortest latency link or closest to the data center will respond. • Least-Connections—Virtual servers with the least connections will respond. • Connection-Limit—Virtual servers will be responded by their connection limit determined by virtual servers' weight: the greater the weight of a virtual server, the more responses it will get.

Settings	Guidelines
	- Bytes-Per-Second—Virtual servers with the lowest traffic will respond. - Server-Performance—Virtual servers with better server-performance in the CPU or Memory (whichever one you give more weight to) will respond.
Alternate	Same as above.
Load Balance Method	Weighted Round Robin
Check Server Status	Enable/disable polling of the local FortiADC SLB. If the server is unresponsive, its virtual servers are not selected for DNS answers.
Check Virtual Server Existence	Enable/disable checks on whether the status of the virtual servers in the virtual server list is known. Virtual servers with unknown status are not selected for DNS answers.
Member	
Server	Select a GLB Servers configuration object.
Server Member	Select the name of the virtual server that is in the servers virtual server list configuration.
Weight	Assigns relative preference among members—higher values are more preferred and are assigned connections more frequently. The default is 1. The valid range is 1-255.
Backup	Enable to designate the member as a backup. Backup members are inactive until all main members are down.

Configuring location lists

A location list configuration consists of a list of locations you select.

To configure a location list:

1. Go to Global Load Balance > FQDN Settings.
2. Click the **Location List** tab.
3. Complete the configuration as described in [Location List settings on page 226](#).
4. Click Save.

Location List settings

Settings	Guidelines
Name	Specify a unique name for the location list.

Settings	Guidelines
GEO IP List	Create a GEO IP list: 1. Click inside the box. 2. Select an option from the drop-down list menu. 3. Repeat Steps 1 and 2 to add more locations to the list. Note: To remove an entry off your list, click the corresponding x sign.

Note: Starting from FortiADC 5.x.x, the GUI shows the full country or region names listed in alphabetical order for location list and data center configuration. The Console uses country or region name abbreviations instead. The abbreviations are done in accordance with the ISO standards. So if you configure a location list or data center from the Console, be sure to consult [ISO-3166-1](#) and/or [ISO 3166-2:CN](#) for the correct abbreviations to use. See the following example commands:

```
config global-load-balance topology
edit "tpl"
set member ZZ US CN65
next
end
```

Where ZZ stands for Reserved, US for United States, and CN65 for China, Xingjiang

Logical Topology

The FortiView>**Global Load Balance**>**Logical Topology** page shows the logical topology of your global load balance configurations.

Adding hosts

To add a host:

1. Click Add Host.
2. Make desired entries or selections as described in [Configuring hosts on page 221](#)
3. Click Save when done.

Note: While in Editor View, you click any component in the logical topology to edit or delete it.

Filtering hosts

The Add Filters button on top of the page allows you to customize the logical topology by:

- Availability
- Host
- Domain Name
- VS Pool
- Server
- Server Member
- Data Center

To add a filter:

1. Click the Add Filters button.
2. Select the desired filter from the drop-down list menu.

Note: You can use the same steps to apply multiple filters. Applied filters appear in front of the Add Filters button in the order they are added. You can remove a filter by clicking the x sign on it.

Configuring a Global DNS policy

The Global DNS policy is a rule base that matches traffic to DNS zones. Traffic that matches both the source and the destination criteria is served by the policy. Traffic that does not match any policy is served by the DNS “general settings” configuration.

Before you begin, you must have:

- A good understanding of DNS and knowledge of the DNS deployment in your network.
- Configured address objects, remote servers, DNS zones, and optional configuration objects you want to specify in your policy.
- Read-Write permission for Global Load Balance settings.

To configure the global DNS policy rule base:

1. Go to Global Load Balance > Zone Tools.
2. Click the **Global DNS Policy** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Global DNS policy configuration on page 228](#).
5. Save the configuration.
6. Reorder rules, as necessary.

Global DNS policy configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Source	Select an address object to specify the source match criteria. See Configuring an address group .
Destination	Select an address object to specify the destination match criteria. See Configuring an address group .
Zone List	Select one or more zone configurations to serve DNS requests from matching traffic. See Configuring DNS zones .
DNS64 List	Select one or more DNS64 configurations to use when resolving IPv6 requests. See Configuring DNS64 .

Settings	Guidelines
Recursion	Enables/disables recursion. If enabled, the DNS server attempts to do all the work required to answer the query. If not enabled, the server returns a referral response when it does not already know the answer.
DNSSEC	Enables/disables DNSSEC.
DNSSEC Validation	Enables/disables DNSSEC validation.
Forward	- First—The DNS server queries the forwarders list before doing its own DNS lookup. - Only—Only queries the forwarders list. Does not perform its own DNS lookups. Note: The internal server caches the results it learns from the forwarders, which optimizes subsequent lookups.
Forwarders	If the DNS server zone has been configured as a forwarder, select the remote DNS server to which it forwards requests. See Configuring remote DNS servers .
Response Rate Limit	Select a rate limit configuration object. See Configuring the response rate limit .
Reordering	
	After you have saved a rule, reorder rules as necessary. The rules table is consulted from top to bottom. The first rule that matches is applied and subsequent rules are not evaluated.

Configuring DNS zones

The DNS zone configuration is the key to the global load balancing solution. This configuration contains the key DNS server settings, including:

- Domain name and name server details.
- Type—Whether the server is the master or a forwarder.
- DNSSEC—Whether to use DNSSEC.
- DNS RR records—The zone configuration contains resource records (RR) used to resolve DNS queries delegated to the domain by the parent zone.

You can specify different DNS server settings for each zone you create. For example, the DNS server can be a master for one zone and a forwarder for another zone.

Before you begin:

- You must have a good understanding of DNS and knowledge of the DNS deployment in your network.
- You must have authority to create authoritative DNS zone records for your network.
- You must have Read-Write permission for Global Load Balance settings.

After you have configured a DNS zone, you can select it in the DNS policy configuration.

To configure the DNS zone:

1. Go to Global Load Balance > Zone Tools.
2. Click the **Zone** tab.

3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [DNS zone configuration on page 230](#).

DNS zone configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference the name in the global DNS policy configuration. Note: - FortiADC supports third-party domain names. - After you initially save the configuration, you cannot edit the name.
Type	- Master—The configuration contains the “master” copy of data for the zone and is the authoritative server for it. - Forward—The configuration allows you to apply DNS forwarding on a per-domain basis, overriding the forwarding settings in the “general” configuration. - FQDN Generate—The zone and its resource record is generated from the global load balancing framework.
Domain Name	The domain name must end with a period. For example: <code>example.com.</code>
DNS policy	Select the DNS policy you want the zone to use.
Forward Options	
Forward	- First—The DNS server queries the forwarder before doing its own DNS lookup. - Only—Only query the forwarder. Do not perform a DNS lookup. - Note: The internal server caches the results it learns from the forwarders, which optimizes subsequent lookups.
Forwarders	Select a remote server configuration object.
Master Options	
TTL	The \$TTL directive at the top of the zone file (before the SOA) gives a default TTL for every RR without a specific TTL set. The default is 86,400. The valid range is 0 to 2,147,483,647.
Negative TTL	The last field in the SOA—the negative caching TTL. This informs other servers how long to cache no-such-domain (NXDOMAIN) responses from you. The default is 3600 seconds. The valid range is 0 to 2,147,483,647.
Responsible Mail	Username of the person responsible for this zone, such as <code>hostmaster.example.com..</code> Note: Format is <code>mailbox-name.domain.com.</code> (remember the trailing dot). The format uses a dot, not the @ sign used in email addresses because @ has other uses in the zone file. Email, however, is sent to <code>hostmaster@example.com</code> .
Primary Server Name	Sets the server name in the SOA record.
Primary Server Address	The IP address of the primary server.
DNSSEC	Enable/Disable DNSSEC

Settings	Guidelines
	Only when a DNS policy has been set, <i>and</i> DNSSC is enabled, will the Back Up DSSET Key appear. It includes the following types of key: • KSK. Type characters for a string key. To regenerate the KSK, disable and re-enable DNSSEC. • ZSK. Type characters for a string key. To regenerate the ZSK, disable and re-enable DNSSEC. • DSSET. It is generated by the system if DNSSEC is enabled for the zone.
DSSET List	Select a DSSET configuration object. See Configuring the DSSET list .
Serial	Set the serial number of the zone. Default 10004. Range 1-4294967295.
Notify Status	Enable/Disable notify status. The IP in "also notify IP list" will be notified only when Notify Status is enabled.
Also Notify IP List	Set a list of IP addresses that will be notified if Notify Status is enabled.
Allow Transfer	Defines a list of IP addresses that are allowed to transfer the DNS zone information. By default there will be "Any" and "None."
FQDN Record	
FQDN Record table	Displays a summary of all DNS RR for the zone, including generated and manually configured RR.
A/AAAA Record	
Hostname	The hostname part of the FQDN, such as <code>www</code>. Note: You can specify the <code>@</code> symbol to denote the zone root. The value substituted for <code>@</code> is the preceding <code>\$ORIGIN</code> directive.
Type	• IPv4 • IPv6
Weight	Assigns relative preference among members—higher values are more preferred and are assigned connections more frequently. The default is 1. The valid range is 1-255.
Address	Specify the IP address of the virtual server.
Method	Weighted Round Robin is the only method supported.
CNAME Record	
Alias	An alias name to another true or canonical domain name (the target). For instance, <code>www.example.com</code> is an alias for <code>example.com</code> .
Target	The true or canonical domain name. For instance, <code>example.com</code> .
NS Record	
Domain Name	The domain for which the name server has authoritative answers, such as <code>example.com</code> .

Settings	Guidelines
	Note: FortiADC supports third-party domain names.
Hostname	The hostname part of the FQDN, such as <code>ns</code> .
Type	• IPv4 • IPv6
Address	Specify the IP address of the name server.
MX Record	
Hostname	The hostname part of the FQDN for a mail exchange server, such as <code>mail</code> .
Priority	Preference given to this RR among others at the same owner. Lower values have greater priority.
Type	• IPv4 • IPv6
Address	Specify the IP address.
TXT Record	
Name	Hostname. TXT records are name-value pairs that contain human readable information about a host. The most common use for TXT records is to store SPF records.
Text	Comma-separated list of name=value pairs. An example SPF record has the following form: <pre>v=spf1 +mx a:colo.example.com/28 -all</pre> If you complete the entry from the the Web UI, do not put the string in quotes. (If you complete the entry from the CLI, you do put the string in quotes.)
SRV Record	
Host Name	The host name part of the FQDN, e.g., <code>www</code> .
Priority	A priority assigned to the target host: the lower the value, the higher the priority.
Weight	A relative weight assigned to a record among records of the same priority: the greater the value, the more weight it carries.
Port	The TCP or UDP port on which the service is provided.
Target Name	The canonical name of the machine providing the service.
PTR Record	
PTR address	A PTR address, such as <code>10.168.192.in-addr.arpa</code> . or <code>1</code>
FQDN	A fully qualified domain name, such as <code>"www.example.com"</code> .

Configuring general settings

The general settings configuration specifies the interfaces that listen for DNS requests. By default, the system listens on the IPv4 and IPv6 addresses of all configured interfaces for DNS requests.

The other settings in the general settings configuration are applied when traffic does not match a Global DNS policy.

Before you begin:

- You must have a good understanding of DNS and knowledge of the DNS deployment in your network.
- You must have Read-Write permission for Global Load Balance settings.

To configure general settings:

1. Go to Global Load Balance > Zone Tools.
2. Click the **General Settings** tab.
3. Complete the configuration as described in [General configuration on page 233](#).
4. Save the configuration.

General configuration

Settings	Guidelines
Global DNS Configuration	Enables/disables this configuration.
Recursion	Enables/disables recursion. If enabled, the DNS server attempts to do all the work required to answer the query. If not enabled, the server returns a referral response when it does not already know the answer.
DNSSEC	Enables/disables DNSSEC.
DNSSEC Validation	Enables/disables DNSSEC validation.
Listen on IPv6	Enables/disables listening for DNS requests on the interface IPv6 address.
Listen on IPv4	Enables/disables listening for DNS requests on the interface IPv4 address.
Traffic Log	Enables/disables traffic log.
Listen on All Interface	Enables listening on all interfaces.
Forward	• First—The DNS server queries the forwarder before doing its own DNS lookup. • Only—Only queries the forwarder. Does not perform its own DNS lookups. Note: The internal server caches the results it learns from forwarders, which optimizes subsequent lookups.
Use System DNS Server	Forwards DNS requests to the system DNS server instead of the forwarders list.
Response Rate Limit	Selects a rate limit configuration object. See Configuring the response rate limit .

Configuring the trust anchor key

DNSSEC validation requires that a DNS name server know the trust anchor key for the root DNS domain in order to validate already signed responses. In general, trust anchor keys do not change often, but they do change occasionally, and might change unexpectedly in the event the keys are compromised.

The FortiADC DNS server is preconfigured with a trust anchor key for the root DNS domain. If you are informed that you must update this key, you can use the configuration editor to paste the new content into the DNS server configuration.

Further reading:

<http://data.iana.org/root-anchors/draft-icann-dnssec-trust-anchor.html>

Before you begin:

- You must have a good understanding of DNSSEC and knowledge of the DNS deployment in your network.
- You must have already obtained the key so that you can copy and paste it into the DNS server configuration.
- You must have Read-Write permission for Global Load Balance settings.

To configure the trust anchor key:

1. Go to Global Load Balance > Zone Tools.
2. Click the **Trust Anchor Key** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Trust anchor key configuration on page 234](#).
5. Save the configuration.

Trust anchor key configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Value	The key value. The key format is a string with the following format: <code>"<domainname>" <num1> <num2> <num3> "<content>"</code> The following is an example: <code>".\" 256 3 5 \"AwEAAbDrWmiIReotvZ6FObgKygZwUxSUJW9z5pjiQMLH0JBGXooHrR16 pdKhI9mNkM8bLUMtwYfgeUOYXIvfagee8rk=\"</code>
Description	Description for the key.

Configuring DNS64

The DNS64 configuration maps IPv4 addresses to AAAA queries when there are no AAAA records. This feature is optional. It can be used in network segments that use NAT64 to support IPv6 client communication with IPv4 backend servers.

Before you begin:

- You must have a good understanding of DNS and knowledge of the DNS deployment in your network.
- You must have configured address objects that specify the network segments for which the DNS64 map applies. See [Configuring an address group](#).
- You must have Read-Write permission for Global Load Balance settings.

After you have created a DNS64 configuration, you can select it a DNS policy configuration.

To configure DNS64:

1. Go to Global Load Balance > Zone Tools.
2. Click the **DNS64** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [DNS64 configuration on page 235](#).

DNS64 configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference the name in the global DNS policy configuration. After you initially save the configuration, you cannot edit the name.
IPv6 Prefix	IP address and netmask that specify the DNS64 prefix. Compatible IPv6 prefixes have lengths of 32, 40, 48, 56, 64 and 96 as per RFC 6052. Each DNS64 configuration has one prefix. Multiple configurations can be defined.
Source Address	Select an address object. Only clients that match the source IP use the DNS64 lookup table.
Mapped Address	Select an address object that specifies the IPv4 addresses that are to be mapped in the corresponding A RR set.
Exclude	Select an address object. Allows specification of a list of IPv6 addresses that can be ignored. Typically, you exclude addresses that do have AAAA records.

Configuring the DSSET list

If you enable DNSSEC, secure communication between the FortiADC DNS server and any child DNS servers is based on keys contained in delegation signer files (DSSET files). In DNSSEC deployments, DSSET files are generated automatically when the zone is signed by DNSSEC.

You use the DSSET list configuration to paste in the content of the DSSET files provided by child domain servers or stub domains.

Note: You use the Global DNS zone configuration to generate the DSSET file for this server. The file generated by the zone configuration editor is the one you give to any parent zone or the registrar of your domain.

Before you begin:

- You must have a good understanding of DNSSEC and knowledge of the DNS deployment in your network.
- You must have used DNSSEC to sign the child domain servers and have downloaded the DSset files to a location you can reach from your management computer.
- You must have Read-Write permission for Global Load Balance settings.

After you have configured a DSSET list, you can select it in DNS zone configuration.

To configure the DSSET list:

1. Go to Global Load Balance > Zone Tools.
2. Click the **DSSET List** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [DSset list configuration on page 236](#).

DSset list configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference the name in the zone configuration (if you enable DNSSEC). After you initially save the configuration, you cannot edit the name.
Filename	Type the filename. The convention is dsset-<domain>, for example, dsset-example.com.
Content	Paste the DSset file content. The content of DSset files is similar to the following: <pre>dns.example.com. IN DS 13447 5 1 A5AD9EFB6840F58CF817F3CC7C24A7ED2DD5559C</pre>

Configuring an address group

An address group is a configuration object that specifies the source and destination IP addresses that are the matching criteria for DNS policies.

Before you begin:

- You must have Read-Write permission for Global Load Balance settings.

After you have configured an address group, you can select it in the DNS policy configuration.

To configure address groups:

1. Go to Global Load Balance > Zone Tools.
2. Click the **Address Group** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration and add members as described in [Address group configuration on page 236](#)

Address group configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the global DNS policy configuration. Note: After you initially save the configuration, you cannot edit the name.

Settings	Guidelines
Member	
Address Type	- IPv4 - IPv6
IP/Netmask	Address/mask notation to match the IP address in the packet header. Create objects to match source IP address and different objects to match destination IP address.
Action	- Include—The rule logic creates an address object that includes addresses matching the specified address block. - Exclude—The rule logic creates an address object that excludes addresses matching the specified address block.

Configuring remote DNS servers

The remote server configuration is used to create a list of DNS forwarders. DNS forwarders are commonly used when you do not want the local DNS server to connect to Internet DNS servers. For example, if the local DNS server is behind a firewall and you do not want to allow DNS through that firewall, you implement DNS forwarding to a remote server that is deployed in a DMZ or similar network region that can contact Internet DNS servers.

Before you begin:

- You must have a good understanding of DNS and knowledge of the remote DNS servers that can be used to communicate with Internet domain servers.
- You must have Read-Write permission for Global Load Balance settings.

After you have configured remote DNS servers, you can select them in DNS zone and DNS policy configurations.

To configure a remote server:

- Go to Global Load Balance > Zone Tools.
- Click the **Remote DNS Server** tab.
- Click **Create New** to display the configuration editor.
- Complete the configuration and add members as described in [Remote DNS server configuration on page 237](#).

Remote DNS server configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the zone configuration (if you use forwarders). Note: After you initially save the configuration, you cannot edit the name.
Member	
Address Type	- IPv4 - IPv6

Settings	Guidelines
Address	IP address of the remote DNS server.
Port	Port number the remote server uses for DNS. The default is 53.

Configuring the response rate limit

The response rate limit keeps the FortiADC authoritative DNS server from being used in amplifying reflection denial of service (DoS) attacks.

Before you begin:

- You must have a good understanding of DNS.
- You must have Read-Write permission for Global Load Balance settings.

After you have created a response rate limit configuration, you can select it in the DNS policy and DNS general settings configurations.

To configure the response rate limit:

1. Go to Global Load Balance > Zone Tools.
2. Click the **Response Rate Limit** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Response rate limit configuration on page 238](#).

Response rate limit configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference the name in the global DNS policy configuration. After you initially save the configuration, you cannot edit the name.
Responses per Second	Maximum number of responses per second. The valid range is 1-2040. The default is 1000.

Chapter 7: Network Security

This chapter includes the following topics:

- [Security features basics on page 239](#)
- [Managing IP Reputation policy settings on page 239](#)
- [Configure IP reputation exception on page 241](#)
- [Configure IP reputation black list on page 241](#)
- [Using the Geo IP block list on page 242](#)
- [Using the Geo IP whitelist on page 244](#)
- [Special Geo codes on page 245](#)
- [Enabling denial of service protection on page 245](#)
- [Configuring an IPv4 firewall policy on page 246](#)
- [Configuring an IPv6 firewall policy on page 247](#)
- [Configuring an IPv4 connection limit policy on page 249](#)
- [Configuring an IPv6 connection limit policy on page 250](#)
- [Anti-virus on page 251](#)
- [Configuring IPS on page 258](#)

Security features basics

In most deployment scenarios, we recommend you deploy FortiGate to secure your network. Fortinet includes security functionality in the FortiADC system to support those cases when deploying FortiGate is impractical. FortiADC includes the following security features:

- Firewall—Drop traffic that matches a source/destination/service tuple you specify.
- Security connection limit—Drop an abnormally high volume of traffic from a source/destination/service match.
- IP Reputation service—Drop or redirect traffic from source IPs that are on the FortiGuard IP Reputation list.
- Geo IP—Drop or redirect traffic from source IPs that correspond with countries in the FortiGuard Geo IP database.
- Web application firewall—Drop or alert when traffic matches web application firewall attack signatures and heuristics.
- Denial of service protection—Drop half-open connections to protect the system from a SYN flood attack.
- IPS—Protect the system based on a robust FortiGuard pattern / signature-based engine.

Managing IP Reputation policy settings

The FortiGuard IP Reputation service provides a database of known compromised or malicious client IP addresses. The database is updated periodically.

The IP Reputation configuration allows you to specify the action the system takes when an SLB virtual server receives traffic from a client with an IP address on the list. [IP Reputation actions on page 240](#) lists limitations for IP Reputation actions.

IP Reputation actions

Action		Profile Limitations
Pass	IPv4 only	Not supported for RADIUS.
Deny	IPv4 only	Not supported for RADIUS.
Redirect	IPv4 only	Not supported for RADIUS, FTP, TCP, UDP.
Send 403 Forbidden	IPv4 only	Not supported for RADIUS, FTP, TCP, UDP.

Note: IP Reputation is also not supported for Layer 4 virtual servers when the Packet Forwarding Mode is Direct Routing.

Basic Steps

1. Configure the connection to FortiGuard so the system can receive periodic IP Reputation Database updates. See [Configuring FortiGuard service settings](#).
2. Optionally, customize the actions you want to take when the system encounters a request from a source IP address that matches the list; and add exceptions. If a source IP address appears on the exceptions list, the system does not look it up on the IP Reputation list. See below.
3. Enable IP Reputation in the profiles you associate with virtual servers. See [Configuring Application profiles](#).

Before you begin:

- You must have Read-Write permission for Firewall settings.

To customize IP Reputation policy rules:

1. Go to Network Security > IP Reputation.
2. Make sure to select the **IP Reputation** tab, which displays all IP reputation policy configuration in FortiADC.
3. Click a policy or the corresponding Edit icon to open the IP Reputation editor.
4. Make the desired changes as described in [IP Reputation policy configuration on page 240](#).
5. Click Save.

IP Reputation policy configuration

Settings	Guidelines
Category	Depending the configuration on FortiGuard IP Reputation service, the IP reputation policy can be one of the following categories: • Botnet • Anonymous Proxy • Phishing • Spam • Others • Black List

Settings	Guidelines
Status	Enable or disable the category.
Action	• Pass • Deny • Redirect • Send 403 Forbidden Note: Layer 4 and TCPS virtual servers do not support Redirect or Send 403 Forbidden. If you apply an IP Reputation configuration that uses these options to a Layer 4 or TCPS virtual server, FortiADC logs the action as Redirect or Send 403 Forbidden but in fact denies the traffic.
Severity	The severity to apply to the event. Severity is useful when you filter and sort logs: • Low • Medium • High
Log	Enable or disable logging.

Configure IP reputation exception

To create an IP Reputation exception:

1. Go to Network Security > IP Reputation.
2. Click the **IP Reputation Exception** tab to add exceptions as described in [IP Reputation exception on page 241](#).
3. Click Save.

IP Reputation exception

Settings	Guidelines
Status	Enable or disable the exception. You might have occasion to toggle the exception off and on.
Type	• IP/netmask: Select this option to allow a specified IP address to pass through. • IP Range: Select this option to allow a specified range of IP addresses to pass through.
IP/Netmask	If IP/netmask is selected in the Type field above, specify a subnet using the address/mask notation.
Start IP / End IP	If IP Range is selected in the Type field above, specify the starting address and ending address of the IP range.

Configure IP reputation black list

Upload the source IP's or CIDRs that you want the ADC to block in the IP reputation black list. When these source IP's try to access the VS, the connection will fail. You can create IP/Netmask or IP Range type black list, back up or restore

files.

The content of IP reputation black list file should be coded in ASCII and every line can be a IP netmask or IP address range. There can be 256 IP netmasks or IP address ranges in the file. It looks like this:

192.168.1.1-192.168.1.10

172.16.1.1-172.16.2.100

10.1.1.0/24

20.1.1.0/24

You use the **Restore** utility to import the file and the **Back Up** utility to export it.

You use the **Clean** utility to erase entries that were imported from the text file. The clean operation does not affect the user-configured entries.

To create an IP Reputation black list:

1. Go to Network Security > IP Reputation
2. Click the **IP Reputation Black List** tab to **Create New** black lists as described in [IP Reputation black list on page 242](#).
3. Click Save.

IP Reputation black list

Settings	Guidelines
Status	Enable or disable the exception. You might have occasion to toggle the exception off and on.
Type	• IP/netmask: Select this option to allow a specified IP address to pass through. • IP Range: Select this option to allow a specified range of IP addresses to pass through.
IP/Netmask	If IP/netmask is selected in the Type field above, specify a subnet using the address/mask notation.
Start IP / End IP	If IP Range is selected in the Type field above, specify the starting address and ending address of the IP range.

Using the Geo IP block list

The FortiGuard Geo IP service provides a database that maps IP addresses to countries, satellite providers, and anonymous proxies. The database is updated periodically.

The Geo IP block list is a policy that takes the action you specify when the virtual server receives requests from IP addresses in the blocked country's IP address space.

For Layer 4 virtual servers, FortiADC blocks access when the first TCP SYN packet arrives. For Layer 7 virtual servers, FortiADC blocks access after the handshake, allowing it to redirect the traffic if you have configured it to do so.

[Geo IP block list actions on page 243](#) lists limitations for Geo IP block list actions.

Geo IP block list actions

Action		Profile Limitations
Pass	IPv4 only	Not supported for HTTP Turbo, RADIUS.
Deny	IPv4 only	Not supported for HTTP Turbo, RADIUS.
Redirect	IPv4 only	Not supported for HTTP Turbo, RADIUS, FTP, TCP, TCPS, UDP.
Send 403 Forbidden	IPv4 only	Not supported for HTTP Turbo, RADIUS, FTP, TCP, TCPS, UDP.

Basic Steps

1. Configure the connection to FortiGuard so the system can receive periodic Geo IP Database updates. See [Configuring FortiGuard service settings](#).
2. Create rules to block traffic from locations.
3. Maintain a whitelist to allow traffic from specified subnets even if they belong to the address space blocked by the Geo IP block list.
4. Select the Geo IP block list and whitelist in the profiles you associate with virtual servers. See [Configuring Application profiles](#).

Before you begin:

- You must have Read-Write permission for Security settings.

To configure a Geo IP block list:

1. Go to Network Security > Geo IP Protection.
2. Click the **Geo IP Protection** tab.
3. Click Create New to create a block list as described in [Geo IP block list configuration on page 243](#).
4. Click Save.
5. Edit your new block list to add members as described in [Geo IP block list configuration on page 243](#).
6. Click **Save** to save your member settings.
7. Click **Save**.

Geo IP block list configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Default Action	• Pass—Allow the traffic. • Deny—Drop the traffic. • Redirect—Send a redirect. You specify the redirect URL on the profile configuration page. • Send 403 Forbidden—Send the HTTP Response code 403. Note: Layer 4 and TCPS virtual servers do not support Redirect or Send 403 Forbidden. If you apply an Geo IP configuration that uses these options to a Layer 4 or TCPS virtual server, FortiADC logs the action as Redirect or Send 403 Forbidden, but in fact denies the traffic.

Settings	Guidelines
Status	Enable or disable the Geo IP block list configuration.
Member	
Log	Enable/disable logging.
Severity	The severity to apply to the event. Severity is useful when you filter and sort logs: • Low • Medium • High
Action	• Pass—Allow the traffic. • Deny—Drop the traffic. • Redirect—Send a redirect. You specify the redirect URL on the profile configuration page. • Send 403 Forbidden—Send the HTTP Response code 403. Note: Layer 4 and TCPS virtual servers do not support Redirect or Send 403 Forbidden. If you apply an Geo IP configuration that uses these options to a Layer 4 or TCPS virtual server, FortiADC logs the action as Redirect or Send 403 Forbidden, but in fact denies the traffic.
Regions	Select a geolocation object. The list includes countries as well as selections for anonymous proxies and satellite providers.

Using the Geo IP whitelist

To configure a Geo IP whitelist:

1. Go to Network Security > Geo IP Protection.
2. Click the **Whitelist** tab to create a whitelist as described in [Geo IP whitelist configuration on page 244](#).
3. Click Save.

Geo IP whitelist configuration

Settings	Guidelines
Name	Configuration name. The name can be up to 35 characters long. Valid characters are A-Z, a-z, 0-9, _, and -. No space is allowed. After you initially save the configuration, you cannot edit the name.
Description	A string to describe the purpose of the configuration, to help you and other administrators more easily identify its use.
Status	Enable/disable the exception. You might have occasion to toggle the exception off and on.
Member	
Type	Select and configure either of the following:

Settings	Guidelines
	IP Subnet—Specify the IP address and CIDR-formatted subnet mask, separated by a forward slash (/), such as 192.0.2.0/24. Dotted quad formatted subnet masks are not accepted. IPv6 addresses are not supported. IP Range—Specify the Start IP and the End IP addresses of the IP range.
Description	Enter a brief description of the IP subnet or IP range, depending on which Type you choose. The description can be up to 1023 characters long. Valid characters are A-Z, a-z, 0-9, _, -, ., and :. No space is allowed.

Special Geo codes

[Special GEO codes and their usage on page 245](#) below describes the usage of special GEO codes.

Special GEO codes and their usage

GEO code	Usage
ZZ	Reserved (IP addresses that are not assigned, e.g., 10.0.0.0/24)
A1	Anonymous Proxy (IP addresses that are defined as anonymous proxy in MaxMind, e.g., 46.19.137.0/24)
A2	Satellite Provider (IP addresses that are defined as satellite provider in MaxMind, e.g., 57.72.6.0/24)
O1	Other Country (Reserved for further use, and no IP address is assigned to this region)

Enabling denial of service protection

You can enable basic denial of service (DoS) prevention to combat [SYN floods](#). When enabled, FortiADC uses the SYN cookie method to track half-open connections. The system maintains a DoS mitigation table for each configured IPv4 virtual server. It times out half-open connections so that they do not deplete system resources.

Note: The DoS feature is supported for traffic to virtual servers only. However, it is not supported for IPv6 traffic or for Layer 4 virtual servers with the Direct Routing packet forwarding mode.

Before you begin:

- You must have Read-Write permission for Firewall settings.

To enable denial of service protection:

- Go to Security > SYN Flood Prevention.
- Enable the SYN Cookie feature.
- Specify a maximum number of half open sockets. The default is 1 (10 connections). The valid range is 1 to 80,000.
- Save the configuration.

Configuring an IPv4 firewall policy

A firewall policy is a filter that allows or denies traffic based on a matching tuple: source address, destination address, and service. By default, firewall policy rules are stateful: if client-to-server traffic is allowed, the session is maintained in a state table, and the response traffic is allowed.

The FortiADC system evaluates firewall policies before other rules. It matches traffic against the firewall policy table, beginning with the first rule. If a rule matches, the specified action is taken. If the session is denied by a firewall policy rule, it is dropped. If the session is accepted, system processing continues.

By default, if firewall rules are not configured, the system does not perform firewall processing; all traffic is processed as if the system were a router, and traffic is forwarded according to routing and other system rules.

Note: You do not need to create firewall rules for routine management traffic associated with the management port or HA ports. The interface “allow access” option enables permitted protocols. The system automatically permits from-self traffic, such as health check traffic, and expected responses.

Before you begin:

- You must have a good understanding and knowledge of firewalls.
- You must have created the address configuration objects and service configuration objects that define the matching tuple in your firewall policy rules.
- You must have Read-Write permission for Firewall settings.

To configure a firewall:

1. Go to Network Security > Firewall > IPv4 Firewall Policy.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in [Firewall policy configuration on page 246](#).
4. Save the configuration.
5. Reorder rules, as necessary.

Firewall policy configuration

Settings	Guidelines
Default Action	Action when no rule matches or no rules are configured: • Deny—Drop the traffic. • Accept—Allow the traffic to pass the firewall.
Rule	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Ingress Interface	Select the interface that receives traffic.
Egress Interface	Select an outgoing interface from the drop-down list <i>if your FortiADC is configured for link load-balancing and/or traffic routing</i> . In both cases, the system will use this interface to forward traffic to its destination.

Settings	Guidelines
	Note: You <i>MUST</i> leave this option blank (default) if your FortiADC is configured for server load-balancing and/or global load-balancing. Otherwise, server load-balancing and/or global load-balancing packets may not match the firewall policy rule.
Source	Select a source address object or address group to use to form the matching tuple.
Destination	Select a destination address object or address group to use to form the matching tuple.
Service	Select a service object to use to form the matching tuple.
Action	- Deny—Drop the traffic. - Accept—Allow the traffic to pass the firewall.
Status	Enabled by default. Note: This button simplifies the implementation of firewall policy/NAT rules, allowing you to turn a policy rule ON or OFF with a click of the button. When a firewall policy rule is disabled, it will be removed from the relevant IP tables, and will be added to the IP table when the rule is enabled.
Reordering	
Reorder	After you have saved a rule, reorder rules as necessary. The rules table is consulted from top to bottom. The first rule that matches is applied and subsequent rules are not evaluated.

Configuring an IPv6 firewall policy

A firewall policy is a filter that allows or denies traffic based on a matching tuple: source address, destination address, and service. By default, firewall policy rules are stateful: if client-to-server traffic is allowed, the session is maintained in a state table, and the response traffic is allowed.

The FortiADC system evaluates firewall policies before other rules. It matches traffic against the firewall policy table, beginning with the first rule. If a rule matches, the specified action is taken. If the session is denied by a firewall policy rule, it is dropped. If the session is accepted, system processing continues.

By default, if firewall rules are not configured, the system does not perform firewall processing; all traffic is processed as if the system were a router, and traffic is forwarded according to routing and other system rules.

Note: You do not need to create firewall rules for routine management traffic associated with the management port or HA ports. The interface “allow access” option enables permitted protocols. The system automatically permits from-self traffic, such as health check traffic, and expected responses.

Before you begin:

- You must have a good understanding and knowledge of firewalls.
- You must have created the address configuration objects and service configuration objects that define the matching tuple in your firewall policy rules.
- You must have Read-Write permission for Firewall settings.

To configure a firewall:

1. Go to Network Security > Firewall > IPv6 Firewall Policy.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in [Firewall policy configuration on page 248](#).
4. Save the configuration.
5. Reorder rules, as necessary.

Firewall policy configuration

Settings	Guidelines
Default Action	Action when no rule matches or no rules are configured: • Deny—Drop the traffic. • Accept—Allow the traffic to pass the firewall.
Rule	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Ingress Interface	Select the interface that receives traffic.
Egress Interface	Select an outgoing interface from the drop-down list <i>if your FortiADC is configured for link load-balancing and/or traffic routing</i> . In both cases, the system will use this interface to forward traffic to its destination. Note: You <i>MUST</i> leave this option blank (default) if your FortiADC is configured for server load-balancing and/or global load-balancing. Otherwise, server load-balancing and/or global load-balancing packets may not match the firewall policy rule.
Source	Select a source address object or address group to use to form the matching tuple.
Destination	Select a destination address object or address group to use to form the matching tuple.
Service	Select a service object to use to form the matching tuple.
Action	• Deny—Drop the traffic. • Accept—Allow the traffic to pass the firewall.
Status	Enabled by default. Note: This button simplifies the implementation of firewall policy/NAT rules, allowing you to turn a policy rule ON or OFF with a click of the button. When a firewall policy rule is disabled, it will be removed from the relevant IP tables, and will be added to the IP table when the rule is enabled.
Reordering	
Reorder	After you have saved a rule, reorder rules as necessary. The rules table is consulted from top to bottom. The first rule that matches is applied and subsequent rules are not evaluated.

Configuring an IPv4 connection limit policy

The firewall connection limit policy allows or denies traffic based on a matching tuple: source address, destination address, and service; and connection count. The purpose is to detect anomalous connection requests.

The limit you specify can be based on the following counts:

- Count of concurrent sessions that match the tuple.
- Count of concurrent sessions from a single host that match the tuple.

The FortiADC system evaluates firewall connection limit policy rules before other rules. It matches traffic against the connection limit table, beginning with the first rule. If no rule matches, the connection is forwarded for further processing. If a rule matches, and the limit has not been reached, the connection is forwarded for further processing. If a rule matches and the limit has been reached, the connection is dropped.

By default, if firewall connection limit rules are not configured, the system does not perform connection limit policy processing. The firewall connection limit can be configured for non-SLB traffic and for Layer 7 SLB traffic, but not Layer 4 SLB traffic.

Note: The purpose of the firewall connection limit is distinct from the virtual server connection limit. The firewall connection limit setting is a security setting; the virtual server connection limit is a capacity setting.

Before you begin:

- You must have a good understanding and knowledge of the capacity of your backend servers.
- You must have created the address configuration objects and service configuration objects that define the matching tuple in your connection limit rules.
- You must have Read-Write permission for Firewall settings.

To configure a firewall connection limit:

1. Click Network Security > Firewall > IPv4 Connection Limit Policy.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in [Connection limit configuration on page 249](#).
4. Save the configuration.
5. Reorder rules, as necessary.

Connection limit configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Ingress Interface	Select the interface that receives traffic.
Egress Interface	Select the interface that forwards traffic.
Source	Select a source address object to use to form the matching tuple.
Destination	Select a destination address object to use to form the matching tuple.
Service	Select a service object to use to form the matching tuple.

Settings	Guidelines
Type	Specify whether the limit is per rule or per host.
Side	When the connection limit is per host, specify whether the connection counter gets incremented when the host IP address appears in: • Source—Only increment the counter if the host is the source address. • Destination—Only increment the counter if the host is the destination address. • Both—Increment the counter if the host is the source or destination address.
Limit	Maximum concurrent sessions. The default is 1,048,576.
Reordering	
	After you have saved a rule, reorder rules as necessary. The rules table is consulted from top to bottom. The first rule that matches is applied and subsequent rules are not evaluated.

Configuring an IPv6 connection limit policy

The firewall connection limit policy allows or denies traffic based on a matching tuple: source address, destination address, and service; and connection count. The purpose is to detect anomalous connection requests.

The limit you specify can be based on the following counts:

- Count of concurrent sessions that match the tuple.
- Count of concurrent sessions from a single host that match the tuple.

The FortiADC system evaluates firewall connection limit policy rules before other rules. It matches traffic against the connection limit table, beginning with the first rule. If no rule matches, the connection is forwarded for further processing. If a rule matches, and the limit has not been reached, the connection is forwarded for further processing. If a rule matches and the limit has been reached, the connection is dropped.

By default, if firewall connection limit rules are not configured, the system does not perform connection limit policy processing. The firewall connection limit can be configured for non-SLB traffic and for Layer 7 SLB traffic, but not Layer 4 SLB traffic.

Note: The purpose of the firewall connection limit is distinct from the virtual server connection limit. The firewall connection limit setting is a security setting; the virtual server connection limit is a capacity setting.

Before you begin:

- You must have a good understanding and knowledge of the capacity of your backend servers.
- You must have created the address configuration objects and service configuration objects that define the matching tuple in your connection limit rules.
- You must have Read-Write permission for Firewall settings.

To configure a firewall connection limit:

1. Click Network Security > Firewall > IPv6 Connection Limit Policy.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in [Connection limit configuration on page 251](#).

4. Save the configuration.
5. Reorder rules, as necessary.

Connection limit configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Ingress Interface	Select the interface that receives traffic.
Egress Interface	Select the interface that forwards traffic.
Source	Select a source address object to use to form the matching tuple.
Destination	Select a destination address object to use to form the matching tuple.
Service	Select a service object to use to form the matching tuple.
Type	Specify whether the limit is per rule or per host.
Side	When the connection limit is per host, specify whether the connection counter gets incremented when the host IP address appears in: • Source—Only increment the counter if the host is the source address. • Destination—Only increment the counter if the host is the destination address. • Both—Increment the counter if the host is the source or destination address.
Limit	Maximum concurrent sessions. The default is 1,048,576.
Reordering	
Reorder	After you have saved a rule, reorder rules as necessary. The rules table is consulted from top to bottom. The first rule that matches is applied and subsequent rules are not evaluated.

Anti-virus


Malware and Advanced Persistent Threats (APT) can cause significant damage to the business of any organization. Malicious codes are commonly used to steal valuable data, gain unauthorized access to networks, or cause products to degrade.

Using a suite of integrated security technologies, Anti-virus (AV) solutions provide protection against a variety of threats, including both known and unknown malicious codes (Malware) and Advanced Targeted Attacks (ATA).

Integrated with the FortiOS AV engine, FortiADC provides an industry-class malware and APT detection and mitigation solution to our customers.

[AV module topology on page 252](#) illustrates how FortiADC's AV module works:

1. Automatically updates the latest attack signatures from FortiGuard to ensure real-time protection.
2. Submits all files, including suspicious files, to an on-premise appliance (FortiSandbox) or cloud-based service (FortiCloud Sandbox) for further analysis after performing the basic AV processing of its own.
3. Malicious files will be dropped or quarantined, and healthy ones will be forwarded to the backend servers.

AV module topology

To use the AV module, you must

- [Creating an AV profile on page 252](#)
- [Setting AV quarantine policies](#)
- [Setting AV service level](#)

Important Notes

- The AV feature does not support HA.
- If FortiADC is in HA mode, you must use the default source-ip for FortiSandbox.
- Try to limit the number of VDOMs when the AV feature is enabled. Otherwise, the capacity of quarantine may become limited.
- All file types are supported by AV feature.

Creating an AV profile

You must configure AV profiles to use the anti-virus service module, which can be done either from the GUI or the Console. Once created, you can include your AV profiles when creating advanced virtual server profiles that use the HTTP or HTTPS protocol. For more information, refer to [Configuring virtual servers](#).

Configure AV profiles from the GUI

To configure an AV profile from the GUI:

1. Click Network Security>Anti Virus.
2. Select the Profile tab.
3. Click the **Create New** button.
4. Make the entries or selections as described in [AV profile configuration on page 253](#).
5. Click Save when done.

AV profile configuration

Settings	Description
Name	A unique name for the AV profile. Note: An AV profile name can contain up to 63 alphanumeric characters.
Comments	A brief description of the profile. Note: A description can be up to 1024 alphanumeric characters long.
Uncomp Size Limit	The maximum size in MB of the memory buffer used to temporarily decompress files. Note: The default is 1 MB. Valid values range from 1 to 10 MB.
Uncomp Nest Limit	The maximum number of levels of nesting (compression) allowed for the system to decompress. Note: The default is 2. Valid values range from 2 to 100.
Scan Bzip2	Scan archives using the bzip2 algorithm. Note: Disabled by default.
Streaming Content Bypass	Enable or disable bypass streaming content (rather than buffering it). Note: Enabled by default.
Oversize Limit	The maximum in-memory file size in KB to be scanned. Note: The default is 1024 KB. Valid values range from 1 to 1024 KB.
Oversize	Select one of the options for the system to handle over-sized files: • Bypass—Ignore oversized files. • Log—Log and block oversized files. • Block—Block oversized files. Note: The default option is Bypass.
Options	Select an option for the system to handle infected files: • AV Monitor—Monitor and log infected files. • Quarantine—Monitor, log, and quarantine infected files. Note: The default is AV Monitor.
Emulator	Enable or disable the Win32 Emulator. Note: Disabled by default to improve throughput.
FSA Analytics	Select an option to submit files to to FortiSandbox. • Disable—No file is submitted. • Suspicious—Only suspicious files are submitted. • All—All files are submitted. Note: The default is Disable.

Settings	Description
Analytics Max Upload	The maximum file size in KB allowed to upload to FortiSandbox. Note: The default is 1024 KB. Valid values range from 1 to 2048 KB.
Analytics DB	Enable or disable supplementing the AV signature databases with the FortiSandbox signature database. Note: Disabled by default.
AV Virus Log	Enable or disable logging for anti-virus scanning. Note: Enabled by default.

Note that FortiADC currently imposes no restriction on the types of files that can be uploaded for AV analysis or evaluation. When scanning files for viruses, it makes no distinction between viruses and Trojans, and submits all suspicious files to FortiSandbox for evaluation. A log is generated whenever a file is uploaded to FortiSandbox.

Configure AV profiles from the Console

To configure an AV profile from the Console, execute the following commands:

```
config security antivirus profile
edit <name_str>
set comment <var-string>
set uncomp-size-limit <limit_int>
set uncomp-nest-limit <limit_int>
set scan-bzip2 {enable | disable}
set streaming-content-bypass {enable | disable}
set oversize-limit <size_int>
set oversize {bypass | log | block}
set options {avmonitor | quarantine}
set emulator {enable | disable}
set fsa-analytics {disable | suspicious | everything}
set analytics-max-upload <integer>
set analytics-db {disable | enable}
set av-virus-log {enable | disable}
end
```

Setting AV quarantine policies


The “quarantined” daemon manages the infected or suspicious files. The quarantine destination can be either the local hard disk.

It’s a multi-process daemon, which receives quarantine requests from the AV daemon and then processes the requests in child processes. It can work in tandem with remote devices to compliment the AV service, such as sending suspicious files to FortiSandbox for deeper inspection or uploading the archive package onto FortiCloud.

In addition, it also manages the use of the storage space, listing the quarantined files, deleting expired files, overriding old files, or dropping new files when there is no enough storage space available.

Note: For the 5.0.0 release, the AV module only supports quarantine on the hard disk and the integration with FortiSandbox, as illustrated in [AV quarantine process flow on page 255](#).

AV quarantine process flow



You can configure AV quarantine policies from the GUI or the Console.

Configuring AV quarantine policies from the GUI

To configure AV quarantine policies from the GUI:

1. Click Network Security>Anti Virus.
2. Click the **Quarantine** tab.
3. Make the entries or selections as described in [AV quarantine policy configuration on page 255](#).
4. Click Save when done.

AV quarantine policy configuration

Settings	Description
Destination	The destination for quarantined files, which could be either of the following: • NULL—Disable quarantine. • Disk—Send quarantined files to the hard disk.
Age Limit	The number of hours that quarantined files are kept on the hard disk. The default is 1 hour. Valid values range from 0 to 336 hours. Note: If the age limit is set to 0 (zero), it means that there is no age limit and quarantined files will remain on the hard disk forever.
Max File Size	The maximum size (in KB) of a single file that can be quarantined. The default is 1024 (KB). Valid values range from 1 to 2048 KB. Note: Files larger than the set Max File Size will not be quarantined. In reality, this value is subject to the available quarantine quota that remains on the hard disk. For example, when there is less than 1024 KB of quarantine quota (disk space reserved for quarantined files) remaining, a file of 1024 KB in size still will not be quarantined even though you've set Max File Size to 1024.
Quarantine Quota	The amount of disk space reserved for quarantining files. The default is 512 MB. Valid values range from 0 to 1024 MB. If the value is set to 0, no files are quarantined.

Settings	Description
Drop Infected	Select either or both of the following: • HTTP • HTTPS • SMTP Note: By default neither option is selected, which means that both types of files are quarantined. If selected, files involving the specified protocol or protocols will be dropped (not quarantined).
Lowspace	Specify the way in which new files are handled when the system disk space is running low, which could be either of the following: • Override Old—Override old quarantine files with new ones. • Drop New—Drop new quarantine files to retain old ones.

Configuring AV quarantine policies from the Console

To configure an AV quarantine policy from the Console, execute the following commands:

```
config security antivirus quarantine
set destination {NULL | disk}
set agelimit <integer>
set maxfilesize <integer>
set quarantine-quota <integer>
set drop-infected { http | https | smtp}
set lowspace {drop-new | ovrw-old}
end
```

Viewing the quarantine monitor

To view the files that have been quarantined according to the policies you set in Network Security > Anti-Virus, go to **Network Security > Quarantine Monitor**.

Quarantined file

File Name	The name of the quarantined file. Format: Checksum+Protocol.
Checksum	The checksum of the file.
Size	The size of the quarantined file.
First Timestamp	The time at which the file was first recorded.
Last Timestamp	The time at which the file was caught.
Service	The protocol of the quarantine file, HTTP, HTTPS, or SMTP.
Status	Infected—the file is infected. Note: no other statuses.
Duplicate Count	How many times the virus was scanned out between the first timestamp and the last timestamp.

Time to Live	Lifetime of the virus in quarantine. 0-336 hours—How many hours the virus has left to live. - Expired—The virus is expired. Note: seldom will you see this, as the expired virus will be removed from the quarantine monitor, since it is no longer relevant. - Forever—The virus will never expire. A copy of the virus is kept here to give notice to the user. Note: The virus is no longer a threat since it is blocked.
Description	The virus type.

Setting AV service level

FortiADC's AV service relies on the system's AV engine and signature databases. The AV engine is upgraded whenever new functions are added. The Updated daemon is responsible for updating the AV engine and the signature databases.

The system offers three types of AV signature databases, namely, Normal, Extended, and Extreme. They represent different levels of AV services. In order for FortiADC to provide you with the level of AV service that you desire, you must choose the right signature database.

Configure AV service level from the GUI

To choose a signature database from the GUI,

1. From the navigation bar, click Network Security>Anti Virus.
2. Click the Settings tab.
3. Select a default DB as described in [Setting AV service level on page 257](#).
4. Click **Save** when done.

Setting AV service level

Settings	Description
Normal	The regular virus database, which includes “In the Wild” viruses and most commonly seen viruses on the network. It provides regular protection.
Extended	The extended virus database, which includes both “In the Wild” viruses and a large collection of zoo viruses that are no longer seen in recent virus studies. It provides enhanced security protection.
Extreme	The extreme virus database, which includes both “In the Wild” viruses and all known zoo viruses that are no longer seen in recent virus studies. It provides the highest level of security protection.

Configure AV service level from the Console

To set the default signature database from Console, execute the following command:

```
config security antivirus settings
set default-db {normal | extended | extreme}
end
```

Configuring IPS

The FortiADC Intrusion Prevention System (IPS) combines signature detection and prevention with low latency and excellent reliability. With intrusion protection, you can create multiple IPS profiles, each containing a complete configuration based on signatures. Then, you can apply any IPS profile to any L4 VS.

Intrusion Prevention System (IPS) technology protects your network from cybercriminal attacks by actively seeking and blocking external threats before they can reach potentially vulnerable network devices.

This section describes how to configure the FortiADC Intrusion Prevention settings.

Predefined Profiles

Every individual IPS Signature takes effect for a particular type of attack, for an effective detection and protection, a well-considered combination of different IPS signatures plays a key role for the whole IPS system. FortiADC has 8 predefined Profiles in respect to: action, application, severity, target, etc. are ready for customers for a fast security-setup

Predefined Profile	Comment
all_default	signatures with default setting
all_default_pass	signatures with PASS action
default	Prevent critical attacks
high_security	Blocks all Critical/High/Medium and some Low severity vulnerabilities
protect_client	Protect against client-side vulnerabilities
protect_email_server	Protect against email server-side vulnerabilities
protect_http_server	Protect against HTTP server-side vulnerabilities
sniffer-profile	Monitor IPS attacks

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 FortiADC unit to detect and stop the attack.

Signatures

IPS signatures are the basis of signature-based intrusion prevention. Every attack can be reduced to a particular string of commands or a sequence of commands and variables. Signatures include this information so your FortiADC 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.

The FortiGuard Intrusion Prevention Service (IPS) provides customers with the latest defenses against stealthy network-level threats through a constantly updated database of known threats and behavior-based signatures.

This update service is backed by a team of threat experts and a close relationship with major application vendors. The best-in-class team also uncovers significant zero-day vulnerabilities continuously, providing FortiADC units with advanced protection ahead of vendor patches.

The IPS Signatures Database is able to be updated automatically or manually by **System > Settings > FortiGuard page**.

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 FortiADC unit conserves resources by looking for attacks only in the protocols used to transmit them. For example, the FortiADC 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 attack signatures. The engine count is configurable by CLI as well. (The recommendation is configuring the engine count as the same count of CPU of the FortiADC has, an ips-engine per CPU)

IPS profiles

The IPS engine does not examine network traffic for all signatures. You must first create an IPS profile and specify which signatures are included. Add signatures to profile individually using signature entries, or in groups using IPS filters.

To view the IPS profiles, go to **Security Profiles > Intrusion Prevention**.

You can group signatures into IPS profiles for easy selection when applying to L4 VS Security. You can define signatures for specific types of traffic in separate IPS profiles, and then select those profiles in profiles designed to handle that type of traffic. For example, you can specify all of the web-server related signatures in an IPS profile, and that the profile can then be applied to a L4 VS Security that controls all of the traffic to and from a web server protected by the unit.

The FortiGuard Service periodically updates the 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 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 profile, 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.

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.

IPS filters

IPS profiles 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 FortiADC 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 profile, go to **Security Profiles > Intrusion Prevention**, select the IPS profile containing the filters you want to view, and select 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 profile is the easiest way. Signature entries are also the only way to include custom signatures in an IPS profile.

Another use for signature entries is to change the settings of individual signatures that are already included in a filter within the same IPS profile. Add a signature entry with the required settings above the filter, and the signature entry will take priority.

Security - L4 VS

To use an IPS profile, you must select it in a L4 VS security options. An IPS profile that it not selected in a policy options will have no effect on network traffic.



IPS does not support NAT46

Session timers for IPS sessions

A session time-to-live (TTL) timer for IPS sessions is available to reduce synchronization problems between the FortiADC Kernel and IPS, and to reduce IPS memory usage.

Creating an IPS Profile

You need to create an IPS profile before specific signatures or filters can be chosen. The signatures can be added to a new profile before it is saved. However, it is good practice to keep in mind that the profile and its included filters are separate things, and that they are created separately. (Predefined Profiles)

To create a new IPS Profile

1. Go to **Security Profiles > Intrusion Prevention**.
2. Select the **Create New** icon in the top of the Edit IPS Profile window.

3. Enter the name of the new IPS Profile.
4. Optionally, enter a comment. The comment will appear in the IPS Profile list.
5. Select **OK**.
6. A newly created Profile is empty and contains no filters or signatures. You need to add one or more filters or signatures before the Profile will be of any use.

Adding IPS signatures to a Profile

1. Go to **Security > Intrusion Prevention**.
2. Select the IPS Profile to which you want to add the signature and click the pencil icon.
3. Under IPS Signatures, select **Add Signature**.
4. Select one or more signatures from the list and click **Apply** to add them to the sensor.
5. After the selected signature has been added to the IPS Signatures, the drop-down list of Action, which is on the right side of the signature, has Default, Pass and Block, is changeable.
6. Click **Apply** on the bottom of the IPS Profile page

Adding an IPS filter to a Profile

While individual signatures can be added to a Profile, a filter allows you to add multiple signatures to a Profile by specifying the characteristics of the signatures to be added.

To create a new pattern based signature and filter

1. Go to **Security Profiles > Intrusion Prevention**.
2. Select the IPS Profile to which you want to add the signature and click the pencil icon.
3. Under IPS Filters, select **Add Filter**.
4. Configure the filter that you require. Signatures matching all of the characteristics you specify in the filter will be included in the filter. Once finished, select **Apply**.

Application refers to the application affected by the attack and filter options include over 25 applications.

OS refers to the Operating System affected by the attack. The options include **BSD, Linux, MacOS, Other, Solaris, and Windows**.

Protocol refers to the protocol that is the vector for the attack; filter options include over 35 protocols, including "other."

Severity refers to the level of threat posed by the attack. The options include **Critical, High, Medium, Low, and Info**.

Target refers to the type of device targeted by the attack. The options include **client** and **server**.

Action	Description
Pass	Select Pass to allow traffic to continue to its destination. Note: to see what the default for a signature is, go to the IPS Signatures page and enable the column Action, then find the row with the signature name in it.
Block	Select Block to drop traffic matching any signatures included in the filter.
Default	Select Default to use the default action of the signature.

5. After the selected signature has been added to the IPS Signatures, the drop-down list of Action, which is on the right side of the Filter, has Default, Pass and Block, is changeable
6. Click **Apply** on the bottom of the IPS Profile page

Adding rate based signatures

These are a subset of the signatures that are found in the database. This group of signatures is for vulnerabilities that are normally only considered a serious threat when the targeted connections come in multiples, a little like DoS attacks.

Adding a rate based signature is straight forward. Select the enable button in the Rate Based Signature table that corresponds with the desired signature.

Predefined IPS Profile

FortiADC has 8 predefined IPS Profiles for the convenience and fast-set-up of users to enable the IPS by an easier way, each predefined profile is created under the attributes of each signature and thoughtful consideration. For users demanding a widely protection but yet ready to create a particular customized one, predefined IPS profiles are highly recommended. They will be kept updated resulted from a periodically database update of the FortiGuard Service. These Profiles are available by directly selecting from Security -> IPS in L4 VS options as well as be considered as a Quick-Enabling-IPS.

Enabling IPS

Currently, the IPS Scanning only supports for the L4VS traffic

- The IPS Profile contains filters, signature entries, or both. These specify which signatures are included in the IPS Profile.

When an IPS Profile is selected in a security option, and all network traffic matching the policy will be checked for the signatures in the IPS Profile.

Configuring Engine Count

For the consideration of varying demands and the performance of different platforms, the Engine-Count of IPS in FortiADC is configurable. The more Engine-Count that a FortiADC has, the better the IPS performs. Every coin has two sides, however, consequently, the more CPU and memory usage will be taken from the whole system.

The default value of the Engine-count is 1, for a better performance accordingly, the configuration could be setting the Engine-Count depends on CPU-Count of the platform has.

Eg: 4-Engine for a 4-Core device. (Refer to the hardware platform reference at the end of this article)

CLI Syntax

```
config global
config system ips
set engine-count {1-256}
next
end
```

Chapter 8: DoS Protection

You use web application firewall policies to scan HTTP requests and responses against known attack signatures and methods and filter matching traffic. This section includes the following topics:

- [Configuring DoS Protection Profile on page 263](#)
- [Configuring HTTP access limit policy on page 265](#)
- [Configuring HTTP connection flood policy on page 266](#)
- [Configuring an HTTP request flood policy on page 266](#)
- [Configuring an IP fragmentation policy on page 267](#)
- [Configuring a TCP SYN flood protection policy on page 268](#)
- [Configuring a TCP slow data flood protection policy on page 268](#)

Configuring DoS Protection Profile

A DoS Protection profile references the DoS policies that are to be enforced.

Before you begin:

- You must have Read-Write permission for Security settings.

After you have configured DoS Protection profile, you can select them in **Server Load Balance > Virtual Server > Security > DoS Protection Profile**.

To configure a DoS Protection Profile:

1. Go to **DoS Protection > DoS Protection Profile**.
2. Click **Create New** to display the configuration editor.

3. Complete the configuration.

Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
HTTP Access Limit	HTTP Access Limit policy. Limit the request number per second from an IP.
HTTP Connection Flood	HTTP Connection Flood policy. Limit the number of connections from a client, which is marked by a cookie.
HTTP Request Flood	HTTP Request Flood policy. Limit the request number per second from a client, which is marked by a cookie.
TCP Slow Data Flood Protection	After the TCP connection is established (the three-way handshake is completed), if FortiADC sends data to the client but the client returns a zero window (a zero window appears when, for example, the client does not take the data out of the TCP receive queue of the client OS when the data sent by the FADC fills up the queue), FortiADC will stop sending data. In this case, FortiADC can actively abort TCP connections and release related resources to avoid occupying its resources for a long time.

4. Save the configuration.

Configuring HTTP access limit policy

HTTP Access Limit policy can limit the speed of HTTP request from a source IP.

Before you begin:

- You must have Read-Write permission for Security settings.

After you have configured HTTP Access Limit policies, you can select them in DoS Protection Profile.

To configure a HTTP Access Limit policy:

- Go to **DoS Protection > Application > HTTP Access Limit**.
- Click **Create New** to display the configuration editor.
- Complete the configuration.

Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces.
Status	Enable Disable. If Enable, this policy will be activated, otherwise it is inactive.
HTTP Request Limit	0-65535. Limits the amount of HTTP requests per second from a certain IP. 0 means no limit for HTTP request.
Action	Pass—Allow the traffic. Deny— Drop the traffic, send a 400 Bad request to the client. Period Block—Deny all the HTTP request from a source IP within a period which specified by Period Block.
Period Block	1-3600 seconds; Default: 60
Log	Enable Disable; If Enable the Action will be log
Severity	High—Log as high severity events. Medium—Log as a medium severity events. Low—Log as low severity events. The default value is High.

- Save** the configuration.

Configuring HTTP connection flood policy

HTTP Connection Flood policy can limit connections from a client which are marked by a cookie.

Before you begin:

- You must have Read-Write permission for Security settings.

After you have configured HTTP Connection Flood policies, you can select them in DoS Protection Profile.

To configure a HTTP Connection Flood policy:

- Go to **DoS Protection > Application > HTTP Connection Flood**.
- Click **Create New** to display the configuration editor.
- Complete the configuration.

Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces.
Status	Enable Disable. If Enable, this policy will be activated, otherwise is in-active.
HTTP Connection Number Limit	1-1024. Limits the number of TCP connections with the same session cookie.
Action	Pass—Allow the traffic. Deny— Drop the traffic, send a 400 Bad request to the client. Period Block—Deny all the HTTP request from a source IP within a period which specified by Period Block.
Period Block	1-3600 seconds; Default: 60
Log	Enable Disable; If Enable the Action will be log
Severity	High—Log as high severity events. Medium—Log as a medium severity events. Low—Log as low severity events. The default value is High.

- Save** the configuration.

Configuring an HTTP request flood policy

HTTP Request Flood policy can limit the speed of HTTP requests from a client which is marked by a cookie.

Before you begin:

- You must have Read-Write permission for Security settings.

After you have configured HTTP Request Flood policies, you can select them in DoS Protection Profile.

To configure a HTTP Request Flood policy:

1. Go to **DoS Protection > Application > HTTP Request Flood**.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration.

Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces.
Status	Enable Disable. If enabled, this policy will be activated, otherwise it is inactive.
HTTP Request Limit	0-65535. Limits the number of HTTP requests per second with the same session cookie. 0 means no limit for HTTP request.
Action	Pass—Allow the traffic. Deny— Drop the traffic, send a 400 Bad request to the client. Period Block—Deny all the HTTP request from a source IP within a period which specified by Period Block.
Period Block	1-3600 seconds; Default: 60
Log	Enable Disable; If Enable the Action will be log
Severity	High—Log as high severity events. Medium—Log as a medium severity events. Low—Log as low severity events. The default value is High.

4. **Save** the configuration.

Configuring an IP fragmentation policy

IP Packet fragmentation assures that IP data grams can flow through any other type of network. It allows data grams created as a single packet to be split into many smaller packets for transmission and reassembled at a receiving host. A DDoS attack can deny services to the network by creating a fragmented data gram of a large enough size to overrun the buffers in your router.

The attack purpose is to consume the system memory and network bandwidth in the shortest possible time. We can limit the maximum usage of memory in each socket, the maximum distance counters between fragmentation packages from the same source IP, and the receiving timeout for an entire package.

Before you begin:

- You must have Read-Write permission for Security settings.

To configure an IP fragmentation policy:

1. Go to **DoS Protection > Networking> IP Fragmentation Protection**.
2. Click **Edit** to display the configuration editor.

3. Complete the configuration.

Max Memory Size Limit	Maximum memory size of the IP fragmentation packet for the vdom. If the limit is reached, FortiADC will stop doing IP fragmentation reassemble.
Min Memory Size Limit	When total IP fragmentation memory size drops to this limit, FortiADC will start to do fragmentation reassemble again.
Timeout	Max life time for each fragmentation queue. All the fragmentation packets in the queue will be dropped if the queue exceed this timeout.

4. **Save** the configuration.

Configuring a TCP SYN flood protection policy

TCP SYN flood protection is a global setting to protect all virtual server traffic from SYN flood attack. After the SYN Cookie option is enabled, each virtual server will monitor SYN rate. If the average SYN rate in 10 seconds exceeds Maximum Half-Open Sockets, it will perform SYN Cookie on all subsequent new connections (SYN packets) of this virtual server until the rate drops to below Maximum Half-Open Sockets.

Before you begin:

- You must have Read-Write permission for Security settings.

To configure a TCP SYN Flood Protection policy:

- Go to **DoS Protection > Networking > TCP SYN Flood Protection**.
- Click **Edit** to display the configuration editor.
- Complete the configuration.

SYN Cookie	Enable/disable syn flood protection.
Maximum Half-Open Sockets	If the average half-open connection rate in 10 seconds for each VS exceeds this setting, it will enable SYN Cookie for all new following TCP connections for this virtual server. If the average rate drops to below this, it will disable SYN Cookie for this virtual server.

4. **Save** the configuration.

Configuring a TCP slow data flood protection policy

A Slow Data attack sends legitimate application layer requests but reads responses very slowly. With that, it may attempt to exhaust the target's connection pool. Slow reading advertises a very small number for the TCP Receive Window size and at the same time empties the client's TCP receive buffers slowly. This ensures a very low data flow rate.

The attack purpose is to consume the system resources (memory, CPU time) slowly. We can disable the connection when sending many probe packages fails in the zero-window timer.

Before you begin:

- You must have Read-Write permission for Security settings.

After you have configured HTTP Request Flood policies, you can select them in DoS Protection Profile.

To configure a HTTP Request Flood policy:

1. Go to **DoS Protection > Networking > HTTP Request Flood**.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration.

Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces.
Status	Enable Disable. If Enable, this policy will be activated, otherwise is inactive.
Probe Interval	Set probe interval time for the TCP zero window timer. After receiving a zero window packet, FortiADC will probe the peer side periodically until it returns with >0 window, or when probe count exceeds the max probe-count.
Probe Count	Max consecutive zero window probe count.
Action	Action after exceed max probe count. Pass—if the probe count exceeds probe-count, stop the probe and pass all the packets in both directions. Deny—deny the connection with RST. Block-period—deny the connection, and block any new connection from the peer side for a period of time.
Severity	High—Log as high severity events. Medium—Log as a medium severity events. Low—Log as low severity events. The default value is High.
Log	Enable or disable log

4. **Save** the configuration.

Chapter 9: Web Application Firewall

You use web application firewall policies to scan HTTP requests and responses against known attack signatures and methods and filter matching traffic. This section includes the following topics:

- [Web application firewall basics on page 270](#)
- [Web application firewall configuration overview on page 271](#)
- [Configuring a WAF Profile on page 272](#)
- [Configuring a Web Attack Signature policy on page 274](#)
- [Configuring a URL Protection policy on page 279](#)
- [Configuring an Advanced Protection policy on page 280](#)
- [Configuring an HTTP Protocol Constraint policy on page 282](#)
- [Configuring CSRF protection on page 285](#)
- [Configuring brute force attack detection on page 286](#)
- [Configuring an SQL/XSS Injection Detection policy on page 288](#)
- [Configuring WAF Exception objects on page 290](#)
- [Configuring a Bot Detection policy on page 290](#)
- [Configuring a Cookie Security policy on page 292](#)
- [Configuring sensitive data protection on page 294](#)
- [Configuring XML Detection on page 296](#)
- [Configuring JSON detection on page 299](#)
- [Importing XML schema on page 301](#)
- [Input Validation on page 303](#)
- [Uploading WSDL files on page 302](#)
- [Web Vulnerability Scanner on page 307](#)

Web application firewall basics

A web application firewall (WAF) is a security policy enforcement point positioned between a client endpoint and a web application. The primary purpose is to prevent attacks against the web servers. A WAF is deployed separately from the web application so that the process overhead required to perform security scanning can be offloaded from the web server, and policies can be administered from one platform to many servers.

A WAF uses methods that complement perimeter security systems, such as the FortiGate next-generation firewall. The FortiADC WAF module applies a set of policies to HTTP scanpoints, which are parsed contexts of an HTTP transaction.

[HTTP scanpoints on page 271](#) illustrates the scanpoints. In the WAF policy configurations, you have options to enable rules to detect attacks at the request line, query string, filename, URI, request headers, request body, response code, or response body.

In particular:

- **Web Attack Signature policy**—The signature database includes signatures that can detect known attacks and exploits that can be found in 22 scanpoints. In your policy configuration, you choose classes of scanpoints to process: HTTP Headers, HTTP Request Body, and HTTP Response Body.


- URL Protection policy—This policy enables you to create rules that detect patterns in the URI or the file extension.
- HTTP Protocol Constraint policy—This policy enables you to create rules that restrict URI, header, and body length; HTTP method, or HTTP response code.
- SQL/XSS Injection Detection policy—This policy includes rules to detect SQL/XSS injection in the HTTP Request URI, HTTP Referer Header, HTTP Cookie Header, or HTTP Request Body.
- Bot Detection—This policy includes rules to detect Bots. A Bot is an application that runs automated tasks over the Internet. The WAF supports two methods for detecting bad Bots: signature detection and behavior detection. You can also use whitelists to exclude known trusted sources (good Bots) from detection.

Policy rules are enforced (action taken) when scanning is completed at four checkpoints:

- HTTP Request Header
- HTTP Request Body
- HTTP Response Header
- HTTP Response Body

If the HTTP Request Header violates a rule, and the action is Deny, the attempted session is dropped and scanning for the transaction stops. If the action is Alert, the event is logged and rules processing continues.


HTTP scanpoints



Web application firewall configuration overview

[WAF configuration overview on page 272](#) shows the relationship between WAF configuration elements. A WAF profile comprises a Web Attack Signature policy, URL Protection policy, HTTP Protocol Constraint policy, SQL/XSS Injection Detection, and Bot Detection policy. The profile is applied to a load balancing virtual server, so all traffic routed to the virtual server is subject to the WAF rules. WAF profiles can be applied to HTTP and HTTPS virtual servers but not HTTP Turbo virtual servers.

WAF configuration overview



Predefined configuration elements

The FortiADC WAF includes many predefined configuration elements to help you get started. It includes predefined WAF profiles, predefined Web Attack Signature policies, predefined HTTP Protocol Constraint policies, and predefined SQL/XSS Injection Detection policies.

Severity

The severity ratings for predefined Web Attack Signatures and the default severity rating for feature options like SQL/XSS Injection Detection are based on the Open Web Application Security Project (OWASP) [Risk Rating Methodology](#). In order to harmonize the significance of severity levels in logs, we recommend you use this methodology to assign severity for any custom elements you create.

Exceptions

You can create exceptions so that traffic to specific hosts or URL patterns is not subject to processing by WAF rules. Exception lists are processed before traffic is inspected. If an exception applies, the traffic bypasses the WAF module.

Basic Steps

1. Create configuration objects that define the exception.
2. Add the exception to a WAF profile configuration or WAF rule configuration.

Configuring a WAF Profile

A WAF profile references the WAF policies that are to be enforced.

[Predefined WAF profiles on page 273](#) describes the predefined profiles. In many cases, you can use predefined profiles to get started.

Predefined WAF profiles

Predefined Profiles	Description
High-Level-Security	- Web Attack Signature policy: High-Level-Security - HTTP Protocol Constraints policy: High-Level-Security - SQL/XSS Injection Detection policy: High-Level-Security
Medium-Level-Security	- Web Attack Signature policy: Medium-Level-Security - HTTP Protocol Constraints policy: Medium-Level-Security - SQL/XSS Injection Detection policy: Medium-Level-Security
Alert-Only	- Web Attack Signature policy: Alert-Only - HTTP Protocol Constraints policy: Alert-Only - SQL/XSS Injection Detection policy: Alert-Only

If desired, you can create user-defined profiles. The maximum number of profiles per VDOM is 255.

Before you begin:

- You can use predefined WAF profiles, create profiles based on predefined feature options, or create profiles based on user-defined configuration objects. If you want to add user-defined configuration objects, you must create them before using this procedure to add them to a WAF profile.
- You must have Read-Write permission for Security settings.

After you have created a WAF profile, you can specify it in a virtual server configuration.

To configure a WAF Profile:

- Go to Web Application Firewall > Web Application Firewall.
- Click the **WAF Profile** tab.
- Click **Create New** to display the configuration editor.
- Complete the configuration as described in [WAF Profile configuration on page 273](#).
- Save the configuration.

WAF Profile configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Description	A string to describe the purpose of the configuration, to help you and other administrators more easily identify its use.
Web Attack Signature	Select a predefined or user-defined Web Attack Signature configuration object.
URL Protection	Select a user-defined URL Protection configuration object.

Settings	Guidelines
HTTP Protocol Constraint	Select a predefined or user-defined HTTP Protocol Constraint configuration object.
SQL/XSS Injection Detection	Select a predefined or user-defined SQL/XSS Injection Detection configuration object.
Exception Name	Select a user-defined exception configuration object. Exceptions identify specific hosts or URL patterns that are not subject to processing by this rule.
Bot Detection	Select a user-defined Bot Detection configuration object.

Configuring a Web Attack Signature policy

The FortiGuard Web Attack Signature service provides a database of attack signatures that is updated periodically to protect against new kinds of attacks. [Web Attack Signature categories and subcategories on page 277](#) summarizes the categories of threats that are detected by the signatures. The categories are reported in logs.

In the Web Attack Signature policy configuration, you can enable/disable the class of scanpoints and the action when traffic matches signatures.

There are three classes of scanpoints:

- HTTP Header—Scans traffic against HTTP header signatures. If you enable a policy at all, you are enabling HTTP header scanning.
- HTTP Request Body—Scans traffic against HTTP request body signatures.
- HTTP Response Body—Scans traffic against HTTP response body signatures.

Header scanning is always a good practice, so enabling a policy always enables header scanning. Body scanning impacts performance, so you have the option of disabling body scanning if system utilization or latency become an issue.

You can specify separate actions for three levels of event severity:

- High—We recommend you deny traffic for high severity events.
- Medium—We recommend you deny or alert, according to your preference. To be strict, deny; otherwise, alert.
- Low—We recommend you allow the traffic and log an alert for low severity events.

[Web Attack Signature predefined policies on page 274](#) describes the predefined policies. You can select the predefined policies in your WAF profiles, or you can create policies that enable a different set of scan classes or a different action. In this release, you cannot exclude individual signatures or create custom signatures. You can enable or disable the scan classes.

Web Attack Signature predefined policies

Policy	Status	Action
High-Level-Security	Scan HTTP header—Enabled. Scan HTTP Request Body—Enabled. Scan HTTP Response Body—Disabled.	High Severity Action—Deny. Medium Severity Action—Deny. Low Severity Action—Alert.
Medium-Level-Security	Scan HTTP header—Enabled. Scan HTTP Request Body—Enabled.	High Severity Action—Deny. Medium Severity Action—Alert.

Policy	Status	Action
Alert-Only	Scan HTTP Response Body—Disabled.	Low Severity Action—Alert.
	Scan HTTP header—Enabled.	High Severity Action—Alert.
	Scan HTTP Request Body—Disabled.	Medium Severity Action—Alert.
	Scan HTTP Response Body—Disabled.	Low Severity Action—Alert.

Basic Steps

1. Configure the connection to FortiGuard so that the system can receive periodic WAF Signature Database updates. See [Configuring FortiGuard service settings](#).
2. Optionally, if you do not want to use the predefined policies, configure Web Attack Signature policies. See below.
3. When configuring the WAF profile, select a policy that you associate with virtual servers. See [Configuring a Web Attack Signature policy](#).

Before you begin:

- You must have read-write permission for security settings.

To configure a Web Attack Signature policy:

1. Go to Web Application Firewall > Known Web Attacks.
2. Click the **Web Attack Signature** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Web Attack Signature configuration on page 275](#).
5. Save the configuration.

Web Attack Signature configuration

Settings	Guidelines
Category	This dialog provides tools for configuring a Web attack signature policy.
Name	Specify a unique name for the Web attack signature policy and click Save. Valid characters are A-Z, a-z, 0-9, _, and -. No space is allowed between characters. Note: Once saved, the policy name cannot be changed.
Category	This section lists the (main) categories of Web attack signatures within the system. Do the following to include the desired categories of Web attack signature in the policy: 1. In the Name column, identify the categories of Web attack signatures of interest. 2. In the Status column, select (check mark) the categories you like to include in the policy. 3. In the Action column, select the action you want to apply to the categories that you select. 4. Double-click the name of a category to view its sub-categories. See Sub-category below.
Sub-category	This section lists the sub-categories of a (main) category of Web attack signature that you have opened (double-clicked) from above. Do the following to enable any of the sub-categories of interest: 1. In the Name column, identify the sub-categories of interest.

Settings	Guidelines
	2. In the Status column, select (check mark) the sub-categories you like to include in the policy.
Signature	This dialog provides tools for searching through and filtering Web attack signatures available within the system.
Search	Use the following options to search for Web attack signatures to display: • Description—Enter a descriptive text string and click Search. • ID—Enter a Web attack signature ID and click Search. • CVE Number—Enter a CVE number related to a Web attack signature and click Search. • Clear Search—Click this button to empty all search fields. Note: Web attack signatures that match your search criterion show up in the Signature section below the moment you click the corresponding Search button.
Filters	Use any or a combination of the following filters to filter the Web attack signatures to be displayed in the Signature section below: • Category—Click the down arrow and select a (main) category of Web attack signatures from the drop-down menu. • Sub-category—Click the down arrow and select a sub-category of the category of Web attack signatures that you have selected. • Status —Click the down arrow and select either (Enable or Disable) from the drop-down menu. • Severity—Click the down arrow and select High, Medium, or Low from the drop-down menu. • With Exception—Click the down arrow and select either (Yes or No) from the drop-down menu. • Clear All—Click this button to clear the existing filters. Note: You can also remove a specific filter by clicking the corresponding x mark.
Signature	This section displays all Web attack signatures that match your search and filter criteria, showing the following information for each Web attack signature: • ID • Status • Name • Severity • Target Application • Exception Name
Signature Detail	This section shows detailed information about the Web attack signature that you've highlighted (clicked) in the Signature section above.
Detail	This tab shows the following information about the selected signature: • Signature ID • Category • Sub-category • Severity • Target Application • Description

Settings	Guidelines
Edit Signature	This tab provides tools for editing a selected Web attack signature. It contains the following fields: - Signature ID—(Read only) Shows the ID of the selected signature. - Status—Click to enable or disable the signature. - Exception Name—Click the down arrow and select an exception from the drop-down menu.

[Web Attack Signature categories and subcategories on page 277](#) summarizes the categories of threats that are detected by the signatures.

Web Attack Signature categories and subcategories

Category (ID)	Subcategory (ID)
Cross Site Scripting (1)	
SQL Injection (2)	
Generic Attacks (3)	OS Command Injection (1)
	Coldfusion Injection (2)
	LDAP Injection (3)
	Command Injection (4)
	Session Fixation (5)
	File Injection (6)
	PHP Injection (7)
	SSI Injection (8)
	UPDF XSS (9)
	Email Injection (10)
	HTTP Response Splitting (11)
	RFI Injection (12)
Trojans (4)	

Category (ID)	Subcategory (ID)
Information Disclosure (5)	Zope Information Leakage (13)
	CF Information Leakage (14)
	PHP Information Leakage (15)
	ISA Server Existence Revealed (16)
	Microsoft Office Document Properties Leakage (17)
	CF Source Code Leakage (18)
	IIS Information Leakage (19)
	Weblogic information leakage (20)
	Generic Filename and Directory leakage (21)
	ASP/JSP Source Code Leakage (22)
	PHP Source Code Leakage (23)
	SQL Error leakage (24)
	HTTP Header Leakage (25)
	WordPress Leakage (26)
	Oracle 9i (27)
Known Exploits (6)	Coppermine Photo Gallery (28)
	Netscape Enterprise Server (29)
	Cisco IOS HTTP Service (30)
	Microsoft SQL Server (31)
	HP OpenView Network Node Manager (32)
	Best Software SalesLogix (33)
	IBM Lotus Domino Web Server (34)
	Microsoft IIS (35)
	Microsoft Windows Media Services (36)
	Dave Carrigan Auth_LDAP (37)
	427BB (38)
	RaXnet Cacti Graph (39)
Credit Card Detection (7)	CHETCPASSWD (40)
	SAP (41)
Bad Robot (8)	

Configuring a URL Protection policy

URL protection policies can filter HTTP requests that match specific character strings and file extensions.

Before you begin:

- You must have Read-Write permission for Security settings.

After you have configured URL protection policies, you can select them in WAF profiles.

To configure a URL Protection policy:

1. Go to Web Application Firewall > Common Attacks Detection.
2. Click the **URL Protection** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [URL Protection configuration on page 279](#).
5. Save the configuration.

URL Protection configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
URL Access Rule	
Full URL Pattern	Matching string. Regular expressions are supported.
Action	• Alert—Allow the traffic and log the event. • Deny—Drop the traffic, send a 403 Forbidden to the client, and log the event. The default is alert.
Severity	• High—Log as high severity events. • Medium—Log as a medium severity events. • Low—Log as low severity events. The default is low.
Exception Name	Select an exception configuration object. Exceptions identify specific hosts or URL patterns that are not subject to processing by this rule.
File Extension Rule	
File Extension Pattern	Matching string. Regular expressions are supported.
Action	• Alert—Allow the traffic and log the event. • Deny—Drop the traffic, send a 403 Forbidden to the client, and log the event. The default is alert.
Severity	• High—Log as high severity events. • Medium—Log as a medium severity events. • Low—Log as low severity events. The default is low.

Settings	Guidelines
Exception Name	Select an exception configuration object. Exceptions identify specific hosts or URL patterns that are not subject to processing by this rule.

Configuring an Advanced Protection policy

The Advanced Protection policy includes the following rules:

- **Content Scraping**—Checks HTTP response header. If the traffic matches the occurrence limit and is over the specified percentage match, it detects web scraping, then executes the relevant actions for the traffic.
- **HTTP Response Code**—Checks HTTP response code. If the traffic matches the occurrence limit and is over the specified percentage match, it detects web scraping, then executes the relevant actions for the traffic.

To configure an Advanced Protection policy:

1. Go to Web Application Firewall>Common Attacks Detection.
2. Click the **Advanced Protection** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Advanced Protection configuration on page 280](#).



If you want to drop a large number of packets when traffic match the rules, you should set action to “block” instead of “deny.”

5. Save the configuration.

Advanced Protection configuration

Settings	Guidelines
Name	Enter a unique Advanced Protection policy name. Valid characters are A-Z, a-z, 0-9, _, and -. No space is allowed. Note: Once saved, the name of an Advanced Protection policy cannot be changed.
Content Scraping	
Content Type	Specify a Content Type for the Content Scraping rule: • text/html • text/plain • text/xml • application/xml • application/soap+xml • application/json
Occurrence Limit	Sets the condition for the limit of the number of responses received from the specified type. If the number of responses received within the time frame (set in Occurrence Within) from the specified type is above this limit, this condition is fulfilled.

Settings	Guidelines
Occurrence Within	Sets the time span during which to count how many times a response is received from the specified type.
Percentage Match	Sets the condition for what percentage of the traffic received is from the specified type, during the given time frame. If the specified type, compared to all traffic, is received above this Percentage Match, this condition is fulfilled. Default is 0, indicating that this condition is disabled by default.
Action	Select which action FortiADC takes when the conditions are fulfilled for Content Scraping. - Alert—Accept the request and generate an alert email, log message, or both. - Deny—Block the request (or reset the connection). - Block—Block subsequent requests from the client for a number of seconds. Also configure Block Period. - Silent-deny—Deny without log. The default value is Alert.
Severity	When FortiADC records violations of this rule in the attack log, each log message contains a Severity Level (severity_level) field. Select which severity level FortiADC uses when using Advanced Protection: - Low - Medium - High The default value is Low.
HTTP Response Code	
Response Code	Specify a Response Code for the HTTP Response Code rule.
Occurrence Limit	Sets the condition for the limit of the number of responses received from the specified type. If the number of responses received within the time frame (set in Occurrence Within) from the specified type is above this limit, this condition is fulfilled.
Occurrence Within	Sets the time span during which to count how many times a response is received from the specified type.
Percentage Match	Sets the condition for what percentage of the traffic received is from the specified type, during the given time frame. If the specified type, compared to all traffic, is received above this Percentage Match, this condition is fulfilled. Default is 0, indicating that this condition is disabled by default.
Action	Select which action FortiADC takes when the conditions are fulfilled for HTTP Response Code. - Alert—Accept the request and generate an alert email, log message, or both. - Deny—Block the request (or reset the connection). - Block—Block subsequent requests from the client for a number of seconds. Also configure Block Period. - Silent-deny—Deny without log. The default value is Alert.
Severity	High—Log as high severity events.

Settings	Guidelines
	• Medium—Log as a medium severity events. • Low—Log as low severity events. The default is low.

Configuring an HTTP Protocol Constraint policy

The HTTP Protocol Constraint policy includes the following rules:

- HTTP request parameters—Limit the length of URIs, headers, and body to prevent several types of attacks, such as buffer overflow and denial of service.
- HTTP request methods—Restrict [HTTP methods](#) allowed in HTTP requests. For example, do not allow the PUT method in HTTP requests to prevent attackers from uploading malicious files.
- HTTP response codes—Drop response traffic containing [HTTP response codes](#) that might contain information attackers can use to craft attacks. For example, some HTTP response codes include fingerprint data like web server version, database version, OS, and so on.

[Predefined HTTP protocol constraint policies on page 282](#) describes the predefined policies.

Predefined HTTP protocol constraint policies

Predefined Rules	Description
High-Level-Security	Protocol constraints enabled with default values. Action is set to deny. Severity is set to high.
Medium-Level-Security	Protocol constraints enabled with default values. Action is set to alert. Severity is set to medium.
Alert-Only	Protocol constraints enabled with default values. Action is set to alert. Severity is set to low.

If desired, you can create user-defined rules to filter traffic with invalid HTTP request methods or drop packets with the specified server response codes.

Before you begin:

- You should have a sense of legitimate URI lengths and HTTP request methods for the destination resources.
- You should know whether your servers include application fingerprint information in HTTP response codes.
- You must have Read-Write permission for Security settings.

To configure an HTTP Protocol Constraint policy:

1. Go to Web Application Firewall > Common Attacks Detection.
2. Click the **HTTP Protocol Constraint** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [HTTP Protocol Constraint configuration on page 283](#).
5. Save the configuration.

HTTP Protocol Constraint configuration

Settings	Guidelines
Name	Enter a unique HTTP protocol constraint policy name. Valid characters are A-Z, a-z, 0-9, _, and -. No space is allowed. Note: Once saved, the name of an HTTP protocol constraint policy cannot be changed.
Request Parameters	
Maximum URI Length	Maximum characters in an HTTP request URI. The default is 2048. The valid range is 1-8192.
Illegal Host Name	Enable/disable hostname checks. A domain name must consist of only the ASCII alphabetic and numeric characters, plus the hyphen. The hostname is checked against the set of characters allowed by the RFC 2616. Disallowed characters, such as non-printable ASCII characters or other special characters (for example, '<', '>', and the like), are a symptom of an attack.
Illegal HTTP Version	Enable/disable the HTTP version check. Well-formed requests include the version of the protocol used by the client, in the form of HTTP/v where v is replaced by the actual version number (one of 0.9, 1.0, 1.1). Malformed requests are a sign of traffic that was not sent from a normal browser and are a symptom of an attack.
Illegal HTTP Multipart	Enable/Disable the HTTP body multipart check. If the content-type is multipart media type, the HTTP body must contain one or more body parts, each preceded by a boundary delimiter line and the last one followed by a closing boundary delimiter line. After its boundary delimiter line, each body part then consists of a header area, a blank line, and a body area. Malformed HTTP requests are a sign of traffic that was not sent from a normal browser and are a symptom of an attack.
Maximum Cookie Number In Request	Maximum number of cookie headers in an HTTP request. The default is 16. The valid range is 1-32.
Maximum Header Number In Request	Maximum number of headers in an HTTP request. The default is 50. Requests with more headers are a symptom of a buffer overflow attack or an attempt to evade detection mechanisms. The valid configuration range is 1-100.
Maximum Request Header Name Length	Maximum characters in an HTTP request header name. The default is 1024. The valid range is 1-8192.
Maximum Request Header Value Length	Maximum characters in an HTTP request header value. The default is 4096. Longer headers might be a symptom of a buffer overflow attack. The valid configuration range is 1-8192.
Maximum URL Parameter Name Length	Maximum characters in a URL parameter name. The default is 1024. The valid range is 1-2048.
Maximum URL Parameter Value Length	Maximum characters in a URL parameter value. The default is 4096. The valid range is 1-8192.
Maximum Request Header Length	Maximum length of the HTTP request header. The default is 8192. The valid range is 1-16384.

Settings	Guidelines
Maximum Request Body Length	Maximum length of the HTTP body. The default is 67108864. The valid range is 1-67108864.
Request Method Rule	
Method	Select one or more methods to match in the HTTP request line: • CONNECT • DELETE • GET • HEAD • OPTIONS • POST • PUT • TRACE • Others Note: The first 8 methods are described in RFC 2616. The group Others contains not commonly used HTTP methods defined by Web Distributed Authoring and Version (WebDAV) extensions.
Action	• Alert—Allow the traffic and log the event. • Deny—Drop the traffic, send a 403 Forbidden to the client, and log the event. The default is alert.
Severity	• High—Log as high severity events. • Medium—Log as a medium severity events. • Low—Log as low severity events. The default is low.
Exception	Select an exception configuration object. Exceptions identify specific hosts or URL patterns that are not subject to processing by this rule.
Response Code Rule	
Minimum Status Code / Maximum Status Code	Start/end of a range of status codes to match. You can specify codes 400 to 599.
Action	• Alert—Allow the traffic and log the event. • Deny—Drop the traffic, send a 403 Forbidden to the client, and log the event. The default is alert.
Severity	• High—Log as high severity events. • Medium—Log as a medium severity events. • Low—Log as low severity events. The default is low.
Exception	Select an exception configuration object. Exceptions identify specific hosts or URL patterns that are not subject to processing by this rule.

Configuring CSRF protection

A cross-site request forgery (CSRF) is an attack that exploits the trust that a site has in a user's browser to transmit unauthorized commands.

Configuration overview

To protect back-end servers from CSRF attacks, you create two lists of items:

- **URL list**—The URL list contains all the URLs that you want to protect. FortiADC will verify the anti-csrf token when you access the URL.
- **Page List**—When FortiADC receives a request for a web page in the page list, it inserts a javascript in the web page. The script runs in the client's web browser and automatically appends a anti-csrf token.



Parameter filters

In some cases, a request for a web page and the requests generated by its links have the same URL. FortiADC cannot distinguish between requests to add javascript to and requests to check for the anti-CSRF parameter.

To avoid this issue, you create unique Page List and URL List items by adding a parameter filter to them. The parameter filter allows you to add additional criteria to match in the URL or HTTP body of a request.

Create your configuration carefully, making sure that all the URLs in the list have corresponding entries in the page list. When FortiADC checks requests for the token but has not added the script to the corresponding web page, it blocks or takes other action against the request.

To configure a CSRF Protection policy:

1. Go to **Web Application Firewall**.
2. Click the **Common Attacks Detection** tab.
3. Click the **CSRF Protection** tab
4. Click **Create New** to display the configuration editor.
5. Fill in the **Name**.
6. Enable the **Status**.
7. Modify the **Action** or **Severity** based on your requirements.
8. Click **Save** to save the configuration.
9. Click **Edit** to display the CSRF Protection.
10. Click **Create New** in CSRF Page to display the configuration editor and fill the Full URL Pattern and enable or disable Parameter Filter based on your security requirements.
11. Click **Create New** in CSRF URL to display the configuration editor and fill the Full URL Pattern and enable or disable Parameter Filter based on your security requirements.
12. Click **Save** to save the configuration.
13. Add the CSRF Protection policy to WAF profile.

CSRF Protection configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces.

Settings	Guidelines
	After you initially save the configuration, you cannot edit the name.
Action	Select which action FortiADC takes when it detects a missing or incorrect anti-CSRF parameter: - Alert—Accept the request and generate an alert email, log message, or both. - Deny—Block the request (or reset the connection). - Block—Block subsequent requests from the client for a number of seconds. Also configure Block Period. - Silent-deny—Deny without log. The default value is Alert.
Severity	When FortiADC records violations of this rule in the attack log, each log message contains a Severity Level (severity_level) field. Select which severity level FortiADC uses when it logs a CSRF attack: - Low - Medium - High The default value is Low.
Full URL Pattern	Supports regular expression.
Parameter Filter	Enable/disable Parameter Filter.
Parameter Name	Name of the parameter.
Parameter Value	Supports regular expression.

Configuring brute force attack detection

Brute Force Attack Detection policies can prevent too many login tests. If an HTTP client tries to log into a server via FortiADC and fails too many times, Brute Force Attack Detection policies can stop it.

Before you begin:

- You must have Read-Write permission for Security settings.

After you have configured Brute Force Attack Detection policies, you can select them in WAF profiles.

To configure a Brute Force Attack Detection policy:

- Go to **Web Application Firewall > Common Attacks Detection**.
- Click the **Brute Force Attack Detection** tab.
- Click **Create New** to display the configuration editor.
- Complete the configuration.

Name	Configuration name. Valid characters are A-Z, a-z, 0-9, __, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Status	On OFF. If On, this policy will be activated, otherwise it is inactive.

Action	Alert—Allow the traffic and log the event. Deny—Drop the traffic, send a 403 Forbidden to the client, and log the event. Block—Drop the traffic, and for a period of time all the traffic from the same IP will be dropped. Slient-deny—Drop the traffic, send a 403 Forbidden to the client. The default is alert.
Severity	High—Log as high severity events. Medium—Log as medium severity events. Low—Log as low severity events. The default is low.
Exception	Select an exception configuration object. Exceptions identify specific hosts or URL patterns that are not subject to processing by this rule.
Comments	A string to describe the purpose of the configuration.

5. **Save** the configuration.
6. **Edit** the new saved configuration.
7. Find "Match Condition" and click **Create New**.
8. Complete the configuration.

Host Status	On OFF; If On, Host Pattern will be shown and needed. The default is OFF.
Host Pattern	Matching string for host name. Regular expressions are supported.
URL Pattern	Matching string. Regular expressions are supported. The input string must start with "/"
Login Failed Code	Matching failed code (HTTP response code). 0 means it does not match this code. The default is 0.
IP Access Limit	1-65535. Specify the number of consecutive login failures. Note: If a pair of HTTP request/response match all the settings above (Host Pattern if Host Status is On, URL Pattern and Login Failed Code if it isn't 0), this is a login failure.

9. **Save** the configuration.



You can add multiple match condition rules by repeating steps 6-9.

10. **Save** the configuration.

Configuring an SQL/XSS Injection Detection policy

SQL/XSS Injection Detection policies detect [SQL injection](#) and [cross-site scripting \(XSS\)](#) attacks. Injection occurs when user-supplied data is sent to an interpreter as part of a command or query. In an SQL injection attack, attackers craft HTTP requests that cause SQL queries to be executed directly against the web application's database. XSS injection attacks cause a web browser to execute a client-side script.

In contrast to signature-based detection, the WAF SQL and XSS injection detector module detects SQL and XSS injection through lexical analysis, which is a complementary method and is faster.

The policy enables/disables scanpoints, the action when traffic matches signatures, and the event severity.

You can enable detection in the following scanpoints:

- SQL Injection: URI—Analyzes content in the URI.
- SQL Injection: Referer—Analyzes content in the HTTP Referer header.
- SQL Injection: Cookie—Analyzes content in the HTTP Cookie header.
- SQL Injection: Body—Analyzes content in the HTTP request body.
- XSS Injection: URI—Analyzes content in the URI.
- XSS Detection: Referer—Analyzes content in the HTTP Referer header.
- XSS Detection: Cookie—Analyzes content in the HTTP Cookie header.
- XSS Detection: Body—Analyzes content in the HTTP request body.

Header scanning is recommended. Body scanning impacts performance, so you have the option of disabling body scanning if system utilization or latency become an issue.

[Predefined SQL injection and XSS detection policies on page 288](#) describes the predefined policies.

Predefined SQL injection and XSS detection policies

Predefined Rules	SQL Injection			XSS		
	Detection	Action	Severity	Detection	Action	Severity
High-Level-Security	All except Body SQL Injection Detection	Deny	High	All except Body XSS Injection Detection	Deny	High
Medium-Level-Security	Only SQL URI SQL Injection Detection	Deny	High	None	Alert	Low
Alert-Only	Only SQL URI SQL Injection Detection	Alert	High	None	Alert	Low

If desired, you can create user-defined policies.

Before you begin:

- You must have Read-Write permission for Security settings.

After you have created an SQL injection/XSS policy, you can specify it in a WAF profile configuration.

To configure an SQL/XSS Injection Detection policy:

1. Go to Web Application Firewall > Common Attacks Detection.
2. Click the **SQL/XSS Injection Detection** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [SQL/XSS Injection Detection configuration on page 289](#).
5. Save the configuration.

SQL/XSS Injection Detection configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
SQL	
SQL Injection Detection	Enable/disable SQL injection detection.
URI Detection	Enable/disable detection in the HTTP request.
Referer Detection	Enable/disable detection in the Referer header.
Cookie Detection	Enable/disable detection in the Cookie header.
Body Detection	Enable/disable detection in the HTTP Body message.
Action	• Alert—Allow the traffic and log the event. • Deny—Drop the traffic, send a 403 Forbidden to the client, and log the event. The default is alert, but we recommend you deny SQL Injection.
Severity	• High—Log as high severity events. • Medium—Log as a medium severity events. • Low—Log as low severity events. The default is low, but we recommend you rate this high or medium.
SQL Exception Name	Select an exception configuration object. Exceptions identify specific hosts or URL patterns that are not subject to processing by this rule.
XSS	
XSS Injection Detection	Enable/disable XSS injection detection.
URI Detection	Enable/disable detection in the HTTP request.
Referer Detection	Enable/disable detection in the Referer header.
Cookie Detection	Enable/disable detection in the Cookie header.
Body Detection	Enable/disable detection in the HTTP Body message.
Action	• Alert—Allow the traffic and log the event. • Deny—Drop the traffic, send a 403 Forbidden to the client, and log the event. The default is alert, but we recommend you deny XSS Injection.

Settings	Guidelines
Severity	- High—Log matches as high severity events. - Medium—Log matches as a medium severity events. - Low—Log matches as low severity events. The default is low, but we recommend you rate this high or medium.
XSS Exception Name	Select an exception configuration object. Exceptions identify specific hosts or URL patterns that are not subject to processing by this rule.

Configuring WAF Exception objects

Exceptions identify specific hosts or URL patterns that are not subject to processing by WAF rules.

Before you begin:

- You must have Read-Write permission for Security settings.

After you have created an exception object, you can specify it in WAF profiles and individual WAF feature rules.

To configure an exception object:

- Go to Web Application Firewall > WAF Profile.
- Click the **Exceptions** tab.
- Click **Create New** to display the configuration editor.
- Complete the configuration as described in [WAF Exception objects on page 290](#).
- Save the configuration.

WAF Exception objects

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Exception Host Status	Enable/disable setting exceptions by host pattern.
Exception Host	Matching string. Regular expressions are supported. For example, you can specify <code>www.example.com</code> , <code>*.example.com</code> , or <code>www.example.*</code> to match a literal host pattern or a wildcard host pattern.
Exception URL	Matching string. Must begin with a URL path separator (/). Regular expressions are supported. For example, you can specify pathnames and files with expressions like <code>\/admin</code> , <code>.*\data\1.html</code> , or <code>\data.*</code> .

Configuring a Bot Detection policy

Bot detection policies use signatures and source behavior tracking to detect client traffic likely to be generated by robots instead of genuine clients. Some bots, such as search engine crawlers, are "good bots" that perform search indexing tasks that can result in more legitimate users being directed to your site. You enable a whitelist to permit those. "Bad

bots" are known to send traffic that has a negative impact on site availability and integrity, such as DDoS attacks or content scrapping. You want to block these.

To get started, you can use predefined whitelists (known good bots) and blacklists (known bad bots). You can also specify a rate limit threshold of HTTP requests/second for sources not matched to either whitelist or blacklist. The rate limit threshold can be useful in detecting "unknown bots".

In the event of false positives, you can use the user-specified whitelist table to fine-tune detection.

Before you begin:

- You must configure the connection to FortiGuard so the system can receive periodic WAF Signature Database updates, including "good bot" and "bad bot" signatures and lists. See [Configuring FortiGuard service settings](#).
- You must have Read-Write permission for Security settings.

After you have configured Bot Detection policies, you can select them in WAF profiles.

To configure a Bot Detection policy:

1. Go to Web Application Firewall > Common Attacks Detection.
2. Click the **Bot Detection** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Bot Detection configuration on page 291](#).
5. Save the configuration.

Bot Detection configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Status	Enable/disable Bot detection.
Search Engine Status	Enable/disable the predefined search engine spider whitelist. The list is included in WAF signature updates from FortiGuard.
Bad Robot Status	Enable/disable the predefined bad robot blacklist. The list is included in WAF signature updates from FortiGuard.
HTTP Request Rate	Specify a threshold (HTTP requests/second/source) to trigger the action. Bots send HTTP request traffic at extraordinarily high rates. The source is tracked by source IP address and User-Agent. The default is 0 (off). The valid range is 0-100,000,000 requests per second.
Action	• Alert—Allow the traffic and log the event. • Deny—Drop the traffic, send a 403 Forbidden to the client, and log the event. The default is alert.
Severity	• High—Log as high severity events. • Medium—Log as a medium severity events. • Low—Log as low severity events. The default is low.
Block Period	The default is 3600 seconds. The valid range is 1-3600.

Settings	Guidelines
	The maximum size of the block IP address table is 100,000 entries. If the table is full, the earliest entry will be deleted.
Whitelist	
IPv4/Netmask	Matching subnet (CIDR format).
URL Pattern	Matching string. Regular expressions are supported.
URL Parameter Name	Matching string. Regular expressions are supported.
Cookie Name	Matching string. Regular expressions are supported.
User Agent	Matching string. Regular expressions are supported.

Configuring a Cookie Security policy

A cookie security policy allows you to configure FortiADC features that prevent cookie-based attacks and apply them in a protection profile. For example, a policy can enable cookie poisoning detection, encrypt the cookies issued by a back-end server, and add security attributes to cookies.

To configure an Cookie Security policy:

1. Go to **Web Application Firewall>Sensitive Data Protection**.
2. Click the **Cookie Security** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Cookie Security configuration on page 292](#).



If you want to drop a large number of packets when traffic match the rules, you should set Action to “block” instead of “deny.”

5. Save the configuration.

Cookie Security configuration

Settings	Guidelines
Name	Enter a unique Cookie Security policy name. Valid characters are A-Z, a-z, 0-9, _, and -. No space is allowed. Note: Once saved, the name of an Cookie Security policy cannot be changed.
Security Mode	No —Does not apply cookie tampering protection or encrypted cookie. Signed —Prevents tampering by tracking the cookie by adding a signature. Encrypted —FortiADC encrypts set-cookie values which have been sent from back-end web server to clients. Clients can only see the encrypted cookies. FortiADC also decrypts cookies which have been submitted by clients before sending them to the back-end server to determine if a cookie attack has been placed.

Settings	Guidelines
Encrypted Cookie Type	All—will encrypt all cookies. List—will encrypt the cookie that matches with the cookie-list. Note: Only applies when Security Mode is set to encrypted.
Cookie Replay	Disable or enable to allow FortiADC to use the IP address of a request to determine the owner of the cookie. If Cookie Replay is enabled, the client IP address will be appended to the set-cookie value before encryption. Once the FortiADC receives the cookie, the cookie will be decrypted and FortiADC will check if the IP matches with the client. Since the public IP of a client is not static in many environments, we recommend that you do not enable cookie-replay. Note: Only applies when Security Mode is set to encrypted. Optional.
Allow Suspicious Cookies	Never—Never allow suspicious cookies. Always—Always allow suspicious cookies. Custom—Don't Block suspicious cookies until the date specified by "Dont_block_until". Select whether or not FortiADC will allow requests that contain unrecognizable cookies or if there are missing cookies. When cookie-replay is enabled, the suspicious cookie is a missing cookie that tracks the client IP address.  In many cases, when you first introduce the cookie security features, the cookies that client browsers have cached earlier will generate false positives. To avoid this problem, either select Never, or select Custom and enter an appropriate date on which to start taking the specified action against suspicious cookies. Note: Only applies when Security Mode is set to encrypted.
Don't Block Until	Specify the date to begin blocking suspicious cookies. Applicable only when Allow Suspicious Cookies is set to custom. Note: Only applies when Security Mode is set to encrypted.
Severity	When FortiADC records violations of this rule in the attack log, each log message contains a Severity Level (severity_level) field. Select which severity level FortiADC uses when using Cookie Security: • Low • Medium • High The default value is Low.
Remove Cookie	Enable so FortiADC will accept the request, but will also remove the cookie before sending it to backend web server. Note: Only applies when Security Mode is set to encrypted or signed.

Settings	Guidelines
HTTP Only	Enable to add "HTTPOnly" flag to cookies. The HttpOnly attribute limits the scope of the cookie to HTTP requests. In particular, the attribute instructs the user agent to omit the cookie when providing access to cookies via "non-HTTP" APIs (such as a web browser API that exposes cookies to scripts). Note: cookie attribute.
Secure	Enable to add the secure flag to cookies. The secure attribute limits the scope of the cookie to "secure" channels (where "secure" is defined by the user agent). When a cookie has the Secure attribute, the user agent will include the cookie in an HTTP request only if the request is transmitted over a secure channel (typically HTTP over Transport Layer Security (TLS)). Note: cookie attribute.
Max Age	Note: cookie attribute. Default value is 0 (do not add max age), range 0- 2147483647. Add the maximum age (in minutes) if the response from the backend server does not already have a "Max-Age" attribute, or does not have an "Expires" attribute.
Exception	See Configuring WAF Exception objects on page 290 .
Action	Select which action FortiADC takes when the conditions are fulfilled for HTTP Response Code. • Alert—Accept the request and generate an alert email, log message, or both. • Deny—Block the request (or reset the connection). • Block—Block subsequent requests from the client for a number of seconds. Also configure Block Period. • Silent-deny—Deny without log. The default value is Alert.
Cookie List	The list of cookies to be encrypted. Note: Only when Security Mode is set to encrypted, and when encrypted_cookie_type is set to "list."

Configuring sensitive data protection

The Data loss prevention (DLP) feature allows Web Application Firewall (WAF) to prevent information leaks, damage and loss. It provides desensitization and warning measures for sensitive information leaks on websites, such as SSN numbers and credit card information, as well as the leakage of sensitive keywords.

- Detects and identifies private and sensitive data generated on the webpage, offering protective measures.
- Provides a built-in illegal and sensitive keyword library.

Before you begin:

- Configure a virtual server with a WAF Profile.

To configure Data Leakage Prevention

1. Go to **Web Application Firewall > Sensitive Data Protection > Sensitive Data Type**.
2. Click **Create New**.

3. Complete the configuration.

Name	Enter the name of the Sensitive Data Type. You will use the name to select the Sensitive Data Type profile in Data Leak Prevention profiles. No spaces.
Description	Comments about this profile. Describe what this profile is used for and what kind of data this regex is used to match.
Regex	Specify the regex string used to match sensitive data. There are two pre-defined regex strings named Credit_Card_Number and US_Social_Security_Number.

4. Click **Save**.5. Go to the **Data Leak Prevention** tab. Click **Create New**.

6. Complete the configuration.

Name	Enter the name of the Data Leak Prevention. You will use the name to select the Data Leak Prevention profile in WAF profiles. No spaces.
Status	Enable or disable this profile. Default is disable.
Masking	Enable masking to replace sensitive data with asterisks(*). Default is disable. Note: When masking is enabled, all target data will be replaced by an asterisk (*) so the threshold value won't take effect here. Masking only works when the action is alert. The connection will be rejected when the action is set as "deny" or "block," so no target data will be replaced.
Action	Sets the action FortiADC will take if a security check detects a potential attack. - Alert—Let the request pass when the profile detects a potential attack, only triggering a WAF log. - Deny—Drop the incoming request and trigger a WAF log. - Block—Block the ip address from the incoming request for 3600 seconds and trigger a WAF log. - silent-deny—Drop the incoming request without triggering a WAF log. Note: You can create a customized action with "Create New."
Severity	Set the severity level in the WAF logs for potential attacks detected by the Data Leak Prevention profile. - High - Medium - Low
URI Pattern	Specified in Data Leak Prevention rules. An empty URI value will prevent this rule from working.
Sensitive Data Type	Specified in Data Leak Prevention rules. Specify a profile created in "Sensitive Data Type."
Threshold	Specified in Data Leak Prevention rules. Set threshold. This means this rule won't take effect until it detects the target data exceeding the specified threshold.

For example, if you set the threshold as 3, the rules will not take effect until the target data is detected three times.

Range 1-10000. Default 1. This value won't work if "masking" is enabled.

7. Click **Create New** to create rules in Data Leak Prevention. Max amount of rules is 256, but detection will stop after matching as many as eight rules.

Example

Create a sensitive-data-type

```
config security waf sensitive-data-type
edit "Credit_Card_Number"
set regex "^3(?:[47]\d{1}|\d{4}(?:\1\d{4}){2}|0[0-5]\d{11}|[68]\d{12})$|^4
(?:\d\d\d\d)?([ -]?)\d{4}(?:\2\d{4}){2}$|^6011([ -]?)\d{4}(?:\3\d{4}){2}$|^5[1-
5]\d\d\d\d([ -]?)\d{4}(?:\4\d{4}){2}$|^2014\d{11}$|^2149\d{11}$|^2131\d{11}$|^1800\d
{11}$|^3\d{15}$"
set description "For credit card numbers from MC, Visa, Amex, Diners/CarteBlanche,
Discover/Novus, Enroute, and JCB. Matches 341-1111-1111-1111 | 5431-1111-1111-1111 |
30569309025904 Non-Matches 30-5693-0902-5904 | 5631-1111-1111-1111 | 31169309025904."
next
End
```

Use it in data-leak-prevention

```
config security waf data-leak-prevention
edit "dlp"
set status enable -> default disable
set action alert -> default alert, means pass with a security log if hit target
config rule
edit 1
set request-uri-pattern / -> default none, means do not scan the content
set sensitive-data-type Credit_Card_Number -> use data-leak-prevention
next
end
next
end
```

Configure the waf profile


```
config security waf profile
edit "WAF"
set data-leak-prevention dlp
next
end
```

Configuring XML Detection

XML is commonly used for data exchange, and hackers sometimes try to exploit security holes in XML code to attack web servers. You can use FortiADC's web application firewall (WAF) to examine client requests for anomalies in XML code. The WAF can also attempt to validate the structure of XML code in client requests using a trusted XML schema file. Configuring XML detection can help to ensure that the content of requests containing XML does not contain any potential attacks.

[XML Check Chain on page 297](#) illustrates how HTTP packets containing XML can be examined when XML detection is configured.

XML Check Chain



XML checks are composed of six parts, and each one carries out a single detection function:

- Format Check—Executes XML format detection.
- XML Schema Validation—Checks to determine whether XML content is well-formed. Must upload an XML schema file.
- Limit Check—Executes XML limit detection sub-module.
- SQL Injection Detection—Executes XML SQL injection detection.
- XSS Feature Library—Executes XML cross-site scripting detection sub-module (XML-SIDM).

Before you begin, you must:

- Configure a virtual server with a WAF Profile. See [Configuring virtual servers on page 76](#) and [Configuring a WAF Profile on page 272](#).

To configure XML Detection:

1. Go to Web Application Firewall > XML & JSON Validation and select the **XML Detection** tab.
2. Click **Create New**.
3. Complete the configuration as described in [XML Detection on page 297](#).
4. Click **Save**.

XML Detection

Settings	Guidelines
Name	Enter the name of the XML Detection profile. You will use the name to select the XML Detection profile in WAF profiles. No spaces.

Settings	Guidelines
XML Format Check	Enable to configure security checks for incoming HTTP requests to determine whether they are well-formed. You can set FortiADC response actions to malformed HTTP requests below.
Soap Format Check	Enable or disable Soap Format Check. Note: When enabled, FortiADC will examine the format of incoming SOAP requests and block those that are ill-formed. This option is disabled by default. If enabled, you can choose to enable or disable WSDL Checks below. FortiADC's Soap format check supports Soap versions 1.1 and 1.2.
WSDL Check	Enable or disable WSDL Check. Note: When enabled, FortiADC will examine the SOAP content in a request against the special characters and OS commands. This option becomes available only when Soap Format Check is enabled above. It is disabled by default. If enabled, you must select a WSDL file below.
WSDL	Select a WSDL file from the list menu, which shows all WSDL files that are shown (uploaded) on the WSDL page. Note: This option allows FortiADC to check the SOAP content in a request against the selected WSDL file, and block the content if it fails the check.
XML Schema Check	Before enabling XML Schema Checks, you must upload an XML schema file to check whether XML content is well-formed. Enable to use XML schema to validate XML content. See Importing XML schema on page 301
XML Schema	Select the XML schema file that you want to use to check whether XML content is valid.
XML Limit Check	Enable to enforce parsing limits to protect web servers from DOS attacks, including XML bombs and transform injections. If enabled, you may change the configuration for the following parameters: • Limit Max Attr • Limit Max Attr Name Len • Limit Max Attr Value Len • Limit Max Cdata Len • Limit Max Elem Child • Limit Max Elem Depth • Limit Max Elem Name Len • Limit Max Namespace • Limit Max Namespace Url Len
Max Attribute	Limits the maximum number of attributes each individual element is allowed to have. The default value is 256. The valid range is 1–256. <i>Available only when XML Limit Checks is enabled.</i>
Max Attribute Name Length	Limits the maximum length of each attribute name. The default value is 128. The valid range is 1–2048. <i>Available only when XML Limit Checks is enabled.</i>
Max Attribute Value Length	Limits the maximum length of each attribute value. The default value is 128. The valid range is 1–2048. <i>Available only when XML Limit Checks is enabled.</i>
Max Cdata Length	Limits the length of the CDATA section for each element. The default value is 65535. The valid

Settings	Guidelines
	range is 1–65535. <i>Available only when XML Limit Checks is enabled.</i>
Max Element Child	Limits the maximum number of children each element is allowed, and includes other elements and character information. The default value is 65535. The valid range is 1–65535. <i>Available only when XML Limit Checks is enabled.</i>
Max Element Depth	Limits the maximum number of nested levels in each element. The default value is 256. The valid range is 1–65535. <i>Available only when XML Limit Checks is enabled.</i>
Max Element Name Length	Limits the maximum length of the name of each element. The default value is 128. The valid range is 1–65535. <i>Available only when XML Limit Checks is enabled.</i>
Max Namespace	Limits the number of namespace declarations in the XML document. The default value is 16. The valid range is 0–256. <i>Available only when XML Limit Checks is enabled.</i>
Max Namespace URL Length	Limits the URL length for each namespace declaration. The default value is 256. The valid range is 0–1024. <i>Available only when XML Limit Checks is enabled.</i>
XML XSS Check	Enable to examine the bodies of incoming XML requests that might indicate possible cross-site scripting attacks. If the request contains a positive match, FortiADC responds with the corresponding action selected below.
XML SQL Injection Check	Enable to examine bodies of incoming requests for inappropriate SQL characters and keywords that might indicate an SQL injection attack. If the request contains a positive match, FortiADC responds with the corresponding action selected below.
Severity	Set the severity level in WAF logs of potential attacks detected by the XML Detection profile. Select one of the following options: • High • Middle • Low
Action	Sets the action FortiADC will take if a security check detects a potential attack. Select one of the following actions: • Alert—Sends an alert when the profile detects a potential attack. • Deny—Blocks the incoming request.
Exception Name	Optional. Select the exception profile that you want to apply to the XML Detection profile. See Configuring WAF Exception objects on page 290 .


Configuring JSON detection

Hackers sometimes try to exploit vulnerabilities in JSON data in HTTP POST operations to attack web servers. You can configure FortiADC's web application firewall (WAF) to enforce security checks that examine client HTTP requests for anomalies in JSON data in HTTP POST operations. This ensures that JSON data reaching web servers is well-formed. Some of the security protections include:

- Running format checks on requests containing JSON data in HTTP POST operations to protect potential security holes.
- Imposing JSON parsing limits to protect against denial-of-service (DOS) attacks.
- Performing JSON cross-site scripting (XSS) checks and JSON SQL Injection checks.

[JSON Check Chain on page 300](#) illustrates how HTTP packets containing JSON can be examined via sequence detection when JSON detection is configured.

JSON Check Chain



JSON checks are composed of four parts, and each one carries out a single detection function:

- Format Check—Executes JSON format detection sub-module (JSON-FDM).
- Limit Check—Executes JSON limit detection sub-module (JSON-LDM).
- SQL Injection Detection—Executes JSON cross-site scripting detection sub-module (JSON-XSSDM).
- XSS Detection—Executes JSON cross-site scripting detection sub-module (JSON-SIDM).

Before you begin, you must:

- Configure a virtual server with a WAF Profile. See [Configuring virtual servers on page 76](#) and [Configuring a WAF Profile on page 272](#).

To configure JSON Detection:

1. Go to Web Application Firewall > XML & JSON Validation and select the **JSON Detection** tab.
2. Click **Create New**.
3. Complete the configuration as described in [JSON Detection on page 300](#).
4. Click **Save**.

JSON Detection

Settings	Guidelines
Name	Enter the name of the JSON Detection profile. You will use the name to select the JSON Detection profile in WAF profiles. No spaces.
JSON Format Checks	Enable to configure security checks for incoming HTTP requests to determine whether they are well-formed. You can set FortiADC response actions to malformed HTTP requests below.
JSON Limit Checks	Enable to enforce parsing limits to protect web servers from attacks such as DOS attacks. If enabled, you may change the configuration for the following parameters: • Limit Max Array Value • Limit Max Depth • Limit Max Object Member • Limit Max String

Settings	Guidelines
Limit Max Array Value	Limits the maximum number of values within a single array. The default value is 256. The valid range is 0–4096. <i>Available only when JSON Limit Checks is enabled.</i>
Limit Max Depth	Limits the maximum depth in a JSON value. The default value is 16. The valid range is 0–4096. <i>Available only when JSON Limit Checks is enabled.</i>
Limit Max Object Member	Limits the number of members in a JSON object. The default value is 64. The valid range is 0–4096. <i>Available only when JSON Limit Checks is enabled.</i>
Limit Max String	Limits the length of a string in a JSON request for a name or a value. The default value is 64. The valid range is 0–4096. <i>Available only when JSON Limit Checks is enabled.</i>
JSON Xss Checks	Enable to examine the bodies of incoming JSON requests that might indicate possible cross-site scripting attacks. If the request contains a positive match, FortiADC responds with the corresponding action selected below.
JSON SQL Injection Checks	Enable to examine the bodies of incoming requests for inappropriate SQL characters and keywords that might indicate an SQL injection attack. If the request contains a positive match, FortiADC responds with the corresponding action selected below.
Severity	Set the severity level in WAF logs of potential attacks detected by the JSON Detection profile. Select from one of the following options: • High • Medium • Low
Action	Sets the action FortiADC will take if a security check detects a potential attack. Select one of the following actions: • Alert—Sends an alert when the profile detects a potential attack. • Deny—Blocks the incoming request.
Exception Name	Optional. Select the exception profile that you want to apply to the JSON Detection profile. See Configuring WAF Exception objects on page 290 .

Importing XML schema

XML schema files specify the acceptable structure of and elements in an XML document. When you use XML schema files to check XML content in HTTP requests, it's easier to describe acceptable content and validate that the content is well-formed.

You can configure FortiADC's web application firewall (WAF) to use trusted XML schema files to validate XML content in HTTP requests that contain XML. Using XML schema files to validate XML content can ensure that client requests to web servers are well-formed and do not contain any potential attacks.

Before you begin, you must:

- Download a trusted XML schema file that you can import to FortiADC. Acceptable file types are `.tar`, `.tar.gz`, or `.zip`.

To import an XML schema file:

1. Go to Web Application Firewall > XML & JSON Validation and select the **XML Schema** tab.
2. Click **Create New**.
3. Enter the name of the XML schema configuration. You will use the name to select the schema file in XML detection profiles. No spaces.
4. Click **Choose File** and select the XML schema file that you want to import.
5. Click **Save**.

Uploading WSDL files

WSDL stands for Web Services Description Language, which is an XML-based interface definition language used to describe the function of Web services. The acronym can also refer to a WSDL file that contains a specific WSDL description of a Web service, as it is in our case. WSDL provides a machine-readable description of how a web service can be called, what parameters it expects, and what data structures it returns.

WSDL is often used in tandem with SOAP and an XML schema to provide Web services. By reading the WSDL file, a client program connecting to a Web service can find out what operations are available on the server. The WSDL file contains all special data types used in the form of XML Schema. The client uses SOAP to call the operations listed in the WSDL file using XML over HTTP.


In FortiADC, WSDL check is an option under Soap Format Check which is part of XML validation. In order to configure this option, you must upload your WSDL file or files to FortiADC.

To upload a WSDL file:

1. On the navigation bar, click Web Application Firewall>XML & JSON Validation.
2. Click the WSDL tab. Click **Create New**. The WSDL dialog opens.
3. Specify a unique name for the WSDL configuration.
4. Click Choose File to browse for and upload the WSDL file.
5. Click **Save**.

Importing JSON schema

JSON Schema describes the structure of a JSON document (for instance, required properties and length limitations). Applications can use this information to validate instances (check that constraints are met), or inform interfaces to collect user input such that the constraints are satisfied.

Software architecture

The schema will validate when the user upload it through CLI/WEB GUI. Only the schema that passes the validation can be saved in ADC.

You can configure FortiADC's web application firewall (WAF) to use trusted JSON schema files to validate JSON content in HTTP requests that contain JSON. Using JSON schema files to validate JSON content can ensure that client requests to web servers are well-formed and do not contain any potential attacks.

Before you begin, you must:

- Download a trusted JSON schema file that you can import to FortiADC. Acceptable file types are `.tar`, `.tar.gz`, or `.zip`.

To import a JSON schema file:

1. Go to Web Application Firewall > XML & JSON Validation and select the **JSON Schema** tab.
2. Click **Create New**.
3. Enter the name of the JSON schema configuration. You will use the name to select the schema file in JSON detection profiles. No spaces.
4. Click **Choose File** and select the JSON schema file that you want to import.
5. Click **Save**.

Input Validation

An Input Validation policy can prevent suspicious HTTP requests. This function will verify the user input from scan points like URL parameter, HTML form, hidden fields, and upload file. If the format isn't correct or FortiADC detects other attacks, the request will be blocked.

To configure an Input Validation policy:

1. Go to Web Application Firewall>Input Validation.
2. Click the **Parameter Validation** tab.
3. Click **Create New** to display the configuration editor. See [Parameter Validation on page 307](#).

Name	Enter a unique Input Validation policy name. Valid characters are A-Z, a-z, 0-9, _, and -. No space is allowed. Note: Once saved, the name of an Input Validation policy cannot be changed.
Host Status	Enable to require that the Host: field of the HTTP request match a protected host name's entry in order to match the URL access rule. Also configure Host.
Host	Select which protected host name's entry (either a web host name or IP address) that the Host: field of the HTTP request must be in to match the URL access rule. Note: Optional. Only available when Host Status is enabled.
Request URL	The HTTP request URL must be start with /. eg./login. This item must be set when configuring the rule. FortiADC will match the other item (rule) when matching the request URL; if the match fails, FortiADC will not attempt to match others.
Action	Select which action FortiADC takes when the conditions are fulfilled for Parameter Validation. • Alert—Accept the request and generate an alert email, log message, or both. • Deny—Block the request (or reset the connection). • Block—Block subsequent requests from the client for a number of seconds. Also configure Block Period. • Silent-deny—Deny without log. The default value is Alert.
Severity	When FortiADC records violations of this rule in the attack log, each log message contains a Severity Level (severity_level) field. Select which severity level FortiADC uses when using Input Validation: • Low • Medium • High The default value is Low .

4. Click **Save**.
5. Edit the newly created Parameter Validation. Under **Parameter Validation Rule Element**, click **Create New**.

Name	Enter a unique Parameter Validation Rule Element name. It must match the value of the name in the input type of the HTML request.
Max Length	The maximum length of the Parameter Validation Rule Element name's value.

6. Click **Save**.
7. Click the **Hidden Field** tab.
8. Click **Create New** to display the configuration editor. See [Hidden Fields on page 307](#).

Name	Enter a unique Hidden Fields policy name. Valid characters are A-Z, a-z, 0-9, _, and -. No space is allowed. Note: Once saved, the name of a Hidden Field policy cannot be changed.
Host Status	Enable to require that the Host: field of the HTTP request match a protected host name's entry in order to match the URL access rule. Also configure Host.
Host	Select which protected host name's entry (either a web host name or IP address) that the Host: field of the HTTP request must be in to match the URL access rule. Note: Optional. Only available when Host Status is enabled.
Request URL	The HTTP request URL must be start with /. eg./login. This item must be set when configuring the rule. FortiADC will match the other item (rule) when matching the request URL; if the match fails, FortiADC will not attempt to match others.
Action	Select which action FortiADC takes when the conditions are fulfilled for Hidden Field. - Alert—Accept the request and generate an alert email, log message, or both. - Deny—Block the request (or reset the connection). - Block—Block subsequent requests from the client for a number of seconds. Also configure Block Period. - Silent-deny—Deny without log. The default value is Alert.
Severity	When FortiADC records violations of this rule in the attack log, each log message contains a Severity Level (severity_level) field. Select which severity level FortiADC uses when using Input Validation: - Low - Medium - High The default value is Low .

9. Click **Save**.
10. Edit the newly created Hidden Field. Under **Post URL**, click **Create New**.

URL	The hidden fields function only works on the configured Post URL.

11. Click **Save**.
12. Edit the newly created Hidden Field. Under **Hidden Fields**, click **Create New**.



To apply this feature, you must enable Session Management in your protection profile.

Name	Enter a unique Parameter Validation Rule Element name. It must match the value of the name in the input type of the HTML request.

13. Click **Save**.
14. Click the **File Restriction** tab.
15. Click **Create New** to display the configuration editor. See [File Restriction on page 307](#)

Name	Enter a unique File Restriction policy name. Valid characters are A-Z, a-z, 0-9, _, and -. No space is allowed. Note: Once saved, the name of a File Restriction policy cannot be changed.
Host Status	Enable to require that the Host: field of the HTTP request match a protected host name's entry in order to match the URL access rule. Also configure Host.
Host	Select which protected host name's entry (either a web host name or IP address) that the Host: field of the HTTP request must be in to match the URL access rule. Note: Optional. Only available when Host Status is enabled.
Request URL	The HTTP request URL must be start with /. eg./login. This item must be set when configuring the rule. FortiADC will match the other item (rule) when matching the request URL; if the match fails, FortiADC will not attempt to match others.
Action	Select which action FortiADC takes when the conditions are fulfilled for File Restriction. - Alert—Accept the request and generate an alert email, log message, or both. - Deny—Block the request (or reset the connection). - Block—Block subsequent requests from the client for a number of seconds. Also configure Block Period. - Silent-deny—Deny without log. The default value is Alert.
Severity	When FortiADC records violations of this rule in the attack log, each log message contains a Severity Level (severity_level) field. Select which severity level FortiADC uses when using Input Validation: - Low - Medium - High The default value is Low .
Upload File Status	Allow: Only allow the selected file type to upload. Block: Block any upload of the selected file type.
Upload File Size	The maximum size of the uploaded file.

16. Click **Save**.

17. Edit the newly created File Restriction. Under **Upload File Type**, click **Create New**.

File Type	The supported file types for the uploaded file.

18. Click save.

19. Go to the **Input Validation Policy** tab. Click **Create New**.

Name	Enter a unique Input Validation policy name. Valid characters are A-Z, a-z, 0-9, _, and -. No space is allowed. Note: Once saved, the name of an Input Validation policy cannot be changed.
Parameter Validation Rule	The Parameter Validation rule created previously.

Hidden Field Rule	The Hidden Field rule created previously.
File Restriction Rule	The File Restriction rule created previously.

20. Click **Save**. You have successfully created an Input Validation policy.

Parameter Validation

Inputs are typically the <input> tags in an HTML form. Input rules define whether or not parameters are required, and their maximum allowed length. Input rules are for visible inputs only, such as buttons and text areas. This function will do the following:

1. Check HOST by simple string or regular expression matching.
2. Check URL by simple string or regular expression matching.
3. Check the parameter name of inputs filed by matching simple string or regular express. Will also restrict the length of the name.

If the conditions are successfully matched, it will execute the specified action.

Hidden Fields

The Hidden Fields rules are for hidden parameters only, from <input type="hidden"> HTML tags. It is often written into an HTML page by the web server when it serves that page to the client, and is not visible on the rendered web page. This function will do the following:

1. Check HOST by simple string or regular expression matching .
2. Check URL by simple string or regular expression matching .
3. Match the configuration of the fetched URL.

If the conditions are successfully matched, it will execute the specified action.

File Restriction

The File Restriction rule is for restricting file uploads based on file type and size. This function will do the following:


1. Check HOST by simple string or regular expression matching.
2. Check URL by simple string or regular expression matching .
3. Check the uploaded file type and file size by simple string or regular expression matching.

If the conditions are successfully matched, it will execute the specified action.

Web Vulnerability Scanner

Web Application Vulnerability Scanner is a set of automated tools which perform black box test on web applications, to look for security vulnerabilities such as Cross-site scripting, SQL injection, command injection, source code disclosure and insecure server configuration.

Scanner



The figure shows the plug-in of the scanner which is configurable in UI/CLI. The user can select which type of vulnerabilities included in each scan. There are 5 types of signatures in our scanner:

1. The **mime signatures** warn about server responses that have an interesting mime. For example anything that is presented as php-source will likely be interesting
2. The **files signatures** will use the content to determine if a response is an interesting file. For example, a SVN file.
3. The **messages signatures** look for interesting server messages. Most are based on errors, such as caused by incorrect SQL queries or PHP execution failures.
4. The **apps signatures** will help to find pages and applications who's functionality is a security risk by default. For example, phpinfo() pages that leak information or CMS admin interfaces.
5. The **context signatures** are linked to injection tests. They look for strings that are relevant to the current injection test and help to highlight potential vulnerabilities.

WVS Task

Configuring WVS Task

1. Go to **Web Application Firewall > Web Vulnerability Scanner**
2. By default you will end up on the **WVS Task** tab.
3. Click **Create New** on the top right. It will open a dialogue box. See the figure WVS Task dialogue below.
4. Complete the configuration as described the table below.
5. **Save** the configuration.
6. Choose the WVS Task to be scanned by clicking the **diamond** in the row. It will turn into a **square** as it scans, and the Task Status will read "Scanning..." or "In Queue" or "Stopped." See the figure Run/Stop below.
7. A report will be generated and WVS Tool will summarize the results in HTML format, zip and store in HD. See [WVS Report on page 312](#)

Notes

- Only one task can run at the same time. If multiple tasks are started, others are added to task queue to wait to run. See the figure Status below.

- If a task is already running it can't be trigger again.
- WVS-task only works for ipv4 pool. ipv6 is not supported.
- It will send a scan according to the pool member port.
- If pool member health-check fails, it will still try to send scan.
- It will **not** send a scan when:
 - there's no pool member.
 - pool member port is 0.
 - pool member status is disable/maintain
- The tasks are **limited to 50**.
- It does not support **HTTP2**
- In **HA**, only the master can start the scanning; it will be triggered only if it is master.
- **Crawl limit**. If, in one task, the refer pool contains multiple real servers, the crawl limits will be dispatched to all the real servers. For example, if the crawl limit is 3000, with 3 servers, the ADC will send 1000 requests to each server.

WVS Task configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Scheduler	Select a scheduler from the schedule group. To configure a scheduler, go to Shared Resources > Schedule Group . See Creating schedule groups on page 346
Profile	Select a profile. Profiles are configured under WVS Profile , the tab to the right of WVS task. See WVS Profile on page 310

WVS Task Dialogue

The screenshot shows the 'WVS Task' configuration dialog. It has a title bar with a hamburger menu icon and the text 'WVS Task'. Below the title bar, there are three configuration fields: 'Name' with a text input field containing the placeholder 'Please input mikey', 'Scheduler' with a dropdown menu showing 'Click to select', and 'Profile' with a dropdown menu showing '1'. At the bottom right of the dialog are two buttons: 'Save' (in blue) and 'Cancel' (in grey).

WVS Task dashboard

Settings	Guidelines
Name	Name of the task
Task Action	Square—Task Status "Scanning..." in process. This task is being scanned. Blank—Task Status "In Queue," waiting to be scanned. Diamond—Task Status "Stopped."
Created Time	Time the task was created.
  	Edit Delete Clone

Run/Stop WVS-task

WVS Task	WVS Profile	WVS Login	Report
Add Filter Create New			
Name	Task Action	Task Status	Created Time
228		Scanning...	2019-01-17 13:23:05
228-p8080		In Queue	2019-01-17 13:20:05
228s		Stopped	2019-01-17 13:20:11
adc		Stopped	2018-09-27 16:58:27
multipool		Stopped	2019-01-17 13:15:03
port0		Stopped	
t1		Stopped	2019-01-17 13:33:26
t2		Stopped	2019-01-17 13:29:15
unreach-ip		Stopped	
vd1-t2		Stopped	2019-01-17 13:31:21
Showing 1 to 10 of 10 entries Show 10 entries Previous 1 Next			

WVS Profile

Creates a WVS Profile that can be selected in [Web Vulnerability Scanner on page 307](#). It gives you the option to select which types of scans you want.

Configuring WVS Profile

1. Go to **Web Application Firewall > Web Vulnerability Scanner**
2. Select the **WVS Profile** tab.
3. Click **Create New** on the top right. It will open a dialogue box.
4. Complete the configuration as described in the table below.
5. **Save** the configuration.
6. The configured profile will appear as an option in the WVS Task dialogue box.

WVS Profile dialogue box

WVS Profile configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Real Server Pool	Select a real sever from the real server pool. To configure a scheduler, you have two options. You can create it from the WVS profile dialogue, or you can go to Server Load Balance > Real Server Pool . See Using real server pools on page 172
HTTP Login Option	Select an HTTP Login Option. Configure a login option in WVS Login on page 311
Scan	You have five scan options. See Web Vulnerability Scanner on page 307
Crawl Limit	Specify a crawl limit.
WVS Exceptions	Specify a WVS Exception. See WVS Exceptions on page 312 .

WVS Login

Configuring WVS Login

1. Go to **Web Application Firewall > Web Vulnerability Scanner**. By default you will end up on the **WVS Login** tab.
2. Click **Create New** on the top right. It will open a dialogue box.
3. **Save** the configuration.

The Login option will now appear in WVS Profile's dialogue box, under **HTTP Login Option**.

WVS Login configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces.

Settings	Guidelines
	After you initially save the configuration, you cannot edit the name.
Login Method	
None	Nothing to specify.
Basic	Username—Specify a username for the login. Password—Specify a password for the login.
Advanced	Username—Specify a username for the login. Password—Specify a password for the login. Auth URL—The full URL in POST for authentication. Auth Target URL—The URL used to POST the form. Auth Verify URL—Used to verify if the username and password authentication failed. Username Field—Field name of the username. Password Field—Field name of the password. Extend Parameter—Extend the parameter for login.

WVS Exceptions

Creates a WVS Exception that can be selected in [Web Vulnerability Scanner on page 307](#)

Configuring WVS Exception

1. Go to **Web Application Firewall > Web Vulnerability Scanner**
2. Select the **WVS Exceptions** tab.
3. Click **Create New** on the top right. It will open a dialogue box.
4. Complete the configuration as described the table below.
5. **Save** the configuration.

WVS Exceptions configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Full URL Pattern	The REGEX pattern used for the exception.

WVS Report

Shows the result of the WVS Task.

WVS Report

Settings	Guidelines
Name	Name of the task.
Created Time	Time the report was created.
  	On the far right. • Download the report • Delete the report • Preview the report
Add Filter	Add Filter—Sort by Name, Created Time

Chapter 10: User Authentication

This chapter includes the following topics:

- [Configuring authentication policies on page 317](#)
- [Configuring user groups on page 319](#)
- [Using the local authentication server on page 320](#)
- [Using an LDAP authentication server on page 321](#)
- [Configuring a RADIUS authentication server on page 323](#)
- [Using Kerberos Authentication Relay on page 325](#)
- [Using HTTP Basic SSO on page 333](#)
- [Configure a SAML service provider on page 335](#)
- [Import IDP Metadata on page 336](#)
- [Configuring AD FS Proxy on page 314](#)

Configuring AD FS Proxy

Microsoft AD FS (Active Directory Federation Services) makes it possible for local users and federated users to use claims-based single sign-on (SSO) to Web sites and services. You can use AD FS to enable your organization to collaborate securely across Active Directory domains with other external organizations by using identity federation. This reduces the need for duplicate accounts, management of multiple log-ons, and other credential management issues that can occur when you establish cross-organizational trusts.

The AD FS Proxy is a service that brokers a connection between external users and your internal AD FS server. It acts as a reverse proxy and typically resides in your organization's perimeter network (aka DMZ). As far as the user is concerned, they do not know they are talking to an AD FS proxy server, as the federation services are accessed by the same URLs.

FortiADC can act as a AD FS Proxy to facilitate the deployment of AD FS. If all the users and applications are internal, there is no need to use FortiADC as AD FS Proxy. If there is a requirement to expose the federation service to the Internet, use FortiADC to replace the AD FS Proxy is helpful.

Adding an AD FS Proxy

1. Click **User Authentication > AD FS Proxy**.
2. Select **Proxy** tab.
3. Click **Create New** to open the AD FS Proxy configuration editor.
4. Make the desired entries or sections, as described in the following table .
5. Save the configuration.

AD FS Proxy

Parameter	Description
Name	Specify a unique name for the AD FS Proxy; Valid characters are A-Z, a-z, 0-9, _, and -. No space is allowed. Note: Once you have saved the configuration, you cannot edit the AD FS Proxy name.
Status	Enable—The proxy can be used by AD FS Publish. Disable—The proxy can't be used anymore. Note: If the proxy is used by at least one AD FS Publish, it can't be disabled.
Method	None: no load balance method will be used, proxy will select the first real server in the AD FS Server Pool. LB METHOD ROUND ROBIN: proxy will select the real server according to Round Robin algorithm.
AD FS Server Pool	Select a real server pool configuration object, which is also an AD FS server farm. See Using real server pools on page 172 . Note: this real server pool must use a SSL profile whose SSL is on, and must also select a local certificate.
Federation Service Name	The FQDN string appointed by the AD FS server.
User Name	A user name used to login to the AD FS server.
Password	The password used to login to the AD FS server.
Server Configuration Update Interval	1-8640000; The time interval of AD FS Proxy to get some configuration from AD FS server. Within the interval, the proxy can only use the cached configuration.
Register Timeout	1-3600; the time of AD FS Proxy waiting for the register response from AD FS server.
Connect Timeout	1-3600; the time of AD FS Proxy setup TCP connection with AD FS server
Response Timeout	1-3600; the time of AD FS Proxy waiting for all the response other than register from AD FS server.
Keepalive Timeout	1-3600; TCP connection keepalive timeout.

Add an AD FS Publish

1. Click User Authentication > AD FS Proxy
2. Select **Publish** tab.
3. Click **Create New** to open the AD FS Publish configuration editor.
4. Make the desired entries or selections, as described in the table below.
5. Save the configuration

AD FS Publish

Parameter	Description
Name	Specify a unique name for the AD FS Proxy; Valid characters are A-Z, a-z, 0-9, _, and -. No space is allowed. Note: Once you have saved the configuration, you cannot edit the AD FS Proxy name.
Status	Enable—The proxy can be used by AD FS Publish. Disable—The proxy can't be used anymore. Note: If the proxy is used by at least one AD FS Publish, it can't be disabled.
AD FS Proxy	Select an AD FS Proxy to publish on it.
Preauthentication Method	Pass Through: ADC will not change the message flow, basically it will only forward the message. AD FS: ADC will do the pre-authentication, if OK, it will forward the following messages.
Relying Party	Relying party trust configuration is received by AD FS Proxy from the AD FS server. This parameter can only be used in the AD FS mode.
External URL	The URL that ADC provide to the external users to serve as the Microsoft Application server such as Exchange server. Example: https://certauth.o365.com/owa/
Backend Server URL	The URL that used for AD FS Proxy to access the Microsoft Application server such as Exchange server. Example: https://certauth.o365.com/owa/

Attach AD FS to a Virtual Server

There are two methods to use the AD FS function for a virtual server.

Attach an AD FS Publish

1. Edit a virtual server.
2. Click **General**.
3. Select a published service for AD FS Published Service.
4. Save the configuration.

Use an AD FS script

1. Complete all the steps in "Attach an AD FS Publish."
2. Click Server Load Balance > **Scripting**.
3. Find the script whose name format is "ADFS_virtual server name_AD FS Publish name." Then clone it.
4. Detach the AD FS Published Service for the virtual server;
5. If the real server pool which was used by the virtual server is different from the AD FS Proxy on which the AD FS Published Service was published, add content routing configuration for the both pools.


6. Attach the content routing created in step 5 to virtual server.
7. Add the cloned script in step 3 into virtual server.
8. Save the configuration.

Configuring authentication policies

Auth policies set the conditions that mandate authentication and reference the user group that has authorization. For example, you can define an auth policy that has the following logic: if the Host header matches example.com and the URI matches /index.html, then the group example-group is authorized. FortiADC supports the Basic Authentication Scheme described in [RFC 2617](#).

[Authorization and authentication on page 317](#) illustrates the client-server communication when authorization is required.

Authorization and authentication



1. The client sends an HTTP request for a URL belonging to a FortiADC virtual server that has an authorization policy.
2. FortiADC replies with an HTTP 401 to require authorization. On the client computer, the user might be prompted with a dialog box to provide credentials.
3. The client reply includes an [Authorization](#) header that gives the credentials.
4. FortiADC sends a request to the server (local, LDAP, or RADIUS) to authenticate the user.
5. The authentication server sends its response, which can be cached according to your user group configuration.
6. If authentication is successful, FortiADC continues processing the traffic and forwards the request to the real server.
7. The real server responds with an HTTP 200 OK.
8. FortiADC processes the traffic and forwards the server response to the client.

Before you begin:

- You must have created the user groups to be authorized with the policy. You also configure users and authentication servers separately. See [Configuring user groups](#).
- You must have read-write permission for Server Load Balance settings.

After you have configured an auth policy, you can select it in the virtual server configuration. Note the following requirements:

- Virtual server type must be Layer 2 or Layer 7.
- Profile type must be HTTP or HTTPS.
- The profile option once-only must be disabled.

To configure an authentication policy:

1. Go to **User Authentication** > Authentication Policy.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in [Authentication policy configuration on page 318](#).
4. Save the configuration.

Authentication policy configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the virtual server configuration. Note: After you initially save the configuration, you cannot edit the name.
Member	
Host Status	If enabled, require authorization only for the specified host. If disabled, ignore hostname in the HTTP request header and require authorization for requests with any Host header. Disabled by default.
Host	Specify the HTTP Host header. If Host Status is enabled, the policy matches only if the Host header matches this value. Complete, exact matching is required. For example, <code>www.example.com</code> matches <code>www.example.com</code> but not <code>www.example.com.hk</code> .
Type	Select either of the following: • Standard • SAML
User Realm	Realm to which the Path URI belongs. The realm is included in the basic authentication header in the HTTP 401 message sent to the client. If a request is authenticated and a realm specified, the same credentials are deemed valid for other requests within this realm.
Path	Require authorization only if the URI of the HTTP request matches this pathname. If none is specified, requests to any URI require authorization. The value is parsed as a match string prefix. For example, <code>/abc</code> matches <code>http://www.example.com/abcd</code> and <code>http://www.example.com/abc/11.html</code> but not <code>http://www.example.com/1abcd</code> .
User Group	Select the user group that is authorized to access the protected resource.

Configuring user groups

User groups are authorized by the virtual server authentication policy. The user group configuration references the authentication servers that contain valid user credentials.

Suggested steps:

1. Configure LDAP and RADIUS servers, if applicable.
2. Configure local users.
3. Configure user groups (reference servers and local users).
4. Configure an authentication policy (reference the user group).
5. Configure the virtual server (reference the authentication policy).

Before you begin:

- You must have created configuration objects for any LDAP and RADIUS servers you want to use, and you must have created user accounts for local users.
- You must have read-write permission for System settings.

After you have created user groups, you can specify them in the server load balancing authentication policy configuration.

To configure a user group:

1. Go to **User Authentication** > User Group.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in [User group configuration on page 319](#).
4. Save the configuration.

User group configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
User Cache	Enable to cache the credentials for the remote users (LDAP, RADIUS) once they are authorized.
Cache Timeout	Timeout for cached user credentials. The default is 300 seconds. The valid range is 1-86,400 seconds.
Authentication Timeout	Timeout for query sent from FortiADC to a remote authentication server. The default is 2,000 milliseconds. The valid range is 1-60,000 milliseconds.
Authentication Log	Specify one of the following logging options for authentication events: • No logging • Log failed attempts • Log successful attempts • Log all (both failed and successful attempts)
Client Authentication Method	• HTML Form • HTTP

Settings	Guidelines
Group Type	- Local—Default. No action is needed. - SSO—Select to enable single sign-on (SSO) and then populate the fields below.
Authentication Relay	Select an authentication relay profile.
Authentication Session Timeout	Specify the authentication session timeout. Valid values range from 1 to 180 minutes. The default is 3 (minutes).
SSO Support	Disabled by default. When enabled, you must specify the SSO domain. See below. Note: Let's suppose that you add two or more virtual servers on FortiADC and they all use the same authentication relay, and then you set the Group Type (above) to SSO and enable SSO Support. When a client visits different services within the defined domain, only in the first request needs to be authenticated. Once authenticated, the client can visit all other services in the same domain.
SSO Domain	Specify the SSO domain.
Log-off URL	Specify the log-off URL.

Using the local authentication server

You can use a local authentication server to authenticate destination server user logins.

Note: The local authentication server does not have user-initiated password management features, so it does not easily scale to large groups of users. For large deployments, we recommend you use RADIUS or LDAP and provide instructions on your website how users can reset, recover, or change their passwords.

Basic steps:

1. Add user accounts to the local authentication server.
2. Select the local authentication server configuration and username when you create user groups.

Before you begin:

- You must have Read-Write permission for System settings.

To use a local authentication server:

1. Go to **User Authentication** > Local User.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in [Local authentication server configuration on page 321](#).
4. Save the configuration.

Local authentication server configuration

Settings	Guidelines
Name	Name of the user account, such as <code>user1</code> or <code>user1@example.com</code> . Do not use spaces or special characters except the 'at' symbol (<code>@</code>) or dot (<code>.</code>). The maximum length is 35 characters. After you initially save the configuration, you cannot edit the name.
Password	Specify a password. The stored password will be encrypted.

Using an LDAP authentication server

Lightweight Directory Access Protocol (LDAP) is an application protocol for accessing and maintaining distributed directory information services over a network. When using LDAP, authentication clients may send "Bind" messages to servers for authentication. Depending on the circumstances, clients may send different kinds of "Bind" messages.

LDAP bind messages

In a server load-balancing client authentication or admin authentication scenario, FortiADC sends binding request to the LDAP server for client authentication. Once a client is successfully authenticated, he or she can then access the LDAP server based on his or her privileges. There are three bind types: simple, anonymous, and regular.

Simple bind

Simple bind means binding with a client's full name. All clients must be located in the same branch specified with the DN.

Anonymous bind

Anonymous bind should be used only if the LDAP server allows it. The LDAP server searches for the client in the entire sub-branches, starting from the specified DN. This bind has two steps: First, FortiADC sends the binding request to specify the search entry point. Then, it sends a search request with the specified scope and filter to the LDAP server to find the given client.

Regular bind

Regular bind can be used when anonymous binding is not allowed on the LDAP server. Regular bind is similar to anonymous bind. The difference is in the initial step. Unlike anonymous bind, regular bind requires that FortiADC get the access privileges on the LDAP server with the specified User DN in the first step. After it has obtained the authorization, FortiADC can then move on to the second step as it does in anonymous bind.

LDAP over SSL (LDAPS) and StartTLS

LDAP over SSL (LDAPS) and startTLS are used to encrypt LDAP messages in the authentication process.

LDAPS is a mechanism for establishing an encrypted SSL/TLS connection for LDAP. It requires the use of a separate port, commonly 636. StartTLS extended operation is LDAPv3 standard mechanism for enabling TLS (SSL) data confidentiality protection. The mechanism uses an LDAPv3 extended operation to establish an encrypted SSL/TLS connection within an already established LDAP connection.

Configuring LDAP binding

You can use an LDAP authentication server to authenticate administrator or destination server user log-ins.

Basic steps:

1. Configure a connection to an LDAP server that can authenticate administrator or user logins.
2. Select the LDAP server configuration when you add administrator users or create user groups.

Before you begin:

- You must know the IP address and port used to access the LDAP server. You must know the CN and DN where user credentials are stored on the LDAP server.
- You must have Read-Write permission for System settings.

To select an LDAP server:

1. Go to **User Authentication** > Remote Server.
2. Select the LDAP Server tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [LDAP server configuration on page 322](#).
5. Save the configuration.

LDAP server configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Server	IP address for the server.
Port	Port number for the server. The commonly used port for LDAP is 389.
Common Name Identifier	Common name (cn) attribute for the LDAP record. For example: cn
Distinguished Name	Distinguished name (dn) attribute for the LDAP record. The dn uniquely identifies a user in the LDAP directory. For example: cn=John%20Doe, dc=example, dc=com
Bind Type	• Simple—bind without user search. It can be used only if all the users belong to the same “branch”.

Settings	Guidelines
	- Anonymous—bind with user search. It can be used when users are in different “branches” and only if the server allows “anonymous search”. - Regular—bind with user search. It can be used when users are in different “branches” and the server does not allow “anonymous search”.
User DN	Available only when Bind Type is "Regular". In that case, enter the user DN.
Password	Available only when Bind Type is "Regular". In that case, enter the user password.
Secure Connection	- Disable - LDAPS - STARTTLS
CA Profile	This field becomes available only when Secure Connection is set to LDAPS or STARTTLS, regardless of the Bind type being selected. In that case, you can either select a CA that has already been provisioned to secure the connection. You may also leave the field blank if you do not want to secure the connection.
Test Connectivity	Tests the connection of the LDAP server.

Configuring a RADIUS authentication server

You can use a RADIUS authentication server to authenticate administrator or destination server user logins.

Basic steps:

1. Configure a connection to a RADIUS server that can authenticate administrator or user logins.
2. Select the RADIUS server configuration when you add administrator users or user groups.

Before you begin:

- You must know the IP address, port, authentication protocol, and shared secret used to access the RADIUS server.
- You must have Read-Write permission for System settings.

To create a RADIUS server configuration:

1. Go to **User Authentication** > Remote Server.
2. Select the RADIUS Server tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [RADIUS server configuration on page 324](#).
5. Save the configuration.

RADIUS server configuration

Settings	Guidelines
Name	Specify a unique name for the RADIUS server configuration. Valid characters are A-Z, a-z, 0-9, _, and -. No space is allowed. After you initially save the configuration, you cannot edit the name.
Server	IP address or DQDN of the remote RADIUS server.
Port	The listening port of the RADIUS server. The commonly used port for a RADIUS server is 1812.
Shared Secret	Shared secret string used when connecting to the server.
Authentication Protocol	- PAP—Password authentication protocol - CHAP—Challenge-Handshake Authentication Protocol. - MS-CHAP—Microsoft version of CHAP. - MS-CHAPv2—Microsoft version of CHAP, version 2.
Timeout	Specify the amount of time that FortiADC must wait for responses from the remote RADIUS server before it times out the connection. Valid values are from 5 to 60 seconds. The default is 5 seconds.
Test Connection	Tests the connectivity of the RADIUS server.

Configuring Duo authentication server support

You can configure FortiADC to support a Duo RADIUS authentication server.

Basic steps:

1. Configure a connection to a RADIUS server that can authenticate administrator or user logins.
2. Select the RADIUS server configuration when you add administrator users or user groups.

Before you begin:

- You must know the IP address, port, authentication protocol, and shared secret used to access the RADIUS server.
- You must have Read-Write permission for System settings.

To configure duo authentication support:

1. Go to **User Authentication** > Remote Server.
2. Select the RADIUS Server tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Configuring Duo authentication server support on page 324](#).
5. Save the configuration.

Configure Duo authentication support

Settings	Guidelines
Name	Name the configuration to something like "Duo RADIUS" to differentiate it from other RADIUS server configurations.
Server	Enter the IP address or DQDN of the Duo RADIUS proxy.
Port	Specify the listening port of the Duo RADIUS proxy.
Shared Secret	Enter the RADIUS secret configured on the Duo RADIUS proxy.
Authentication Protocol	Be sure to select PAP for Duo RADIUS support.
Timeout	Specify the amount of time that FortiADC must wait for responses from the remote RADIUS server before it times out the connection. Valid values are from 5 to 60 seconds. For Duo RADIUS support, we recommend using 60 seconds.

You can also configure a Duo RADIUS server using the following commands from the Console:

```
config user radius
    edit <name>
        set auth-type {chap|ms_chap|ms_chapv2|pap}
        set port <integer>
        set secret <password>
        set server <string>
        set timeout <integer>
        set vdom <datasource>
    next
end
```

Using Kerberos Authentication Relay

Kerberos authentication is a computer authentication protocol that works on the basis of tickets (i.e., credentials). It provides several authentication choices, allowing nodes communicating over a non-secure network to verify each others' identity securely via a Key Distribution Center (KDC) and Service Tickets (STs). It is primarily used for client-server authentication model and provides mutual authentication by which both the client and the server verify each others' identity.

Kerberos authentication is built upon symmetric key cryptography and requires a trusted third party, and may also resort to the use of public-key cryptography in certain phases of the authentication process. By default, Kerberos Authentication Relay uses UDP port 88.

The Kerberos authentication consists of the following logical components:

- Client
- Authentication Server (AS)
- Ticket Granting Server (TGS)
- Service Server (SS)

Often, the AS and TGS are located on the same physical server, i.e., the KDC.

Authentication Workflow

The following paragraphs highlight the workflow of Kerberos authentication.

Step 1: Client authentication

The client sends a cleartext (i.e., unencrypted) message of the user ID to the Authentication Server (AS) requesting services that the user wants to use. The client does not send either the secret key or the password to the AS. The AS generates the secret key by hashing the password of the user found at the database, e.g., Active Directory in Windows Server. The AS then checks to see if the client is in its database. If it is in the database, the AS sends back the following two messages to the client:

- Message A: Client/TGS Session Key encrypted using the secret key of the client/user.
- Message B: Ticket Granting Ticket (TGT) which includes the client ID, client network address, ticket validity period, and the client/TGS session key) encrypted using the secret key of the TGS.

Once the client receives Messages A and B, it attempts to decrypt Message A with the secret key generated from the password entered by the user. If the user entered password does not match the password in the AS database, the client's secret key will be different and thus unable to decrypt message A. With a valid password and secret key, the client decrypts Message A to obtain the Client/TGS Session Key. This session key is used for further communications with the TGS. Note that the client cannot decrypt Message B, as it is encrypted using TGS's secret key. At this point, the client has enough information to authenticate itself to the TGS.

Step 2: Client service authorization

When requesting services, the client sends the following messages to the TGS:

- Message C: Composed of the TGT from Message B and the ID of the requested service.
- Message D: Authenticator, which is composed of the client ID and the time-stamp, encrypted using the Client/TGS Session Key.

Upon receiving Messages C and D, the TGS retrieves Message B out of Message C. It decrypts Message B using the TGS secret key. This gives the TGS the "client/TGS session key". Using this key, the TGS decrypts Message D (Authenticator) and sends the following two messages to the client:

- Message E: Client-to-server ticket, which includes the client ID, client network address, validity period, and Client/Server Session Key, encrypted using the service's secret key.
- Message F: Client/Server Session Key encrypted with the Client/TGS Session Key.

Step 3: Client service request

Upon receiving Messages E and F from TGS, the client has enough information to authenticate itself to the SS. The client connects to the SS and sends the following two messages:

- Message E: From the previous step (the client-to-server ticket, encrypted using service's secret key).
- Message G: A new Authenticator, which includes the client ID and time-stamp encrypted using the Client/Server Session Key.

The SS decrypts the ticket using its own secret key to retrieve the Client/Server Session Key. Using the sessions key, the SS decrypts the Authenticator and sends the following message to the client to confirm its true identity and willingness to serve the client:

- Message H: The time-stamp found in client's Authenticator, plus 1 in version 4, but not necessary in version 5[2][3]), encrypted using the Client/Server Session Key.

The client decrypts the confirmation using the Client/Server Session Key and checks whether the time-stamp is correct. If it is correct, then the client can trust the server and start issuing service requests to the server.

The server provides the requested services to the client.

FortiADC Kerberos authentication implementation

Implementation of Kerberos authentication involves the following configurations in FortiADC:

- Authentication Relay. See the following paragraph.
- User Group. See [Configuring user groups on page 319](#).
- Authentication Policy. See [Configuring authentication policies on page 317](#)
- Virtual Server. See [Configuring virtual servers on page 76](#)

Configure Authentication Relay (Kerberos)

Use the following steps to configure Kerberos authentication:


1. Click **User Authentication** > Authentication Relay.
2. Click Create New to open the configuration editor dialog.
3. Make the desired entries or selections as described in [Kerberos authentication configuration on page 327](#).
4. Click Save when done.

Kerberos authentication configuration


Settings	Guidelines
Name	Specify the name of the configuration.
Delegation Type	• Kerberos (Be sure to select this option.) • HTTP Basic
KDC IP	Enter the IP address of the KDC.
KDC Port	88
Realm	Specify the realm in all upper-case characters.
Delegator Account	Specify the delegator account. Required.
Delegator Password	Specify the delegator password. Required.
Authorization	• HTTP Error 404 • Always
Delegated SPN	Specify the delegated SPN. Required.
Add Default Domain	Disabled by default. When selected, specify the default domain below.
Default Domain	Enter the default domain.

Kerberos Connectivity Test


After creating a Kerberos Authentication relay, the **Test** function will appear:




Click on the **Test** symbol, and it will open the **User Principle** dialogue. Set a user name.



If you do the Kerberos relay configuration and the user principle was corrected, you will get this response.



If it failed, you will get an error prompt.



Two-factor authentication

Normally, you are required to use your user name and password to log into your account on a system or network. In this single-factor authentication, your password is the only piece of information you need to access your account. In this case, you are presenting to the system or network a shared secret, which is your password, to authenticate your identity. Had a hacker obtained or figured out your password, your password would be compromised.

Two-factor authentication is a means for authenticating a user's identity using two different pieces of information or factors. The primary advantage of two-factor authentication is that it provides a greater level of security than single-factor authentication does. Generally, the two factors are something you must know (password) and something you must have (e.g., a token). This makes it harder for a hacker to gain access to your account because the hacker would have to have both your password and the security token.

FortiADC works in tandem with FortiAuthenticator to provide two-factor authentication. With this integration, you are required to provide your password and the security token generated by FortiAuthenticator and delivered to a specified email address to gain access to FortiADC.

To take advantage of this feature, you must

- On FortiAuthenticator, create an administrator user account, a user group, and set FortiADC as a RADIUS client.
- On FortiADC, set FortiAuthenticator as the RADIUS server.

You do not have to perform these two tasks in any specific order, but you do need to have administrator access to both FortiADC and FortiAuthenticator, which allow you to carry out the configurations.

Note: Keep in mind that, for the current release, two-factor authentication works with RADIUS server (FortiAuthenticator) only; it does not work with any other remote server.

Configuring FortiAuthenticator for two-factor authentication

FortiADC uses FortiAuthenticator as the remote authentication server, which provides the security token needed for two-factor authentication on FortiADC. If you wanted to require that all FortiADC users of your organization use two-factor authentication to log into the appliance, you must first configuring FortiAuthenticator, which involves the following tasks:

1. Creating user accounts
2. Create a user group and add users to it.
3. Designate FortiADC as a RADIUS service client

Note: The following instructions assume that you have FortiAuthenticator installed on your network and you have administrator access to it.

Creating user accounts on FortiAuthenticator

To create a user account on FortiAuthenticator:

1. From the menu bar on the left, select Authentication > User Management > Local User.
2. Click Create New to open the Create New Local User page.
3. Make all the required entries or selections as highlighted in [FortiAuthenticator configuration on page 330](#).

4. Click OK when done.
5. Repeat Steps 1 through 4 to create as many user accounts as needed.

FortiAuthenticator configuration

The screenshot displays the FortiAuthenticator configuration page for a user account. The left sidebar shows the navigation tree under 'System' > 'Authentication' > 'User Management' > 'Local Users'. The main content area is divided into three sections:

- User Account Policies:**
 - Username: john_doe_ii
 - ☐ Disabled
 - ☒ Password-based authentication [\[Change Password\]](#)
 - ☒ Token-based authentication
 - Deliver token code by: ☐ FortiToken ☒ Email ☐ SMS
 - ☒ Allow RADIUS authentication
- User Role:**
 - Role: ☒ Administrator ☐ Sponsor ☐ User
 - ☒ Full permission
 - ☐ Web service access
 - ☐ Restrict admin login from trusted management subnets only
- User Information:**
 - First name: Last name:
 - Email address: Phone number:
 - Mobile number: SMS gateway:
 - Street address:
 - City: State/Province:

Configuring FortiADC a user group

Once you have created all the local user accounts, you need to create a user group and add the users to it.

To configure a user group:

1. From the menu bar on the left, select Authentication > User Management > User Groups.
2. Click Create New to open the Create New User Group page.
3. Specify a unique name for the user group.
4. Make sure the Local radio button is selected.

5. Add all the users to the user group.
6. Click OK when done.

Set FortiADC as a RADIUS Service client

As a remote authentication server, FortiAuthenticator serves as a RADIUS server, whereas FortiADC functions as a RADIUS client. Therefore, upon setting up the user group, the next thing you need to do is to set your FortiADC appliance as the RADIUS service client, and link the user group to it.

To set your FortiADC as a RADIUS service client:

1. From the menu bar on the left, select Authentication > RADIUS Service > Clients.
2. Click Create New to open the Add RADIUS Client page.
3. In the Name field, specify a unique name for the RADIUS Service Client configuration.
4. For Client Address, select the IP/Hostname radio button, and enter your FortiADC appliance's IP address or hostname.
5. For Secret, enter the shared secret between FortiAuthenticator and FortiADC, making sure that it matches the Shared Secret you specify when configuring the RADIUS server on your FortiADC appliance.
6. For Authentication method, select Enforce two-factor authentication.
7. For User input format, select realm\username.
8. In the Realm column, click the down arrow in the Realm column and select Local | Local users.
9. In the Groups column, check the Filter check box and select the user group you have configured earlier.
10. Click Save.
11. Click OK when done.

Note: Figure xxx highlights the required fields for configuring RADIUS service client.

Configuring FortiADC for two-factor authentication

In the preceding section, we've stated that, in the two-factor authentication process, FortiAuthenticator serves as the RADIUS server that provides services to FortiADC. We discussed, among other things, how to set FortiADC as a client of FortiAuthenticator.

In this section, we talk about how to configure FortiADC as FortiAuthenticator's client, which involves the following tasks:

1. Create RADIUS server configuration using FortiAuthenticator.
2. Create admin user accounts with RADIUS authentication.

The following instructions assume you have administrator access to FortiADC.

Creating a RADIUS server configuration using FortiAuthenticator

In order to let FortiAuthenticator provide authentication services for FortiADC, you need to choose FortiAuthenticator as the remote server from the FortiADC side.

To configure a RADIUS configuration using FortiAuthenticator:

1. On FortiADC's main navigation bar, click User Authentication > Remote Server.
2. Select the RADIUS Server tab.
3. Click Create New to open the RADIUS dialog box.
4. In the Name field, specify a unique name for the RADIUS server configuration.
5. In the Server field, enter the IP address of the FortiAuthenticator that you've configured earlier.
6. In the Port field, accept the default port number, which is 1812.
7. In the Shared Secret field, enter the secret key that you specified when configuring FortiAuthenticator.
8. In the Authentication Protocol field, accept the default value or click the down arrow to select another option from the list menu.
9. Click Save when done.

Adding admin user accounts with RADIUS authentication

Once you have set FortiAuthenticator as the RADIUS server to provide authentication service to FortiADC, you must then associate FortiADC user accounts with FortiAuthenticator.

It is important to note that the user names you choose on FortiADC must match those that you have added on FortiAuthenticator. Otherwise, the two-factor authentication will not work.

To add admin user using RADIUS authentication:

1. On FortiADC's main navigation bar, click System > Administrator.
2. Click Create New to open the Admin dialog box.
3. In the Name field, specify the user name of the admin account, making sure that it matches one the users names you specified on FortiAuthenticator.
4. In the Trusted Hosts field, leave it as is or specify the IP address of a specific host. (Note: If left as is, a user can manage FortiADC via this admin account from any host; if the IP address of a specific host is specified, then a user can manage FortiADC via this admin account from that host only.)
5. In the Global Admin field, accept the default (No) or select Yes. (Note: If left as is, you must select Profile and the VDOM or VDOMs that the admin account can manage; If Yes is selected, then this admin account becomes a global administrator and can manage all VDOMs on this FortiADC appliance.)
6. In the Authentication Type field, be sure to select RADIUS.
7. In the RADIUS Server field, select the RADIUS server configuration you've created on FortiADC, as discussed in the preceding paragraph.
8. In the Wildcard field, leave as is (OFF) or turn it ON. (Note: Once the Wildcard feature is enabled, in addition to the admin user configured on FortiADC, any users configured on the RADIUS server (i.e., FortiAuthenticator) can log into FortiADC and still be mapped to the specific admin profile.)
9. Click Save when done.
10. Repeat the above steps to create as many admin user accounts as needed.

Two-factor authentication in action

In the preceding two sections, we talked about how to configure FortiAuthenticator and FortiADC for two-factor authentication. The following shows the general work flow in which two-factor authentication works when you are trying to log into FortiADC:

1. On FortiADC's login page, you enter your username and password, and click Log In.
2. FortiADC presents your login credentials to FortiAuthenticator.

3. After verifying your user name and password, FortiAuthenticator generates a security token and sends it to the email address that you specified when setting up your account on FortiAuthenticator. At the same time, the Token field pops up on FortiADC's login page, right below the password field.
4. You retrieve the token from your email, copy and paste it into the Token field on FortiADC's login page, and click Log In.
5. FortiADC sends your login information, along with the token, to FortiAuthenticator for authentication.
6. After verifying that the your have the correct token, FortiAuthenticator lets you log into FortiADC.

Using HTTP Basic SSO

When an application uses a Credentials Management API to prompt for user credentials, you must enter the required information that can be validated either by the operating system or by the web application. You can specify your domain credentials information in either of the following formats:

- User Principal Name (UPN)
- Down-Level Logon Name

The UPN format is used to specify an Internet-style name, such as `UserName@Example.Fortinet.com`. [Anatomy of a UPN on page 333](#) presents an anatomy of a UPN:

Anatomy of a UPN

Component	Comment	Example
User name	The name of an account	JohnDoeII
Separator	The at sign (@)	@
UPN suffix	Also known as the domain name	Example.Fortinet.com

The down-level logon name format specifies a domain and a user account in that domain, for example, `DOMAIN\UserName`. [Anatomy of a down-level logon name on page 333](#) highlights the components of a down-level logon name:

Anatomy of a down-level logon name

Component	Description	Example
NetBIOS domain name	Domain name	Domain
Separator	The backslash (\)	\
User account name	Also known as the login name	User name

FortiADC supports HTTP basic SSO when Client Authentication Method is set to be either HTML Form Authentication or HTML Basic Authentication.

For HTTP basic SSO, FortiADC forwards the client's credentials to the web application via the HTTP "Authorization" header. For example, `username/password "user1/fortinet"` from a client is added to the HTTP header in the format `"Authorization: Basic dXNlcjE6Zm9ydGluZXQ="`, and then forwarded to the back-end web application.

You can use either UPN or down-level logon name to log into a web application, and FortiADC adds the domain offload of your logon name for your convenience. Automatically adding the default domain prefix enables you to log in using your user name alone in environments where both user name and domain name are required for the same purpose. This feature comes in handy when you forget your domain name while trying to log into a web application..

Configure HTTP Basic SSO

Use the following steps to configure HTTP basic SSO authentication:

1. Click **User Authentication** > Authentication Relay.
2. Click Create New to open the configuration editor dialog.
3. Make the desired entries or selections as described in [HTTP Basic SSO authentication configuration on page 334](#).
4. Click Save when done.

HTTP Basic SSO authentication configuration

Settings	Guidelines
Name	Specify the name of the authentication relay configuration.
Delegation Type	Select HTTP Basic
Authorization	Select either of the following: • HTTP Error 401—If selected, FortiADC relays the authentication credentials only when it encounters an HTTP 401 error from the back-end server. • Always—If selected, FortiADC relays the authentication credentials all the time.
Domain Prefix Support	This is a switch to enable or disable the default domain prefix function. Sometimes the domain controller requires the user to log in with the user name format "domain\username" such as 'KFOR\user1' When this option is enabled, the user can also successfully log in by only entering 'user1' because FortiADC is able to automatically add the prefix 'KFOR' and then send 'KFOR\user1' to the server.
Domain Prefix	The value will be added as the domain prefix when the Domain Prefix Support is enabled (above), and when the user inputs the username without the domain. Note: The value of this domain prefix MUST be a valid NetBIOS domain name.

SAML and SSO

Web Single Sign-on (SSO) is an approach that allows single sign-on (SSO) for multiple web applications that have established a common agreement on how to exchange user information. End users provide their credentials only once and are recognized by all of the Web applications, even if they are deployed in different domains and use different identity stores. Web SSO also allows the use of a single identity store by all of the Web apps.

Security Assertion Markup Language (SAML) defines an XML-based framework for describing and exchanging security information among online business entities. It is the most popular protocol for implementing Web SSO.

The SAML protocol has two components—the Service Provider (SP) and the Identify Provider (IDP). They use SAML-defined formatted XML to talk to each other and deliver the identity information called Authentication Assertion.

FortiADC support SAML 2.0, which offers the following benefits:

- Provides support for service provider (SP) and Identity Provider (IDP) Metadata
- Provides single sign-on (SSO) experience for all virtual server resources linked with the user log-in

Functioning as an SP, FortiADC supports the following IDPs:

- FortiAuthenticator (Factory default)
- Shibboleth
- OpenAM/OpenSSO

Configure a SAML service provider

You must configure your SPs in order to use SAML authentication. To configure an SP, you must have the required IDP metadata file imported into FortiADC ahead of time. See [Import IDP Metadata on page 336](#) for more information.

Once you have imported the needed IDP metadata file into FortiADC, you can use the following steps to configure a SAML service provider:

1. Click **User Authentication** > SAML.
2. Select the SAML Service Providers tab, if it is not selected.
3. Click **Create New** to open the SAML Service Providers configuration editor.
4. Make the desired entries or selections, as described in [Configure a SAML service provider on page 335](#).
5. Click Save when done.

Configure a SAML service provider

Parameter	Description
SAML Service Provider	Use this page to configure an SAML service provider.
Name	Specify a unique name for the SAML service provider.
Entity ID	Specify the SAML service provider's entity ID, which is the SAML service provider's URL.
Local Certification	Select an option. The default is Factory.
Service URL	/SSO
Assertion Consuming Service Binding Type	Post.
Assertion Consuming Service Path	/SAML2/Post
Single Logout Binding Type	Post
Single Logout Path	/SLO/Logout
IDP Metadata	Select an IDP metadata file. Note: You must have the IDP metadata file imported into FortiADC ahead of time.
Metadata Export Service Location	/Metadata

Parameter	Description
Authentication Session Lifetime	28800
Authentication Session Timeout	3600
SSO Status	Enable(d) by default, which allows FortiADC to forward SSO information to the real server, which in turn gets the authentication information and implements the SSO function.
Export Assertion Status	Enable(d) by default, which allows FortiADC to send to the real server the URL where the Authentication Assertion (.i.e., identity information) can be fetched.
Export Assertion Path	/GetAssertion
Export Cookie Status	Enable(d) by default, which allows FortiADC to send to the real server the cookie of a site that the user last visited.
Export Assertion ACL	
IP Netmask	Enter the IP address of the real server (or the IP Netmask if the real server is one of a group of real servers) that requests authentication assertions.

Import IDP Metadata

A SAML metadata file provides the information of a client, such as its entity ID, credential, and so on. It also contains a couple of URLs so that the server knows where to send different requests, e.g., log-in requests, attribute query requests, etc. You need to import this metadata to your SAML component so that it knows which client it should talk to.

Another purpose is to establish a trust relationship between the Service Provider (SP) and Identity Provider (IdP). In this case, SAML metadata is used to exchange configuration information between the SP and the IdP, and viceversa. The metadata can be signed and encrypted so that the data is transferred securely. The other side may need the corresponding public key to validate and decrypt it and then can be used to understand and establish the connection with the SP or IdP

To import a SAML IDP metadata file:

1. Click **User Authentication** > SAML.
2. Select the IDP Metadata tab.
3. Click Import.
4. Follow the instructions onscreen to import the IDP metadata file.

Note: With the 5.0.0. release, FortiADC has enhanced its SAML IDP file parsing and SP metadata format. For IDP files, it can accept any XML with or without the default namespace set to 'md'. For SP metadata, the SP metadata no longer uses the default namespace 'md' and has removed the non-standard extension. In addition, metadata is required in SP metadata, signing, and encrypt, which is also a required setting for some IDPs.

This enhancement has modified the SP metadata XML file. So if you have an existing SAML configuration in an earlier version and would like to upgrade to 5.x.x, you **MUST** upon the upgrade reconfigure your SAML service providers and import the new SP metadata XML file.

Chapter 11: Shared Resources

This chapter includes the following topics:

- [Configuring health checks on page 337](#)
- [Monitoring health check status on page 345](#)
- [Creating schedule groups on page 346](#)
- [Creating IPv4 address objects on page 346](#)
- [Configuring IPv4 address groups on page 347](#)
- [Creating IPv6 address objects on page 348](#)
- [Configuring IPv6 address groups on page 349](#)
- [Managing ISP address books on page 350](#)
- [Creating service objects on page 352](#)
- [Creating service groups on page 354](#)
- [Configuring WCCP on page 355](#)

Configuring health checks

In server load balancing deployments, the system uses health checks to poll the members of the real server pool to test whether an application is available. You can also configure additional health checks to poll related servers, and you can include results for both in the health check rule. For example, you can configure an HTTP health check test and a RADIUS health check test. In a web application that requires user authentication, the web server is deemed available only if the web server and the related RADIUS server pass the health check.

In link load balancing deployments, the health check can poll either the ISP link group member itself or a “beacon” server that is deployed on the other side of the ISP link. A beacon is an IP address that must be reachable in order for the link to be deemed available. A beacon can be any IP address, such as a main office, core router, or virtual server at another data center.



If you expect a backend server is going to be unavailable for a long period, such as when it is undergoing hardware repair, it is experiencing extended down time, or when you have removed it from the server farm, you can improve the performance of the FortiADC system by setting the status of the pool member to Disabled, rather than allowing the system to continue to attempt health checks.

[Predefined health check configuration objects on page 338](#) describes the predefined health checks. You can get started with these or create custom objects.

Predefined health check configuration objects

Predefined	Description
LB_HLTHCK_HTTP	Sends a HEAD request to the server port 80. Expects the server to return an HTTP 200.
LB_HLTHCK_HTTPS	Sends a HEAD request to the server port 443. Expects the server to return an HTTP 200.
LB_HLTHCK_ICMP	Pings the server.
LB_HLTHCK_TCP_ECHO	Sends a TCP echo to server port 7. Expects the server to respond with the corresponding TCP echo.

Before you begin:

- You must have a good understanding of TCP/IP and knowledge of the services running on your backend servers.
- You must know the IP address, port, and configuration details for the applications running on backend servers. For some application protocol checks, you must specify user credentials.
- You must have Read-Write permission for Load Balance settings.

After you have configured a health check, you can select it in the SLB server pool, LLB link group, or GLB server configuration.

To configure a health check:

1. Go to Shared Resources > Health Check.
2. Click **Create New** to display the configuration editor.
3. Select one of the following options:
 - ICMP
 - TCP Echo
 - TCP
 - HTTP
 - HTTPS
 - DNS
 - RADIUS
 - SMTP
 - POP3
 - IMAP4
 - RADIUS Accounting
 - FTP
 - Oracle
 - TCP Half Open Connection
 - TCP SSL
 - SNMP
 - SSH
 - L2 Detection
 - UDP
 - SIP
 - SIP-TCP
 - SNMP-Custom
 - RSTP
 - MySQL
 - Diameter
4. Complete the configuration as described in [Health check configuration on page 339](#).
5. Save the configuration.



You can clone a predefined configuration object to help you get started with a user-defined configuration.

To clone a configuration object, click the clone icon  that appears in the tools column on the configuration summary page.

Health check configuration

Settings	Guidelines
General	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Type	Select a type of health check.
Destination Address Type	- IPv4 - IPv6
Destination Address	IP address to send health check traffic. In server load balancing deployments, if you do not specify an IP address, the real server IP address is used. You might configure IP address for a health check if you are configuring a combination of health checks to poll related servers. In link load balancing deployments, if you do not specify an IP address, the destination IP address is the address of the gateway. You can configure IP address if you want to test connectivity to a beacon on the other side of the gateway, or if you want to test whether service traffic is allowed to pass through the link.
Hostname	For HTTP or HTTPS health checks, you can specify the hostname (FQDN) instead of the destination IP address. This is useful in VM environments where multiple applications have the same IP address.
Interval	Seconds between each health check. Should be more than the timeout to prevent overlapping health checks. The default is 10.
Timeout	Seconds to wait for a reply before assuming that the health check has failed. The default is 5.
Up Retry	Attempts to retry the health check to see if a down server has become available. The default is 1.
Down Retry	Attempts to retry the health check to see if an up server has become unavailable. The default is 1.
Specifics	
ICMP	
No specific options	Simple ping to test connectivity.
TCP Echo	
No specific options	Simple ping to test connectivity.

Settings	Guidelines
TCP / TCP Half Open Connection / UDP	
Port	Listening port number of the backend server. Usually HTTP is 80, FTP is 21, DNS is 53, POP3 is 110, IMAP4 is 143, RADIUS is 1812, and SNMP is 161.
TCP SSL	
Port	Listening port number of the backend server. Usually HTTP is 80, FTP is 21, DNS is 53, POP3 is 110, IMAP4 is 143, RADIUS is 1812, and SNMP is 161.
SSL Ciphers	Default selections are recommended.
Local Cert	For TCP SSL only. Click the down arrow and select a local SSL Health Check Client certificate from the list menu. The certificate titled "Factory" is the default certificate shipped with your FortiADC. The rest, if any, are the custom certificates that you have created.
HTTP/HTTPS	
Port	Listening port number of the backend server. Usually HTTP is 80. If testing an HTTP proxy server, specify the proxy port.
SSL Ciphers	For HTTPS only. Default selections are recommended.
Local Cert	For HTTPS only. See TCP / TCP Half Open Connection / TCP SSL / UDP above.
HTTP CONNECT	If the real server pool members are HTTP proxy servers, specify an HTTP CONNECT option: - Local CONNECT—Use HTTP CONNECT to test the tunnel connection through the proxy to the remote server. The member is deemed available if the request returns status code 200 (OK). - Remote CONNECT—Use HTTP CONNECT to test both the proxy server response and remote server application availability. If you select this option, you can configure an HTTP request within the tunnel. For example, you can configure an HTTP GET/HEAD request to the specified URL and the expected response. - No CONNECT—Do not use the HTTP CONNECT method. This option is the default. The HTTP CONNECT option is useful to test the availability of proxy servers only. See the FortiADC Deployment Guide for FortiCache for an example that uses this health check.
Remote Host	If you use HTTP CONNECT to test proxy servers, specify the remote server IP address.
Remote Port	If you use HTTP CONNECT to test proxy servers, specify the remote server port.
Method Type	HTTP method for the test traffic: - HTTP GET—Send an HTTP GET request to the server. A response to an HTTP GET request includes HTTP headers and HTTP body. - HTTP HEAD—Send an HTTP HEAD request. A response to an HTTP HEAD request includes HTTP headers only.
Send String	The request URL, such as /contact.php.
Receive String	A string expected in return when the HTTP GET request is successful.

Settings	Guidelines
Status Code	The health check sends an HTTP request to the server. Specify the HTTP status code in the server reply that indicates a successful test. Typically, you use status code 200 (OK). Other status codes indicate errors.
Match Type	What determines a failed health check? • Match String • Match Status • Match All (match both string and status) Not applicable when using HTTP HEAD. HTTP HEAD requests test status code only.
DNS	
Domain Name	The FQDN, such as www.example.com, to use in the DNS A/AAAA record health check.
Address Type	• IPv4 • IPv6
Host Address	IP address that matches the FQDN, indicating a successful health check.
RADIUS / RADIUS Accounting	
Port	Listening port number of the backend server. Usually RADIUS is 1812 and RADIUS accounting is 1813.
Username	User name of an account on the backend server.
Password	The corresponding password.
Password Type	• User—If the backend server does not use CHAP, select this option. • CHAP—If the backend server uses CHAP and does not require a secret key, select this option.
Secret Key	The secret set on the backend server.
NAS IP Address	NAS IP address RADIUS attribute (if the RADIUS server requires this attribute to make a connection).
SIP / SIP-TCP	
Port	Specify the port number. Valid values range from 0 to 65535.
SIP Request Type	Specify the SIP request type to be used for health checks: • SIP Options • SIP Register
Status Code	The expected response code. If not set, response code 200 is expected. Specify 0 if any reply should indicate the server is available.
SMTP	
Port	Listening port number of the backend server. Usually SMTP is 25.
Domain Name	The FQDN, such as www.example.com, to use in the SMTP HELO request used for health checks.

Settings	Guidelines
	If the response is OK (250), the server is considered as up. If there is error response (501) or no response at all, the server is considered down.
POP3	
Port	Listening port number of the backend server. Usually POP3 is 110.
Username	User name of an account on the backend server.
Password	The corresponding password.
IMAP4	
Port	Listening port number of the backend server. Usually IMAP4 is 143.
Username	User name of an account on the backend server.
Password	The corresponding password.
Folder	Select an email mailbox to use in the health check. If the mailbox does not exist or is not accessible, the health check fails. The default is INBOX.
FTP	
Port	Listening port number of the backend server. Usually FTP is 21.
User name	User name of an account on the backend server.
Password	The corresponding password.
File	Specify a file that exists on the backend server. Path is relative to the initial login path. If the file does not exist or is not accessible, the health check fails.
Passive	Select this option if the backend server uses passive FTP.
SNMP	
Port	Listening port number of the backend server. Usually SNMP is 161 or 162.
CPU	Maximum normal CPU usage. If overburdened, the health check fails.
Memory	Maximum normal RAM usage. If overburdened, the health check fails.
Disk	Maximum normal disk usage. If the disk is too full, the health check fails.
Agent type	• UCD • Windows 2000
Community	Must match the SNMP community string set on the backend server. If this does not match, all SNMP health checks fail.
Version	SNMP v1 or v2c.
CPU Weight	100
Memory Weight	100
Disk Weight	100

Settings	Guidelines
SNMP-Custom	
Port	Listening port number of the backend server. Usually SNMP is 161 or 162.
Community	Must match the SNMP community string set on the backend server. If this does not match, all SNMP health checks fail.
Version	SNMP v1 or v2c.
OID	String specifying the OID to query
Value Type	Abstract syntax notation (ASN) value type: • ASN_INTEGER • ASN_OCTET_STR • ASN_OBJECT_ID • ASN_COUNTER • ASN_INTEGER
Compare Type	• Equal • Less • Greater
Counter Value	Specify the value for the evaluation.
SSH	
Port	Listening port number of the backend server. Usually SSH is 22.
Username	Username for test login.
Password	Corresponding password.
L2 Detection	
No specific options	Link Layer health checker. Sends ARP (IPv4) or NDP (IPv6) packets to test whether a physically connected system is available.
RTSP	
Port	Specify the listening port number. Valid values range from 0 to 65535.
RTSP Method Type	RTSP Options
Status Code	200
MySQL	
Port	Specify the listening port number of the MySQL server. Valid values range from 0 to 65535.
Username	Specify the database user name. (Optional)
Password	Specify the database password, if applicable.
MySQL Server Type	Select either of the following: • Master (Default) • Slave

Settings	Guidelines
Diameter	
Origin Host	Specify the FortiADC appliance that originates the Diameter message. The value is in FQDN format and used to uniquely identify a Diameter node for duplicate connection and routing loop detection. Note: Some Diameter servers do not accept multiple connections from the same origin host. If you set the origin host the same as the origin host (Identity) of the Diameter load-balance profile and use the health check and Diameter load balance profile in the same virtual server, the health check or the Diameter load-balance profile may run into certain undefined problems.
Origin Realm	Specify the realm of the FortiADC appliance that originates the Diameter message. The value is in FQDN format.
Vendor ID	Specify the type Unsigned32 vendor ID which contains the IANA "SMI Network Management Private Enterprise Codes" value assigned to the vendor of a Diameter application. The default is 12356.
Product Name	Specify the type UTF8String product name which contains the vendor assigned name for the product.
Host IPv4 Address	Specify the type IPv4 address used to inform a Diameter peer of the sender's IP address when the destination address type is IPv4. The default is blank, meaning that it is the address of the FortiADC's outgoing interface.
Host IPv6 Address	Specify the type IPv6 address used to inform a Diameter peer of the sender's IP address when the destination address type is IPv6. The default is blank, meaning that it is the address of the FortiADC's outgoing interface.
Auth Application ID	Specify the type Unsigned32 authentication application ID used to advertise support of the authentication and authorization portion of an application. This field is optional; the default is 0 (zero).
Acct Application ID	Specify the type Unsigned32 accounting application ID used to advertise support of the accounting portion of an application. This field is optional; the default is 0 (zero).
Oracle	Note: Oracle DB HC only supports Hardware models in 5.1.0
Port	Listening port number of the OracleDB server.
Username	Specify the database username
Password	Specify the database password
Connect type	Select one of the following: • Service name • SID • Connect string
Service name	Use this to specify the service name.
SID	Use this to specify the SID

Settings	Guidelines
Connect String	Use this to specify the connect string
Oracle-send-string	Send a string (command) to the OracleDb server
Oracle-receive-string	The string we accept in order to receive
Row	The row in which the send string (command) takes effect
Column	The column in which the send string (command) takes effect
Script	
Port	Specify the port that the script uses
Script	Specify the script which we create or which we have pre-defined



In SLB deployments, a health check port configuration specifying port 0 acts as a wildcard. The port for health check traffic is imputed from the real server pool member.

In LLB and GLB deployments, specifying port 0 is invalid because there is no associated configuration to impute a proper port. If your health check port configuration specifies port 0, you will not be able to use it in an LLB or GLB configuration.

Monitoring health check status

FortiADC enables you to monitor the health of server in real time directly from your desktop, as described below.

1. Click Shared Resources > Health Check.
2. Click the Health Check Monitor tab.
3. Configure the health check monitor as described in [Checking server health on page 345](#).
4. Click Start to perform the health check. The result will show in the Monitor Information.

Checking server health

Parameter	Description
IP Address	Enter the IP address of the remote server.
Health Check	Select the health check configuration.
Port	Enter the port number, if applicable. Note: This field is available only for health check configurations that require port numbers.

Creating schedule groups

You create schedule objects to use in link load balancing policies. A policy rule can be time-bound: one time, daily, weekly, or monthly.

Basic Steps

1. Create a schedule object.
2. Select the schedule when you configure the link policy.

Before you begin:

- You must have Read-Write permission for System settings.

To create schedule objects:

1. Go to Shared Resources > Schedule Group.
2. Click Create New to display the configuration editor.
3. Give the schedule a name, save it, and add schedule members as described in [Schedule member configuration on page 346](#).
4. Save the configuration.

Schedule member configuration

Settings	Guidelines
Name	Unique group name. Valid characters are A-Z, a-z, 0-9, __, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Member	
Name	Unique member name. Valid characters are A-Z, a-z, 0-9, __, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Type	• One Time • Daily • Weekly • Monthly
Start Date	YYYY/MM/DD.
End Date	YYYY/MM/DD.
Start Time	HH:MM.
End Time	HH:MM.

Creating IPv4 address objects

You create address objects to specify matching source and destination addresses in policies.

The following policies use address objects:

- Firewall policies
- QoS policies
- Connection limit policies
- Link load balancing policies

Note: For link load balancing, you can also add address objects to address groups, which can then be used in link load balance policies.

Basic Steps

1. Create address objects.
2. Select them when you configure address groups or policies.

Note: Before you begin, you must have Read-Write permission for System settings.

To create an address object:

1. Click Shared Resources > Address.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in [Address object configuration on page 347](#).
4. Click Save.

Address object configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Type	• IPv4/Netmask • Address Range
IPv4/Netmask (or IPv6/Netmask)	Specify a subnet using the IP address/mask notation.
Address Range	Specify the start and end of an address range.

Configuring IPv4 address groups

You configure address group objects when you have more than one address object you want to specify in rules that match source or destination addresses. For example, if you subscribe customer 1 and customer 2 to a group of links, then you can create rules that match the customer 1 OR customer 2 address space and load balance the set of gateways assigned to them.

The following policies use address groups:

- Link load balancing policies

Basic Steps

1. Create address objects.
2. Configure address group objects. You can add up to 256 members in a group.
3. Select the address groups when you configure your policies.

Before you begin:

- You must have Read-Write permission for System settings.

To configure an address group:

1. Click Shared Resources > Address.
2. Click the **Address Group** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Address Group configuration on page 348](#).
5. Click Save.

Address Group configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Member List	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Address	Select an address object.

Creating IPv6 address objects

You create address objects to specify matching source and destination addresses in policies.

The following policies use address objects:

- Firewall policies
- QoS policies
- Connection limit policies
- Link load balancing policies

Note: For link load balancing, you can also add address objects to address groups, which can then be used in link load balance policies.

Basic Steps

1. Create address objects.
2. Select them when you configure address groups or policies.

Note: Before you begin, you must have Read-Write permission for System settings.

To create an address object:

1. Click Shared Resources > IPv6 Address.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in [IPv6 Address object configuration on page 349](#).
4. Click Save.

IPv6 Address object configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Type	• IPv6/Netmask • Address Range
IPv4/Netmask (or IPv6/Netmask)	Specify a subnet using the IP address/mask notation.
Address Range	Specify the start and end of an address range.

Configuring IPv6 address groups

You configure address group objects when you have more than one address object you want to specify in rules that match source or destination addresses. For example, if you subscribe customer 1 and customer 2 to a group of links, then you can create rules that match the customer 1 OR customer 2 address space and load balance the set of gateways assigned to them.

The following policies use address groups:

- Link load balancing policies

Basic Steps

1. Create address objects.
2. Configure address group objects. You can add up to 256 members in a group.
3. Select the address groups when you configure your policies.

Before you begin:

- You must have Read-Write permission for System settings.

To configure an address group:

1. Click Shared Resources > Address.
2. Click the **IPv6 Address Group** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Address Group configuration on page 350](#).
5. Click Save.

Address Group configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Member List	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Address	Select an address object.

Managing ISP address books

ISP address books contain IP subnet addresses and associated province location settings for ISP links.

The following policies use the ISP address book objects:

- ISP routes
- LLB proximity routes
- LLB policies
- GLB data center configuration

The province setting is used in GLB deployments in China to enable location awareness that is province-specific. For example, a user can be directed to a data center in specific location inside the country, such as Beijing or Guangdong, rather than simply China.

[ISP address book types on page 350](#) shows the three types of address book entries:

- **Predefined**—Addresses and associated province location settings for China Mobile, China Telecom, and China Unicom. The IP subnet addresses in the predefined address books are not exposed in the user interface. The predefined package is provided to make it easier for you to configure a route when all you know and all you need to know is the name of the ISP that hosts the link.
- **Restored**—Addresses imported from a text file. The IP subnet addresses in the restored address books are not exposed in the user interface. “Restored” addresses can help you rapidly build an ISP address book configuration. “Restored” addresses can help you rapidly build an ISP address book configuration.
- **User-defined**—In the ISP address configuration, you can modify the predefined and restored address books by specifying subnets to add or exclude from them. This gives you flexibility in case you encounter address conflicts or the ISP instructs you to add a subnet address manually.


You can also create new user-defined entries for other ISPs.

Note: In systems with multiple VDOMs, these commands apply to the current VDOM only. In other words, if you configure an exclusion, it is applicable to the current VDOM only; it does not change the predefined address book.

You can use the **Inquire** utility to see whether an IP address belongs to any of the address books. If an address can be found in more than one address book, the results are returned in the following priority:

1. User-defined
2. Restored
3. Predefined

ISP address book types



The text file for the Restored entries has the following format:

```
#this is a comment line
ISP name:ABC
Province:Beijing
1.1.1.0/24
Province:Unknown
2.2.0.0 255.255.0.0
#this is a comment line too
3.3.3.3/32
ISP name:DEF
Province:Shanghai
4.4.4.0 255.255.255.0
5.5.0.0/16
```

You use the **Restore** utility to import the file and the **Back Up** utility to export it.

You use the **Clean** utility to erase entries that were imported from the text file. The clean operation does not affect the predefined addresses or user-configured entries. If a restored entry has user-configured elements (for example, an exclude list), the clean operation clears the addresses but preserves the configuration and converts it to a user-defined type.

Basic Steps

1. Create ISP address objects.
2. Select them when you configure your policies.

Note: Before you begin, you must have read-write permission for System settings.

Create an ISP address book object

To create an ISP address book object:

1. Click Shared Resource > Address.
2. Click the **ISP Address** tab.
3. Click **Create New**. The ISP Address dialog opens.
4. Complete the configuration as described in [ISP address object configuration on page 352](#).
5. Click Save.

ISP address object configuration

Settings	Guidelines																																	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.																																	
Address	Address/mask notation specifying a subnet to add it to the address book entry.																																	
Excluded Address	Address/mask notation specifying a subnet to be excluded from the address book entry. Create exclusions to predefined and restored address books only. Note: This field applies to predefined and restored address books only; it is not applicable or available for user-defined address books.																																	
Province	Select the associated province location. The configuration supports the following selections: <table><tr><td>Anhui</td><td>Henan</td><td>Shanxi (Taiyuan)</td></tr><tr><td>Beijing</td><td>Hubei</td><td>Shanxi (Xian)</td></tr><tr><td>Chongqing</td><td>Hunan</td><td>Sichuan</td></tr><tr><td>Fujian</td><td>Jiangsu</td><td>Tianjin</td></tr><tr><td>Gansu</td><td>Jiangxi</td><td>Xianggang</td></tr><tr><td>Guangdong</td><td>Jilin Liaoning</td><td>Xinjiang</td></tr><tr><td>Guangxi</td><td>Neimenggu</td><td>Xizang</td></tr><tr><td>Guizhou</td><td>Ningxia</td><td>Yunnan</td></tr><tr><td>Hainan</td><td>Qinghai</td><td>Zhejiang</td></tr><tr><td>Hebei</td><td>Shandong</td><td>Unknown</td></tr><tr><td>Heilongjiang</td><td>Shanghai</td><td></td></tr></table>	Anhui	Henan	Shanxi (Taiyuan)	Beijing	Hubei	Shanxi (Xian)	Chongqing	Hunan	Sichuan	Fujian	Jiangsu	Tianjin	Gansu	Jiangxi	Xianggang	Guangdong	Jilin Liaoning	Xinjiang	Guangxi	Neimenggu	Xizang	Guizhou	Ningxia	Yunnan	Hainan	Qinghai	Zhejiang	Hebei	Shandong	Unknown	Heilongjiang	Shanghai	
Anhui	Henan	Shanxi (Taiyuan)																																
Beijing	Hubei	Shanxi (Xian)																																
Chongqing	Hunan	Sichuan																																
Fujian	Jiangsu	Tianjin																																
Gansu	Jiangxi	Xianggang																																
Guangdong	Jilin Liaoning	Xinjiang																																
Guangxi	Neimenggu	Xizang																																
Guizhou	Ningxia	Yunnan																																
Hainan	Qinghai	Zhejiang																																
Hebei	Shandong	Unknown																																
Heilongjiang	Shanghai																																	

Creating service objects

FortiADC provides more than two dozen predefined services, as shown on the Shared Resources > Service > Service page. In addition, it allows you to create your service objects as well. Service objects are an important part of the following policy configurations:

- Firewall policies
- QoS policies

- Connection limit policies
- Link load balancing policies

Note: For link load-balancing, you can also add service objects to service groups; then use service groups in LLB policies.

Basic Steps

1. Create service objects.
2. Select them when you configure service groups or policies.

Before you begin:

- You must have Read-Write permission for System settings.

To create a service object:

1. Go to Shared Resources > Service.
2. Select the Service tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Service object configuration on page 353](#).
5. Save the configuration.

Service object configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Note: Once created, the name cannot be changed.
Protocol Type	Select one of the following: • ip (default) • icmp • tcp • udp • tcp-and-udp • sctp
Protocol	1 Note: This applies only when Protocol Type is to set to IP. In that case, it displays the protocol number without port.
Specify Source Port	This option becomes available when TCP, UDP, SCTP, or TCP-AND-UDP is selected as the protocol type. When selected, you also need to specify the Minimum Source Port and Maximum Source Port below.
Minimum Source Port	1
Maximum Source Port	65535

Settings	Guidelines
Minimum Destination Port	1
Maximum Destination Port	-65535

Creating service groups

You configure service group objects when you have more than one service you want to specify in a rule that matches service. You can group all Web services and group all mail services, for example, if you want to have rules that treat those as groups.

The following policies use service groups:

- Link load balancing policies

Basic Steps

1. Create service objects.
2. Configure service group objects. You can add up to 256 members in a group.
3. Select the service groups when you configure your policies.

Before you begin:

- You must have Read-Write permission for System settings.

To configure a service group:

1. Go to Shared Resources > Service.
2. Click **Service Group**.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Service Group configuration on page 354](#).
5. Save the configuration.

Service Group configuration

Settings	Guidelines
Name	Specify a unique name for the service group configuration. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Member List	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Service	Select a service object.

Configuring WCCP

Web Cache Communication Protocol (WCCP) is a Cisco-developed content-routing protocol that provides a mechanism to redirect traffic flows in real-time. The FortiADC supports Version 2 (WCCPv2).

With WCCP, the FortiADC can forward client traffic to WCCP compatible devices, where additional actions will be performed (that are not native to the FortiADC), and then, after undergoing these processes, the traffic will be sent back to the FortiADC.

To configure a WCCP object:

1. Go to System > **WCCP**.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration according to the table below.
4. Click **save**.
5. Go to Networking > Interface. Select an interface and open the dialogue.
6. Under Mode Specifics, find the **WCCP button**, and click On. (Default is off).
7. Click **save**.
8. Go to Server Load Balance > Virtual Server.
9. Select a virtual server. Go to **Monitoring**.
10. Enable the WCCP button, click on.



Only Layer 7 Virtual Servers are supported.

WCCP configuration

Settings	Description
Service ID	Name of the service group. Range 0-255.
Authentication	• Disable—No password is required. Default. • Enable—Opens up a text box. Specify the password.
Forward Method	• GRE—Encapsulates the intercepted packet in an IP GRE header with a source IP address of the WCCP server and a destination IP address of the target WCCP client. This allows the WCCP server to be multiple Layer 3 hops away from the WCCP client. • L2—Rewrites the destination MAC address of the intercepted packet to equal the MAC address of the target WCCP client. L2 forwarding requires that the WCCP server is Layer 2 adjacent to the WCCP client. • any—Cache server determines the method.
Return Method	Defines how a cache server declines a redirected packet, and returns it to the FortiADC (see forward-method above for option descriptions).
Assignment Method	Defines which assignment method the FortiADC prefers:

Settings	Description
	• HASH—A hash key based on any combination of the source and destination IP and port of the packet. • MASK—A mask value specified with a maximum of 7 bits and, like the hash key, can be configured to cover both the source and destination address space. • any—Cache server determines the method.
Group Address	IP multicast address used by the cache routers. The default, 0.0.0.0, means the FortiADC will ignore multicast WCCP traffic. Otherwise, set the address between 244.0.0.0 to 239.255.255.255.
Router ID	IP address known to all cache engines, and identifies an interface on the FortiADC to the cache engines. If all cache engines connect to the same FortiADC interface, use the default address of 0.0.0.0. However, if the cache engines can connect to different FortiADC interfaces, you must set router-id to a specific IP address, which must then be added to the configuration of the cache engines that connect to that interface.
Server List	IP address and netmask for up to four cache servers.

Chapter 12: Basic Networking

This chapter includes the following topics:

- [Configuring network interfaces on page 357](#)
- [Configuring static routes on page 366](#)
- [Configuring policy routes on page 368](#)

See [Chapter 18: Advanced Networking](#) for advanced topics.

Configuring network interfaces

This section covers the following topics:

- [Physical interface](#)
- [VLAN interface](#)
- [Aggregate interface](#)
- [Loopback interface](#)
- [Softswitch](#)
- [Configuring network interfaces](#)
- [Configuring management interface on page 364](#)

Physical interfaces


Each physical network port (or vNIC on FortiADC-VM) has a network interface that directly corresponds to it—that is, a “physical network interface.”

Physical ports have three uses:

- **Management**—The network interface named port1 is typically used as the management interface.
- **HA**—If you plan to deploy HA, you must reserve a physical port for HA heartbeat and synchronization traffic. Do *not* configure the network interface that will be used for HA; instead, leave it unconfigured or “reserved” for HA.
- **Traffic**—The remaining physical ports can be used for your target traffic—these are your “traffic interfaces.”

Traffic interfaces can be associated with logical interfaces. The system supports two types of logical interfaces: VLAN and aggregate. [Physical and logical interfaces on page 357](#) illustrates how physical ports are associated with physical and logic interfaces.

Physical and logical interfaces



With VLANs, multiple VLAN logical interfaces are associated with a single physical port. With link aggregation, it is the reverse: multiple physical interfaces are associated with a single aggregate logical interface.

[Physical network interfaces on page 358](#) lists factory default IP addresses for physical network interfaces.

Physical network interfaces

Network Interface*	IPv4 Address/Netmask	IPv6 Address/Netmask
port1	192.168.1.99/24	::/0
port2	0.0.0.0/0	::/0
port3	0.0.0.0/0	::/0
port4	0.0.0.0/0	::/0
...		
* The number of physical network interfaces varies by model.		

VLAN interface

You can use [IEEE 802.1q](#) VLAN to reduce the size of a broadcast domain, thereby reducing the amount of broadcast traffic received by network hosts and improving network performance.

Unlike physical LANs, VLANs do not require you to install separate hardware switches and routers to achieve this effect. Instead, VLAN-compliant switches restrict broadcast traffic based upon whether its VLAN ID matches that of the destination network. As such, VLAN trunks can be used to join physically distant broadcast domains as if they were close.

The VLAN ID is part of the tag that is inserted into each Ethernet frame in order to identify traffic for a specific VLAN. FortiADC appliances handle VLAN header addition automatically, so you do not need to adjust the maximum transmission unit (MTU). Depending on whether the device receiving a packet operates at Layer 2 or Layer 3 of the network, a VLAN tag might be added, removed, or rewritten before forwarding to other nodes on the network. For example, a Layer 2 switch typically adds or removes a tag when forwarding traffic among members of the VLAN, but does not route tagged traffic to a different VLAN ID. In contrast, a FortiADC content-based routing policy might forward traffic between different VLAN IDs (also known as inter-VLAN routing).

Note: VLANs are not designed to be a security measure, and should not be used where untrusted devices and/or individuals outside of your organization have access to the equipment. VLAN tags are not authenticated, and can be ignored or modified by attackers. VLAN tags rely on the voluntary compliance of the receiving host or switch.

Aggregate interface

Link aggregation (also called NIC teaming/bonding or link bundling) forms a network interface that queues and transmits over multiple wires (also called a port channel), instead of only a single wire (as FortiADC would normally do with a single network interface per physical port). This multiplies the bandwidth that is available to the network interface, and therefore is useful if FortiADC is deployed inline with your network backbone.

Link aggregation on FortiADC complies with [IEEE 802.1ax](#) and [IEEE 802.3ad](#) and distributes Ethernet frames using a modified round-robin behavior. If a port in the aggregation fails, traffic is redistributed automatically to the remaining ports with the only noticeable effect being a reduced bandwidth. When broadcast or multicast traffic is received on a port in the aggregation, reverse traffic will return on the same port.

When link aggregation uses a round-robin that considers only Layer 2, Ethernet frames that belong to an HTTP request can sometimes arrive out of order. Because network protocols at higher layers often do not gracefully handle this (especially TCP, which may decrease network performance by requesting retransmission when the expected segment does not arrive), FortiADC's frame distribution algorithm is configurable. For example, if you notice that performance with link aggregation is not as high as you expect, you could try configuring FortiADC to queue related frames consistently to the same port by considering the IP session (Layer 3) and TCP connection (Layer 4), not simply the MAC address (Layer 2).

You must also configure the router, switch, or other link aggregation control protocol (LACP)-compatible device to which FortiADC is connected with the same speed/duplex settings, and it must have ports that can be aggregated. In a deployment like this, the two devices use the cables between the ports to form a trunk, not an accidental Layer 2 (link) network loop. FortiADC uses LACP to detect the following conditions:

- Suitable links between itself and the other device, and form a single logical link.
- Individual port failure so that the aggregate interface can redistribute queuing to avoid a failed port.

Loopback interface

A loopback interface is a virtual interface. Like any other interface, a loopback interface can be assigned an address of its own. Unlike any other interface, a loopback interface, once configured, is always up and available. Because a loopback interface never goes down, it is often used for troubleshooting, i.e., the FortiADC appliance, in our case.

In addition, loopback interfaces are also used by BGP and OSPF protocols to determine properties specific to the protocols for a device or network.

Softswitch

A softswitch, or software switch, is a virtual switch that is implemented at the software or firmware level rather than the hardware level. It can be used to simplify communication between devices connected to different FortiADC interfaces. For example, using a softswitch, you can place the FortiADC interface connected to an internal network on the same subnet as your wireless interfaces. This allows devices on the internal network to communicate with devices on the wireless network without any additional configuration.

A softswitch can also be useful if you require more hardware ports for the switch on a FortiADC unit. For example, if your FortiADC has a 4-port switch, WAN1, WAN2, and DMZ interfaces, and you need one more port, you can create a softswitch that includes the 4-port switch and the DMZ interface all on the same subnet. Such applications also apply to wireless interfaces, virtual wireless interfaces, and physical interfaces.

Similar to a hardware switch, a softswitch functions like a single interface. It has one IP address, and all interfaces in the softswitch are on the same subnet. Traffic between devices connected to each interface is not regulated by security policies, and traffic passing in and out of the switch is affected by the same policy. For more information, see the [FortiADC Transparent Mode Configuration Guide](#).

Configuring network interfaces

You can edit the physical interface configuration. You cannot create or delete a physical interface configuration.

Before you begin:

- You must have Read-Write permission for System settings.

To configure a network interface:

- Go to Networking > Interface.
- Double-click the row for a physical interface to edit its configuration or click **Create New** if you want to configure an aggregate or VLAN interface.
- Complete the configuration as described in [Network interface configuration on page 360](#).
- Save the configuration.

Network interface configuration

Settings	Guidelines
Common Settings	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name
Status	The Status column is not the detected physical link status; it is the administrative status (Up/Down) that indicates whether you permit the network interface to receive and/or transmit packets.
Allow Access	Allow inbound service traffic. Select from the following options: - HTTP—Enables connections to the web UI. We recommend this option only for network interfaces connected to a trusted private network, or directly to your management computer. - HTTPS—Enables secure connections to the web UI. We recommend this option instead of HTTP. - Ping—Enables ping and traceroute to be received on this network interface. When it receives an ECHO_REQUEST ("ping"), FortiADC will reply with ICMP type 0 (ECHO_RESPONSE or "pong"). - SNMP—Enables SNMP queries to this network interface. - SSH—Enables SSH connections to the CLI. We recommend this option instead of Telnet. - Telnet—Enables Telnet connections to the CLI. We recommend this option only for network interfaces connected to a trusted private network, or directly to your management computer.
Dedicated HA management IP	Note: Starting from the v. 4.8.1 release, this option is replaced by "Management Interface". Therefore, it is removed from the GUI though it still remains on the Console. For more information, see Configuring management interface on page 364 .

Settings	Guidelines
Virtual Domain	If applicable, select the virtual domain to which the configuration applies.
Mode	- Static—Specify a static IP address. The IP address must be on the same subnet as the network to which the interface connects. Two network interfaces cannot have IP addresses on the same subnet (i.e. overlapping subnets). - PPPoE—Use PPPoE to retrieve a configuration for the IP address, gateway, and DNS server. For example, if this interface uses a DSL connection to the Internet, your ISP may require this option.
Static	
Traffic Group	Select either of the following: - Default - Create New
Floating	Enable/Disable floating IP.
Floating IP	Enter the floating IP.
IPv4/Netmask	Specify the IP address and CIDR-formatted subnet mask, separated by a forward slash (/), such as 192.0.2.5/24. Dotted quad formatted subnet masks are not accepted.
IPv6/Netmask	Specify the IP address and CIDR-formatted subnet mask, separated by a forward slash (/), such as 2001:0db8:85a3::8a2e:0370:7334/64. Dotted quad formatted subnet masks are not accepted.
Secondary IP Address	Secondary IP addresses can be used when you deploy the system so that it belongs to multiple logical subnets. If you assign multiple IP addresses to an interface, you must assign them static addresses. To add secondary IP addresses, enable the feature and save the configuration. After you have saved it the first time, you can edit it to add secondary IP addresses and enable inbound traffic to that address.
PPPoE	
Username	PPPoE account user name.
Password	PPPoE account password.
Discovery Retry Timeout	Seconds the system waits before it retries to discover the PPPoE server. The default is 5 seconds. The valid range is 1-255.
DNS Server Override	Use the DNS addresses retrieved from the PPPoE server instead of the one configured in the FortiADC system settings.
Retrieve Default Gateway	Use the default gateway retrieved from the PPPoE server instead of the one configured in the FortiADC system settings.
Type	If you are editing the configuration for a physical interface, you cannot set the type. If you are configuring a logical interface, you can select from the following options: - Aggregate—A logical interface you create to support the aggregation of multiple physical interfaces. - VLAN—A logical interface you create to VLAN subinterfaces on a single physical interface.

Settings	Guidelines
Aggregate	
Member	Select the physical interfaces that are included in the aggregation.
Aggregate Mode	Link aggregation type: • 802.3ad • Balance-alb • Balance-rr • Balance-tlb • Balance-xor • Broadcast
Aggregate Algorithm	Connectivity layers that will be considered when distributing frames among the aggregated physical ports: • Layer 2 • Layer 2-3 • Layer 3-4
VLAN	
VLAN ID	VLAN ID of packets that belong to this VLAN. If one physical network port (that is, a VLAN trunk) will handle multiple VLANs, create multiple VLAN subinterfaces on that port, one for each VLAN ID that will be received. If multiple different physical network ports will handle the same VLANs, on each of the ports, create VLAN subinterfaces that have the same VLAN IDs. The valid range is between 1 and 4094. The value you specify must match the VLAN ID added by the IEEE 802.1q-compliant router or switch connected to the VLAN subinterface.
Interface	Physical interface associated with the VLAN; for example, port2.
Secondary IP List	
IP Address	Secondary IP addresses can be used when you deploy the system so that it belongs to multiple logical subnets. If you assign multiple IP addresses to an interface, you must assign them static addresses. To add secondary IP addresses, enable the feature and save the configuration. After you have saved it the first time, you can edit it to add secondary IP addresses and enable inbound traffic to that address. For each address, specify an IP address using the CIDR-formatted subnet mask, separated by a forward slash (/), such as 192.0.2.5/24.
Allow Access	Select the services that are allowed to send inbound traffic.
HA Node IP List	
IP Address	You use the HA node IP list configuration in an HA active-active deployment. On each HA cluster node, add an HA node IP list that includes an entry for each cluster node. When the

Settings	Guidelines
	appliance is in standalone mode, it uses the physical port IP address; when it is in HA mode, it uses the HA node IP list address. For each address, specify an IP address using the CIDR-formatted subnet mask, separated by a forward slash (/), such as 192.0.2.5/24.
Node ID	ID of the corresponding node.
Allow Access	Select the services that are allowed to send inbound traffic.



In an HA active-active deployment, if an interface uses secondary IP addresses, you must use the CLI to enable the HA node secondary IP address list, and then configure the list:

```
FADC # config system interface
```

```
FADC (interface) # edit port3
```

```
FADC (port3) # set ha-node-secondary-ip enable
```

```
FADC (port3) # config ha-node-secondary-ip-list
```

```
FADC (ha-node-second~r) # edit 1
```

```
Add new entry '1' for node 2221
```

```
FADC (1) # set ip 192.168.1.100
```

```
FADC (1) # set allowaccess https http ping snmp ssh
```

```
FADC (1) # end
```

```
FADC (port3) # end
```

To configure a physical interface using the CLI:

```
config system interface
edit <port_name>
set ip <ip&netmask>
set allowaccess {http https ping snmp ssh telnet}
end
```

To configure an aggregate interface using the CLI:

```
config system interface
edit <specified_name>
set type agg
set aggregate-mode {802.3ad | balance-alb | balance-rr | balance-
    tlb | balance-xor | broadcast}
set aggregate-algorithm {layer2 | layer2_3 | layer3_4}
set member <port_name> <port_name>
set ip <ip&netmask>
end
```

To configure a VLAN interface using the CLI:

```
config system interface
edit <specified_name>
set type vlan
set vlanid <number>
set interface <port_name>
set ip <ip&netmask>
end
```

Configuring management interface

The management interface should be used exclusively by the FortiADC administrator to manage the devices, physical or virtual, (such as configuring or debugging it). It should be an interface through which FortiADC's management traffic (such as license authenticating) can traverse at any time without affecting normal network traffic. It is especially useful for slave devices in HA active-passive mode. The management interface has the highest access permissions, and the FortiADC administrator should make sure that it is used for management traffic only, and avoid using it for normal traffic.

You can configure the management interface from either the GUI or the CLI. This section discusses how to configure the management interface from the GUI. For instructions on how to configure management interface using the CLI, see the section "Moving from 'Dedicated HA Management IP' to 'Management Interface'" at the end of this section.

Note:

- It must be noted that, because the management interface is a global configuration, it must and can only be configured from the "global" system interface and used by the "global" administrator. Therefore, the option is NOT available on any VDOM.
- This "management interface" is a virtual interface, which is quite different from the default, factory-set, "physical" management interface used to set up the appliance for the first time, as discussed in [Step 2: Configure the management interface on page 54](#), Chapter 3: "Getting Started", of this Handbook.

To configure the management interface:

1. From FortiADC's global interface, click Networking > Interface to open the interface configuration page.
2. In the Management Interface section, click the edit button, the pencil, in the top right corner to enable the management interface. The fields for management interface configuration appear on the page.
3. Make the desired selections and entries as described in [Management interface configuration on page 365](#).
4. Click Save when done.

Management interface configuration

Option	Guidelines
Management Status	Enable this option.
Management Interface	Select an interface (port) from the list menu. Note: The management interface handles all incoming and outgoing management traffic. <i>It must and can only work in promiscuous mode.</i>
Management IP	Enter the IP address of the management interface. Note: Once enabled, the management network IP becomes active in all each modes (i.e., standalone, active-passive, active-active, and VRRP). Therefore, the management interface IP address must be unique and must NOT be used in regular functions, such as the virtual server IP addresses, source NAT pool IP addresses, source NAT pool trans-to IP addresses, 1-to-1 NAT external/mapped IP addresses, and all the other IP addresses configured on the interface. Otherwise, it will conflict with the HA functions.
Management IP Allow Access	Select the type or types of management traffic that are allows to access the Management interface.
Management MAC Address	Specify the MAC address of the management interface. Note: If you do not specify a management MAC address, FortiADC will automatically populate the field with a random MAC address when you click the Save button

"Dedicated HA Management IP" vs. "Management Interface"

In pre-FortiADC 4.8.1 releases, the GUI had an option in interface configuration (Networking > Interface > Add) which allows you to set an interface as the "Dedicated HA Management IP", which functions exactly the same as the "Management Interface" in 4.8.1. With the 4.8.1 release, that option is removed from the GUI (even though it is still available in the Console) is replaced by the "Management Interface". If you have a dedicated HA management IP configured on a pre-4.8.1 version of FortiADC, we highly recommend that you delete it, and then configure a management interface instead, after you've upgraded to 4.8.1. This will help streamline your interface configuration and make system management easier.

All this can be done through FortiADC's Console only. The following instructions show how to delete your old "Dedicated HA Management IP" and configure the "Management Interface" using the Console in FortiADC 4.8.1:

Step 1: Remove the "Dedicate HA Management IP"

Execute the following commands:

```
config system interface
edit "port1"
set dedicate-to-mgmt disable
unset ip
next
end
```

Step 2: Configure the "Management Interface":

Execute the following commands:

```
config system ha
set mgmt-status enable
set mgmt-interface port1
set mgmt-ip 10.106.129.120/24
set mgmt-ip-allowaccess https ping ssh snmp http telnet
set mgmt-mac-addr fe:02:98:41:93:f8
end
```

Configuring static routes

Network systems maintain route tables to determine where to forward TCP/IP packets. Routes for outbound traffic are chosen according to the following priorities:

- Link local routes—Self-traffic uses link local routes.
- LLB Link Policy route—Configured policy routes have priority over default routes.
- Policy route—Configured policy routes have priority over default routes.
- Static route / ISP route / OSPF route—Priority is based on the distance metric. By default, distance for static routes is 10, for ISP is 20, for OSPF is 110, for EBGp is 20, and for IBGP is 200. The distance metric is configurable for static routes and OSPF routes, but not for ISP routes.
- Default LLB Link Policy route—Default routes have lower priority than configured routes.
- Default static route / OSPF route—Default routes have lower priority than configured routes.

The system evaluates content route rules first, then policy routes, then static routes. The packets are routed to the first route that matches. The static route table, therefore, is the one that must include a "default route" to be used when no more specific route has been determined.

Static routes specify the IP address of a next-hop router that is reachable from that network interface. Routers are aware of which IP addresses are reachable through various network pathways, and can forward those packets along pathways capable of reaching the packets' ultimate destinations. The FortiADC system itself does not need to know the full route, as long as the routers can pass along the packet.

You must configure at least one static route that points to a router, often a router that is the gateway to the Internet. You might need to configure multiple static routes if you have multiple gateway routers, redundant ISP links, or other special routing cases.

Before you begin:

- You must have Read-Write permission for System settings.

To configure a static route:

1. Go to Networking > Routing.
The configuration page displays the Static tab.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in [Static route configuration on page 367](#).
4. Save the configuration.

Static route configuration

Settings	Guidelines
Destination	Address/mask notation to match the destination IP in the packet header. It is a best practice to include a default route. If there is no other, more specific static route defined for a packet's destination IP address, a default route will match the packet, and pass it to a gateway router so that any packet can reach its destination. If you do not define a default route, and if there is a gap in your routes where no route matches a packet's destination IP address, packets passing through the FortiADC towards those IP addresses will, in effect, be null routed. While this can help to ensure that unintentional traffic cannot leave your FortiADC and therefore can be a type of security measure, the result is that you must modify your routes every time that a new valid destination is added to your network. Otherwise, it will be unreachable. A default route ensures that this kind of locally-caused "destination unreachable" problem does not occur. Specify 0.0.0.0/0 or ::/0 to set a default route for all packets.
Gateway	Specify the IP address of the next-hop router where the FortiADC system will forward packets for this static route. This router must know how to route packets to the destination IP addresses that you have specified, or forward packets to another router with this information. For a direct Internet connection, this will be the router that forwards traffic towards the Internet, and could belong to your ISP. The gateway must be in the same subnet as the interface used to reach it.
Distance	The default administrative distance is 10, which makes it preferred to OSPF routes that have a default of 110. We recommend you do not change these settings unless your deployment has exceptional requirements.

To configure a static route using the CLI:



```
config router static
edit 1
set destination <ip address/netmask>
set gateway <ip address>
set distance <value>
end
```

Configuring policy routes

Network systems maintain route tables to determine where to forward TCP/IP packets. Policy routes set the gateway for traffic with a source and destination that match the policy.

Routes for outbound traffic are chosen according to the following priorities:

1. Link local routes—Self-traffic uses link local routes.
2. LLB Link Policy route—Configured policy routes have priority over default routes.
3. Policy route—Configured policy routes have priority over default routes.
4. Static route / ISP route / OSPF route—Priority is based on the distance metric. By default, distance for static routes is 10, for ISP routes is 20, and for OSPF routes is 110. The distance metric is configurable for static routes and OSPF routes, but not ISP routes.
5. Default LLB Link Policy route—Default routes have lower priority than configured routes.
6. Default static route / OSPF route—Default routes have lower priority than configured routes.

The system evaluates policy routes, then static routes. The packets are routed to the first route that matches. The policy route table, therefore, need not include a “default route” for packets that do not match your policy because those packets can be forwarded to the default route set in the static route table.

Most policy route settings are optional, so a matching route might not provide enough information to forward the packet. In that case, the FortiADC appliance 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 destination address is the only match criteria in the policy route, the FortiADC appliance looks up the IP address of the next-hop router in its routing table. This situation could occur when interfaces are dynamic (such as DHCP or PPPoE) and you do not want or are unable to specify a static IP address of the next-hop router.

Before you begin:

- You must have Read-Write permission for System settings.

To configure a policy route:

1. Go to Networking > Routing.
2. Click the **Policy** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Policy route configuration on page 368](#).
5. Save the configuration.

Policy route configuration

Settings	Guidelines
Source	Address/mask notation to match the source IP in the packet header. To match any value, either leave it blank or enter 0.0.0.0/32.
Destination	Address/mask notation to match the destination IP in the packet header. To match any value, leave it blank or enter 0.0.0.0/32.
Gateway	IP address of the next-hop router where the FortiADC system will forward packets for this policy route. This router must know how to route packets to the destination subnet, or forward packets to another router with this information.

Chapter 13: System Management

This chapter includes the following topics:

- [Configuring basic system settings on page 371](#)
- [Configuring system time on page 372](#)
- [Updating firmware on page 373](#)
- [Configuring an SMTP mail server on page 377](#)
- [Configuring FortiGuard service settings on page 377](#)
- [Pushing/pulling configurations on page 379](#)
- [Configuring FortiSandbox service on page 381](#)
- [Backing up and restoring configuration on page 382](#)
- [SCP support for configuration backup on page 386](#)
- [Rebooting, resetting, and shutting down the system on page 386](#)
- [Create a traffic group on page 387](#)
- [Manage administrator users on page 389](#)
- [Create administrator users on page 390](#)
- [Configure access profiles on page 391](#)
- [Enable password policies on page 394](#)
- [Two-factor authentication on page 329](#)
- [Configuring SNMP on page 394](#)
- [Download SNMP MIBs on page 396](#)
- [Configure SNMP threshold on page 396](#)
- [Configure SNMP v1/v2 on page 397](#)
- [Configure SNMP v3 on page 398](#)
- [Configuring central management on page 399](#)
- [Manage and validate certificates on page 402](#)
- [Generating a certificate signing request on page 405](#)
- [Creating a local certificate group on page 408](#)
- [Importing intermediate CAs on page 410](#)
- [Creating an intermediate CA group on page 411](#)
- [OCSP stapling on page 412](#)
- [Validating certificates on page 413](#)
- [Importing CRLs on page 415](#)
- [Adding OCSPs on page 416](#)
- [Importing OCSP signing certificates on page 420](#)
- [Importing CAs on page 421](#)
- [Creating a CA group on page 422](#)
- [System alerts on page 423](#)
- [Configuring alert policies on page 424](#)
- [Creating alert configurations on page 425](#)
- [Configuring alert actions on page 423](#)
- [Configuring a syslog object on page 429](#)

- [Configuring an email alert object on page 429](#)
- [Configuring SNMP trap servers on page 428](#)
- [HSM Integration on page 430](#)

Configuring basic system settings

The basic system settings page includes configuration options for the following settings and features:

- Hostname
- Web UI language
- Management service ports
- DNS
- Virtual domain

Before you begin:

- You must have Read-Write permission for System settings.

To configure basic system settings:

1. Click System > Settings.
The configuration page displays the Basic tab.
2. Complete the configuration as described in [Basic settings configuration on page 371](#).
3. Save the configuration.

Basic settings configuration

Settings	Guidelines
Hostname	You can configure a hostname to facilitate system management. If you use SNMP, for example, the SNMP system name is derived from the configured hostname. The hostname can be up to 35 characters in length. It can include US-ASCII letters, numbers, hyphens, and underscores, but not spaces and special characters. The System Information widget and the <code>get system status</code> CLI command display the full hostname. If the hostname is longer than 16 characters, the name is truncated and ends with a tilde (~) to indicate that additional characters exist, but are not displayed.
Language	English or Simplified Chinese.
Idle Timeout	Log out an idle administrator session. The default is 30 minutes.
HTTP Port	Specify the port for the HTTP service. Usually, HTTP uses port 80.
HTTPS Port	Specify the port for the HTTPS service. Usually, HTTPS uses port 443.
Telnet Port	Specify the port for the Telnet service. Usually, Telnet uses port 25.
SSH Port	Specify the port for the SSH service. Usually, SSH uses port 22.
Primary DNS	The system must be able to contact DNS servers to resolve IP addresses and fully qualified domain names. Your Internet service provider (ISP) might supply IP addresses of DNS servers, or you might want to use the IP addresses of your own DNS servers. You must provide unicast,

Settings	Guidelines
	non-local addresses for your DNS servers. Localhost and broadcast addresses are not accepted.
	Incorrect DNS settings or unreliable DNS connectivity can cause issues with other features, such as FortiGuard services and NTP system time.
Secondary DNS	IPv4/IPv6 address of the secondary DNS server for your local network.
Virtual Domain	Enables the virtual domain feature. Before you enable it, make sure you understand how the system implements virtual domains. See Chapter 16: Virtual Domains .
Config Sync Enable	Enable/disable the configuration synchronization feature. This feature is related to Pushing/pulling configurations , not HA synchronization. Disabled by default.

Configuring system time

The system time must be accurate for many features to work, including scheduling, logging, and SSL/TLS-related features.

We recommend that you use Network Time Protocol (NTP) to maintain the system time. As an alternative when NTP is not available or is impractical, you can set the system time manually.

You can change the system time with the web UI or the CLI.

Before you begin:

- You must have Read-Write permission for System settings.

To configure the system time:

- Go to System > Settings.
- Click the **Maintenance** tab.
- Complete the configuration as described in [System time configuration on page 372](#).
- Save your changes.

System time configuration

Setting	Guidelines
System Time	Displays the system time. You can use NTP to set the system time, or use the controls to set the system time manually. Specify time in HH:MM:SS format.
Daylight Saving Time	Enable if you want the system to adjust its own clock when its time zone changes between daylight saving time (DST) and standard time.
Time Zone	Select the time zone where the appliance is located.
NTP	
NTP	Select to use NTP.
NTP Server	Specify a space-separated list of IP addresses or FQDNs for an NTP server or pool, such as <code>pool.ntp.org</code> .

Setting	Guidelines
	To find an NTP server, go to http://www.ntp.org .
Synchronizing Interval	Specify how often the system synchronizes its time with the NTP server. The default is 60 minutes. The valid range is 1-1440.

To configure NTP using the CLI:

```
config system time ntp
set ntpsync enable
set ntpserver {<server_fqdn> | <server_ipv4>}
set syncinterval <minutes_int>
end
```



To configure the system time manually:

```
config system time ntp
set ntpsync disable
end
config system time manual
set zone <timezone_index>
set daylight-saving-time {enable|disable}
end
execute date <MM/DD/YY> <HH:MM:SS>
```

Updating firmware

This topic includes the following information:

- [Upgrade considerations](#)
- [Updating firmware using the web UI](#)
- [Updating firmware using the CLI](#)

Upgrade considerations

The following considerations help you determine whether to follow a standard or non-standard upgrade procedure:

- HA—Updating firmware on an HA cluster requires some additions to the usual steps for a standalone appliance. For details, see [Updating firmware for an HA cluster](#).
- Re-imaging—If you are installing a firmware version that requires a different size of system partition, you might be required to re-image the boot device. Consult the release notes. In that case, do *not* install the firmware using this procedure. Instead, see [Restoring firmware \(“clean install”\)](#).
- Downgrades—If you are downgrading the firmware to a previous version, and the settings are not fully backwards compatible, the system might remove incompatible settings or use the default values for that version of the firmware. You might need to reconfigure some settings.

Important: Read the release notes for release-specific upgrade considerations.

Updating firmware using the web UI

User interface for managing firmware updates on page 374 shows the user interface for managing firmware (either upgrades or downgrades). Firmware can be loaded on two disk partitions: the active partition and the alternate partition. The upgrade procedure:

- Updates the firmware on the inactive partition and then makes it the active partition.
- Copies the firmware on the active partition, upgrades it, and installs it in place of the configuration on the inactive partition.

For example, if partition 1 is active, and you perform the upgrade procedure:

- Partition 2 is upgraded and becomes the active partition; partition 1 becomes the alternate partition.
- The configuration on partition 1 remains in place; it is copied, upgraded, and installed in place of the configuration on partition 2.

The reason for this is to preserve the working system state in the event upgrade fails or is aborted.

User interface for managing firmware updates

The screenshot displays the FortiADC web UI. On the left is a navigation menu with 'System' selected, showing sub-items like 'Settings', 'High Availability', 'Traffic Group', 'Administrator', 'SNMP', 'Debug', 'Certificate', 'Manage Certificates', 'Verify', and 'Alert'. The main content area has tabs for 'System' and 'Settings', with 'Maintenance' selected under 'System'. The 'Time' section shows: System Date Time (2018-09-27 08:16:57), Daylight Saving Time (disable), Time Zone ((GMT+9:00)Osaka,Sapporo,Tokyo,Seoul), NTP (enable), NTP Server (pool.ntp.org), and Synchronizing Interval (60). The 'Firmware' section features an 'Upgrade Firmware' button and a table:

Partition	Active	Last Upgrade	Firmware Version
1	Disable	Tue Sep 18 04:25:35 2018	FA-HYV-5.01.00-FW-build0227
2	Enable	Tue Sep 25 07:11:51 2018	FA-HYV-5.01.01-FW-build0239

At the bottom right of the firmware section is a 'Boot Alternate Firmware' button.

Before you begin:

- Download the firmware file from the Fortinet Customer Service & Support website: <https://support.fortinet.com/>
- Read the release notes for the version you plan to install.
- Back up your configuration before beginning this procedure. Reverting to an earlier firmware version could reset settings that are not compatible with the new firmware.
- You must have super user permission (user **admin**) to upgrade firmware.

To boot the firmware on the alternate partition:

- Click **Boot Alternate Firmware**.

The system reboots, the alternate becomes the active firmware, and the active becomes the alternate firmware.

To update firmware:

1. Go to System > Settings.
2. Click the **Maintenance** tab.
3. Scroll to the Upgrade section.
4. Click **Choose File** to locate and select the file.
5. Click ⓘ to upload the firmware and reboot.

The system replaces the firmware on the alternate partition and reboots. The alternate (upgraded) partition becomes the active, and the active becomes the alternate.



When you update software, you are also updating the web UI. To ensure the web UI displays the updated pages correctly:

- Clear your browser cache.
- Refresh the page.

In most environments, press Ctrl-F5 to force the browser to get a new copy of the content from the web application. See the Wikipedia article on browser caching issues for a summary of tips for many environments:

https://en.wikipedia.org/wiki/Wikipedia:Bypass_your_cache.

Updating firmware using the CLI

The CLI upgrade procedure replaces the firmware on the alternate partition and reboots. The alternate (upgraded) partition becomes the active, and the active becomes the alternate.

Note: The CLI does not have an equivalent of the web UI **Boot Alternative Firmware** command.

Before you begin:

- Read the release notes for the version you plan to install. If information in the release notes is different from this documentation, follow the instructions in the release notes.
- You must be able to use TFTP to transfer the firmware file to the FortiADC. Download and install a TFTP server, like `tftpd` ([Windows](#), [Mac OS X](#), or [Linux](#)), on a server on the same subnet as the FortiADC.
- Download the firmware file from the Fortinet Customer Service & Support website: <https://support.fortinet.com/>
- Copy the firmware image file to the root directory of the TFTP server.
- Back up your configuration before beginning this procedure.
- You must have super user permission (user **admin**) to upgrade firmware.



TFTP is not secure, and it does not support authentication. You should run it only on trusted administrator-only networks, and never on computers directly connected to the Internet. Turn off `tftpd` off immediately after completing this procedure.

To install firmware via the CLI:

1. Connect your management computer to the FortiADC console port using an RJ-45-to-DB-9 serial cable or a null-modem cable.
2. Initiate a connection to the CLI and log in as the user **admin**.
3. Use an Ethernet cable to connect FortiADC port1 to the TFTP server directly, or connect it to the same subnet as the TFTP server.
4. If necessary, start the TFTP server.
5. Use the following command to transfer the firmware image to the FortiADC system:

```
execute restore image tftp <filename> <tftp_ipv4>
```

The following example shows an upgrade:

```
FortiADC-VM # execute restore image tftp FAD_VM-v400-build0308-FORTINET.out 192.0.2.1
This operation will replace the current firmware version!
Do you want to continue? (y/n)y
Connect to tftp server 192.0.2.1 ...
Please wait...
```

```
#####
```

```
Get image from tftp server OK.
```

```
Check image trailer OK.
```

```
Check image OK.
```

```
FortiADC-VM #
```

The following example shows a downgrade:

```
FortiADC-VM # execute restore image tftp FAD_VM-v400-build0307-FORTINET.out 192.0.2.1
This operation will replace the current firmware version!
Do you want to continue? (y/n)y
Connect to tftp server 192.0.2.1 ...
Please wait...
```

```
#####
```

```
Get image from tftp server OK.
```

```
Check image trailer OK.
```

```
This operation will downgrade the current firmware version!
```

```
Do you want to continue? (y/n)y
```

```
FortiADC-VM #
```

6. To verify the upgrade, display the system version number:

```
FortiADC-VM # get system status
Version: FortiADC-VM v4.2.0,build0307,150209
VM Registration: Valid: License has been successfully authenticated with registration servers.
VM License File: License file and resources are valid.
VM Resources: 1 CPU/1 allowed, 1620 MB RAM/2048 MB allowed, 23 GB Disk/1024 GB allowed
...
```



If the download fails after the integrity check with the error message `invalid compressed format (err=1)`, but the firmware matches the integrity checksum on the Fortinet Customer Service & Support website, try a different TFTP server.

Configuring an SMTP mail server

You can configure an SMTP email server if you want to send alerts by email. See [Configuring report email](#) for information on alerts.

Before you begin:

- You must have Read-Write permission for System settings.

To configure SMTP:

1. Go to System > Settings.
2. Click the **Services** tab.
3. Complete the configuration as described in [SMTP configuration on page 377](#).
4. Save the configuration.

SMTP configuration

Settings	Guidelines
Address	IP address or FQDN of an SMTP server (such as FortiMail) or email server that the appliance can connect to in order to send alerts and/or generated reports.
Port	Listening port number of the server. Usually, SMTP is 25.
Authentication	Enable if the SMTP server requires authentication.
Security	STARTTLS is an extension to plain text communication protocols. It enables a plain text connection to be upgraded to an encrypted (TLS or SSL) connection instead of using a separate port for encrypted communication. Specify this option if you have implemented STARTTLS for your mailserver; otherwise, select none .
Username	Username for authentication to the SMTP server.
Password	Password for authentication to the SMTP server.

Configuring FortiGuard service settings

FortiGuard periodically updates the WAF Signature Database, IP Reputation Database, and Geo IP Database. You can go to the FortiGuard website to download the update packages that you can upload to FortiADC, or you can schedule automatic updates.

Before you begin:

- If you want to perform a manual update, you must download the update file from the FortiGuard website.

You must have Read-Write permission for System settings.

To configure FortiGuard service settings:

1. Go to System > Settings.
2. Click the **FortiGuard** tab.

3. Complete the configuration as described in [FortiGuard service configuration on page 378](#).
4. Save the configuration.

FortiGuard service configuration

Settings	Guidelines
Support Contract	
Registration	Review your registration and license information. If you need to update your registration or renew your license, click Login Now to open the login page for the Fortinet Service & Support website. Note: If your license is invalid, FortiGuard does not send updates to your FortiADC. The functionality on your FortiADC unit remains intact and useful even though it is out of date.
Hardware	Shows the hardware model of your FortiADC unit.
Firmware	Shows the firmware version on your FortiADC unit.
Enhanced Support	Shows the status of Enhanced Support of your FortiADC unit. .
Comprehensive Support	Shows the status of Comprehensive Support of your FortiADC unit.
FortiGuard Services	
WAF Signature	Shows the version of the Web Application Firewall Signature file on your FortiADC unit. To manually update the file, click Update to display controls that enable you to select and upload the latest WAF Signature file.
IP Reputation	Shows the version of the IP Reputation file on your FortiADC unit. To manually update the file, click Update to display controls that enable you to select and upload the latest IP reputation file.
Geo IP	Shows the version and region of the Geo IP file on your FortiADC unit. To manually update the file, click Update to display controls that enable you to select and upload the latest Geo IP file.
Web Filter	Shows the status of the Web Filter on your FortiADC unit.
Update Schedule	
Scheduled Update	Click the button to enable or disable the Scheduled Update feature. Note: If enabled, you must set the frequency, date, or time of the update schedule. See below.
Scheduled Update Frequency	• Every—Schedule periodic updates. Specify the update interval to perform the scheduled update. • Daily—Schedule daily updates. Specify the time of the day to perform the scheduled update. • Weekly—Schedule weekly updates. Specify the day and time to perform the scheduled update.
Scheduled Update Day	Select the day of the week for the scheduled update.
Scheduled Update Time	Specify the time (hour and minute) for the scheduled update.
Override Server	Click the button to enable or disable the Override Server feature.

Settings	Guidelines
	Note: This feature provides another option for your FortiADC to connect to FortiGuard when it (FortiADC) is unable to connect to FortiGuard via the default FortiGuard server IP address. If enabled, you must enter the Override Server Address that you have obtained from the Fortinet Service and Support team. See below.
Override Server Address	Enter the Override Server Address provided by the Fortinet Service and Support team.
Tunneling	Click the button to enable or disable tunneling. If enabled, you must configure all the settings for the tunneling function. See below. Note: Tunneling, or port forwarding, is a way of transmitting private (usually corporate) data through a public network in a disguised way — the routing nodes in the public network are unaware that the transmission is part of a private network.
Tunneling Address	Enter the Tunneling Address that was provided to you.
Tunneling Port	Enter the Tunneling Port number that was provided to you.
Tunneling Username	Specify your user name for the tunneling configuration.
Tunneling Password	Specify your password for the tunneling configuration.
Save	Click the Save button to save your FortiGuard service configuration.
Web Filter	
Cache Status	Click the button to enable or disable caching of the categorical lists of websites. Note: FortiGuard maintains massive lists of web sites classified into categories so that you can enforce categorical decisions in your rules, like "do not do SSL forward proxy for sites belonging to the Personal Privacy category."
Cache TTL	Specify a cache expiration value. The default is 3600. The valid range is from 10 to 86,400. When the cache expires, FortiADC initiates an update from FortiGuard.
FDS Port	Specify the port to receive updates. The default is 53. An alternative is 8888.
Save	Click Save to save your Web Filter configuration.

Pushing/pulling configurations

You can use the sync list configuration page to push or pull sets of configuration objects to or from a target FortiADC appliance. The push/pull operation is a manual operation. It is not repeated automatically.

Before you begin:

- Configuration synchronization must be enabled on the appliances. Go to System > Settings > Basic.
- You must plan for the impact the configuration push/pull has on the target deployment.
- You must have Read-Write permission for System settings.

To push or pull a configuration:

1. Click System > Settings.
2. Click the **Sync List** tab.

3. Click **Create New** and complete the configuration as described in [Table 126](#).
After you have saved the configuration, it is added to the configuration table.
4. To execute the push/pull operation, select the configuration from the table, select **From** or **To**, and click **Sync**.
5. Check the Status column in the table to see the result of the push/pull operation.
6. Log into the target appliance and check the configuration logs (Log & Report > Log Browsing > Event Log > Configuration). Notice the log entries for each configuration change resulting from the push/pull operation.

Sync List configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Server IP	IP address of the remote appliance.
Password	Password for the admin account on the remote appliance.
Type	• System—Includes <code>config config</code>, <code>config system</code> (except <code>config system mailserver</code>), <code>config user</code>, and <code>config vdom</code> commands. • Networking—Includes <code>config router</code> commands. • LB—Includes <code>config load-balance</code> commands. • Log—Includes <code>config log</code> commands and <code>config system mailserver</code>. • LLB—Includes <code>config link-load-balance</code> commands. • GDS—Includes <code>config global-load-balance</code> and <code>config global-dns-server</code> commands. • Security—Includes <code>config security waf</code> commands. • User—Includes <code>config user</code> commands. Note: For each of the above settings, there are certain parameters that cannot be synchronized through the Sync List feature. For details,

[Table 127](#) highlights the commands that cannot be synced using the Sync List feature, and must be handled manually on a per appliance basis..

Commands that cannot be synced via the Sync List feature

Module	Commands
System	• <code>system global</code> • <code>system tcpdump</code> • <code>system accprofile</code> • <code>system admin</code> • <code>system ha</code> • <code>system snmp sysinfo</code> • <code>system snmp community</code> • <code>system snmp user</code> • <code>system alert-snmp-trap</code> • <code>system fortiguard</code> • <code>system hsm info</code> • <code>system hsm partition</code> • <code>config sync-list</code>

Module	Commands
Networking	• firewall qos-filter • firewall qos-filter6 • router policy • router isp • router setting • firewall nat-snat • firewall vip • router md5-ospf • router ospf • router bgp • system interface • router static
LLB	• link-load-balance virtual-tunnel • link-load-balance flow-policy
Security	• firewall policy • firewall policy6 • firewall connlimit • firewall connlimit6
SLB	• load-balance ippool • load-balance virtual-server
GLB	• global-load-balance link • global-load-balance virtual-server-pool • global-load-balance host • global-load-balance analytic • global-dns-server general • global-dns-server policy
Log & Report	• system mailserver

Configuring FortiSandbox service

FortiADC is integrated with FortiSandbox to enhance its anti-virus capabilities. Upon detecting suspicious traffic segments, FortiADC first conducts some basic analysis of its own and then forwards them to FortiSandbox for further analysis. The latter will then drop or quarantine the malicious traffic segments and forward healthy traffic segments to the back-end servers.

To configure FortiSandbox services:

1. From the navigation bar, click System>Settings.
2. Click the FortiSandbox tab.
3. Make the selection or entries as described in [FortiSandbox service configuration on page 382](#).
4. Click Save when done.

FortiSandbox service configuration

Settings	Description
Type	Select either of the following: FSA—FortiSandbox appliance.
Status	Click the button to enable or disable FortiSandbox service. Note: FortiSandbox is disabled by default.
Server	Enter the IP address of the FortiSandbox appliance. Note: This option applies if you want to use a on-premise FortiSandbox appliance for service.
Email	The email address of the person to be notified.
Source IP	The IP address of the source interface on the FortiADC appliance.

FortiCloud Sandbox file upload limits

[FortiCloud Sandbox file upload limit on page 382](#) shows the maximum number of files per minute that you can upload to FortiCloud Sandbox from various FortiADC platforms.

FortiCloud Sandbox file upload limit

Platform	Number of files uploaded per minute
FortiADC 60F/VM01	5
FortiADC100—400/VM02	10
FortiADC 700D/VM04	20
FortiADC 1000—2000/VM08	50
FortiADC 4000	100

Backing up and restoring configuration

You use the backup procedure to save a copy of your system configuration. A full backup is a zip file.

The backup feature has a few basic uses:

- Saving the configuration as CLI commands that a co-worker or Fortinet support can use to help you resolve issues with misconfiguration.
- Restoring the system to a known functional configuration.
- Creating a template configuration you can edit and then load into another system using the restore procedure.

A complete configuration backup is a zip file that includes the complete configuration files, plus any files you have imported, including error page files, script files, and ISP address book files.

In the event that FortiADC experiences hardware failure, being able to restore the entire backup configuration minimizes the time to reconfigure the system.

All backup files follow the same file-naming convention: `hostname_date_time`. For example, a backup file named "FortiADC-VM_20171214_0830.txt" means that the backup is made of a system whose hostname is "FortiADC-VM", the backup is made at 08:30 on December 14, 2017. It must be noted that the date and time in the backup file name reflects the date and time in your FortiADC's system settings when the backup is performed.

Note: Configuration backups do *not* include data such as logs and reports.



Back up files can include sensitive information, such as HTTPS certificate private keys. We strongly recommend that you password-encrypt your backup files and store them in a secure location.

Before you begin:

- If you are restoring a configuration, you must know its management interface configuration in order to access the web UI after the restore procedure is completed. Open the configuration file and make note of the IP address and network requirements for the management interface (port1). You must also know the administrator username and password.
- You must have Read-Write permission for system settings.

To backup or restore your system configuration:

1. From navigation bar, click System > Settings.
2. Click the **Backup & Restore** tab.
3. Select the desired action and storage location, as described in [Backup and restore configuration on page 383](#).
4. Follow the instructions in the following paragraphs to back up or restore your configuration, or schedule auto backups.

Backup and restore configuration

Actions	Guidelines
Mode	Select one of the options: • Back Up—Use this option to back up the current configuration. Note: The backup is saved to a text file. • Restore—Use this option to restore a previous configuration. The restore file must be a text file. • Auto Backup—Use this option to let FortiADC automatically back up its configuration as scheduled.
Storage	Select one of the storage locations: • Local PC/Server—The local PC or server. (Note: When scheduling auto backups, this refers to the SFTP server.) • ADC—Your FortiADC device.
Entire Configuration	Enable this option to include error page files, script files, and ISP address book files in the backup file. Note: The backup is saved to a tar file. ADC

Run a manual backup

You can back up your FortiADC system configuration at any time from the System>Settings>Backup & Restore page using the following procedures.

1. Select Back Up.
2. Select a storage location for the backup file, Local PC/Server or ADC.
3. Specify a name.
4. Add a password if you want.
5. The maximum total backup file size differs by model. For more information, see [Maximum total backup file size by hardware model on page 386](#).
6. Click Back Up.

Note: If you've chosen to back up your configuration to the local PC or server, the backup file will appear in the lower-left corner of the GUI. The configuration backup file can be found on the PC or server where all downloaded files are stored. When backing up to a local PC or server, you have the option to use a password to protect the backup file. The option is disabled by default. To use this option, you must enable it first, and then create a password for the configuration backup you are going to do. Be sure to remember the password because it is required when you restore the configuration backup file.

If you've chosen to back up to FortiADC device, the backup file will show up in the table on the Backup & Restore page, where you can either download or upload the backup file using the Download or Upload icon to the far-right column of the same row.

Restore a backup configuration

Use the following procedures to restore a backup of a previous configuration.

1. Select Restore.
2. Select the storage location where the backup file resides.
3. To restore from the Local PC/Server, click Choose File, **then upload the desired file**.
4. To restore from FortiADC, select the backup from the table, and click the corresponding **Restore** icon, on the far right.

Note: The time required to restore a backup file varies, depending on the size of the file and the speed of your network connection. Your web UI session is terminated when the system restarts. To continue using the web UI, refresh the web page and log in again.

If the restored system has a different management interface configuration than the previous configuration, you must access the web UI using the new management interface IP address.

Schedule auto backups

FortiADC's auto backup feature allows you to conveniently set up configuration backup schedules so that it can perform the backups for you automatically according to the schedule. Backup files can be saved on your FortiADC or a local device via SFTP. It must be noted that you can only store up to 10 backup files on FortiADC at any given time and that the size of all backup files combined must not exceed the limit allowed on your hardware model, as stipulated in the table below.

The Auto Backup configuration page also comes with an Overwrite Config check box, which (if enabled) will let the system automatically delete backup files when the number or the size of saved backup files exceeds either limit. Removal of backup files is done in a FIFO (first-in, first-out) fashion, starting with the oldest backup. If Overwrite Config is not enabled, the system will generate error log messages when the backup files exceed the limits.

Schedule auto backups onto FortiADC:

1. Select Auto Backup.
2. Select ADC as the storage location where the backup files will be saved.
3. Enable the scheduled backup radio button.
4. Specify the scheduled backup frequency, and set the schedule accordingly.
5. Select the Overwrite Config radio button (recommended).
6. Click Save.

Schedule auto backups onto an SFTP server:

To schedule auto backups onto an SFTP server, you must have a user account on the server and provide the information required about the server, such as its IP address, port number, backup location, and your account user name and password.

1. Select Auto Backup.
2. Select Local PC/Server (SFTP server) as the storage location where the backup files will be saved.
3. Select the Scheduled Backup radio button.
4. Specify the scheduled backup frequency, and set the schedule accordingly.
5. Enter the IP address of the SFTP server.
6. Enter the port of the SFTP server.
7. Specify the backup file path on the SFTP server, in Folder.
8. Enter your username for the SFTP server.
9. Enter your password for the SFTP server.
10. Click Save.

Schedule auto backups from the Console

Use the following commands to set up auto backup from the Console:

```
config system auto-backup
set storage {sftp| disk}
set address <ip>
set port <port>
set username <name>
set password <password>
set folder <local directory>
set overwrite {enable|disable}
set schedule-backup-day {Sunday | Monday | Tuesday | Wednesday | Thursday | Friday |
Saturday}
set schedule-update-frequency {daily|weekly|every}
set schedule-update-time <hh:mm>
set backup-status {enable|disable}
```

end

Maximum total backup file size by hardware model

Hardware model	Maximum total backup file size
FortiADC 60F	50 MB
FortiADC 100F	50 MB
FortiADC200D	50 MB
FortiADC 200F	50 MB
FortiADC 300D	100 MB
FortiADC 400D	100 MB
FortiADC 700D	100 MB
FortiADC 1000F	100 MB
FortiADC 1500D	100 MB
FortiADC 2000D	100 MB
FortiADC 2000F	200 MB
FortiADC 4000D	200 MB
FortiADC 4000F	200 MB
All FortiADC VMs	100 MB

SCP support for configuration backup

This feature provides a secure method to transfer FortiADC configuration files from FortiADC to your host, using the Secure Shell (SSH) protocol.

To send your configuration backup file to your host, execute the following command:

```
execute backup {config|config-file} scp <server user name> <server password> <directory>  
    <filename> <server ip>[:port]
```

Example:

```
execute backup config-file scp fortinet fadc /etc/home/ backup_config 192.1.2.3
```

Rebooting, resetting, and shutting down the system

The following items have the indicated usage:

- **Reboot**—Reboots the operating system.
- **Reset**—Resets the configuration to the default factory values.
- **Shut Down**—Shuts down the system. When the system is shut down, it is unavailable to forward traffic.



Do not unplug or switch off the FortiADC appliance without first shutting down the operating system. The shutdown process enables the system to finish writing any buffered data, and to correctly spin down and park the hard disks. Failure to do so could cause data loss and hardware problems.

To reboot the system:

Do one of the following:

- Go to the dashboard, and in the System Information widget, click **Reboot**.
- From the CLI console, enter the following command:

```
execute reboot
```

To perform a factory reset:

Do one of the following:

- Go to the dashboard, and in the System Information widget, click **Reset**.
- From the CLI console, enter the following command:

```
execute factoryreset
```

To power off the system:

To shut down the system:

- Go to the dashboard, and in the System Information widget, click **Shut Down**.
- From the CLI console, enter the following command:

```
execute shutdown
```

The system does not emit disk activity noise when shutdown is complete.

To completely power off:


- For hardware appliances, press the power button if there is one. Power supplies and switches vary by hardware model. On some, you press the power button; on others, you flip the switch to either the off (O) or on (I) position.
- For FortiADC-VM, power off the virtual machine.

Create a traffic group

A *traffic group* is a set of VRIDs. Each VRID keeps talking with its peers using 'hello' packets via its heartbeat interface so that each VRID can be aware of its peers (master or slave) operating state and perform a VRRP fail-over in case the master node fails. The different VRIDs have no relationship with each other.

In [Traffic group on page 388](#), both VRID1 and VRID2 use Device1 as the master. When Port2 on Device1 fails, all traffic between the client and the server can't pass through the device

Traffic group



To solve this problem, you can create a traffic group and add both VRID1 and VRID2 as its members, and set the rule that the whole traffic group to fail over to the success device when either VRID fails. In this case, if Device1's Port2 fails, the whole traffic group will fail over to Device2.

Using the VRID concept, FortiADC allows you to add objects with floating IP address, such as interface, virtual server, IP pool, and SNA T pool, etc. to a traffic-group. With this configuration, it will trigger the whole traffic group to switch over when a resource fails.

Normally, the number of traffic groups should be same as the number of devices in an HA group for HA active-active configurations. FortiADC comes with a predefined traffic group named 'default'. You can use this default traffic group if you only need an HA active-passive deployment. Otherwise, you must configure your own traffic groups before making HA active-active configurations, using the instructions discussed in the following paragraphs.

Create a traffic group via the command line interface

Use the following commands to create a new traffic group:

```
config system traffic-group
edit traffic-group-1
set preempt enable
set network-failover enable
set failover-order 1 3 5
next
end
```

Note: The failover sequence must be configured according to the order of node IDs. This means that if a node is dead, the next node in queue will take over handling the traffic. If you want to decide when a node should retake the traffic over from power-down to start-up, you **MUST** enable the Preempt option.

Create a traffic group from the Web GUI

Use the following steps to configure a traffic group from FortiADC's web interface:

1. Click System > Traffic Group.
2. Click Create New to open the Traffic Group dialog.

3. Make the desired entries or selections as described in [Traffic-group parameters on page 389](#).
4. Click Save when done.

Traffic-group parameters

Parameter	Description
Traffic Group Name	Specify a unique name for the traffic group.
Preempt	Disabled by default. If enabled, the node will retake control of traffic from power-down to start-up.
Remote IP Monitor	Disabled by default. When enabled, the system will actively monitor the remote beacon IP addresses to determine the available network path.
Failover Order	Follow the hint onscreen to set the failover sequence among the ports.

Manage administrator users

This topic includes the following information:

- [Administrator user overview](#)
- [Create administrator users](#)
- [Configure access profiles](#)
- [Enable password policies](#)

Administrator user overview

In its factory default configuration, FortiADC has one administrator account named **admin**. The user of this account has permissions that grant read-write access to all system functions.

Unlike other administrator accounts, this default **admin** cannot be deleted. The **admin** account is similar to a root administrator account. This account always has full permission to view and change all system configuration options, including viewing and changing *all* other administrator accounts. You cannot alter the name and permissions of this default admin account.

To prevent accidental changes to the configuration, it is best that only network administrators, and if possible, only a single person, use the **admin** account.

You can use the **admin** account to configure more administrator accounts for other users. Accounts can be created with different levels of access. If you require such role-based access control (RBAC) restrictions, or if you simply want to harden security or prevent inadvertent changes to other administrators' areas, you can do so using access profiles. For example, you can create an account for a security auditor who must only be able to view the configuration and logs, but *not* change them.

Basic steps

1. Create administrator user accounts with permissions provisioned by the profiles.
2. Configure access profiles to provision permissions to roles.
3. Enable password policies.

Create administrator users

We recommend that only network administrators—and if possible, only a single person—use the **admin** account. You can configure accounts that provision different scopes of access. For example, you can create an account for a security auditor who must only be able to view the configuration and logs, but *not* change them.

Before you begin:

- If you want to use RADIUS or LDAP authentication, you must have already have created the RADIUS server or LDAP server configuration.
- You must have Read-Write permission for System settings.

To create an administrator user account:

1. Go to System > Administrator.
2. Make sure the Admin tab is selected.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Administrator user configuration on page 390](#).
5. Click Save.

Administrator user configuration

Settings	Guidelines
Name	Name of the administrator account, such as <code>admin1</code> or <code>admin@example.com</code>. Do not use spaces or special characters except the 'at' symbol (<code>@</code>). The maximum length is 35 characters. If you use LDAP or RADIUS, specify the LDAP or RADIUS username. This is the user name that the administrator must provide when logging in to the CLI or web UI. The users are authenticated against the associated LDAP or RADIUS server. After you initially save the configuration, you cannot edit the name.
Trusted Hosts	Source IP address and netmask from which the administrator is allowed to log in. For multiple addresses, separate each entry with a space. You can specify up to three trusted areas. They can be single hosts, subnets, or a mixture. Configuring trusted hosts hardens the security of the system. In addition to knowing the password, an administrator must connect only from the computer or subnets you specify. Trusted host definitions apply both to the web UI and to the CLI when accessed through Telnet, SSH, or the CLI console widget. Local console access is <i>not</i> affected by trusted hosts, as the local console is by definition not remote, and does not occur through the network. If ping is enabled, the address you specify here is also a source IP address to which the system will respond when it receives a ping or traceroute signal. To allow logins only from <i>one</i> computer, enter only its IP address and 32- or 128-bit netmask: <pre>192.0.2.1/32 2001:0db8:85a3::8a2e:0370:7334/128</pre> To allow login attempts from any IP address (not recommended), enter: <pre>0.0.0.0/0</pre>

Settings	Guidelines
	Caution: If you restrict trusted hosts, do so for <i>all</i> administrator accounts. Failure to do so means that all accounts are still exposed to the risk of brute force login attacks. This is because if you leave even <i>one</i> administrator account unrestricted (i.e. 0.0.0.0/0), the system must allow login attempts on all network interfaces where remote administrative protocols are enabled, and wait until <i>after</i> a login attempt has been received in order to check that user name's trusted hosts list. Tip: If you allow login from the Internet, set a longer and more complex New Password, and enable only secure administrative access protocols. We also recommend that you restrict trusted hosts to IPs in your administrator's geographical area. Tip: For improved security, restrict all trusted host addresses to single IP addresses of computer(s) from which <i>only</i> this administrator will log in.
Global Admin	- No —Default. If selected, the account can access the virtual domain specified in this configuration only. - Yes—If selected, the account can access all virtual domains.
Profile	Select a user-defined or predefined profile. The predefined profile named super_admin_prof is a special access profile used by the admin account. However, selecting this access profile will <i>not</i> confer all permissions of the admin account. For example, the new administrator would not be able to reset lost administrator passwords. Note: This option does not appear for the admin administrator account, which by definition always uses the super_admin_prof access profile.
Virtual Domain	Optional. If you have enabled the virtual domain feature, select the virtual domain that this administrator can view and manage.
Authentication Type	- Local—Use the local administrator authentication server. - RADIUS—Use a RADIUS authentication server. Select the RADIUS server configuration. - LDAP—Use an LDAP authentication server. Select the LDAP server configuration. Note: This option does not apply to a global admin account.
Password	Set a strong password for all administrator accounts. The password should be at least eight characters long, be sufficiently complex, and be changed regularly. To check the strength of your password, you can use a utility such as Microsoft's password strength meter.
Confirm Password	Re-enter the same password.

Configure access profiles

Access profiles provision permissions to roles. The following permissions can be assigned:

- Read (view access)
- Read-Write (view, change, and execute access)
- No access

When an administrator has only read access to a feature, the administrator can access the web UI page for that feature, and can use the `get` and `show` CLI command for that feature, but cannot make changes to the configuration.

In larger companies where multiple administrators divide the share of work, access profiles often reflect the specific job that each administrator does (“role”), such as account creation or log auditing. Access profiles can limit each administrator account to their assigned role. This is sometimes called role-based access control (RBAC).

[Areas of control in access profiles on page 392](#) lists the administrative areas that can be provisioned. If you provision read access, the role can view the web UI menu (or issue a CLI get command). If you provision read-write access, the role can save configuration changes (or issue a CLI set command).

For complete access to *all* commands and abilities, you must log in with the administrator account named **admin**.

Areas of control in access profiles

Web UI Menus	CLI Commands
System	config system diagnose hardware diagnose sniffer diagnose system execute date execute ping execute ping-options execute traceroute
Router	config router
Server Load Balance	config load-balance
Link Load Balance	config link-load-balance
Global Load Balance	config global-dns-server config global-load-balance
Security	config firewall config security waf
Log & Report	config log config report execute rebuild-db
* For each <code>config</code> command, there is an equivalent <code>get/show</code> command. The <code>config</code> commands require write permission. The <code>get/show</code> commands require read permission.	

Before you begin:

- You must have Read-Write permission for System settings.

To configure administrator profiles:

- Click System > Administrator.
- Click the **Access Profile** tab.
- Click **Create New** to display the configuration editor.
- Complete the configuration as described in [Configure access profiles on page 391](#).
- Click Save.

Access profile configuration

Settings	Guidelines
Name	Specify a name for the access profile configuration. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces.
System	Select one of the following: • None—Do not provision access for the menu. • Read Only—Provision ready-only access. • Read-Write—Enable the role to make changes to the configuration.
Networking	Select one of the following: • None—Do not provision access for the menu. • Read Only—Provision ready-only access. • Read-Write—Enable the role to make changes to the configuration.
User	Select one of the following: • None—Do not provision access for the menu. • Read Only—Provision ready-only access. • Read-Write—Enable the role to make changes to the configuration.
Server Load Balance	Select one of the following: • None—Do not provision access for the menu. • Read Only—Provision ready-only access. • Read-Write—Enable the role to make changes to the configuration.
Link Load Balance	Select one of the following: • None—Do not provision access for the menu. • Read Only—Provision ready-only access. • Read-Write—Enable the role to make changes to the configuration.
Global Load Balance	Select one of the following: • None—Do not provision access for the menu. • Read Only—Provision ready-only access. • Read-Write—Enable the role to make changes to the configuration.
Security	Select one of the following: • None—Do not provision access for the menu. • Read Only—Provision ready-only access. • Read-Write—Enable the role to make changes to the configuration.
Log & Report	Select one of the following: • None—Do not provision access for the menu. • Read Only—Provision ready-only access. • Read-Write—Enable the role to make changes to the configuration.
Shared Resource	For each category, set the permission: • None—Do not provision access for the menu. • Read Only—Provision ready-only access. • Read-Write—Enable the role to make changes to the configuration.

The **super_admin_prof** access profile, a special access profile assigned to the **admin** account and required by it, appears in the list of access profiles. It exists by default and cannot be changed or deleted. The profile has permissions similar to the UNIX root account.

Enable password policies

A password policy is a set of rules designed to enhance computer security. A good password policy encourages users to create strong passwords and use them properly. For your network and data security and integrity, we strongly recommend the enforcement of strong password policies when using FortiADC.

To enable password policy:

1. Go to System > Administrator.
2. Select the Password Policy tab.
3. Complete the configuration as described in [Password policy configuration](#).
4. Click Save.

Password policy configuration

Settings	Guidelines
Password Policy	Enabled by default.
Minimum Length	Specify the minimum length requirement of passwords, which can be from 8 (default) to 32 characters in length.
Must Contain	Select the restrictions you want to impose on passwords: • Upper Case Letter—If selected, passwords must contain upper-case letters. • Lower Case Letter—If selected, passwords must contain lower-case letters. • Number—If selected, passwords must contain numbers. • Non-alphanumeric —If selected, passwords must contain non-alphanumeric characters.

Configuring SNMP


Many organizations use *SNMP* (simple network management protocol) to track the health of their systems. FortiADC supports SNMP v1, v2c, and v3.

SNMP depends on network devices that maintain standard management information bases (MIBs). *MIBs* describe the structure of the management data maintained on the device. Some MIB definitions are standard for all network devices, and some are vendor and product-family specific.

The FortiADC system runs an *SNMP agent* to communicate with the *SNMP manager*. The agent enables the system to respond to *SNMP queries* for system information to the SNMP manager.

[SNMP communication on page 394](#) illustrates the basic communication.

SNMP communication



With SNMP v1 and v2c managers, you configure *SNMP communities* to connect FortiADC and the SNMP manager. The SNMP Manager sends the community string along with all SNMP requests. If the community string is correct, the device responds with the requested information. If the community string is incorrect, the device simply discards the request and does not respond.



Fortinet strongly recommends that you do not add FortiADC to the community named `public`. This default name is well-known, and attackers that attempt to gain access to your network often try this name first.

With SNMPv3 managers, you configure *SNMP users* to connect FortiADC and the SNMP manager. Queries and traps include username/password authentication, along with an encryption key. FortiADC implements the user security model described in [RFC 3414](#).

Before you begin:

- On the SNMP manager, you must verify that the SNMP manager is a member of the community to which the FortiADC system belongs, and you must compile the necessary Fortinet-proprietary management information blocks (MIBs) and Fortinet-supported standard MIBs. For information on Fortinet MIBs, see [Appendix A: Fortinet MIBs](#).
- In the FortiADC interface settings, you must enable SNMP access on the network interface through which the SNMP manager connects.
- You must have Read-Write permission for System settings.

To configure SNMP system information:

1. Go to System > SNMP.
2. Click the System Information tab.
3. Complete the configuration as described in [SNMP settings on page 395](#).
4. Save the configuration.

SNMP settings

Settings	Guidelines
SNMP Agent	Disabled by default. Enable to activate the SNMP agent so that the system can receive SNMP queries.
Description	A description or comment about the system, such as <code>dont-reboot</code> . The description can be up to 35 characters long, and can contain only letters (a-z, A-Z), numbers, hyphens (-) and underscores (_).
Contact	Contact information for the administrator or other person responsible for this system, such as a phone number (555-5555) or name (j.doe). The contact information can be up to 35 characters long, and can contain only letters (a-z, A-Z), numbers, hyphens (-) and underscores

Settings	Guidelines
	(_).
Location	Physical location of the appliance, such as <code>floor2</code> . The location can be up to 35 characters long, and can contain only letters (a-z, A-Z), numbers, hyphens (-) and underscores (_).

Downloading SNMP MIB files

You can download the FortiADC SNMP MIB file or the Fortinet core MIB file using the links at the bottom of the page.

For more information, refer to [Appendix A: Fortinet MIBs on page 597](#).

Download SNMP MIBs

FortiADC allows you to download full FortiADC and Fortinet Core MIB files, which provides more options for server load balance, global server load balance, link load balance, and firewall with SNMP traps.

To download an SNMP MIB file:

1. Click System > SNMP.
2. Click the **System Information tab**.
3. In the FortiADC SNMP MIB section, click Download FortiADC MIB File or Download Fortinet Core MIB File.
4. Follow the instructions onscreen to complete the download.

Configure SNMP threshold

To configure SNMP threshold:

1. Go to System > SNMP.
2. Click the Threshold tab.
3. Complete the configuration as described in [SNMP threshold on page 396](#).
4. Save the configuration.

SNMP threshold

Settings	Guidelines
CPU	• Trigger—The default is 80% utilization. • Threshold—The default is 3, meaning the event is reported when the condition has been triggered 3 times in a short period. • Sample Period—The default is 600 seconds. • Sample Frequency—The default is 30 seconds.
Memory	• Trigger—The default is 80% utilization. • Threshold—The default is 3, meaning the event is reported when the condition has been triggered 3 times in a short period. • Sample Period—The default is 600 seconds.

Settings	Guidelines
	Sample Frequency—The default is 30 seconds.
Disk	- Trigger—The default is 90% utilization. - Threshold—The default is 1, meaning the event is reported each time the condition is triggered. - Sample Period—The default is 7200 seconds. - Sample Frequency—The default is 3600 seconds.

Configure SNMP v1/v2

To configure SNMP v1/v2:

1. Go to System > SNMP.
2. Click the SNMPv1/v2 tab.
3. Complete the configuration as described in [SNMP settings on page 397](#).
4. Save the configuration.

SNMP settings

Settings	Guidelines
SNMPv1/v2	
Name	Name of the SNMP community to which the FortiADC system and at least one SNMP manager belongs, such as <code>management</code>. You must configure the FortiADC system to belong to at least one SNMP community so that community's SNMP managers can query system information. You can add up to three SNMP communities. Each community can have a different configuration for queries and traps. You can also add the IP addresses of up to eight SNMP managers to each community to which IP addresses are permitted to query the FortiADC system.
SNMP v1 Status	Select to enable the SNMP v1 configuration.
SNMP v1 Port	Enter the port number on which the system listens for SNMP v1 queries from the SNMP managers in this community. The default is 161.
SNMP v2 Status	Select to enable the SNMP v2 configuration.
SNMP v2 Port	Enter the port number on which the system listens for SNMP v2 queries from the SNMP managers in this community. The default is 161.
Host	
IP Address	Enter the subnet address for the SNMP manager to be permitted to query the FortiADC system. SNMP managers have read-only access. You can add up to 8 SNMP managers to each community. To allow any IP address using this SNMP community name to query the FortiADC system, enter <code>0.0.0.0/0</code> . For security best practice reasons, however, this is not recommended.



Test both traps and queries (assuming you have enabled both). Traps and queries typically occur on different port numbers, and therefore verifying one does not necessarily verify that the other is also functional.

To test queries, from your SNMP manager, query the FortiADC appliance.

To test traps, cause one of the events that should trigger a trap.

Configure SNMP v3

To configure SNMP v3:

1. Go to System > SNMP.
2. Click the SNMPv3 tab.
3. Complete the configuration as described in [SNMP v3 on page 398](#).
4. Save the configuration.

SNMP v3

Settings	Guidelines
SNMP v3	
Name	User name that the SNMP Manager uses to communicate with the SNMP Agent. After you initially save the configuration, you cannot edit the name.
Status	Enable/disable the configuration.
Security Level	• No Auth And No Privacy—Do not require authentication or encryption. • Auth But No Privacy—Authentication based on MD5 or SHA algorithms. Select an algorithm and specify a password. • Auth And Privacy—Authentication based on MD5 or SHA algorithms, and encryption based on AES or DES algorithms. Select an Auth Algorithm and specify an Auth Password; and select a Private Algorithm and specify a Private Password.
SNMP v3 Port	Enter the port number on which the system listens for SNMP v3 queries from the SNMP managers. The default is 161.
Host	
IP Address	Enter the subnet address for the SNMP manager to be permitted to query the FortiADC system. SNMP managers have read-only access. You can add up to 8 SNMP managers to each community. To allow any IP address using this SNMP community name to query the FortiADC system, enter 0.0.0.0/0. For security best practice reasons, however, this is not recommended.



Test both traps and queries (assuming you have enabled both). Traps and queries typically occur on different port numbers, and therefore verifying one does not necessarily verify that the other is also functional.

To test queries, from your SNMP manager, query the FortiADC appliance.

To test traps, cause one of the events that should trigger a trap.

Configuring central management


Central Management allows the ADC to be connected to an ADC Manager. Multiple ADC's can be managed by the ADC Manager. If you have large networks with multiple ADC's, with the ADC Manager you can simplify the configuration of these ADC's (for example, setting multiple ADC's to the same configuration), and view all of their logs and statistics together.

The ADC Manager is a powerful tool that gives you more effective control over your ADC's.

This guide will show you how to enable central management on your particular ADC. You will enter the IP address of your manager, then enable Central Management, therefore allowing the ADC Manager to manage your ADC.

See the [ADC Manager Administration Guide](#).

Basic configuration of two ADC's linked to a Manager



Configuring central management

1. Global > System > Central Management
2. Click the edit function at the far right. At first, the IP address should be empty, and the Type should be none.
3. Configure according to the table below.

Central Management settings

Settings	Guidelines
Type	The type of the Central Management None—Initial State of CM Agent. FortiADC Manager—The ADC is connected to the Manager.
Address	The IP address or hostname of FortiADC-Manager.
Interval	How often the ADC tries to connect to the Manager. Default 10 seconds. Range 10- 120.
Register	Enable/disable register to ADC Manager. That is to say, cut off or attach connection to the ADC Manager. Default is disable
Management Status	The connection status of the ADC. • Online—ADC Manager successfully connects to CM Server. • Offline—ADC Manager failed to connect CM Server. It can happen at the first connection trial or if ADC Manager lost the connection. Note: ADC Manager updates info to CM Server every minute and will make state as Offline when it does not get response 2 times. • Reject—Occurs when ADC Manager tries to connect with 'State is not None' and CM Server does not have the record of this ADC(identify by license). The connection will be rejected by CM Server.

Note: **When register is enabled**, modifying other central management settings is forbidden. Other central management settings are grayed out. Furthermore, a warning message is shown when on login.

This FortiADC is currently managed by a FortiCM device


 All changes should be performed from a FortiCM to avoid conflict. How would you like to proceed?

Logout


Login Read-Only

Login Read-Write

When ADC is set to 'Read Only mode', it means all configurations can only be viewed, although the admin access profile has Read-Write permission.



When trying to write configurations in Read Only, the error message is shown:



The CM Agent state change log can be found in **System Logs**.

Log & Report									
Log Browsing									
Event Log									
Security Log									
Traffic Log									
Script Log									
Aggregate Log									
☐ Configuration ☒ System ☐ Admin ☐ User ☐ Health Check ☐ SLB ☐ LLB ☐ GLB ☐ Firewall									
Filter Setting Download Refresh									
Date	Time	Log Level	Submod	User	Action	Status	Message	Description	
2018-10-21	20:15:50	information	none	none	none	success	CM agent state changes to ONLINE	CM agent state changed	
Date		2018-10-21		Time		20:15:50			
Log ID		0003011000		Log Level		information			
Message ID		10806		Submod		none			
User		none		UI		none			
Action		none		Status		success			
Description		CM agent state changed		Message		CM agent state changes to ONLINE			
Type		event		Sub Type		system			
Vdom		root							
2018-10-21	20:11:05	information	none	none	none	success	CM agent state changes to OFFLINE	CM agent state changed	

Manage and validate certificates

This section includes the following topics:

- [Overview](#)
- [Prerequisite tasks](#)
- [Manage certificates](#)
- [Validate certificates](#)

Overview

The FortiADC system is able to process the following two types of TLS/SSL traffic:

- **System administration**—Administrators connect to the web UI (HTTPS connections only). When you connect to the web UI, the system presents its own default “Factory” certificate. This certificate is used only for connections to the web UI. It cannot be removed. Do not use this certificate for server load balancing traffic.
- **Server load balancing**—Clients use SSL or TLS to connect to a virtual server. When you use FortiADC as a proxy for SSL operations normally performed on the backend real servers, you must import the X.509 v3 server certificates and private keys that the backend servers would ordinarily use, as well as any certificate authority (CA) or intermediate CA certificates that are used to complete the chain of trust between your clients and your servers.

The FortiADC system supports all of the TLS/SSL administration methods commonly used by HTTPS servers, including:

- **Server name indication (SNI)**—You can require clients to use the TLS extension to include the server hostname in the TLS client hello message. Then, the FortiADC system can select the appropriate local server certificate to present to the client.
- **Local certificate store**—A certificate store for the X.509 v3 server certificates and private keys that the backend servers would ordinarily use.
- **Intermediate CAs store**—A store for Intermediate CAs that the backend servers would ordinarily use to complete the chain of server certificates. HTTPS transactions use intermediate CAs when the server certificate is signed by an intermediate certificate authority (CA) rather than a root CA.
- **Certificate Authorities (CAs) store**—A store for CA certificates that the back-end servers would ordinarily use to verify the CA signature in client certificates or the signature of an OCSP Responder.
- **OCSP**—Use Online Certificate Status Protocol (OCSP) to obtain the revocation status of certificates.
- **CRL**—Use a Certificate Revocation List (CRL) to obtain the revocation status of certificates.
- **Certificate validation policy**—You can configure certificate validation policies that use OCSP or CRL. These policies can be associated with load balancing profiles.
- **All digital certificates of RSA and ECDSA key types**—whether they are local, remote, intermediate, or CA root certificates.
- **Multiple CA, CRL, and OCSP configurations.**
- **Client certificate forwarding.**
- **SNI forwarding.**
- **Email alert on certificate expiration, CRL expiration, and OCSP stapling expiration.**

Note: The factory certificate is the default certificate for any application over SSL/TSL. It is a unique certificate that presents the credentials of your FortiADC. Upon system start, FortiADC automatically generates a self-signed factory certificate with its identifier (i.e., common name) which is your FortiADC's serial number. For example, if a trial license is

in use, then the common name (CN) for the factory.cer would be FADV0000000TRIAL; if the license is imported, the factory.cer would be FADV080000072226.

Certificates and their domains

You can generate or import certificates in the global domain (i.e., FortiADC appliance) and individual VDOM domains (i.e., virtual machines). The visibility and use of certificates or certificate groups may vary, depending where (the domain) they are created. Below are the general guidelines regarding the availability and use of certificates or certificate groups.

- **Local Certificates/intermediate Certificates**—If generated or imported in a specific VDOM domain, they can be viewed and deleted in that VDOM only, but not visible in any other VDOM or the global domain; if generated or imported in the global domain, they can be viewed and downloaded by all VDOMS, but can be deleted only in the global domain.
- **Local Certificate Groups/Intermediate CA Groups**—If added in a specific VDOM domain, they can be viewed, edited, or referenced in that VDOM domain only, but not visible in any other VDOMs or the global domain; if added in the global domain, they can be visible to all VDOM domains, but can be edited only in the global domain.
- **CA/CRL/OCSP Signing Certificates**—If imported in a specific VDOM domain, they can be viewed or deleted only in that VDOM, but not visible in any other VDOM domain or the global domain; if imported in the global domain, they can be viewed or downloaded in all VDOM domains, but can be deleted only in the global domain.
- **Verify/CA Group/OCSP**—If added in a specific VDOM domain, they can be viewed or edited or referenced to in that VDOM domain only, but not visible in any other VDOM domain or the global domain; if added in the global domain, they can be viewed or referenced to in all VDOMs, but can be edited only in the global domain.

Prerequisite tasks

You must download the certificates from your backend servers so that you can import them into the FortiADC system. This example shows how to download a CA certificate from Microsoft Windows 2003.

To download a CA certificate from Microsoft Windows 2003 Server:

1. Go to `https://<ca-server_ipv4>/certsrv/`.
where `<ca-server_ipv4>` is the IP address of your CA server.
2. Log in as Administrator. Other accounts may not have sufficient privileges.
The Microsoft Certificate Services home page appears. [Welcome page on page 403](#) is an example of this page.
Welcome page

Welcome

Use this Web site to request a certificate for your Web browser, e-mail client, or other program. By using a certificate, you can verify your identity to people you communicate with over the Web, sign and encrypt messages, and, depending upon the type of certificate you request, perform other security tasks.

You can also use this Web site to download a certificate authority (CA) certificate, certificate chain, or certificate revocation list (CRL), or to view the status of a pending request.

For more information about Certificate Services, see [Certificate Services Documentation](#).

Select a task:

[Request a certificate](#)

[View the status of a pending certificate request](#)

[Download a CA certificate, certificate chain, or CRL](#)

3. Click the **Download CA certificate, certificate chain, or CRL** link to display the Download a CA Certificate, Certificate Chain, or CRL page. [Download a CA Certificate, Certificate Chain, or CRL page on page 404](#) is an example of this page.
4. From Encoding Method, select **Base64**.
5. Click **Download CA certificate**.

Download a CA Certificate, Certificate Chain, or CRL page

Download a CA Certificate, Certificate Chain, or CRL

To trust certificates issued from this certification authority, [install this CA certificate chain](#).

To download a CA certificate, certificate chain, or CRL, select the certificate and encoding method.

CA certificates:

Current [myca]

Encoding method:

☐ DER
☒ Base 64

[Download CA certificate](#)
[Download CA certificate chain](#)
[Download latest base CRL](#)
[Download latest delta CRL](#)

Manage certificates

This section discusses the following tasks you can perform on the System > Certificate > Manage Certificates page:

- [Generating a certificate signing request](#)
- [Importing local certificates](#)
- [Importing intermediate CAs](#)

- [Creating an intermediate CA group](#)
- [Creating a local certificate group](#)
- [OCSP stapling on page 412](#)

Generating a certificate signing request

Many commercial certificate authorities (CAs) provide websites where you can generate your own certificate signing request (CSR). A CSR is an unsigned certificate file that the CA will sign. When a CSR is generated, the associated private key that the appliance will use to sign and/or encrypt connections with clients is also generated.

If your CA does not provide this service, or if you have your own private CA such as a Linux server with OpenSSL, you can use FortiADC to generate a CSR and private key. This CSR can then be submitted for verification and signing by the CA.

Before you begin:

- You must have Read-Write permission for System settings.

To generate a certificate signing request:

1. Go to System > Certificate > Manage Certificates.
2. Click the **Local Certificate** tab.
3. Click **Generate** to display the configuration editor.
4. Complete the configuration as described in [CSR configuration on page 405](#).
5. Click Save when done.

The system creates a private and public key pair. The generated request includes the public key of the FortiADC appliance and information such as the IP address, domain name, or email address. The FortiADC appliance private key remains confidential on the FortiADC appliance. The Status column of the new CSR entry is **Pending**.

6. Select the row that corresponds to the certificate request.
7. Click **Download**.
Standard dialogs appear with buttons to save the file at a location you select. Your web browser downloads the certificate request (.csr) file.
8. Upload the certificate request to your CA.
After you submit the request to a CA, the CA will verify the information in the certificate, give it a serial number, an expiration date, and sign it with the public key of the CA.
9. If you are not using a commercial CA whose root certificate is already installed by default on web browsers, download your CA's root certificate, and then install it on all computers that will be connecting to your FortiADC appliance. Otherwise, those computers might not trust your new certificate.
10. After you've received the signed certificate from the CA, import the certificate into the FortiADC system.

CSR configuration

Settings	Guidelines
Generate Certificate Signing Request	
Certification Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. The maximum length is 35 characters. Note: This is the name of the CSR file, not the host name/IP contained in the certificate's <code>Subject</code> line.

Settings	Guidelines
Subject Information	
ID Type	Select the type of identifier to use in the certificate to identify the virtual server: - Host IP—The <i>static</i> public IP address of the FortiADC virtual server in the IP Address field. If the FortiADC appliance does not have a static public IP address, use the email or domain name options instead. Note: Do NOT use this option if your network has a dynamic public IP address. Your web browser will display the “Unable to verify certificate” or similar error message when your public IP address changes. - Domain Name—The fully qualified domain name (FQDN) of the FortiADC virtual server, such as <code>www.example.com</code>. This does not require that the IP address be static, and may be useful if, for example, your network has a dynamic public IP address and therefore clients connect to it via dynamic DNS. Do not include the protocol specification (<code>http://</code>) or any port number or path names. - E-Mail—The email address of the owner of the FortiADC virtual server. Use this if the virtual server does not require either a static IP address or a domain name. Depending on your choice for ID Type, related options appear.
IP Address	Enter the static IP address of the FortiADC appliance, such as <code>10.0.0.1</code>. The IP address should be the one that is visible to clients. Usually, this should be its public IP address on the Internet, or a virtual IP that you use NAT to map to the appliance’s IP address on your private network. This option appears only if ID Type is Host IP.
Domain Name	Enter the FQDN of the FortiADC appliance, such as <code>www.example.com</code>. The domain name <i>must</i> resolve to the IP address of the FortiADC appliance or backend server according to the DNS server used by clients. (If it does not, the clients’ browsers will display a <code>Host name mismatch</code> or similar error message.) This option appears only if ID Type is Domain Name.
Email	Enter the email address of the owner of the FortiADC appliance, such as <code>admin@example.com</code>. This option appears only if ID Type is E-Mail.
Distinguished Information	
Organization Unit	Name of organizational unit (OU), such as the name of your department. This is optional. To enter more than one OU name, click the + icon, and enter each OU separately in each field.
Organization	Legal name of your organization.
Locality (City)	City or town where the FortiADC appliance is located.
State/Province	State or province where the FortiADC appliance is located.
Country/Region	Country where the FortiADC appliance is located.
Email	E-mail address that may be used for contact purposes, such as <code>admin@example.com</code> .
Key Information	
Key Type	Select either of the following: - RSA - ECDSA

Settings	Guidelines
Key Size/ Curve Name	For RSA key, select one of the following key sizes: • 512 Bit • 1024 Bit • 1536 Bit • 2048 Bit • 4096 Bit. Note: Larger keys use more computing resources, but provide better security. For ECDSA, select one of the following curve names: • prime256v1 • secp384r1 • secp521r1
Enrollment Information	
Enrollment Method	• File-Based—You must manually download and submit the resulting certificate request file to a CA for signing. Once signed, upload the local certificate. • Online SCEP—The FortiADC appliance automatically uses HTTP to submit the request to the simple certificate enrollment protocol (SCEP) server of a CA, which will validate and sign the certificate. For this selection, two options appear. Enter the CA Server URL and the Challenge Password.

Importing local certificates

You can import (upload) the following types of X.509 server certificates and private keys into the FortiADC system:

- Base64-encoded
- PKCS #12 RSA-encrypted

Before you begin:

- You must have Read-Write permission for System settings.
- You must have downloaded the certificate and key files and be able to browse to them
- so that you can upload them.

To import a local certificate:

1. Go to System > Certificate > Manage Certificates.
2. Click the **Local Certificate** tab.
3. Click **Import** to display the configuration editor.
4. Complete the configuration as described in [Local certificate import configuration on page 407](#).
5. Click Save when done.

Local certificate import configuration

Settings	Guidelines
Type	Click the down arrow and select one of the following options from the drop-down menu: • Local Certificate—Use this option only if you have a CA-signed certificate that was

Settings	Guidelines
	originated from a CSR generated in FortiADC . See Generating a certificate signing request on page 405. Note: It is important to make sure that the load-balancer (FortiADC appliance) you use to import a local certificate is the same appliance where the CSR was generated because it is where the key matching the certificate resides. The import operation will fail without the matching key on the same hardware system. • PKCS12 Certificate—Use this option only if you have a PKCS #12 password-encrypted certificate with its key in the same file. • Certificate—Use this option only if you have a certificate and its key in separate files. Note: Additional fields are displayed depending on your selection.
Local Certificate	
Certificate File	Browse for and upload the certificate file that you want to use.
PKCS12 Certificate	
Certificate Name	Specify the certificate name that can be referenced by other parts of the configuration, such as <code>www_example_com</code> . The maximum length is 35 characters. Do not use spaces or special characters.
Certificate File	Browse for and upload the certificate file that you want to use.
Password	Specify the password to decrypt the file. If the file was encrypted by a password when generated, the same password must be provided when the file is imported to FortiADC. If the file was generated without a password, there is no need to specify a password when importing the file to FortiADC.
Certificate	
Certificate Name	Specify the name that can be referenced by other parts of the configuration, such as <code>www_example_com</code> . The maximum length is 35 characters. Do not use spaces or special characters.
Certificate File	Browse for and upload the certificate file that you want to use.
Key File	Browse for and upload the corresponding key file.
Password	Specify the password to decrypt the file. If the file was encrypted by a password when generated, the same password must be provided when the file is imported to FortiADC. If the file was generated without a password, there is no need to specify a password when importing the file to FortiADC.

Creating a local certificate group

Local certificate groups are used to facilitate the configuration of profiles that are associated with a virtual server.

Before you begin, you must:

- Have Read-Write permission for System settings.
- Have already added the certificates to the local certificate store and intermediate CAs to the intermediate certificate store, and created an intermediate CA group.
- Optionally, create an OCSP Stapling configuration.

To create a local certificate group:

1. Go to System > Certificate > Manage Certificates.
The configuration page displays the Local Certificate Group tab.
2. Click **Create New** to display the configuration editor.
3. Enter the Group Name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. The maximum length is 35 characters. After you initially save the configuration, you cannot edit the name.
4. Click **Save**.
5. To add Group Members to a Local Certificate Group, double-click the group or click the (edit) icon in the row of the group that you want to modify.
6. Click **Create New**.
7. Complete the configuration as described in [Local certificate group configuration on page 409](#).
8. Click Save.

Local certificate group configuration

Settings	Guidelines
Default	Check this check box only if you want to make this local certificate the default for the group. Note: Only one local certificate can be set as the default in a group. If one local certificate has already been set as the default, you must disable (uncheck) it in order to set another one as the default. By default, the first local certificate in the group becomes the default if no other local certificate is set as the default.
Local Certificate	Select a local certificate to add to the group.
OCSP Stapling	Select an OCSP Stapling configuration. The local certificate in the OCSP Stapling configuration must match the local certificate in the local certificate group member. See OCSP stapling on page 412 .
Intermediate CA group	Select an intermediate CA group to add to the local group. (Optional)
Extra Certificate	FortiADC supports dual SSL certificates, one for an RSA-based SSL certificate and the other for an ECDSA-based SSL certificate. This option allows you to add an additional local certificate along with an additional OCSP stapling and intermediate CA group to a local certificate group configuration. Note: This extra local certificate, which is optional, must be of a different format from the local certificate you selected in the first place. In other words, if the local certificate is RSA-based, then this extra local certificate must be ECDSA-based, or vice versa.
Extra Local Certificate	Select an extra local certificate which is different from the local certificate.
Extra OCSP Stapling	Select an extra OCSP stapling configuration. The extra local certificate in the extra OCSP stapling configuration must match the extra local certificate in the extra local certificate group member. (Optional) Note: This option is available only when the Extra Local Certificate has already been set.
Extra Intermediate CA Group	Select an extra intermediate CA group to add to the extra local certificate group. (Optional) Note: This option is available only when the Extra Local Certificate is set.

Note: In general, ECDSA certificates are a good choice for both client and server because they require less time and fewer resources to process. However, for some old web browsers that do not support ECDSA certificates, RSA is the only choice. So, having both an RSA certificate and an ECDSA certificate in the same local certificate group configuration allows FortiADC to take full advantage of the benefits that they offer.

You can also assign two certificates to a local certificate group from the Console, as illustrated in the following example commands:

```
config system certificate local_cert_group
edit "dual"
config group_member
edit 1
set local-cert intermediate02-leafCA-leaf-Serve-RSA
set OCSP-stapling intermediate02-leafCA-leaf-Serve-RSA
set intermediate-ca-group RSA-intermediate02-leaf
set local-cert-extra intermediate02-leafCA-leaf-Serve-ECC
set OCSP-stapling-extra intermediate02-leafCA-leaf-Serve-ECC
set intermediate-ca-group-extra RSA-intermediate02-leaf
next
end
next
end
```

Importing intermediate CAs

An intermediate CA store is for the intermediate CA certificates that back-end servers would normally use to complete the chain of server certificates, if any. HTTPS transactions use intermediate CAs when the server certificate is signed by an intermediate certificate authority (CA) rather than a root CA.

In FortiADC, a root CA can be imported as an "intermediate CA".

Before you begin, you must:

- Have Read-Write permission for System settings.
- Know the URL of an SCEP server or have downloaded the certificate and key files and be able to browse to them so that you can upload them.

To import an intermediate CA:

1. Go to System > Certificate > Manage Certificates.
2. Click the **Intermediate CA** tab.
3. Click **Import** to display the configuration editor.
4. Complete the configuration as described in [Intermediate CA import configuration on page 410](#).
5. Click Save when done.
6. Repeat Steps 3 through 5 to import as many intermediate CAs as needed.

Intermediate CA import configuration

Settings	Guidelines
Certificate Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. The maximum length is 35 characters. After you initially save the configuration, you cannot edit the name.
Import Method	• SCEP—Use Simple Certificate Enrollment Protocol. SCEP allows routers and other

Settings	Guidelines
	intermediary network devices to obtain certificates. File—Upload a file.
SCEP	
SCEP URL	Specify the URL of the SCEP Server.
CA Identifier	Enter the identifier of the CA on the SCEP server, if applicable.
File	
Certificate File	Browse for and upload the the certificate file on the local machine.
Key File	Browse for the corresponding PEM key file that you want to upload. Note: Both a certificate file and key file are required for the intermediate CA used in SSL decryption by the forward proxy.
Password	Password to encrypt the files in local storage.

Creating an intermediate CA group

You select an intermediate CA group configuration object in the local certificate group, so you should configure in the group all the Intermediate CAs that would be needed by the backend servers that belong to a single virtual server.

Before you begin:

- You must have Read-Write permission for System settings.
- You must have already added the Intermediate CAs to the Intermediate CA certificate store.

To create an Intermediate CA group:

- Go to System > Certificate > Manage Certificates.
- Click the **Intermediate CA Group** tab.
- Click **Create New** to display the configuration editor.
- Complete the configuration as described in [Intermediate CA group configuration on page 411](#).
- Save the configuration.

Intermediate CA group configuration

Settings	Guidelines
Group Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. The maximum length is 35 characters. After you initially save the configuration, you cannot edit the name.
Group Member	
Intermediate CA	Select the Intermediate CA to add to the group,
Default	Check this check box only if you want to make this intermediate CA the default for the group. Note: Only one intermediate CA can be set as the default in an intermediate CA group. If one intermediate CA has already been set as the default, you must disable (uncheck) it in order to set another one as the default. By default, the first intermediate CA in a group becomes the default if no intermediate CA is set as the default,

OCSP stapling

OCSP stapling is an improved approach to OCSP for verifying the revocation status of certificates. Rather than having the client contact the OCSP server to validate the certificate status each time it makes a request, FortiADC can be configured to periodically query the OCSP server and cache a time-stamped OCSP response for a set period. The cached response is then included, or "stapled," with the TLS/SSL handshake so that the client can validate the certificate status when it makes a request.

This method of verifying the revocation status of certificates shifts the resource cost in providing OCSP responses from the client to the presenter of a certificate. In addition, because fewer overall queries to the OCSP responder will be made when OCSP stapling is configured, the total resource cost in verifying the revocation status of certificates is also reduced. FortiADC allows you to upload an OCSP response file, configure an OCSP to let FortiADC download the OCSP response from the OCSP server, or both.

Before you begin, you must:

- Have Read-Write permission for System settings.
- Add a local certificate. See [Importing local certificates](#).
- Add a CA certificate. See [Importing intermediate CAs](#).
- Add an OCSP configuration or have an OCSP response file. See [Adding OCSPs](#).

To configure OCSP stapling:

1. Go to System > Certificate > Manage Certificates.
2. Click the **OCSP Stapling** tab.
3. Click **+ Import** to display the configuration editor.
4. Complete the configuration as described in [OCSP stapling configuration on page 412](#).
5. Click Save.

OCSP stapling configuration

Settings	Guidelines
Name	Enter the mkey.
Local Certificate	Select the local certificate to add to the OCSP stapling configuration.
Issuer Certificate	Select the CA certificate to add to the OCSP stapling configuration.
OCSP	Select the OCSP configuration to add to the OCSP stapling configuration. If an OCSP configuration is not selected, import an OCSP Response from a file (see below). You can both select an OCSP configuration and upload an OCSP response file; in this case, FortiADC will first use the OCSP response file and then automatically update using the OCSP configuration.
Response Update Ahead Time	Available only when you select an OCSP configuration. This option is meaningful only when the next update field in the OCSP response is present in a selected OCSP stapling response. Enter the time before the next scheduled update at which FortiADC will start the download for the next update. The default value is 1 hour.
Response Update Interval	Available only when you select an OCSP configuration. Enter the next update interval if the downloaded OCSP response is the same or FortiADC fails to download the new OCSP response. The default value is 5 minutes.

Settings	Guidelines
	If the next update field in the OCSP response is not present, FortiADC will attempt to download the next update periodically according to this parameter.
OCSP Response	Enable to import an OCSP response from a file. PEM and DER formats are supported.



To configure OCSP stapling using the CLI:

```
config system certificate OCSP_stapling
edit <ocsp_stapling_name>
set OCSP
set OCSP-response
set issuer-certificate
set local-certificate
set response-update-ahead-time
set response-update-interval
```

Note: When configuring OCSP stapling in the CLI, only PEM format file types are supported.

Validating certificates

This section discusses the ways to validate client certificates and real server certificates from within the FortiADC system. It covers the following topics:

- [Importing CAs](#)
- [Creating a CA group](#)
- [Importing remote certificates](#)
- [Importing CRLs](#)
- [Adding OCSPs](#)
- [Validating certificates](#)

Configure a certificate verification object

To be valid, a client certificate must meet the following criteria:

- Must not be expired or not yet valid
- Must not be revoked by either certificate revocation list (CRL) or, if enabled, online certificate status protocol (OCSP)
- Must be signed by a certificate authority (CA) whose certificate you have imported into the FortiADC appliance

Certificate verification rules specify the CA certificates to use when validating client certificates, and they specify a CRL and/or OCSP server, if any, to use for certificate revocation checking.

You select a certificate verification configuration object in the profile configuration for a virtual server or in a real-server-SSL profile. If the client presents an invalid certificate during the authentication phase of a SSL/TLS session initiation, the FortiADC system will not allow the connection.

Before you begin:

- You must have Read-Write permission for System settings.
- You must have already created CA, OCSP or CRL configuration.

After you have configured a certificate verification object, you can include it in a virtual server profile or a Real Server SSL Profile, and it will be used to validate certificates presented to FortiADC.



Note: For the same certificate object you can configure multiple CRL files.

To configure a certificate verification object:

1. Go to System > Certificate > Verify.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in [Certificate verify configuration on page 414](#).
4. Click Save when done. The newly certificate verification object appears on the Verify page.
5. Click the Edit icon in the far-right column (or double-click the entry) to open the configuration editor.
6. In the Group Member panel, select the CA, OCSP, or CRL of interest.
7. Click Save when done.

Certificate verify configuration

Settings	Guidelines
Name	Enter a unique name for the certificate verification object that you are creating. Valid characters are A-Z, a-z, 0-9, _, and -. The maximum length is 35 characters. No space is allowed.
verify-depth	Note: CLI only. The default value is 1, but you may select any value from 0 to 255.
customize-error-ignore	Note: This option is available from the CLI only. Enable or disable <code>customize-error-ignore</code> . The option is disabled by default. If it's enabled, you are required to select the <code>ca-ignore-errors</code> and <code>cert-ignore-errors</code> , as described below.
ca-ignore-errors	Note: CLI only. When <code>customize-error-ignore</code> is enabled, the following options become available for you to choose from: • <code>UNABLE_TO_GET_ISSUER_CERT</code> • <code>UNABLE_TO_GET_CRL</code> • <code>CERT_NOT_YET_VALID</code> • <code>CERT_HAS_EXPIRED</code> • <code>CRL_NOT_YET_VALID</code> • <code>CRL_HAS_EXPIRED</code> • <code>DEPTH_ZERO_SELF_SIGNED_CERT</code>

Settings	Guidelines
	- SELF_SIGNED_CERT_IN_CHAIN - UNABLE_TO_GET_ISSUER_CERT_LOCALLY - UNABLE_TO_VERIFY_LEAF_SIGNATURE - CERT_CHAIN_TOO_LONG - INVALID_CA - INVALID_PURPOSE - CERT_UNTRUSTED - CERT_REJECTED Note: If <code>customize-error-ignore</code> is disabled (by default), the CLI shows the following: <code>ca-ignore-errors: UNABLE_TO_GET_ISSUER_CERT UNABLE_TO_GET_CRL CERT_UNTRUSTED</code>
<code>cert-ignore-errors</code>	Note: CLI only. When <code>customize-error-ignore</code> is enabled, the following options become available for you to choose from: - UNABLE_TO_GET_ISSUER_CERT - UNABLE_TO_GET_CRL - CERT_NOT_YET_VALID - CERT_HAS_EXPIRED - CRL_NOT_YET_VALID - CRL_HAS_EXPIRED - DEPTH_ZERO_SELF_SIGNED_CERT - SELF_SIGNED_CERT_IN_CHAIN - UNABLE_TO_GET_ISSUER_CERT_LOCALLY - UNABLE_TO_VERIFY_LEAF_SIGNATURE - CERT_CHAIN_TOO_LONG - INVALID_CA - INVALID_PURPOSE - CERT_UNTRUSTED - CERT_REJECTED Note: If <code>customize-error-ignore</code> is disabled (by default), the CLI shows the following: <code>cert-ignore-errors: UNABLE_TO_GET_CRL</code>
Group Member	
CA	Select a CA (Required).
OCSP	Select an OCSP (Optional).
CRL	Select a CRL (Optional).

Importing CRLs

A certificate revocation list (CRL) is a file that contains a list of revoked certificates with their serial numbers and their revocation dates. The file also contains the name of the issuer of the CRL, the effective date, and the next update date. By default, the shortest validity period of a CRL is one hour.

Some potential reasons for certificates to be revoked include:

- A CA server was hacked and its certificates are no longer trustworthy.
- A single certificate was compromised and is no longer trustworthy.

- A certificate has expired and is not supposed to be used past its lifetime.

You can either upload a CRL file from your local machine or specify the URL of the CRL file



Online Certificate Status Protocol (OCSP) is an alternative to CRL. OCSP is useful when you do not want to deploy CRL files, for example, or want to avoid the public exposure of your PKI structure. For more information, see [Adding OCSPs](#).

Before you begin, you must:

- Have Read-Write permission for System settings.
- Know the URL of a CRL server or have the CRL files downloaded onto your local machine.

To import a CRL file:

1. Go to System > Certificate > Verify.
2. Click the **CRL** tab.
3. Click **Import** to display the configuration editor.
4. Complete the configuration as described in [CRL configuration on page 416](#).
5. Click Save when done.
6. Repeat Steps 3 through 5 to import as many CRLs as needed.

CRL configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. The maximum length is 35 characters. After you initially save the configuration, you cannot edit the name.
Import Method	
HTTP	If selected, FortiADC will download the CRL file from an HTTP server. You must specify the HTTP URL.
SCEP	If selected, FortiADC will download the CRL file from an SCEP server. You must specify the SCEP URL.
File	If selected, you will need to browse for the CRL file on your local machine and upload it into FortiADC.
LDAP	If selected, FortiADC will download the CRL file from the LDAP server (User Authentication > Remote Server > LDAP Server).
CRLDP	If selected, FortiADC will get the address of the CRL file from the extension ("CRL Distribution Points") stored in the client certificate.

Adding OCSPs

FortiADC supports the validation of client digital certificates using Online Certificate Status Protocol (OCSP). In such a configuration, FortiADC contacts the OCSP Responder (i.e., the certificate management system), which maintains the current revocation status information of client certificates or backend server certificates, to determine the current status of digital certificate presented to it. It can then decide whether to allow or block the TLS/SSL connections, based on the status of the client certificates provided by the OCSP Responder.

OCSP enables you to validate certificate status by real-time online query, rather than by importing certificate revocation list (CRL) files. Since distributing and installing CRL files can be a considerable burden in large organizations, and because delay between the release and install of the CRL represents a vulnerability window, this can often be preferable.

During the process of TLS/SSL handshake, FortiADC will send an OCSP status request for the client certificate or backend server certificate to the OCSP Responder. The OCSP Responder then verifies whether the status request contains the information required to identify the certificate and returns a signed response with the status of the inquired certificate, which could be one of the following:

- Good = The certificate has not yet been revoked.
- Revoked = The certificate has been revoked.
- Unknown = The OCSP Responder has no information about the requested certificate, and therefore is able to determine its status.

Note: FortiADC only accepts client certificates in "Good" status as determined by the OCSP Responder as valid.

To use OCSP queries, you must first install the certificates of trusted OCSP servers.

Before you begin, you must:

- Have Read-Write permission for System settings.
- Know the URL of an OCSP server
- Have downloaded the certificate and key files and be able to browse to them so that you can upload them.
- Have already imported the OCSP signing certificates into FortiADC. See [Importing remote certificates](#) and [Creating a CA group](#).

To add an OCSP verify object:

1. Go to System > Verify.
2. Click the **OCSP** tab.
3. Click **Create New** to display the OCSP configuration editor.
4. Complete the configuration as described in [OCSP certificate configuration on page 417](#).
5. Click Save when done.
6. Repeat Steps 3 through 5 to add as many OCSP verify objects as needed.

OCSP certificate configuration

Settings	Guidelines
Name	Enter a unique name for the client certificate validation object that uses OCSP. Valid characters are A-Z, a-z, 0-9, _, and -. The maximum length is 35 characters. No space is allowed.
OCSP URL	Specify the URL of the OCSP Responder.
Verify Others	Upon receiving the OCSP response from the OCSP server, FortiADC first performs OCSP basic verify to validate the OCSP responder's signature. Enable (default)—When Verify Others is enabled, you must select a OCSP Signing Certificate (see OCSP Signing Certificates below). The OCSP basic verify succeeds when the selected OCSP signing certificate matches the OCSP response signature. Otherwise, the OCSP basic verify will fail and the TLS/SSL connection will be terminated. Disable—When Verify Others is disabled, you must select a CA chain. The OCSP basic verify will be carried out in the following sequence:

Settings	Guidelines
	1. The OCSP response signing certificate must be one of the certificates in the CA group or a certificate issued by one of the certificates in the CA group. Also, the certificates must form a chain from the OCSP signing certificate all the way to a self-signed root CA. Otherwise, the OCSP basic verify will fail. 2. If Step 1 (above) is successful, the validation will proceed like this: If the Issuer Criteria Check field is selected (enabled by default), then the OCSP signing certificate can be either the issuing CA of the certificate whose status FortiADC must validate, or a dedicated OCSP signing certificate issued by this issuing CA. The validation succeeds if this criterion is met. Otherwise, the validation process will move onto Step 3 (below). 3. If the OCSP signing certificate is issued by one of the certificates in the CA group, but is not a dedicated OCSP signing certificate, then the validation will proceed like this: If the root CA of this OCSP signing certificate is a trusted self-signed root CA and the "Accept Trusted Root CA" field is selected (enabled by default), then the validation will succeed. Otherwise, the validation will fail.
OCSP Signing Certificates	Select the client certificate of which you'd like to verify the signature of the OCSP Responder that signs it. Note: This option is applicable only when Verify Others is enabled. You MUST select a OCSP signing certificate which must have been imported into FortiADC in advance. See .
CA Chain	Click the down arrow and select a CA group from the list menu. Note: This becomes available only when Verify Others is disabled. In that case, you must select a CA chain (i.e., CA group). It's highly recommended that you have CA groups configured in advance to use this option. See Creating a CA group .
Issuer Criteria Check	Enable/Disable issuer-criteria check. Note: This option comes in hand in hand with CA Chain, and is only available when Verify Others is disabled (see Verify Others above). It is enabled by default, but you can uncheck it if you do not want to validate the certificate issuer's identity.
Accept Trusted Root CA	Enable/Disable accept trusted root CA. Note: This option becomes available only when Criteria Check is enabled (see Criteria Check above). It is enabled by default, in which case FortiADC will accept trusted root CA in the validation process. Uncheck it if you do not want to use this feature.
Timeout	Specify the amount of time in milliseconds (from 1 to 2147483647) the OCSP responder must wait before it times out. The default is 200.
Max age	Specify the maximum amount of time in seconds (from -1 to 214748364) the OCSP responder must check. Note: Setting it to -1 disables max-age check.
Host Header	Specify the host name (Optional).
Reject OCSP Response with Missing Nextupdate	By default, this option is disabled (unselected). In that case, FortiADC accepts all OCSP responses, including those without the nextupdate field. This may have some potential security repercussions, especially if the max-age filed in the OCSP response is not set. To minimize the security risk, you can enable this option so that FortiADC will automatically reject OCSP responses that do not have the nextupdate field.

Settings	Guidelines
	Note: As a good practice, we recommend that, if this option is enabled, you should set an acceptable max-age value (see above) as well so that FortiADC can also check the max-age of the OCSP response. It must be noted that max-age check is an extra, user-enforced check, and that it has nothing to do with the OCSP server's behavior. In other words, once a max-age is set, then FortiADC will enforce the max-age check no matter whether or not the SCSP server sets the nextupdate field in OCSP response.
Caching	Enable or disable OCSP caching. Note: Enabled by default. For a detailed discussion about the function of OCSP caching, see OCSP caching .
Caching Thisupd Extra Maxage	Specify the number of seconds before the this-update-time. The cache will be discarded if the current timestamp is behind the this-update-time in OCSP response. Note: The default is -1, which means that the existing cache will always be used. The smaller value will be used if the max-age and the caching-thisupd-extra-maxage both exist. If one of them is -1, the other one will be used.
Caching Nextupd Ahead Time	Specify the number of seconds before the next-update-time. The cache will be discarded when the current timestamp is ahead of the next-update-time in OCSP response. Note: The default is -1, which means that the existing cache will always be used. Setting the value to 0 means that the cache will expire after the next-update-time, and setting it to 2147483647 makes the cache always expired so that FortiADC always needs to get the latest result from an OCSP server. Warning: There is a default leeway of 60 seconds. So when you set "Caching Nextupd Ahead Time" to x, it means the cache will expire at "x" before "next-update-time", plus 60 seconds.
Nonce Check	Enable or disable nonce check. Note: This option is enabled by default.
Tunneling	Click the button to enable or disable tunneling. If enabled, you must configure all the settings for the tunneling function. See below. Note: Tunneling, or port forwarding, is a way of transmitting private (usually corporate) data through a public network in a disguised way — the routing nodes in the public network are unaware that the transmission is part of a private network.
Tunneling Address	Enter the Tunneling Address that was provided to you.
Tunneling Port	Enter the Tunneling Port number that was provided to you.
Tunneling Password	Specify your password for the tunneling configuration.
Tunneling Username	Specify your user name for the tunneling configuration.
Save	Click the Save button to save your OCSP service configuration.

OCSP caching

OCSP caching is a technique used to speed up OCSP checking. When a client accesses FortiADC or FortiADC accesses a real server for the first time, it (FortiADC) queries the certificate's status using OCSP and caches the response. In

subsequent accesses, the same client or real server will get verified directly from cache, if available.

OCSP caching essentially caches the result of an OCSP verification, not the whole OCSP response. It keeps the certificate status in the buffer for a specified period of time. OCSP verification results can be either obtained by querying an OCSP server or from an OCSP stapling response received from backend real servers.

It must be noted that configuration of OCSP caching is done on a per-VDOM basis and in *rlimit*.

Each OCSP configuration has a flag to let you decide whether to enable OCSP caching or not. Each haproxy process has one and only one OCSP cache which is shared among all OCSP servers.

If OCSP caching is enabled, FortiADC will search its cache first. If no OCSP response result is found in the cache or the cached result has expired (expired OCSP result will be removed from cache), it will query the OCSP server for an updated one. FortiADC uses issuer and serial number hash as key, and also store some extra information (e.g., subject name hash) as extra key. It also implements LRU (least recently used) caching policy. It forms two links: one is to search using key (as an eb-tree) and the other is to implement the LRU caching scheme. You can configure how much memory to use and the maximum period of time to cache (which is useful if the next-update is missing) and cache the nextupdate ahead time.

Implementation of the LRU caching scheme means that frequently used cache would not expire because it will get updated itself upon expiration (replacing itself with a new one) and the least recently used cache may be removed even though it is far from expiration.

When system configuration has changed, FortiADC either restarts the process of haproxy, or performs dynamic reload. In case of a restart, the cache is cleared. In case of dynamic reload, the cache is kept. Modification of cache memory size will restart the haproxy process. Changing other OCSP parameters will trigger dynamic reload.

You can use the existing OCSP max-age to control the lifespan of a cached item, or the "cache-thisupd-extra-maxage" and the "cache-nextupd-ahead-time" to manipulate the caching behavior.

Configure OCSP caching from the Console

```
Config system certificate ocs
Edit "ocsp"
Set caching-flag enable/disable
Set caching-thisupd-extra-maxage 2
Set caching-nextupd-ahead-time 10
End
```

```
Config system vdom
Edit "root"
Set OCSP-caching-maximum-memory 4M
End
```

Importing OCSP signing certificates

OCSP signing certificates are certificates with no private keys. For dynamic certification revocation, you must verify them through an OCSP server. This option allows you to import remote (OCSP) certificates into FortiADC and use them to verify the OCSP response signature.

Before you begin, you must:

- Have Read-Write permission for System settings.
- Have the remote certificates downloaded onto your local machine so that you can upload it to FortiADC.

To import an OCSP-signing certificate:

1. Go to System > Certificate > verify.
2. Click the **OCSP Signing Certificate** tab.
3. Click **Import** to display the configuration editor.
4. Complete the configuration as described in [Importing an OCSP signing certificate on page 421](#).
5. Click Save when done.
6. Repeat Steps 3 through 5 to import as many remote certificates as needed.

Importing an OCSP signing certificate

Settings	Guidelines
Name	Enter a unique name for the remote certificate you want to import. Valid characters are A-Z, a-z, 0-9, _, and -. The maximum length is 35 characters. No space is allowed.
OCSP Signing Certificates	Browse for and upload the remote certificate file of interest.

Once an OCSP signing certificate has been uploaded into FortiADC, the name of the certificate file shows up under the Remote tab. You can view or remove the certificate from this page using the corresponding icons in the far-right column of the page.

Importing CAs

The certificate authority (CA) store is used to authenticate the certificates of other devices. When the FortiADC system is presented with a certificate, it examines the CA's signature, comparing it with the copy of the CA's certificate already imported into the CA store. If the public key matches the private key, the client's or device's certificate is considered legitimate.

In web browsers, the CA store includes trusted root CAs that can be used to establish trust with servers that have certificates signed by the issuing CAs. In an SSL forward proxy deployment, FortiADC acts as a proxy for the client, so you might want to import client browser CAs, create a CA group, and create a certificate verification policy to verify server certificates against this group. You can examine the CA store in common web browsers to come up with a good list of CAs to download and then import. The following list has links for some common web browsers:

- Apple iOS: <https://support.apple.com/en-us/HT204132>
- Google Chrome and Mozilla Firefox: <https://wiki.mozilla.org/CA:IncludedCAs>
- Microsoft Internet Explorer: <https://technet.microsoft.com/en-us/library/dn265983.aspx>

You must do one of the following:

- Import the certificates of the signing CA and all intermediate CAs to FortiADC's store of CA certificates.
- In *all* personal certificates, include the full signing chain up to a CA that FortiADC knows in order to prove that the clients' certificates should be trusted.
- If the signing CA is not known, that CA's own certificate must likewise be signed by one or more other intermediary CAs, until both the FortiADC appliance and the client or device can demonstrate a signing chain that ultimately leads to a mutually trusted (shared "root") CA that they have in common. Like a direct signature by a known CA, this proves that the certificate can be trusted.

Before you begin, you must:

- Have Read-Write permission for System settings.
- Know the URL of an SCEP server or have downloaded the certificate and key files and be able to browse to them so that you can upload them.

To import a CA:

1. Go to System > Certificate > Verify.
2. Click the **CA** tab.
3. Click **Import** to display the configuration editor.
4. Complete the configuration as described in [CA import configuration on page 422](#).
5. Click **Save** when done.
6. Repeat Steps 3 through 5 to import as many CAs as needed.

CA import configuration

Settings	Guidelines
Certificate Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. The maximum length is 35 characters. No space is allowed.
Import Method	• SCEP—Use Simple Certificate Enrollment Protocol. SCEP allows routers and other intermediary network devices to obtain certificates. • File—Upload a file.
SCEP	
SCEP URL	Enter the URL of the SCEP server.
CA Identifier	Enter the identifier for a specific CA on the SCEP server.
File	
Local PC	Browse for the certificate file on the local machine and upload it to FortiADC.

Creating a CA group

CA groups are only used to verify the signature of the OCSP Responder.

Include in the CA group all of the CAs for the pool of backend servers to be associated with a single virtual server.

Before you begin, you must:

- Have Read-Write permission for System settings.
- Have already added the CAs to the CA certificate store.

To create a CA group:

1. Go to System > Certificate > Verify.
2. Click the **CA Group** tab.
3. Click **Create New** to display the configuration editor.
4. Name the CA group and click Save when done. The new CA group appears on the CA Group page.
5. Click the Edit icon in the far-right column (or double-click the CA group) to bring up the configuration editor.

6. Click Create New.
7. Complete the configuration as described in [CA group configuration on page 423](#).
8. Click Save when done.
9. Repeat Steps 6 through 8 to add as many CAs to the group as needed.

CA group configuration

Settings	Guidelines
Group Name	Specify a unique name for the CA group that you are creating. Valid characters are A-Z, a-z, 0-9, _, and -. The maximum length is 35 characters. No space is allowed.
Group Member	
CA	Click the down arrow and select the desired CA from the list menu to add to the group.

System alerts

This section provides a description of the alert system and instructions to configure it to monitor important system events and metrics. You can create alert policies that monitor and provide alerts for some of the following:

- User authentication events
- Security events
- HA information
- Server Load Balance information
- Link Load Balance information
- Global Load Balance information
- Appliance information, including temperatures and fan speeds
- System information and metrics, including CPU, memory, and disk usage

For more information about events and metrics that you can monitor in alert policies, see [Creating alert configurations](#).

The chapter includes the following topics:

- [Creating alert configurations](#)
- [Configuring alert actions](#)
- [Configuring alert policies](#)

Configuring alert actions

Alert actions define how FortiADC responds to triggered alert configurations in an alert policy. You can configure FortiADC to send logs of alerts to syslog servers, email recipients, and SNMP managers.

Before you begin, keep the following in mind:

- If you want to send logs of alerts to a syslog server, you must configure a remote syslog server. See [Configuring syslog settings](#).
- If you want to send messages to recipients via email, you must configure alert email settings. See [Configuring report email](#).

- If you want to send SNMP traps, you must configure an SNMP trap server. See [Configuring SNMP trap servers on page 428](#)

To configure alert actions:

1. Go to System > Alert > Alert and select the **Alert Actions** tab.
2. Click **Create New**.
3. Complete the configuration as described in [Alert Actions on page 424](#).
4. Click **Save**.

Alert Actions

Settings	Guidelines
Name	Specify the name of the alert action. You will use the name to select the alert action in alert policies.
Syslog	Select the remote syslog server(s) you want to include.
Email	Select the email recipient(s) you want to include.
SNMP Trap	Select the SNMP manager(s) you want to include.

Configuring alert policies

Alert policies allow you to select groups of alert configurations to monitor. If alert configurations in the policy are triggered, FortiADC will send alert messages according to the alert actions selected in the policy.

Before you begin:

- You must configure alert actions. See [Configuring alert actions](#).
- If you want to use custom alert configurations, you must create new ones. See [Creating alert configurations](#).

To configure alert policies:

1. Go to System > Alert > Alert and select the **Alert Policy** tab.
2. Click **Create New**.
3. Complete the configuration as described in [Alert Policy on page 424](#).
4. Click **Save**.
5. To add alert configurations to the policy, see [To add alert configurations in an alert policy](#).

Alert Policy

Settings	Guidelines
Name	Specify the name of the alert policy. No spaces.
Status	Select Enable so that FortiADC will generate alerts according to the policy.
Action	Select the alert action to determine how FortiADC will respond to alert configurations in the policy.
Comments	Enter comments or a description of the policy for your records.

To add alert configurations in an alert policy:

1. Go to System > Alert > Alert and select the **Alert Policy** tab.
2. Select the policy to which you want to add alert configurations.
3. In the **Alert Member** section, click **Create New**.
4. Complete the configuration as described in [Alert configurations in a policy on page 425](#).
5. Click **Save**.
6. Repeat Steps 3 through 5 to add as many alert configurations as needed.

Alert configurations in a policy

Settings	Guidelines
Name	Specify the name of the alert configuration for the policy. No spaces.
Status	Enable so that the alert policy will monitor the alert configuration and generate alerts when it's triggered.
Alert	Select the alert configuration you want to include.
Inherit	Enable so that the alert action set in the alert policy will determine how FortiADC responds when the alert configuration is triggered. Disable if you want to set a custom alert action for the alert configuration.
Action	Only available when Inherit is disabled. Select the custom alert action that FortiADC will use to respond to the alert configuration when it's triggered.

Creating alert configurations

Alert configurations are specific events or metrics that you can monitor. If the alert configurations are triggered, you can define alert actions for them in alert policies. FortiADC comes equipped with a number of default alert configurations; you can further configure these to fit your environment's particular needs.

To create alert configurations:

1. Go to System > Alert > Alert and select the **Alert Config** tab.
2. Click **Create New**.
3. Complete the configuration as described in [Alert Config on page 425](#).
4. Click **Save**.

Alert Config

Settings	Guidelines
Name	Specify the name of the alert configuration. You will use the name to select the alert configuration in alert policies.
Priority	Set the alert level of the alert configuration: • High • Middle • Low The alert level is color-coded and denotes the severity of the alert configuration.

Settings	Guidelines
Rolling Window	Enable to define a Rolling Window Time and Number of Occurrence (see below). The Rolling Window Time sets a period of time in which a number of events must take place before an alert is triggered. The number of events that must take place within this period of time is set in the Number of Occurrences option. Note: the Throttle Alert option may override and suppress alerts defined by the rolling window.
Rolling Window Time	Available only when Rolling Window is enabled (see above). Specify the range of time (in seconds) for the rolling window. The valid range is 1–3600.
Alert Expiry Time	Specify the time (in seconds) until the alert is no longer active in the web interface. Once the alert expires and is no longer active, it is still visible, but will be grayed out. The valid range is 3600–7776000. The default value is 86400
Number of Occurrences	Available only when Rolling Window is enabled (see above). Specify the number of events that must take place before FortiADC will trigger the alert. The valid range is 1–3600. The default value is 1.
Throttle Alert	Specify a range of time (in seconds) in which FortiADC will trigger an alert. Within the range of time, only one alert will trigger after any number of events in the alert configuration occur. The valid range is 1–3600. The default value is 300.
Description	Enter a comments or description of the alert configuration as needed.
Source Type	Select either of the options: - Event—Select this option to choose an event that triggers the alert. - Metric—Select this option to specify the metric that triggers the alert. To use this option, you must configure the Alert Metric Expire Member as described at the end of this section.
Event Occurs	Note: This option is available only when <code>Event</code> is selected in the Source Type field. Select the event to be monitored in the alert configuration. Note: A brief description of the selected event appears below the drop-down menu box.
Object	Note: This option is available when <code>Metric</code> is selected in the Source Type field. Select one of the following options: - System - Virtual Server - Interface
Duration	Note: This option is available when <code>Metric</code> is selected in the Source Type field. Specify the length of time (in seconds) required for a selected "metric" to exist before an alert is triggered.
Instance	Note: This option is available only when either <code>Virtual Sever</code> or <code>Interface</code> is selected in the Object field. - Virtual Server—Select a virtual server (name) from the drop-down menu. - Interface—Select a network interface (port) from the drop-down menu.

To modify default alert configurations:

You cannot edit or delete default alert configurations, but you can clone them and create custom alert configurations.

1. Go to System > Alert > Alert and select the **Alert Config** tab.
2. Click the  (clone) icon in the row of the default alert configuration that you want to modify.

3. Complete the configuration as described in [Alert Config on page 425](#).
4. Click **Save**.

To add metrics to alert configurations:

Before you begin, you must create and save an alert configuration in which the Source Type is `Metric` and the **Duration** is defined.

1. Go to System > Alert > Alert.
2. Select the **Alert Config** tab
3. Double-click the alert configuration in the row of the alert configuration that you want to modify.
4. Set **Rolling Window** to Off.
5. Set **Source Type** to Metric.
6. Set **Object** to System. This will open up the **Alert Metric Expire Member** section all the way below.
7. In that section, click **Create New**.
8. Complete the configuration as described in [Add metrics to alert configurations on page 427](#).
9. Click **Save**.
10. Complete Steps 3 through 5 for as many metrics as you want to monitor in an alert configuration.

Add metrics to alert configurations

Settings	Guidelines
Name	Specify a name for the metric.
Metric Occurs	Select among the following metrics that the event configuration will monitor: - dev_stats.avg_cpu_usage—total average CPU usage as a percentage of CPU available to the server - dev_stats.avg_mem_usage—total average memory usage as a percentage of memory available to the server - dev_stats.avg_disk_usage—virtual disk1 capacity usage
Comparator	The metric is compared to the Value field according to the selected option: - Ge—greater than - Le—less than - Eq—equal to The alert configuration will trigger if the specified value satisfies the selected option.
Value	Specify the metric value that the Comparator uses to determine if the metric triggers an alert. Enter the scalar portion of the value. For example, if you want to specify 2 milliseconds, 2 is the scalar and milliseconds is the unit of measure. Once the scalar portion of the value is defined, the Vantage web interface will auto-populate the unit portion of the value field based on the metric selected.

Note: After a metric-object-instance has been specified in an alert, the system will not prohibit you from deleting it in other part of the system configuration. For instance, a port named `vlan1` is added in network configuration and then used as a metric-object-instance in an alert. If `vlan1` is deleted later on in network configuration, FortiADC will not generate an error message for this action.

Configuring SNMP trap servers

Simple Network Management Protocol (SNMP) allows you to collect and exchange hardware and software information about devices on your network. You can configure an SNMP trap server so that FortiADC's alert system is able to send SNMP traps about important events and metrics. For details about events and metrics that FortiADC can monitor, see [Creating alert configurations on page 425](#).

Before you begin:

- You must have Read-Write permission for System settings.
- In the FortiADC interface settings, you must enable SNMP access on the network interface through which the SNMP manager connects.
- On the SNMP manager, you must verify that the SNMP manager is a member of the community to which the FortiADC system belongs, and you must compile the necessary Fortinet-proprietary management information blocks (MIBs) and Fortinet-supported standard MIBs. For information on Fortinet MIBs, see [Configuring SNMP trap servers on page 428](#) and [Appendix A: Fortinet MIBs](#).

To configure an SNMP manager:

1. Go to System > Alert > Alert Resource and select the **SNMP Trap Server** tab.
2. Click **Create New**.
3. Complete the configuration as described in [SNMP trap server configuration on page 428](#).
4. Click **Save**.

SNMP trap server configuration

Settings	Guidelines
Name	Enter the name of the trap server. No spaces. You will use this name to select the trap server in an Alert Actions profile. See Configuring alert actions on page 423 .
Hosts	Enter the IP address of the SNMP manager(s) that will receive traps.
Version	Enter the version of SNMP that you want to utilize for the trap server.
Local Port	Enter the source port number for trap packets sent to the SNMP manager(s) for the trap server. The default port is 162.
Remote Port	Enter the destination port number for trap packets sent to the SNMP manager(s) for the trap server. The default port is 162.
Note:	The following options apply to SNMP v3 only.
Security Level	Choose one of the following three security levels: • No Auth and No Privacy—Enables no additional authentication or encryption compared to SNMP v1 and v2 • Auth But No Privacy—Enables authentication only. The SNMP manager needs to supply the password specified in this community configuration. Also specify Auth Algorithm and the associated Auth Password below. • Auth and Privacy—Enables both authentication and encryption. Also specify Auth Algorithm, Auth Password, Private Algorithm, and Private Password below. Note: This option is available only for v3 SNMP managers.

Settings	Guidelines
Auth Algorithm	Specify the authentication algorithm. Ensure that the SNMP manager and FortiADC use the same algorithm. Note: Available only when the selected Security Level is Auth But No Privacy or Auth and Privacy.
Auth Password	Specify the password for the authentication algorithm. Ensure that the SNMP manager and FortiADC use the same password. Note: Available only when the selected Security Level is Auth But No Privacy or Auth and Privacy.
Private Algorithm	Specify the encryption algorithm. Ensure that the SNMP manager and FortiADC use the same algorithm. Note: Available only when the selected Security Level is Auth and Privacy.
Private Password	Specify the password for the encryption algorithm. Ensure that the SNMP manager and FortiADC use the same password. Note: Available only when the selected Security Level is Auth and Privacy.

Configuring an email alert object

FortiADC is able to send email alerts based on your specification. To use this feature, you must configure your own email alert objects to keep track the emails

To configure an email object:

1. Go to System > Alert > Alert Resource and select the Email tab.
2. Click **Create New**.
3. Complete the configuration as described in [Syslog server configuration on page 429](#).
4. Click **Save**.

Syslog server configuration

Settings	Guidelines
Name	Enter a name for the email alert object, e.g., Accounting. No spaces. You will use this name to select the email alerts in the Alert Actions profile.
Mail From	Enter the email address of the email sender.
Mail To	Enter the email address of the email recipient.

Configuring a syslog object

Syslog is an industry standard for sending log messages across a network. Because the syslog protocol provides a wide range of system information, syslog monitoring has been an important part of network monitoring.

A syslog server receives and analyzes syslog messages, stored in a high performance database. It checks the content of received syslog messages and trigger alarms depending on the content and severity. To enable FortiADC to track syslog alerts (i.e., syslog messages), you must configure a syslog object.

To configure a syslog object:

1. Go to System > Alert > Alert Resource and select the Syslog tab.
2. Click **Create New**.
3. Complete the configuration as described in [Syslog server configuration on page 430](#).
4. Click **Save**.

Syslog server configuration

Settings	Guidelines
Name	Enter a name for the syslog message object. No spaces. You will use this name to select the syslog in an Alert Actions profile.
Syslog Server	Enter the IP address of the syslog server that will receive syslog messages.
Port	Enter the port of the syslog server. The default is 514.

HSM Integration

A hardware security module (HSM) is a dedicated device for managing digital keys and performing cryptographic operations. An HSM can be a plug-in card or an external device directly connected to a computer or network server. Purposefully designed to protect the crypto-key life cycle, HSMs have been used by some of the world's most security-conscious entities to protect their cryptographic infrastructure by securely managing, processing, and storing cryptographic keys inside a hardened, tamper-resistant device.

Because of their strengths in securing cryptographic keys and provisioning encryption, decryption, authentication, and digital signing services for a wide range of applications, HSMs have been used by enterprises worldwide to safeguard their online transactions, identities, and applications.

Integrating FortiADC with SafeNet Network HSM

Starting from Version 4.7.2, FortiADC has integrated with SafeNet Network HSM. It enables you to retrieve a per-connection, SSL session key from the HSM server instead of loading the private key and certificate stored on FortiADC.

The integration requires specific configuration steps on both the FortiADC and the HSM appliances, as outlined below:

On the HSM appliance:

- Create one or more HSM partitions for FortiADC
- Send the FortiADC client certificate to the HSM server
- Register the FortiADC HSM client to the partition(s)
- Retrieve the HSM server certificate

On the FortiADC appliance:

- Configure communication with the HSM server, including using the server and client certificates to register FortiADC as a client of the HSM server
- Generate a certificate-signing request (CSR) that includes the HSM's configuration information
- Upload the signed certificate to FortiADC

It must be noted that

- Currently, FortiADC supports the SafeNet Network HSM only.
- HSM support is disabled on FortiADC by default. You must enable it via the CLI for the feature to become available on the FortiADC GUI. To enable HSM support from the CLI, execute the following commands:

```
config system global
set hsm enable
```

- You must have the HSM server certificate available on your local PC or a network drive.
- HSM integration supports all HA modes, i.e., active-active, active-passive, and VRRP.
- HSM partition is a global configuration that can be used from individual VDOMs.
- HSM integration does not support configuration synchronization (config-sync), but local certificate using HSM can be synchronized to peer FortiADC appliances. Keep in mind that this local certificate may NOT function properly on peer FortiADC appliances.
- Network Trust Links (NTLs) IP check (ntls ipcheck) must be disabled on the HSM server for HA configuration.

The following instructions assume that you have (1) HSM support enabled on FortiADC and (2) access to the HSM server certificate from your PC.

Preparing the HSM appliance

Before starting to configure FortiADC-HSM integration, you must configure the SafeNet Network HSM first using the following steps:

1. On the SafeNet Network HSM, use the `partition create` command to create and initialize a new HSM partition that uses password authentication.
Note: This is the partition FortiADC uses on the HSM server. You can create more than one partition, but all the partitions are assigned to the same client. For more information, see HSM-related documentation.
2. Use the SCP utility and the following command to send the FortiADC client certificate to the HSM:
`scp <fortiadc_ip>.pem admin@<hsm_ip>:`
3. Using SSH, connect to the HSM server using the admin account. Then, use the following command to register a client for FortiADC on the HSM server:
`lunash:> client register -c <client_name> -ip <fortiadc_ip>`, where `<client_name>` is the name you specify that identifies the client.
4. Use the following command to assign the client you registered to the partition you've created in Step 1 above:
`lunash:> client assignPartition -client <client_name> -partition <partition_name>`
You can verify the assignment using the following command:
`lunash:> client show -client <client_name>`
5. Repeat the client assignment process for any additional partitions you've created for FortiADC.
6. Use the SCP utility and the following command to retrieve the server certificate file from the HSM server:
`scp <hsm_username>@<hsm_ip>:server.pem /usr/lunasa/bin/server_<hsm_ip>.pem`
7. On the FortiADC GUI, navigate to System>HSM to bring up the HSM configuration page.
8. Complete the HSM configuration as described in [HSM Configuration Parameters on page 431](#). Then move on to [Generating a certificate-signing request on FortiADC on page 432](#).

HSM Configuration Parameters

Parameter	Description
Client Certificate	

Parameter	Description
Client IP	Enter the IP address of the interface (i.e., port) which FortiADC uses to generate the client certificate. Note: This IP address is the common name of client certificate. FortiADC is the client of the HSM server. The client and server certificates are used in SSL connection between FortiADC and the HSM server.
Generate	Click this button to generate the client certificate that you've specified above. Note: Use this option only if you do not have an existing client certificate on FortiADC.
Download	Click this button retrieve the client certificate that you have just generated or stored on FortiADC. Note: You must generate a client certificate if you do not have one already residing on FortiADC. See above.
Configuration	Complete the following entries or selections to configure the FortiADC-HSM integration.
Server IP	Enter the IP address of the HSM server.
Port	Specify the port via which FortiADC establishes an NTLS connection with the HSM server. The default value is 1792.
Timeout	Specify a timeout value for the connection between FortiADC and the HSM server. The default is 20000. Valid values range from 5000 to 20000 milliseconds.
Upload Server Certificate File	Click Browse to browse for the server certificate file that you retrieved earlier.
Register	Click this button to register FortiADC as a client of the HSM sever using the specified server and client certificates. Note: This action generates a config file, e.g., <code>/example.conf</code>
Unregister	Click this button to clear all HSM-related configurations on the back-end.
Partition	Click Create New to create partition or Delete to remove a selected partition. Note: FortiADC can accept only one partition. Once a partition is added, the Register and Unregister buttons become dimmed out, meaning you cannot make any change to the HSM configuration. To edit the HSM configuration, you must delete the partition first.
Partition Name	Specify the name of a partition to which the FortiADC HSM client is assigned.
Password	Specify the password for the partition.

Note: When configure your CSR to work with an HSM, the CSR generation process creates a private key on both the HSM and the FortiADC. The private key on the HSM is the "real" key that secures communication when FortiADC uses the signed certificate. The key found on the FortiADC is used when you upload the certificate to FortiADC.

Generating a certificate-signing request on FortiADC

Once you have completed configuring the HSM server, you must generate a certificate-signing request which references the HSM connection and partition from inside FortiADC.

To generate a certificate-signing request:

1. On the FortiADC GUI, navigate to System > Manage Certificates > Local Certificate.
2. Click Generate to bring up the Local Certificate configuration page.
3. Configure the certificate-signing request as described in [Generating a certificate-signing request on page 433](#). Then move on to [Downloading and uploading the certificate request \(.csr\) file on page 435](#).

Generating a certificate-signing request

Parameter	Description
Generate Certificate Signing Request	Complete the following entries or selections to configure the FortiADC-HSM integration.
Certificate Name	Specify a name for the certificate request, e.g., www.example.com. This can be the name of your web site.
Subject Information	Specify the information that the certificate is required to contain in order to uniquely identify the FortiADC appliance. This area varies depending on the ID Type you select.
ID Type	Select the type of identifier to use in the certificate to identify the FortiADC appliance: • Host IP — Select this option if the FortiADC appliance has a static IP address, and then enter the public IP address of the FortiADC appliance in the IP field. If the FortiADC appliance does not have a public IP address, use Domain Name or Email instead. See below. • Domain Name — Select this option if the FortiADC appliance has a static IP address and subscribes to a dynamic DNS service. Enter the FQDN of the FortiADC appliance, such as www.example.com, in the Domain Name field, but do NOT include the protocol specification (http://) or any port number or path names. • Email — Select this option if the FortiADC appliance does not require either a static IP address or a domain name. Enter the email address of the owner of the FortiADC appliance in the Email field. The ID type you can select varies by whether or not your FortiADC appliance has a static IP address, a fully-qualified domain name (FQDN), and by the primarily intended use of the certificate. For example, if your FortiADC appliance has both a static IP address and a domain name, but you will primarily use the local certificate for HTTPS connections to the web UI by the domain name of the FortiADC appliance, you might prefer to generate a certificate based upon the domain name of the FortiADC appliance rather than its IP address. Depending on your choice for ID Type, the other options may vary.
IP	Note: This option appears only if the ID Type is Host IP. Enter the static IP address of the FortiADC appliance, such as 10.0.0.1. The IP address must be the one visible to clients. Usually, this should be its public IP address on the Internet, or a virtual IP that you use NAT to map to the appliance's IP address on your private network.
Domain Name	Note: This option appears only if the ID Type is Domain Name. Enter the fully qualified domain name (FQDN) of the FortiADC appliance, such as www.example.com. The domain name must resolve to the static IP address of the FortiADC appliance or a protected server.

Parameter	Description
Email	Note: This option appears only if the ID Type is Email. Enter the email address of the owner/user of the FortiADC appliance, such as admin@example.com.
Distinguished Information	The following information is OPTIONAL in the certificate; it is NOT required.
Organization unit	Enter the name of your organizational unit (OU), such as the name of your department. To enter more than one OU name, click the + icon, and enter each OU in each separate field.
Organization	Enter the legal name of your organization.
Locality(City)	Enter the name of the city or town where the FortiADC appliance is deployed.
State/Province	Enter the name of the state or province where the FortiADC appliance is deployed.
Country/Region	Select the name of the country where the FortiADC appliance is deployed.
Email	Enter an email address that may be used for contact purposes, such as admin@example.com.
Key Information	Enter the information pertinent to the key.
Key Type	This field shows the type of algorithm used to generate the key. Note: It's read-only and cannot be changed. FortiADC 4.7.2 supports RSA key type only.
Key Size	Select one of the following key sizes: • 512 bit • 1024 bit • 1536 bit • 2048 bit • 4096 bit Note: Larger keys may take longer to generate, but provide better security.
HSM	Select this option if the private key for the connections is provided by an HSM appliance instead of FortiADC. Note: This option is available only if you have enabled HSM via the CLI using the <code>config system global</code> command. For more information, see Integrating FortiADC with SafeNet Network HSM on page 430 .
Partition Name	Enter the name of the partition where the private key for this certificate is located on the HSM server. Note: This option becomes available only when HSM is selected. See above.
Enrollment Information	
Enrollment Method	Select either of the following: • File Based —If selected, you must manually download and submit the resulting certificate signing request (.csr) file to a certificate authority (CA) for signing. Once signed, you need to upload the local certificate. <i>This is the only enrollment method if HSM is selected.</i>

Parameter	Description
	Online SCEP — If selected, the FortiADC appliance will automatically use HTTP to submit the certificate-signing request to the simple certificate enrollment protocol (SCEP) server of a CA, which will validate and sign the certificate. Note: For this selection, two more options appear: CA Server URL and Challenge Password. <i>This option is not available if HSM is selected.</i>

Downloading and uploading the certificate request (.csr) file

Normally, when generating a certificate-signing request, the FortiADC appliance creates a private and public key pair. The generated request includes the public key of the FortiADC appliance and information such as the FortiADC appliance's IP address, domain name, or email address. The FortiADC appliance's private key remains confidential on the FortiADC appliance. The Status column of the entry is PENDING.

If you configured your CSR to work with the FortiADC-HSM integration, the CSR generation process creates a private key both on the HSM and on FortiADC appliances. The private key on the HSM is used to secure communication when FortiADC uses the certificate. The FortiADC private key is used when you upload the certificate to FortiADC.

After you have submitted a certificate-signing request from inside FortiADC as discussed above, you must go back to the System > Management Certificates > Local Certificate page to download the certificate request (.csr) file, and then upload that file to your certificate authority (CA) by taking the following steps:

1. On the System > Manage Certificates > Local Certificate page, locate the entry of the certificate request.
2. Click the Download icon.
Note: The time it takes to download the certificate request (.csr) file varies, depending on the size of the file and the speed of your network connection. After the file is downloaded, save it at a location on your machine.
3. Upload the certificate request (.csr) file to your CA.
Note: Upon receiving the certificate request file, the CA will verify the information in the certificate, give it a serial number and an expiration date, and sign it with the public key of the CA.
4. If you are not using a commercial CA whose root certificate is already installed by default on web browsers, download your CA's root certificate, and then install it on all computers that will be connecting to your FortiADC appliance.
Note: You must have the certificate installed on the computers. Otherwise, they may not trust your new certificate. After you have received the signed certificate from the CA, upload it to FortiADC, as discussed below.

Uploading the server certificate to FortiADC

You must have the Read and Write permission to upload server certificates to the FortiADC appliance.

To upload the server certificate to FortiADC:

1. On the FortiADC GUI, navigate to the System > Manage Certificates > Local Certificate page.
2. Click Import.
3. Make the selections as described in [Uploading a server certificate on page 436](#), and click Save.

Uploading a server certificate

Parameter	Description
Type	Click the down arrow and select one of the following options from the drop-down menu: - Local Certificate—Use this option only if you have a CA-signed certificate that was originated from a CSR generated in FortiADC . See HSM Integration on page 430. Note: It is important to make sure that the load-balancer (FortiADC appliance) you use to import a local certificate is the same appliance where the CSR was generated because it is where the key matching the certificate resides. The import operation will fail without the matching key on the same hardware system. - PKCS12 Certificate—Use this option only if you have a PKCS #12 password-encrypted certificate with its key in the same file. - Certificate—Use this option only if you have a certificate and its key in separate files. Note: Additional fields are displayed depending on your selection.
Certificate File	Click Browse to locate the certificate file that you want to upload.
Certificate Name	The name of the certificate. Note: This field applies when Type is Certificate or PKCS12.
Key File	Click Browse to locate the key file that you want to upload with the certificate. Note: This option is available only if Type is Certificate.
Password	Enter the password used to encrypt the server certificate file. Note: This enables FortiADC to decrypt and install the certificate. This option is available only if Type is Certificate or PKCS12 Certificate.

Once a certificate is uploaded to FortiADC, you can use it in a policy or server pool configuration.

Chapter 14: Logging and Reporting

This chapter includes the following topics:

- [Downloading logs on page 437](#)
- [Using the security log on page 438](#)
- [Using the traffic log on page 445](#)
- [Using the script log on page 452](#)
- [Configuring local log settings on page 452](#)
- [Configuring syslog settings on page 454](#)
- [Configuring fast stats log settings on page 456](#)
- [Enabling real-time statistics on page 456](#)
- [Configuring report email on page 457](#)
- [Configuring reports on page 457](#)
- [Configuring Report Queries on page 459](#)
- [Configuring fast reports on page 461](#)

Downloading logs

You can download the local collection of raw log files. You might do this if you are following manual procedures for storing log data or performing ad hoc analysis or troubleshooting.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To download logs:

1. Go to Log & Report > Log Browsing.
2. Complete the configuration as described in [Download logs on page 437](#).
3. Click the **Download** button.

Download logs

Settings	Guidelines
Log/Sublog	Event Log: • Configuration • System • Admin • User • Health Check • SLB • LLB

Settings	Guidelines
	• GLB • Firewall Security Log: • IP Reputation • DoS • WAF • GEO Traffic Log: • SLB Layer 4 • SLB HTTP • SLB TCPS • SLB RADIUS • GLB • SLB SIP • SLB DRP • SLB DNS • SLB RTSP • SLB SMTP • SLB RTMP • SLB DIAMETER • SLB MySQL Script Log: • SLB Aggregate Log • Syncflood • GEO • IP Reputation • WAF
Filters	Configure the filters. Note: Filter options may vary, depending on the type and/or sub-type of the log that you select.

Using the security log

The Security Log table displays logs related to security features.

By default, the log is filtered to display IP Reputation logs, and the table lists the most recent records first.

You can use the following category filters to review logs of interest:

- IP Reputation—Traffic logged by the IP Reputation feature
- DoS—Traffic logged by the SYN Flood feature
- WAF—Traffic logged by the WAF feature
- Geo—Traffic logged by the Geo IP block list feature
- AV—Traffic logged by the AV feature

Within each category, you can use Filter Setting controls to filter the table based on the values of matching data:

- Date
- Time
- Proto
- Service
- Src
- Src_port
- Dst
- Dst_port
- Vs Name
- Action

The last column in each table includes a link to log details.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To view and filter the log:

1. Go to Log & Report > Log Browsing.
2. Click the **Security Logs** tab to display the attack log.
3. Click **Filter Settings** to display the filter tools.
4. Use the tools to filter on key columns and values.
5. Click **OK** to apply the filter and redisplay the log.

[IP Reputation log on page 439](#) to [Geo IP log on page 442](#) list the log columns in the order in which they appear in the log.

IP Reputation log

Column	Example	Description
date	date=2014-12-02	Log date.
time	time=10:27:01	Log time.
log_id	log_id=0200004230	Log ID.
type	type=attack	Log type: attack.
subtype	subtype=ip_reputation	Log subtype: ip_reputation.
pri	pri=warning	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=13065998	Message ID.
count	count=1	For IP reputation, count=1.
severity	severity=high	Rule severity.
proto	proto=6	Protocol.
service	service=http	Service.

Column	Example	Description
src	src=4.4.4.4	Source IP address.
src_port	src_port=49301	Source port.
dst	dst=2.2.2.2	Destination IP address.
dst_port	dst_port=80	Destination port.
policy	policy=vs1	Virtual server name.
action	action=deny	Policy action.
srccountry	srccountry=cn	Location of the source IP address.
dstcountry	dstcountry=us	Location of the destination IP address.
msg	msg=msg	Security rule name, category, subcategory, and description of the attack.

DoS log

Column	Example	Description
date	date=2014-12-02	Log date.
time	time=10:27:01	Log time.
log_id	log_id=0200004230	Log ID.
type	type=attack	Log type: attack.
subtype	subtype=synflood	Log subtype: synflood.
pri	pri=warning	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=13065998	Message ID.
count	count=1	For DoS, number of timeouts sent per destination.
severity	severity=high	Always "high" for DoS.
proto	proto=0	Protocol.
service	service=http	Service.
src	src=173.177.99.94	Source IP address.
src_port	src_port=49301	Source port.
dst	dst=10.61.2.100	Destination IP address.
dst_port	dst_port=80	Destination port.
policy	policy=unknown	For DoS, policy=unknown.
action	action=deny	Policy action.

Column	Example	Description
srccountry	srccountry=cn	Location of the source IP address.
dstcountry	dstcountry=us	Location of the destination IP address.
msg	msg=msg	Security rule name, category, subcategory, and description of the attack.

WAF log

Column	Example	Description
date	date=2015-07-22	Log date.
time	time=10:27:01	Log time.
log_id	log_id=0202008074	Log ID.
type	type=attack	Log type: attack.
subtype	subtype=waf	Log subtype: waf.
pri	pri=alert	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=1512	Message ID.
count	count=1	Rule match count.
severity	severity=low	Rule severity.
proto	proto=6	Protocol.
service	service=http	Service.
src	src=1.1.1.1	Source IP address.
src_port	src_port=34352	Source port.
dst	dst=2.2.2.2	Destination IP address.
dst_port	dst_port=80	Destination port.
policy	policy=vs1	Virtual server name.
action	action=pass	Policy action.
sigid	sigid=1	Attack signature ID.
subcat	subcat=waf_subtype	WAF module: waf_web_attack_signature, waf_url_access, waf_http_protocol_cont and waf_sql_xss_injection_detect.
http_host	http_host=192.168.1.140:8080	HTTP Host header in HTTP request. Maximum length is 64. Longer URIs are truncated and appended with
http_url	http_url=/bigdata	URI in HTTP request. Maximum length is 128. Longer URIs are truncated and appended with
pkt_hdr	pkt_hdr=header	Contents of the packet header that matched the attack signature.
srccountry	srccountry=Australia	Location of the source IP address.

Column	Example	Description
dstcountry	dstcountry=France	Location of the destination IP address.
msg	msg="Find Attack ID: 1010010001 NAME: "HTTP Method Violation" CATEGORY: "HTTP Protocol Constraint" SUB_CATEGORY: "Request Method Rule""	Security rule name, category, subcategory, and description of the attack.

Geo IP log

Column	Example	Description
date	date=2014-12-02	Log date.
time	time=10:27:01	Log time.
log_id	log_id=0200004230	Log ID.
type	type=attack	Log type: attack.
subtype	subtype=geo	Log subtype: geo.
pri	pri=warning	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=13065998	Message ID.
count	count=1	Rule match count.
severity	severity=high	Rule severity.
proto	proto=0	Protocol.
service	service=http	Service.
src	src=173.177.99.94	Source IP address.
src_port	src_port=49301	Source port.
dst	dst=10.61.2.100	Destination IP address.
dst_port	dst_port=80	Destination port.
policy	policy=vs1	Virtual server name.
action	action=deny	Policy action.
srccountry	srccountry=cn	Location of the source IP address.
dstcountry	dstcountry=us	Location of the destination IP address.
msg	msg=msg	Security rule name, category, subcategory, and description of the attack.

AV log

Column	Example	Description
date	date=2014-12-02	Log date.
time	time=10:27:01	Log time.
log_id	log_id=0200004230	Log ID.
msg_id	message id=362301459	Message ID
virus category	virus category=N/A	Virus Category.
count	count=1	Rule match count.
severity	severity=high	Rule severity.
proto	proto=0	Protocol.
service	service=http	Service.
src	src=173.177.99.94	Source IP address.
src_port	src_port=49301	Source port.
dst	dst=10.61.2.100	Destination IP address.
dst_port	dst_port=80	Destination port.
type	type=attack	Type
subtype	subtype=av	Sub Type
action	action=deny	Policy action.
srccountry	srccountry=cn	Location of the source IP address.
dstcountry	dstcountry=us	Location of the destination IP address.
msg	msg=msg	Security rule name, category, subcategory, and description

Column	Example	Description
		of the attack.
sign_id	sign_id=0	Signature ID
virus_id	virus_id=0	Virus ID
av_anatype	av_anatype=analytics	AV AnaType
url	url=none	URL
virus/botnet	virus/botnet=N/A	Virus/Botnet
Submitted to FortiSandbox	Submitted_to_Fortisandbox=no	Submitted to FortiSandBox
quar file name	quar_file_name=N/A	Quar File Name
Proto Method	proto_method=none	Proto Method
AV Profile	av_profile=AV1	AV Profile
FortiSandbox Checksum	B08663FD9FC147D6ADBB3D70DCEC1271A4288C71D887D44811D93E366D91AD2C	

Using the traffic log

The Traffic Log table displays logs related to traffic served by the FortiADC deployment.

By default, the log is filtered to display Server Load Balancing - Layer 4 traffic logs, and the table lists the most recent records first.

You can use the following category filters to review logs of interest:

- SLB Layer 4—Traffic served by Layer-4 virtual servers
- SLB HTTP—Traffic served by virtual servers with HTTP profiles
- SLB TCPS—Traffic served by virtual servers with TCPS profiles
- SLB RADIUS—Traffic served by virtual servers with RADIUS profiles
- GLB—Traffic served by global load balancing policies
- SLB SIP—Traffic served by virtual servers with SIP profiles
- SLB RDP—Traffic served by virtual servers with RDP profiles
- SLB DNS —Traffic served by virtual servers with DNS profiles
- SLB RTSP —Traffic served by virtual servers with RTSP profiles
- SLB SMTP —Traffic served by virtual servers with SMTP profiles
- SLB RTMP—Traffic served by virtual servers with RTMP profiles
- SLB DIAMETER—Traffic served by Diameter profiles
- SLB MySQL—Traffic served by MySQL profiles.
- LLB — Traffic served by LLB profiles.

Within each category, you can use Filter Setting controls to filter the table based on the values of matching data:

- Date
- Time
- Proto
- Service
- Src
- Src_port
- Dst
- Dst_port
- Policy
- Action

The last column in each table includes a link to log details.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To view and filter the log:

1. Go to Log & Report > Log Access > Traffic Logs to display the traffic log.
2. Click **Filter Settings** to display the filter tools.

3. Use the tools to filter on key columns and values.
4. Click **Apply** to apply the filter and redisplay the log.

[SLB Layer 4 and SLB TCPS logs on page 446](#) to [GLB log on page 451](#) list the log columns in the order in which they appear in the log.

SLB Layer 4 and SLB TCPS logs

Column	Example	Description
date	date=2014-12-01	Log date.
time	time=07:50:36	Log time.
log_id	log_id=0102007810	Log ID.
type	type=traffic	Log type.
subtype	subtype=slb_tcps	Log subtype: slb_layer4, slb_tcps.
pri	pri=information	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=522030	Message ID.
duration	duration=55	Session duration.
ibytes	ibytes=138	Bytes in.
obytes	obytes=303	Bytes out.
proto	proto=6	Protocol.
service	service=tcps	Service.
src	src=31.1.1.103	Source IP address in traffic received by FortiADC.
src_port	src_port=5534	Source port.
dst	dst=21.1.1.101	Destination IP address in traffic received by FortiADC (IP address of the virtual server).
dst_port	dst_port=443	Destination port.
trans_src	trans_src=31.1.1.103	Source IP address in packet sent from FortiADC. Address might have been translated.
trans_src_port	trans_src_port=5534	Source port in packet sent from FortiADC.
trans_dst	trans_dst=21.1.1.101	Destination IP address in packet sent from FortiADC (IP address of the real server).
trans_dst_port	trans_dst_port=443	Destination port in packet sent from FortiADC.
policy	policy=L7vs	Virtual server name.
action	action=none	For most logs, action=none.
srccountry	srccountry=Reserved	Location of the source IP address.
dstcountry	dstcountry=Reserved	Location of the destination IP address.
real_server	real_server=2_2_2_10	Real server configured name.

SLB HTTP logs

Column	Example	Description
date	date=2014-12-01	Log date.
time	time=07:50:36	Log time.
log_id	log_id=0102007810	Log ID.
type	type=traffic	Log type.
subtype	subtype=slb_http	Log subtype: slb_http.
pri	pri=information	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=522030	Message ID.
duration	duration=55	Session duration.
ibytes	ibytes=138	Bytes in.
obytes	obytes=303	Bytes out.
proto	proto=6	Protocol.
service	service=http	Service.
src	src=31.1.1.103	Source IP address in traffic received by FortiADC.
src_port	src_port=5534	Source port.
dst	dst=21.1.1.101	Destination IP address in traffic received by FortiADC (IP address of the virtual server).
dst_port	dst_port=443	Destination port.
trans_src	trans_src=31.1.1.103	Source IP address in packet sent from FortiADC. Address might have been translated.
trans_src_port	trans_src_port=5534	Source port in packet sent from FortiADC.
trans_dst	trans_dst=21.1.1.101	Destination IP address in packet sent from FortiADC (IP address of the real server).
trans_dst_port	trans_dst_port=443	Destination port in packet sent from FortiADC.
policy	policy=L7vs	Virtual server name.
action	action=none	For most logs, action=none.
http_method	http_method=get	HTTP method.
http_host	http_host=10.61.2.100	Host IP address.
http_agent	http_agent=curl/7.29.0	HTTP agent.
http_url=	http_url=/ip.php	Base URL.
http_qry	http_qry=unknown	URL parameters after the base URL.
http_cookie	http_cookie=unknown	Cookie name.
http_retcode	http_retcode=200	HTTP return code.

Column	Example	Description
user	user=user1	User name.
usergrp	usergrp=companyABC	User group.
auth_status	auth_status=success	Authentication success/failure.
srccountry	srccountry=Reserved	Location of the source IP address.
dstcountry	dstcountry=Reserved	Location of the destination IP address.
real_server	real_server=2_2_2_10	Real server configured name.

SLB RADIUS log

Column	Example	Description
date	date=2014-12-01	Log date.
time	time=07:50:36	Log time.
log_id	log_id=0102007810	Log ID.
type	type=traffic	Log type.
subtype	subtype=slb_radius.	Log subtype: slb_radius.
pri	pri=information	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=522030	Message ID.
duration	duration=55	Session duration.
ibytes	ibytes=138	Bytes in.
obytes	obytes=303	Bytes out.
proto	proto=6	Protocol.
service	service=radius	Service.
src	src=31.1.1.103	Source IP address in traffic received by FortiADC.
src_port	src_port=5534	Source port.
dst	dst=21.1.1.101	Destination IP address in traffic received by FortiADC (IP address of the virtual server).
dst_port	dst_port=443	Destination port.
trans_src	trans_src=31.1.1.103	Source IP address in packet sent from FortiADC. Address might have been translated.
trans_src_port	trans_src_port=5534	Source port in packet sent from FortiADC.
trans_dst	trans_dst=21.1.1.101	Destination IP address in packet sent from FortiADC (IP address of the real server).
trans_dst_port	trans_dst_port=443	Destination port in packet sent from FortiADC.
policy	policy=L7vs	Virtual server name.

Column	Example	Description
action	action=none	For RADIUS, action=auth or acct.
user	user=user1	RADIUS accounting username.
srccountry	srccountry=Reserved	Location of the source IP address.
dstcountry	dstcountry=Reserved	Location of the destination IP address.
real_server	real_server=2_2_2_10	Real server configured name.

SLB RDP logs

Column	Example	Description
date	date=2016-03-18	Log date.
time	time=11:48:29	Log time.
log_id	log_id=107005800	Log ID.
type	type=traffic	Log type.
subtype	subtype=slb_rdp	Log subtype: slb_rdp.
pri	pri=information	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=1321705	Message ID.
duration	duration=2	Session duration.
ibytes	ibytes=92	Bytes in.
obytes	obytes=400	Bytes out.
proto	proto=6	Protocol.
service	service=http	Service.
src	src=192.168.1.1	Source IP address in traffic received by FortiADC.
src_port	src_port=37869	Source port.
dst	dst=192.168.1.142	Destination IP address in traffic received by FortiADC (IP address of the virtual server).
dst_port	dst_port=8080	Destination port.
trans_src	trans_src=2.2.2.2	Source IP address in packet sent from FortiADC. Address might have been translated.
trans_src_port	trans_src_port=58661	Source port in packet sent from FortiADC.
trans_dst	trans_dst=2.2.2.10	Destination IP address in packet sent from FortiADC (IP address of the real server).
trans_dst_port	trans_dst_port=80	Destination port in packet sent from FortiADC.
policy	policy=vs-l7	Virtual server name.
action	action=none	For most logs, action=none.

Column	Example	Description
srccountry	srccountry=Reserved	Location of the source IP address.
dstcountry	dstcountry=Reserved	Location of the destination IP address.
real_server	real_server=r_22210	Real server configured name.

SLB SIP logs

Column	Example	Description
date	date=2016-01-29	Log date.
time	time=18:06:48	Log time.
log_id	log_id=0106001134	Log ID.
type	type=traffic	Log type.
subtype	subtype=slb_sip	Log subtype: slb_sip.
pri	pri=information	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=154799	Message ID.
duration	duration=1	Session duration.
ibytes	ibytes=44346	Bytes in.
obytes	obytes=2.2.2.10	Bytes out.
proto	proto=6	Protocol.
service	service=http	Service.
src	src=N/A	Source IP address in traffic received by FortiADC.
src_port	src_port=43672	Source port.
dst	dst=192.168.1.142	Destination IP address in traffic received by FortiADC (IP address of the virtual server).
dst_port	dst_port=8080	Destination port.
trans_src	trans_src=2.2.2.2	Source IP address in packet sent from FortiADC. Address might have been translated.
trans_src_port	trans_src_port=80	Source port in packet sent from FortiADC.
trans_dst	trans_dst=N/A	Destination IP address in packet sent from FortiADC (IP address of the real server).
trans_dst_port	trans_dst_port=none	Destination port in packet sent from FortiADC.
policy	policy=invite	Virtual server name.
action	action=sip: bob@1.1.1.1 v2.0	Invite sent to.
sip_method	sip_method=from: alice@2.2.2.2	Invite sent from.

Column	Example	Description
sip_uri	sip_uri=to: server@3.3.3.3	SIP server IP address.
sip_from	sip_from=callid:1111111	SIP call ID.
sip_to	sip_to=200	
sip_callid	sip_callid=Reserved	Reserved.
sip_retcode	sip_retcode=Reserved	Reserved.
srccountry	srccountry=Reserved	Location of the source IP address.
dstcountry	dstcountry=Reserved	Location of the destination IP address.
real_server	real_server=2_2_2_10	Real server configured name.

GLB log

Column	Example	Description
date	date=2014-12-01	Log date.
time	time=07:50:36	Log time.
log_id	log_id=0102007810	Log ID.
type	type=traffic	Log type.
subtype	subtype=dns	Log subtype: dns.
pri	pri=information	Log severity.
vd	vd=root	Virtual domain.
msg_id	msg_id=522030	Message ID.
proto	proto=6	Protocol.
src	src=31.1.1.103	Source IP address.
src_port	src_port=5534	Source port.
dst	dst=21.1.1.101	Destination IP address.
dst_port	dst_port=443	Destination port.
policy	policy=policy	Global load balancing policy name.
action	action=none	For most logs, action=none.
fqdn	fqdn=pool.ntp.org	FQDN from client request.
resip	resip=4.53.160.75	DNS response IP address.
srccountry	srccountry=Reserved	Location of the source IP address.
dstcountry	dstcountry=Reserved	Location of the destination IP address.

LLB log

Column	Example	Description
date	date=2014-12-01	Log date.
time	time=07:50:36	Log time.
log_id	log_id=0114000000	Log ID.
type	type=traffic	Log type.
subtype	subtype=llb	Log subtype: llb
pri	pri=information	Log severity.
vd	vd=root	Virtual domain.
msg_id	msg_id=522030	Message ID.
duration	duration=120	Session duration
ibytes	ibytes=1131	Bytes in
obytes	obytes=492	Bytes out
proto	proto=6	Protocol.
src	src=31.1.1.103	Source IP address.
src_port	src_port=5534	Source port.
dst	dst=21.1.1.101	Destination IP address.
dst_port	dst_port=443	Destination port.
policy	policy=Link_Policy	Link Policy.
action	action=vtunnel	Group Type (Link Group or Virtual Tunnel) in Link Group
srrcountry	srrcountry=Japan	Location of the source IP address
dstcountry	dstcountry=France	location of the destination IP address
gateway	gateway=none	Gateway in Link Group

Using the script log

The Script Log table shows all the scripts.

Note: This feature is available for the SLB (server load balance) module only.

Configuring local log settings

The local log is a datastore hosted on the FortiADC system.

Typically, you use the local log to capture information about system health and system administration activities. We recommend that you use local logging during evaluation and verification of your initial deployment, and then configure remote logging to send logs to a log management repository where they can be stored long term and analyzed using preferred analytic tools.

Local log disk settings are configurable. You can select a subset of system events, traffic, and security logs.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To configure local log settings:

1. Go to Log & Report > Log Setting.
The configuration page displays the Local Log tab.
2. Complete the configuration as described in [Local logging configuration on page 453](#).
3. Save the configuration.

Local logging configuration

Settings	Guidelines
Status	Select to enable local logging.
File Size	Maximum disk space for a local log file. The default is 200 MB. When the current log file reaches this size, a new file is created.
Log Level	Select the lowest severity to log from the following choices: • Emergency—The system has become unstable. • Alert—Immediate action is required. • Critical—Functionality is affected. • Error—An error condition exists and functionality could be affected. • Warning—Functionality might be affected. • Notification—Information about normal events. • Information—General information about system operations. • Debug—Detailed information about the system that can be used to troubleshoot unexpected behavior. For example, if you select Error, the system collects logs with level Error, Critical, Alert, and Emergency. If you select Alert, the system collects logs with level Alert and Emergency.
Disk Full	Select log behavior when the maximum disk space for local logs (30% of total disk space) is reached: • Overwrite—Continue logging. Overwrite the earliest logs. • No Log—Stop logging.
Event	Select to enable logging for events.
Event Category	This option becomes available only when the Event check box is selected. In that case, select the types of events to collect in the local log:

Settings	Guidelines
	• Configuration—Configuration changes. • Admin—Administrator actions. • System—System operations, warnings, and errors. • User—Authentication results logs. • Health Check—Health check results and client certificate validation check results. • SLB—Notifications, such as connection limit reached. • LLB—Notifications, such as bandwidth thresholds reached. • GLB—Notifications, such as the status of associated local SLB and virtual servers. • Firewall—Notifications for the "firewall" module, such as SNAT source IP pool is using all of its addresses.
Traffic	Select to enable logging for traffic processed by the load balancing modules.
Traffic Category	The following options become available only when the Traffic check-box is selected. See above. • SLB—Server Load Balancing traffic logs related to sessions and throughput. • GLB—Global Load Balancing traffic logs related to DNS requests. • LLB—Link Load Balancing traffic logs related to session and throughput.
Security	Select to enable logging for traffic processed by the security modules.
Security Category	• DoS—SYN flood protection logs. • IP Reputation—IP Reputation logs. • WAF—WAF logs. • GEO—Geo IP blocking logs. • AV—AV logs. • Enable All—All types of log mentioned above.
Script	Select to enable scripting.
Script Category	SLB is selected by default and required.

Configuring syslog settings

A remote syslog server is a system provisioned specifically to collect logs for long term storage and analysis with preferred analytic tools.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To configure syslog settings:

1. Go to Log & Report > Log Setting.
2. Click the **Syslog Server** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Syslog configuration on page 455](#).
5. Save the configuration.

Syslog configuration

Settings	Guidelines
Status	Select to enable the configuration.
Address	IP address of the syslog server.
Port	Listening port number of the syslog server. Usually this is UDP port 514.
Log Level	Select the lowest severity to log from the following choices: - Emergency—The system has become unstable. - Alert—Immediate action is required. - Critical—Functionality is affected. - Error—An error condition exists and functionality could be affected. - Warning—Functionality might be affected. - Notification—Information about normal events. - Information—General information about system operations. - Debug—Detailed information about the system that can be used to troubleshoot unexpected behavior. For example, if you select Error, the system sends the syslog server logs with level Error, Critical, Alert, and Emergency. If you select Alert, the system collects logs with level Alert and Emergency.
CSV	Send logs in CSV format. Do not use with FortiAnalyzer.
Facility	Identifier that is not used by any other device on your network when sending logs to FortiAnalyzer/syslog.
Event	Select to enable logging for events.
Event Category	Select the types of events to send to the syslog server: - Configuration—Configuration changes. - Admin—Administrator actions. - System—System operations, warnings, and errors. - User—Authentication results logs. - Health Check—Health check results and client certificate validation check results. - SLB—Notifications, such as connection limit reached. - LLB—Notifications, such as bandwidth thresholds reached. - GLB—Notifications, such as the status of associated local SLB and virtual servers. - Firewall—Notifications for the "firewall" module, such as SNAT source IP pool is using all of its addresses.
Traffic	Select to enable logging for traffic processed by the load balancing modules.
Traffic Category	- SLB—Server Load Balancing traffic logs related to sessions and throughput. - GLB—Global Load Balancing traffic logs related to DNS requests. - LLB—Link Load Balancing traffic logs related to session and throughput
Security	Select to enable logging for traffic processed by the security modules.
Security Category	- DoS—SYN flood protection logs. - IP Reputation—IP Reputation logs.

Settings	Guidelines
	- WAF—WAF logs. - GEO—Geo IP blocking logs.
Script	Select to enable scripting.
Script Category	SLB is elected by default.

Configuring fast stats log settings

The fast stats log feature enables real-time statistics collection for fast reports. By default, the feature is enabled, but you can disable it if you like.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To enable or disable the fast stats log feature:

- Go to Log & Report > Log Setting.
- Click the Fast Stats tab.
- Complete the configuration as described in [Fast stats log configuration on page 456](#).
- Save the configuration.

Fast stats log configuration

Settings	Guidelines
Status	Enable/disable fast statistics. The feature is enabled by default.
Traffic	Enable/disable fast statistics for traffic logs. The feature is enabled by default.
Traffic Category	Enable/disable fast statistics for traffic categories. SLB is enabled by default..
Security	Select to enable logging for traffic processed by the security modules. Disabled by default.
Security Category	- DoS—SYN flood protection logs. - IP Reputation—IP Reputation logs. - WAF—WAF logs. - GEO—Geo IP blocking logs.

Enabling real-time statistics

The fast statistics feature enables real-time statistics collection for fast reports. Enabled by default. You can disable fast statistics if you encounter issues.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To enable/disable real-time statistics:

1. Go to Log & Report > Log Setting.
2. Click the **Fast Stats** tab.
3. Complete the configuration as described in [Fast Statistics configuration on page 457](#).
4. Save the configuration.

Fast Statistics configuration

Settings	Guidelines
Status	Enable/disable fast statistics. Enabled by default.
Traffic	Enable/disable fast statistics for traffic logs. Enabled by default.
Traffic Category	Enable/disable fast statistics for traffic categories. SLB is enabled by default.

Configuring report email

You can configure report email objects to work with an SMTP mail server. See [Configuring an SMTP mail server](#) for information on how to set up the connection to the mail server.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To configure report email objects:

1. Click Log & Report > Report Email.
2. Click the **Create New** tab.
3. Complete the configuration as described in [Report mail configuration on page 457](#).
4. Click Save.

Report mail configuration

Settings	Guidelines
Name	Enter a name for the report email configuration object, e.g., Accounting. No spaces.
Mail To	Enter the email address of the report email recipient.
Mail From	Enter the email address of the report email sender.

Configuring reports

You can configure on-demand or scheduled reports.

Before you begin:

- If you want reports to include user-defined queries, you must configure the queries before you configure the report.
- You must have Read-Write permission for Log & Report settings.

To configure a report:

1. Go to Log & Report > Report Config.
The Report tab is displayed.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in [Report configuration on page 458](#).
4. Save the configuration.

To run an on-demand report:

- In the report table, the final column for has a "run report" icon (▶). Click it.

To view a generated report:

- Go to Log & Report > Report > Overall.

Report configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the zone configuration (if you use forwarders). Note: After you initially save the configuration, you cannot edit the name.
On Schedule	Enable/disable reporting on schedule.
Period	Select a report period. If you select absolute or last N-hours, last N-days, or last N-weeks, additional controls are displayed for you to set these variables.
Schedule Type	Daily or on specified days.
Schedule Weekdays	If you do not schedule the report daily, specify the days on which to run it.
Schedule Hour	0-23.
Email Format	Attachment format. Only PDF is supported. If you schedule reports and set this option, the report is sent on schedule to all addresses in the Log & Report > Report Email > Recipient list.
Email Subject	Message subject.
Email Body	Message body.
Email Attachname	Filename for attachment.
Email Compress	Enable/disable compression of the attachment.
Query List	Select queries to include in the report.

Configuring Report Queries

The predefined list of queries covers the most common administrator and stakeholder interests. It includes the following:

- SLB-Top-Policy-By-Bytes
- SLB-Top-Source-By-Bytes
- SLB-Top-Source-Country-By-Bytes
- SLB-History-Flow-By-Bytes (total traffic over time)
- LLB-Top-Link-by-Bytes
- LLB-History-Flow-By-Bytes (total traffic over time)
- DNS-Top-Policy-by-Count
- DNS-Top-Source-by-Count
- Attack-Top-Destination-For-IPReputation-By-Count
- Attack-Top-Source-For-IPReputation-By-Count
- Attack-Top-Source-Country-For-IPReputation-By-Count
- Attack-Top-Destination-For-Synflood-By-Count
- Attack-Top-Destination-For-GEO-By-Count
- Attack-Top-Source-For-GEO-By-Count
- Attack-Top-Source-Country-For-GEO-By-Count
- Attack-Top-Destination-For-WAF-By-Count
- Attack-Top-Source-For-WAF-By-Count
- Attack-Top-Source-Country-For-WAF-By-Count
- Event-Top-Admin-Login-By-Count
- Event-Top-Failed-Admin-Login-By-Count
- Event-Top-Admin-Config-By-Count

If necessary, you can create your own query configuration objects.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

After you have created a query configuration object, you can select it in the report configuration.

To configure report queries:

1. Go to Log & Report > Report Config.
The Report tab is displayed.
2. Click the **Query Set** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [Query configuration on page 460](#).
5. Save the configuration.

Query configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the zone configuration (if you use forwarders). Note: After you initially save the configuration, you cannot edit the name.
Module	• SLB • LLB • DNS • Attack • Event
SLB	
SLB Submodule	• All—Queryset will include all SLB queries • HTTP—Queryset will include only HTTP queries
SubType	Submodule All has the following subtypes: • top_policy (virtual server) • top_source • top_source_country • slb_history_flow (total traffic over time) Submodule HTTP has the following subtypes: • top_policy (virtual server) • top_pool_member
Traffic Sort Type	Submodule All has the following Traffic Sort Types • sessions • bytes Submodule HTTP has the following Traffic Sort Types: • sessions • bytes • CPS • RPS • BPS • Average Session Duration • Transaction Latency
LLB	
Traffic Sort Type	• sessions • bytes
LLB Subtype	• top_link • slb_history_flow (total traffic over time)
DNS	
DNS Sort Type	Only count is applicable.
DNS Subtype	• Top_Policy

Settings	Guidelines
	top_source
Attack	
Attack Sort Type	Only count is applicable.
Attack Subtype	- top_destip_for_geo - top_destip_for_ipreputation - top_destip_for_sysflood - top_destip_for_waf - top_source_country_for_geo - top_source_country_for_ipreputation - top_source_country_for_waf - top_source_for_geo - top_source_for_ipreputation - top_source_for_waf
Event	
Event Sort Type	Only count is applicable.
Event Subtype	- top_admin_login - top_failed_admin_login - top_admin_config

Configuring fast reports

Fast reports are real time statistics displayed on the Dashboard > Data Analytics page.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

After you have created a query configuration object, you can select it in the report configuration.

There are two ways to configure a fast report.

The Log & Report route:

1. Go to Log & Report > Report Config
2. Click the **Fast Report** tab.
3. Click **Create New**.
4. Complete the configuration as described in [Fast report configuration on page 462](#).
5. Save the configuration.

The Fortiview Route:

1. Fortiview > Data Analytics
2. Click the **+ Add Widget** button in the far right. The **Fast Report** dialogue will open up. It will be the same as in the Log & Report route.
3. Complete the configuration as described in [Fast report configuration on page 462](#).
4. Save the configuration.

Fast report configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the zone configuration (if you use forwarders). Note: After you initially save the configuration, you cannot edit the name.
Module	Select one of the following options: • SLB • Security
SLB SubType	Select an option from the list menu: • Top Src • Top Dest • Top Browser • Top OS • Top Dev • Top Domain • Top URL • Top Referrer • Top Source Country • Top Session
Security SubType	Select an option from the list menu: • Top Attack Type for All • Top Attack Type by VS for All • Top VS for DDoS • Top Destination Country for DDoS • Top VS for GEO • Top Source for GEO • Top Destination for GEO • Top Source Country for GEO • Top Destination Country for GEO • Top Action by Source for GEO • Top Action by Source Country for GEO • Top Category by VS for IP Reputation • Top Source for IP Reputation • Top Destination for IP Reputation • Top Source Country for IP Reputation • Top Destination Country for IP Reputation • Top Attack Type by VS for WAF

Settings	Guidelines
	• Top Attack Type by Source Country for WAF • Top Attack Type by Source for WAF • Top Attack by Destination Country for WAF • Top Attack by Destination for WAF
History Chart	Enable/Disable.
Time Range	Select an option from the list menu: • 10MINS • 1HOUR • 1DAY • 1WEEK • 1MONTH
Data Type	Select either of the following: • Bandwidth • Session

Display logs via CLI

FortiADC allows you to display logs using the CLI, with filtering functions.

```
FortiADC-VM # execute log
delete-file      delete-file
delete-type      delete-type
display          display the log message
filter           set filter for log browsing
list-type        list-type
rebuild-db       rebuild-db
```

```
FortiADC-VM # execute log display

<startline integer >=0 >    show log from startline

FortiADC-VM # execute log filter

<type|subtype|field|clear|show>    set ,clear,show filters
```

Where:

-
- type <event|traffic|attack>
 - subtype <subtype_value> ex:slb_http
 - field <field_name> <field_value_list>

Chapter 15: High Availability Deployments

This chapter includes the following topics:

- [HA feature overview on page 465](#)
- [HA system requirements on page 470](#)
- [HA configuration synchronization on page 470](#)
- [Configuring HA settings on page 471](#)
- [Monitoring an HA cluster on page 475](#)
- [Updating firmware for an HA cluster on page 476](#)
- [Deploying an active-passive cluster on page 478](#)
- [Deploying an active-active cluster on page 481](#)
- [Deploying an active-active-VRRP cluster on page 493](#)

HA feature overview

FortiADC appliances can be deployed as standalone units or as high availability (HA) clusters.

A *cluster* is two or more nodes. A *node* is an instance of the appliance/system. In a cluster, one node is the *primary node*, also called the *master node*. The other members of the cluster are *secondary nodes*, also called *slave nodes*.

The primary node has a special role. It has a one-to-many relationship with member nodes. Both configuration updates and software updates are initiated by the primary node and pushed to member nodes.

The system selects the primary node based on the following criteria:

- Link health (if monitor ports links are down, the node is considered down)
- Remote IP monitor health check results
- Override setting (prefers priority to uptime)
- Most available ports
- Highest uptime value
- Lowest device priority number (1 has greater priority than 2)
- Highest-sorting serial number—Serial numbers are sorted by comparing each character from left to right, where 9 and z are the greatest values. The system gives preference to higher values over lower values.

HA solutions depend on two types of communication among cluster members:

- Synchronization—During initialization, the primary node pushes its configuration (with noted exceptions) to member nodes. After initialization has completed, the nodes synchronize their session tables.
- Heartbeats—A cluster node indicates to other nodes in the cluster that it is up and available. The absence of heartbeat traffic indicates the node is not up and is unavailable.

There are three types of HA clusters:

- Active-Passive—Only the primary node is active, so it is the only node that receives traffic from adjacent routers. Typically, there is one other node that is in standby mode. It assumes active status if the primary node undergoes maintenance or otherwise becomes unavailable.

- **Active-Active**—All nodes receive traffic. Active-Active deployments support load balancing and failover among up to eight cluster members.
- **Active-Active-VRRP** —FortiADC's Active-Active-VRRP mode uses a VRRP-like protocol, and can function in both HA Active-Passive mode and HA Active-Active mode, depending on the number of traffic groups used in the configuration. When only one traffic group is used, it actually functions in Active-Passive mode; when two or more traffic groups are used, it works in Active-Active mode.

In an Active-Passive cluster, only the management IP address for the primary node is active. In an active-passive cluster, you can log into a node only when it has primary node status and its IP address is active. To access the user interface of an appliance in standby status (the active-passive slave), you must use a console port connection.


In an Active-Active cluster, the IP addresses for all interfaces are unique, including the management interface. When the appliance is in standalone mode, the physical port IP address is active; when it is in HA mode, the address assigned to it in the HA node IP list address is active. You can log into any node using the active IP address for its management port.

In an Active-Active-VRRP cluster, FortiADC uses hbdev for members status communication. It also allows you to configure sync+session, persistence sync, and image sync functions via hbdev and dataport, which is essentially the same as the HA-AA/AP mode. Note that FortiADC is unable to communicate with third-party VRRP devices because it actually doesn't use the VRRP protocol at all.

Tip: You can use the `execute ha manage` command to log into the console of a member node. See the CLI reference.


[Basic active-passive cluster on page 466](#) shows an *active-passive* cluster in a single network path. In an active-passive cluster, the primary node is the active node that handles all traffic. In the event that the primary node experiences hardware failure or system maintenance, *failover* takes place. In failover, the standby node becomes the primary node and processes the traffic that is forwarded along the network path. The new primary node sends gratuitous ARP to notify the network to direct traffic for the virtual MAC addresses (vMAC) to its network interfaces. It takes the IP addresses of the unresponsive node.

Basic active-passive cluster



[Redundant path active-passive cluster on page 467](#) shows an active-passive cluster in a *redundant path*. A topology like this is a best practice because it is fully redundant, with no single point of failure. If the gateway, load balancer, or switch were to fail, the failover path is chosen.

Redundant path active-passive cluster




[Basic active-active cluster on page 469](#) shows an *active-active* cluster. An active-active cluster supports load-balancing and failover among up to eight member nodes. The routers on either side of the cluster must be configured to use equal cost multipath (ECMP) to distribute traffic to the FortiADC cluster nodes. All nodes actively receive and forward traffic.

The primary node has a special role. It handles all FTP and firewall traffic, and it acts as the failover node for all of the other nodes in the cluster.

The failover mechanism is the same as an active-passive deployment, with the primary node acting as the standby node for all other cluster members. If a member node fails, the primary node takes the IP addresses of the unresponsive node and notifies the network via ARP to redirect traffic for that vMAC to its own network interfaces. For example, in [Basic active-active cluster on page 469](#), node1 is the primary node. If node2 were to fail, its traffic would failover to node1. If node3 were to fail, its traffic would also failover to node1. If the primary node were to fail, a new primary node would be elected, and it would function as the master in all respects, including its role as the new standby node for failover from all other cluster members.

Basic active-active cluster



HA system requirements

- Appliances must have the same hardware model and same firmware version.
- Redundant network topology: if an active node fails, physical network cabling and routes must be able to redirect traffic to the other member nodes.
- At least one physical port on both HA appliances to be used for heartbeat and data traffic between cluster members. For active-passive failover pairs, you can connect the ports directly via a crossover cable. For active-active clusters with more than two members, you can connect the nodes via the same Layer 2 switch.
- Heartbeat and synchronization traffic between cluster nodes occur over the physical network ports that you designate. If switches are used to connect the nodes, the interfaces must be reachable by Layer 2 multicast.
- Each appliance must be licensed. If using FortiADC-VM, the license must be paid; trial licenses will not function.



FortiADC-VM supports HA. However, if you do not want to use the native HA, you can use your hypervisor or VM environment manager to install your virtual appliances over a hardware cluster to improve availability. For example, VMware clusters can use vMotion or VMware HA.

HA configuration synchronization

Normally in an HA configuration, the master node pushes most of its configuration to the other member nodes. This is known as HA configuration synchronization. If automatic synchronization is enabled, synchronization occurs automatically when an appliance joins the cluster, and it repeats every 30 seconds thereafter. If synchronization is not enabled, you must initiate synchronization manually.

HA configuration synchronization includes:

- Core CLI-style configuration file (fadc_system.conf)
- X.509 certificates, certificate signing request files (CSR), and private keys
- Layer-7 virtual server error message files
- Layer-4 TCP connection state, Layer-4 persistence table, and Layer-7 persistence table (Source Address Persistence table only)
- Health check status (active-passive deployments only)

For most settings, you configure only the primary node, and its settings are pushed to other members.

[HA settings that are not synchronized on page 470](#) summarizes the configuration settings that are not synchronized. All other settings are synchronized.

HA settings that are not synchronized

Setting	Explanation
Hostname	The hostnames are not synchronized to enable you to use unique names.
SNMP system information	Each member node has its own SNMP system information so that you can maintain accurate, separate data in SNMP collections. However, the network interfaces of a standby node are not active, so they cannot be actively monitored with SNMP.

Setting	Explanation
RAID level	RAID settings are hardware-dependent and determined at boot time by looking at the drives (for software RAID) or the controller (hardware RAID), and are not stored in the system configuration. Therefore, they are not synchronized.
HA settings	Most of the HA configuration is not synchronized in order to support HA system operations. In particular: • Priority and Override settings—These settings are used to elect a primary node, so they are not synchronized to enable differentiation. • Group ID—Nodes with the same Group ID join a cluster. The setting precedes and determines group membership, so it is set manually. • HA mode—Many administrators prefer to be able to switch the primary node from an HA mode to standalone mode without the other nodes following suit, or to switch a secondary node to standalone mode and have that setting not overwritten by periodic synchronization, so the HA mode setting is not pushed from the primary node to the member nodes. • Node list and Local Node ID—These settings are for active-active mode only. They identify a node uniquely within an active-active load balancing group, so they are not synchronized to enable differentiation.

In addition to HA settings, the following data is *not* synchronized either:

- Log messages—These describe events that happened on a specific appliance. After a fail-over, you might notice that there is a gap in the original active appliance's log files that corresponds to the period of its down time. Log messages created during the time when the standby was acting as the active appliance (if you have configured local log storage) are stored there, on the original standby appliance.
- Generated reports—Like the log messages that they are based upon, reports also describe events that happened on that specific appliance. As such, report settings are synchronized, but report output is not.

You can view the status of cluster members from the dashboard of the primary node. You might need to log into the system for a non-primary member node in the following situations:

- To configure settings that are not synchronized.
- To view log messages recorded about the member node itself on its own hard disk.
- To view traffic reports for traffic processed by the member node.

Configuring HA settings

Note: Currently, FortiADC only supports HA configurations for IPv4 address mode; HA is not supported on IPv6.

Before you begin:

- You must have Read-Write permission to items in the System category.

To configure HA settings:

1. Go to System > High Availability.
2. Complete the configuration as described in [High availability configuration on page 472](#).
3. Save the configuration.

After you have saved the configuration, cluster members begin to send heartbeat traffic to each other. Members with the same Group ID join the cluster. They send synchronization traffic through their data links.

High availability configuration

Settings	Guidelines
Cluster Mode	- Standalone - Active-Passive - Active-Active - Active-Active-VRRP
Basic Settings	
Active-Passive	
Group Name	Name to identify the HA cluster if you have more than one. This setting is optional, and does not affect HA function. The maximum length is 63 characters.
Group ID	Number that identifies the HA cluster. Nodes with the same group ID join the cluster. If you have more than one HA cluster on the same network, each cluster must have a different group ID. The group ID is used in the virtual MAC address that is sent in broadcast ARP messages. The valid range is 0 to 31. The default value is 0.
Config Priority	The default value is 100, but you can specify any numeric value ranging from 0 to 255. Note: FortiADC 4.7.x has introduced a new parameter called config-priority for HA configuration. It allows you to determine which configuration the system uses when synchronizing the configuration between the HA nodes. Therefore, upon upgrading to FortiADC 4.7.x, it is highly recommended that you use this option to manually set different HA configuration priority values on the nodes. Otherwise, you'll have no control over the system's master-slave configuration sync behavior. When the configuration priority values are identical on both nodes (whether by default or by configuration), the system uses the configuration of the appliance with the larger serial number to override that of the appliance with the smaller serial number. When the configuration priority values on the nodes are different, the configuration of the appliance with the lower configuration priority will prevail.
Active-Active	
Group Name	Same as Active-Passive. See above.
Group ID	Same as Active-Passive. See above.
Config Priority	Same as Active-Passive. See above.
Local Node ID	A number that uniquely identifies the member within the cluster. The valid range is from 0 to 7. This number is used in the virtual MAC address that is sent in ARP responses.
Node List	Select the node IDs for the nodes in the cluster. An active-active cluster can have up to eight members.
Active-Active-VRRP	
Group Name	Same as Active-Passive. See above.
Group ID	Same as Active-Passive. See above.
Config Priority	Same as Active-Passive. See above.
Local Node ID	Same as Active-Active. See above.


Settings	Guidelines
Synchronization	
Layer 7 Persistence Synchronization	Enable to synchronize Layer 7 session data used for persistence to backend servers. When enabled, the Source Address Persistence table is synchronized between HA members. When not enabled, a node that receives traffic due to failover would not know that a session had been created already, so it will be treated as a new session. Synchronization of the persistence table is not required for cookie-based or hash-based persistence methods to get the desired result. Client traffic will be routed to the same backend server. Synchronization of the persistence table is not possible for SSL session ID. When the session via the first node is terminated, the client must re-establish an SSL connection via the second node. When a client requests a new SSL connection with an SSL server, the initial TCP connection has an SSL Session ID of 0. This zero value tells the server that it needs to set up a new SSL session and to generate an SSL Session ID. The server sends the new SSL Session ID in its response to the client as part of the SSL handshake.
Layer 4 Persistence Synchronization	Enable to synchronize Layer 4 session data used for persistence to backend servers. When enabled, the Source Address Persistence table is synchronized between HA members. When not enabled, a node that receives traffic because of load balancing or failover would not know that a session had been created already, so it will be treated as a new session. Synchronization of the persistence table is not required for hash-based persistence methods to get the desired result. Client traffic will be routed to the same backend server.
Layer 4 Connection Synchronization	Enable to synchronize Layer 4 connection state data. When enabled, the TCP session table is synchronized. If subsequent traffic for the connection is distributed through a different cluster node because of failover, the TCP sessions can resume without interruption. When not enabled, a node that receives traffic because of failover would not know that a session had been created already, and the client will be required to re-initialize the connection.
Advanced Settings	
Priority	Number indicating priority of the member node when electing the cluster primary node. This setting is optional. The smaller the number, the higher the priority. The default is 5. The valid range is from 0 to 9. Note: By default, up time is more important than this setting unless Override is enabled. See below.
Override	Enabled by default. This makes device priority (see above) a more important factor than up time when selecting the primary node.
Heartbeat Interval	Number of 100-millisecond intervals at which heartbeat packets are sent. This is also the interval at which a node expects to receive heartbeat packets. This part of the configuration is pushed from the primary node to member nodes. The default is 2. The valid range is 1 to 20 (that is, between 100 and 2,000 milliseconds).

Settings	Guidelines
	Note: Although this setting is pushed from the primary node to member nodes, you should initially configure all nodes with the same Detection Interval to prevent inadvertent failover from occurring before the initial synchronization.
Lost Heartbeat Threshold	Number of times a node retries the heartbeat and waits to receive HA heartbeat packets from the other nodes before concluding the other node is down. This part of the configuration is pushed from the primary node to member nodes. Normally, you do not need to change this setting. Exceptions include: • Increase the failure detection threshold if a failure is detected when none has actually occurred. For example, in an active-passive deployment, if the primary node is very busy during peak traffic times, it might not respond to heartbeat packets in time, and a standby node might assume that the primary node has failed. • Decrease the failure detection threshold or detection interval if administrators and HTTP clients have to wait too long before being able to connect through the primary node, resulting in noticeable down time. The valid range is from 1 to 60. Note: Although this setting is pushed from the primary node to member nodes, you should initially configure all nodes with the same HB Lost Threshold to prevent inadvertent failover from occurring before the initial synchronization.
ARP Times	Number of times that the cluster member broadcasts extra address resolution protocol (ARP) packets when it takes on the primary role. (Even though a new NIC has not actually been connected to the network, the member does this to notify the network that a new physical port has become associated with the IP address and virtual MAC of the HA cluster.) This is sometimes called “using gratuitous ARP packets to train the network,” and can occur when the primary node is starting up, or during a failover. Also configure ARP Packet Interval. Normally, you do not need to change this setting. Exceptions include: • Increase the number of times the primary node sends gratuitous ARP packets if an active-passive cluster takes a long time to fail over or to train the network. Sending more gratuitous ARP packets may help the failover to happen faster. • Decrease the number of times the primary node sends gratuitous ARP packets if the cluster has a large number of VLAN interfaces and virtual domains. Because gratuitous ARP packets are broadcast, sending them might generate a large amount of network traffic. As long as the active-passive cluster fails over successfully, you can reduce the number of times gratuitous ARP packets are sent to reduce the amount of traffic produced by a failover. The valid range is 1 to 60. The default is 5.
ARP Interval	Number of seconds to wait between each broadcast of ARP packets. Normally, you do not need to change this setting. Exceptions include: • Decrease the interval if an active-passive cluster takes a long time to fail over or to train the network. Sending ARP packets more frequently may help the failover to happen faster.


Settings	Guidelines
	Increase the interval if the cluster has a large number of VLAN interfaces and virtual domains. Because gratuitous ARP packets are broadcast, sending them might generate a large amount of network traffic. As long as the active-passive cluster fails over successfully, you can increase the interval between when gratuitous ARP packets are sent to reduce the rate of traffic produced by a failover. The valid range is from 1 to 20. The default is 6 seconds.
Data	Set the network interface to be used for data synchronization among cluster nodes. You can configure up to two data ports. If one data port fails, its traffic fails over to the next data port. If all data ports fail, data synchronization traffic fails over to the heartbeat port. If you do not configure a data port, the heartbeat port is used for synchronization. Use the same port numbers for all cluster members. For example, if you select port3 on the primary node, select port3 as the data port interface on the other member nodes.
Remote IP Monitor	Enable or disable active monitoring of remote beacon IP addresses to determine if the network path is available. Note: This option is disabled by default. If enabled, you must specify the Failover Threshold and Failover Hold Time described below.
Failover Threshold	Number of unreachable remote-ip-monitor-list to indicate failure. The default is 5. The valid range is 1-64.
Failover Hold Time	If failover occurs due to a remote IP monitor test, and this node's role changes (to master or slave), it cannot change again until the hold time elapses. The hold time can be used to prevent looping. The default hold time is 120 seconds. The valid range is from 60 to 86400.

Monitoring an HA cluster

You can view HA status from the system dashboard. Go to System > Dashboard > Main and hover over the **Standalone** tab on the top right. **See Detail** will appear. Click it.



HA Status page




You can also use log messages, alert emails, and SNMP to monitor HA events, such as when failover has occurred. The system logs HA node status changes as follows:

- When HA is initialized: HA device Init
- When a member joins a group: Member (FAD2HD3A12000003) join to the HA group
- When the HA configuration is changed from standalone to an active-passive or active-active cluster mode: HA device into Slave mode

The following figure shows FortiADC HA event objects in an SNMP manager.

FortiADC HA event objects in an SNMP manager



Updating firmware for an HA cluster

You can upgrade firmware on all nodes in a cluster from the primary node.

The following process occurs when you perform the HA upgrade procedure:

1. The primary node pushes the firmware image to the member nodes.
2. The primary node notifies the member nodes of the upgrade, and it takes their user traffic during the upgrade.
3. The upgrade command is run on the member nodes, the systems are rebooted, and the member nodes send the primary node an acknowledgment that upgrade has been completed.
4. The upgrade command is run on the primary node, and it reboots. When the system is rebooting, a member node assumes primary status, and the traffic fails over from the former primary node to the new primary node.

After the upgrade process is completed, the system determines whether the original node becomes the primary node, according to the HA Override setting:

- If Override is *enabled*, the cluster considers the Device Priority setting. Both nodes usually make a second failover in order to resume their original roles.
- If Override is *disabled*, the cluster considers uptime first. The original primary node will have a smaller uptime due to the order of reboots during the firmware upgrade. Therefore it will *not* resume its active role; instead, the node with the greatest uptime will remain the new primary node. A second failover will *not* occur.

Reboot times vary by the appliance model, and also by differences between the original firmware version and the firmware version you are installing.

The administrator procedure for an HA cluster is similar to the procedure for installing firmware on a standalone appliance. To ensure minimal interruption of service to clients, use the following steps. The same procedure applies to both active-active and active-passive clusters.



If *downgrading* to a previous version, do *not* use this procedure. The HA daemon on a member node might detect that the primary node has older firmware, and attempt to upgrade it to bring it into sync, undoing your downgrade.

Instead, switch out of HA, downgrade each node individually, then switch them back into HA mode.

Before you begin:

- Download the firmware file from the Fortinet Customer Service & Support website: <https://support.fortinet.com/>
- Read the release notes for the version you plan to install.
- Back up your configuration before beginning this procedure. Reverting to an earlier firmware version could reset settings that are not compatible with the new firmware.
- You must have super user permission (user **admin**) to upgrade firmware.
- Verify that the cluster node members are powered on and available on *all* of the network interfaces that you have configured. If required ports are not available, HA port monitoring could inadvertently trigger an additional failover, resulting in traffic interruption during the firmware update.

To upgrade the firmware for an HA cluster:

1. Log into the web UI of the *primary* node as the `admin` administrator.
2. Go to System > Settings.
3. Click the **Maintenance** tab.
4. Scroll to the **Upgrade Firmware** button.
5. Click **Choose File** to locate and select the file.
6. Enable the HA Sync.
7. Click  to upload the firmware and start the upgrade process.

After the new firmware has been installed, the system reboots.



When you update software, you are also updating the web UI. To ensure the web UI displays the updated pages correctly:

- Clear your browser cache.
- Refresh the page.

In most environments, press Ctrl-F5 to force the browser to get a new copy of the content from the web application. See the Wikipedia article on browser caching issues for a summary of tips for many environments:

https://en.wikipedia.org/wiki/Wikipedia:Bypass_your_cache.

Deploying an active-passive cluster

This topic includes the following information:


- [Overview](#)
- [Basic steps](#)
- [Best practice tips](#)

Overview

In an active-passive cluster, one node is the active appliance; it processes traffic. The other node is passive; it is ready to assume the role of the active appliance if the primary node is unavailable.

You configure the system to send heartbeat packets between the pair to monitor availability. The system continually polls the activity of the heartbeat packets. If the active appliance becomes unresponsive, failover occurs: the standby becomes active. [An active-passive cluster at failover—IP address transfer to the new active member on page 478](#) illustrates the process: (1) the standby node sends gratuitous ARP to notify adjacent routers to direct traffic for the virtual MAC addresses (vMAC) to its network interfaces; (2) It takes the IP addresses of the unresponsive node.

An active-passive cluster at failover—IP address transfer to the new active member



When the former active appliance comes back online, it might or might not assume its former active role. The system selects the active member based on the following criteria:

- Link health (if monitor ports links are down, the node is considered down)
- Remote IP monitor health check results
- Override setting (prefers priority to uptime)
- Most available ports
- Highest uptime value
- Lowest device priority number (1 has greater priority than 2)
- Highest-sorting serial number—Serial numbers are sorted by comparing each character from left to right, where 9 and z are the greatest values. The system gives preference to higher values over lower values.

Basic steps

To deploy an active-passive cluster:

1. License all FortiADC appliances in the HA cluster, and register them, including FortiGuard services, with the Fortinet Customer Service & Support website:
<https://support.fortinet.com/>
2. Physically link the FortiADC appliances that make up the HA cluster.
You must link at least one of their ports (for example, port4 to port4) for heartbeat and synchronization traffic between members of the cluster. You can do either of the following:
 - Connect the two appliances directly with a crossover cable.
 - Link the appliances through a switch. If connected through a switch, the heartbeat interfaces must be reachable by Layer 2 multicast.
3. Configure the secondary node:
 - a. Log into the secondary appliance as the **admin** user.
 - b. Complete the HA settings as described in [Configuring HA settings](#).**Important:** Set the Device Priority to a higher number than the preferred primary node; for example, set it to 2.
4. Configure the primary node:
 - a. Log into the primary appliance as the **admin** user.
 - b. Complete the configuration for all features, as well as the HA configuration.**Important:** Set the Device Priority to a lower number than the secondary node; for example, set it to 1.

Note: After you have saved the HA configuration changes, cluster members join or rejoin the cluster. After you have saved configuration changes on the primary node, it automatically pushes its configuration to the secondary node.

Best practice tips

The following tips are best practices:

- Be careful to maintain the heartbeat link(s). If the heartbeat is accidentally interrupted, such as when a network cable is temporarily disconnected, the other nodes assume that the primary node has failed. In an active-passive deployment, failover occurs. If no failure has actually occurred, both nodes can be operating as the active node simultaneously.
- If you link HA appliances through switches, to improve fault tolerance and reliability, link the ports through two *separate* switches. Also, do *not* connect these switches to your overall network, which could introduce a potential attack point, and could also allow network load to cause latency in the heartbeat, which could cause an unintentional failover.

Deploying an active-active cluster

This topic includes the following information:

- [Configuration overview](#)
- [Basic steps](#)
- [Expected behavior](#)
- [Best practice tips](#)

Configuration overview

[HA active-active deployment on page 482](#) shows an example of an active-active cluster. In an active-active cluster, traffic from the upstream router can be load-balanced among up to eight member nodes.

Load balancing depends on the equal cost multipath (ECMP) configuration on adjacent routers. The routers on either side of the cluster must be configured to use ECMP to distribute traffic to the FortiADC cluster nodes. In the example, assume that the FortiADC configuration includes virtual servers belonging to subnet 10.61.0.0/24. On Router A, you configure equal cost routes as follows:

destination: 10.61.0.0/24 gateway: 10.61.51.1

destination: 10.61.0.0/24 gateway: 10.61.51.2

destination: 10.61.0.0/24 gateway: 10.61.51.3

Likewise, on Router B, you configure equal cost routes for server-to-client traffic:

destination: 0.0.0.0/0 gateway: 10.65.51.1

destination: 0.0.0.0/0 gateway: 10.65.51.2

destination: 0.0.0.0/0 gateway: 10.65.51.3

Active-active clusters also support *failover*. The primary node is the backup node for each of the other nodes in the cluster. If a member node fails, the primary node takes its IP address and sends gratuitous ARP to adjacent routers to direct traffic for that virtual MAC address (vMAC) to its own network interfaces.

The FortiADC configuration involves the following components:

- Primary node system and feature configuration
- Interface configuration (HA node IP list)
- HA configuration

In an active-active cluster, one of the nodes is selected as the *primary node*, and the others are *member nodes*. In this example, node1 is the primary node and node2 and node3 are member nodes. When the cluster is formed, the configuration for node1 is pushed to node2 and node3.

When you configure the network interfaces for nodes in an active-active cluster, in addition to the interface primary IP address, you configure an HA node IP list that specifies special HA IP addresses of each node in the cluster. The HA node IP list for port2 in the example has the following values:

10.61.51.1/16 node1

10.61.51.2/16 node2


10.61.51.3/16 node3

Likewise, the HA node IP list for port3 has the following values:

10.65.51.1/16 node1
 10.65.51.2/16 node2
 10.65.51.3/16 node3

Finally, you log into each node when it is in standalone mode to configure its HA settings. When you are ready to form the cluster, change the setting to HA active-active. The system state changes when a node joins a cluster.

HA active-active deployment



Note: The example shows routers on both sides of the FortiADC cluster. Your deployment might not have a router between the FortiADC cluster and the real server pool. In this case, if your real servers support load balancing methods like ECMP, the expected behavior is the same as what is described here. If not, it is expected that the real servers route reply traffic to the cluster node that sent them the client traffic.

Basic steps

To deploy an active-active cluster:

1. License all FortiADC appliances in the HA cluster, and register them, including FortiGuard services, with the Fortinet Customer Service & Support website: <https://support.fortinet.com/>.
2. Physically link the FortiADC appliances that make up the HA cluster.
You must link at least one of their ports (for example, port4 to port4) for heartbeat and synchronization traffic between members of the cluster. You can do either of the following:
 - If only two nodes, connect the two appliances directly with a crossover cable.
 - If more than two nodes, link the appliances through a switch. If connected through a switch, the interfaces must be reachable by Layer 2 multicast.
3. Configure member nodes:
 - a. Log into the member nodes as the **admin** user.
 - b. Complete the HA configuration as described in [Configuring HA settings](#).
Important: Set the Device Priority to a higher number than the preferred primary node; for example, set it to 2.
4. Configure the preferred primary node:
 - a. Log into the primary node as the **admin** user.
 - b. Configure network interfaces so that each traffic interface has an HA node IP address list in addition to its physical port IP address. See [Configuring network interfaces](#).
When HA is set to standalone, the system uses the physical port IP address. When HA is set to active-active, the system uses the HA node IP address.
 - c. Complete the configuration for all features, as well as the HA configuration.
Important: Set Device Priority to a lower number than the member nodes; for example, set it to 1.

Note: After you have saved the HA configuration changes, cluster members join or rejoin the cluster. After you have saved configuration changes on the primary node, it automatically pushes its configuration to the member nodes.

Expected behavior

In active-active deployments, be sure to enable data synchronization. In particular, enable the following settings:

- Layer 4 Connection Synchronization—Synchronizes TCP connection state data.
- Layer 4 Session Synchronization—Synchronizes the source IP address table used for persistence to backend servers.
- Layer 7 Session Synchronization—Synchronizes the source IP address table used for persistence to backend servers.


The sections that follow describe how the cluster uses synchronized data.

Traffic to TCP virtual servers

When Layer 4 synchronization is enabled, the cluster nodes share TCP connection state and Layer 4 source IP address data for traffic to Layer 4 virtual servers (and Layer 2 TCP and Turbo HTTP virtual servers, which are packet-based). The node that receives the first SYN packet forwards the traffic to the real server, and, at the same time, multicasts the session data to the other nodes in the cluster.


[TCP traffic flow when ECMP results in forwarding through same node on page 484](#) illustrates the sequence of the traffic flow when client-to-server and server-to-client session traffic are routed through the same node.

TCP traffic flow when ECMP results in forwarding through same node



1. Router A uses ECMP to select a cluster node to which to forward a client connection request—in this case, node1.
2. The cluster node forwards the traffic to a real server and multicasts the session data to the cluster via the data port.
3. Router B uses ECMP to select a cluster node to which to forward the server response traffic—also node1.
4. The cluster node forwards the traffic to the client and multicasts the session data to the cluster.


[TCP traffic flow when synchronization has occurred on page 485](#) illustrates the sequence of the traffic flow when client-to-server and server-to-client session traffic are routed through different nodes and synchronization has occurred before the second node receives the response traffic.

TCP traffic flow when synchronization has occurred

1. Router A uses ECMP to select a cluster node to which to forward a client connection request—in this case, node1.
2. The cluster node forwards the traffic to a real server and multicasts the session data to the cluster via the data port.
3. Router B uses ECMP to select a cluster node to which to forward the server response traffic. In this case, it selects node2.
4. If the session has already been synchronized between node1 and node2, node2 forwards the traffic to the client and multicasts the session data to the cluster.

TCP traffic flow when synchronization has not yet occurred on page 485 illustrates the sequence of the traffic flow when client-to-server and server-to-client session traffic are routed through different nodes and synchronization has not yet occurred when the second node receives the response traffic.

TCP traffic flow when synchronization has not yet occurred




1. Router A uses ECMP to select a cluster node to which to forward a client connection request—in this case, node1.
2. The cluster node forwards the traffic to a real server and multicasts the session data.
3. Router B uses ECMP to select a cluster node to which to forward the server response traffic. In this case, it selects node2.
4. Because the session has not yet been synchronized between node1 and node2, node2 multicasts the traffic to the cluster.
5. When node1 receives traffic from node2, it forwards the traffic to the client and multicasts the session data.

Traffic to HTTP virtual servers

When Layer 7 synchronization is enabled, the cluster nodes share source IP data for traffic to HTTP virtual servers differently when the virtual server profile Source option is enabled. When the Source option is enabled, the traffic FortiADC forwards to the real server has the client source IP address; when disabled, it has the FortiADC HA cluster node IP address.

[HTTP traffic flow when the Source profile option is not enabled on page 487](#) illustrates the sequence of the traffic flow when the Source option is not enabled.


HTTP traffic flow when the Source profile option is not enabled



1. Router A uses ECMP to select a cluster node to which to forward a client connection request—in this case, node1.
2. The cluster node forwards the traffic to a real server. Because the Source option was not enabled, the source IP address in the FortiADC-to-real-server traffic is the node1 HA cluster node IP address, and this becomes the destination IP address for the response traffic.
3. Router B does not use ECMP; instead, it forwards the traffic to the node1 HA cluster IP address.
4. The cluster node finds the real client IP address in its session table and forwards the traffic to the client.

[HTTP traffic flow when the Source profile option is enabled on page 488](#) illustrates the sequence of the traffic flow when the Source option is enabled.

HTTP traffic flow when the Source profile option is enabled



1. Router A uses ECMP to select a cluster node to which to forward a client connection request—in this case, node1.

2. The cluster node forwards the traffic to a real server. Because the Source option is enabled, the source IP address in the FortiADC-to-real-server traffic is the true client IP address, and this becomes the destination IP address for the server-to-client traffic.
3. Router B uses ECMP and might forward the traffic to any node in the cluster. In this example, it forwards the traffic to node2.
4. Because the server-to-client response was not expected by node2, it multicasts the traffic to the cluster.
5. When node1 receives the server-to-client response data from node2, it forwards the response to the client.

Note: In an active-active deployment, the virtual server profile Source option adds latency to the transaction. To reduce latency, use an alternative to the Source option, such as the X-Forwarded-For option, if you have a requirement that the original client IP be logged by the real server.

FTP traffic and traffic processed by firewall rules


In an active-active deployment, FTP traffic and firewall traffic are always forwarded through the primary node only.

FTP has both a control connection and a data connection associated with client-server communication. The two “channels” make it difficult to support asymmetric routes in an active-active cluster.

In addition, traffic processed by the stateful firewall rules is also not load-balanced.

[FTP or firewall traffic flow when ECMP selects the primary node on page 489](#) illustrates the sequence of the traffic flow when ECMP results in traffic being forwarded through the primary node.


FTP or firewall traffic flow when ECMP selects the primary node



1. Router A uses ECMP to select a cluster node to which to forward a client connection request. In this case, it selects the primary node, node1.
2. The primary node forwards the traffic to a real server.
3. Router B uses ECMP to select a cluster node to which to forward the server response traffic—also node1.
4. The primary node forwards the traffic to the client.

[FTP or firewall traffic flow when ECMP results in an asymmetric route on page 490](#) illustrates the sequence of the traffic flow when ECMP results in an asymmetric route.


FTP or firewall traffic flow when ECMP results in an asymmetric route



1. Router A uses ECMP to select a cluster node to which to forward a client connection request. In this case, it selects the primary node, node1.
2. The cluster node forwards the traffic to a real server.
3. Router B uses ECMP to select a cluster node to which to forward the server response traffic—in this case, node2.
4. Because the server-to-client response was not expected by node2, it forwards traffic to the cluster.
5. When the primary node receives traffic from node2, it forwards it to the client.

FTP or firewall traffic flow when ECMP results in traffic sent to a non-primary node on page 491 illustrates the sequence of the traffic flow when ECMP results in client-to-server traffic sent to a non-primary node.

FTP or firewall traffic flow when ECMP results in traffic sent to a non-primary node



1. Router A uses ECMP to select a cluster node to which to forward a client connection request to a real server destination IP address. In this case, it selects a member node, node3.
2. Firewall traffic is forwarded by the primary node only, so node3 multicasts the session data to the cluster.
3. The primary node forwards the traffic to a real server.
4. Router B uses ECMP to select a cluster node to which to forward the server response traffic—in this case, node2.
5. Because the server-to-client response was not expected by node2, it forwards traffic to the cluster.
6. When the primary node receives traffic from node2, it forwards it to the client.

Best practice tips

The following tips are best practices:

- Be careful to maintain the heartbeat link(s). If the heartbeat is accidentally interrupted, such as when a network cable is temporarily disconnected, the other nodes assume that the primary node has failed. In an active-active deployment, a new primary node is elected among member nodes. If no failure has actually occurred, both nodes can be operating as primary nodes simultaneously.
- If you link HA appliances through switches, to improve fault tolerance and reliability, link the ports through two *separate* switches. Also, do *not* connect these switches to your overall network, which could introduce a potential attack point, and could also allow network load to cause latency in the heartbeat, which could cause an unintentional failover.

Advantages of HA Active-Active-VRRP

Compared with HA Active-Passive or Active-Active clusters, an HA Active-Active-VRRP cluster offers the following advantages:

- The HA Active-Active mode is an device-based HA mode, in which the HA fail over will switch over the whole failed device even in cases where only one monitor port fails.
- In FortiADC HA Active-Active-VRRP mode, you can manually assign a virtual server to a traffic group, enabling you to do traffic load design based on virtual servers.
- In HA Active-Active-VRRP mode, FortiADC only synchronizes the session table/persistence table to the next available device in the same traffic group using the “failover-order” command. In cases where you have more than two devices in the cluster, this synchronization mechanism can turn out to be more efficient than HA Active-Passive or Active-Active mode because the session/persistence table will be synced to the whole HA group. In this sense, FortiADC actually supports the N+M hot-backup function.
- HA Active-Active mode must work together with an external router with the ECMP route configured to distribute traffic to different Active-Active nodes; HA Active-Active-VRRP mode does not need this external router to do ECMP traffic distribution — Both sides can simply point their respective gateway to the VRRP floating IP.
- In HA Active-Active-VRRP mode, different devices in the same traffic group have the same HA status. Once you have pointed both the client and the server side gateways to the floating IP in the same traffic, the incoming/outgoing traffic will going to the same device. As a result, HA Active-Active-VRRP mode doesn't need to multicast the traffic itself to the HA group, which should offer the best network performance and efficiency.
- In HA Active-Active mode, the AA-Master will take over all AA-NotWorking nodes' traffic. If multiple AA devices have failed, the AA-Master will have to process much more traffic than the AA-Slave nodes, which may exhibit some unexpected behavior under abnormal high traffic stress.
- In terms of sync session, you are unable to access the real server's IP address from the client directly in HA Active-Active mode, but you don't have this limitation in HA Active-Active-VRRP mode.

Deploying an active-active-VRRP cluster

This topic includes the following information:

- [Configuration overview](#)
- [Basic steps](#)
- [Best practice tips](#)

Configuration overview

The Virtual Router Redundancy Protocol (VRRP) is designed to eliminate the single point of failure inherent in the static default routed environment. VRRP specifies an election protocol that dynamically assigns responsibility for a virtual router to one of the VRRP routers on a LAN. The VRRP router controlling the IP address(es) associated with a virtual router is called the Master, and forwards packets sent to these IP addresses. The election process provides dynamic fail-over in the forwarding responsibility should the Master become unavailable. Any of the virtual router's IP addresses on a LAN can then be used as the default first hop router by end-hosts. The advantage of VRRP is a higher availability default path without requiring configuration of dynamic routing or router discovery protocols on every end-host.

A virtual router is defined by its virtual router identifier (VRID) and a set of IP addresses. A VRRP router may associate a virtual router with its real address on an interface, and may also be configured with additional virtual router mappings and priority that the virtual router can back up. The mapping between VRID and addresses must be coordinated among all VRRP routers on a LAN.

FortiADC only adopts the VRRP concept, but not the exact VRRP protocol itself. For this reason, its HA Active-Active VRRP mode can only be called a VRRP-like HA mode.

VRRP configurations can be used as a high availability (HA) solution to ensure that your network maintains connectivity with the Internet (or with other networks) even if the default router for your network fails. Using VRRP, you can assign VRRP routers as master or backup routers. The master router processes traffic, while the backup routers monitor the master router and start forwarding traffic the moment the master router fails.

VRRP is described in RFC 3768.

FortiADC 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. In a VRRP configuration, when a FortiADC unit operating as the master unit fails, a backup unit automatically takes its place and continues processing network traffic. In such a situation, all traffic to the failed unit transparently fails over to the backup unit that takes over the role of the failed master FortiADC unit. When the failed FortiADC unit is restored, it will once again take over processing traffic for the network. See [An active-active-VRRP cluster configuration using two FortiADC units on page 494](#).

An active-active-VRRP cluster configuration using two FortiADC units

In an active-active-VRRP cluster, one of the nodes is selected as the primary node of a traffic group, and the rest of the nodes are member nodes of the traffic group. Traffic from the upstream can be load-balanced among up to eight member nodes. Active-active-VRRP clusters also support failover. If the primary node fails, the traffic group work on this node will fail over to one of the backup nodes which will send gratuitous ARP to adjacent devices to redirect traffic for its own MAC address to all network interfaces within the traffic group.

The FortiADC VRRP configuration involves the following:

- Traffic group and their features
- Interface and virtual server (pertinent floating IP and traffic group)
- HA

Note: It is important to note that FortiADC only supports VRRP configuration between two or more FortiADC units. It can NOT be integrated into a VRRP group formed with any third-party VRRP devices.

Basic steps

To deploy an active-active-VRRP cluster:

1. Configure the HA active-active-VRRP cluster.

<https://support.fortinet.com/>

For example:

```
config system ha
set mode active-active-vrrp
set hbdev port2
set group-id 14
set local-node-id 1
end
```

2. Configure the traffic group.

Configure the traffic group and set its parameters. The failover sequence must be configured according to the order of node IDs. This means that if a node is dead, the next node in queue will take over handling the traffic. If you want to decide when a node should retake the traffic over from power-down to start-up, you can enable the preempt.

If only two nodes, connect the two appliances directly with a crossover cable.

If more than two nodes, link the appliances through a switch. If connected through a switch, the interfaces must be reachable by Layer 2 multicast.

```
config system traffic-group
edit "traffic-group-1"
set failover-order 1 2
next
end
```

3. Configure applications and relate them with the traffic group

Relate applications with the traffic group in the virtual server configuration and interface + IP configuration. If no traffic group is related, the "default" traffic group will be used.

For example (Relate a virtual server to a traffic group):

```
config load-balance virtual-server
edit "vs1"
set packet-forwarding-method FullNAT
set interface port1
set ip 10.128.3.4
set load-balance-profile LB_PROF_HTTP
set load-balance-method LB_METHOD_DEST_IP_HASH
set load-balance-pool rs1
set ippool-list vs1-pool vs1-pool-1
set traffic-group traffic-group-1
next
```

For example (Relate an interface and IP address with a traffic group):

```
edit "port1"
set vdom root
set ip 10.128.3.1/16
set allowaccess https ping ssh snmp http telnet
set traffic-group traffic-group-1
set floating enable
set floating-ip 10.128.3.3
next
end
```

Best practice tips

The following tips are best practices:

Note: After you have saved the HA configuration changes, cluster members join or rejoin the cluster. After you have saved configuration changes on the primary node, it automatically pushes its configuration to the member nodes.

Chapter 16: Virtual Domains

This chapter includes the following topics:

- [Virtual domain basics on page 497](#)
- [Enabling the virtual domain feature on page 497](#)
- [Creating a virtual domain on page 498](#)
- [Assigning network interfaces and admin users to VDOMs on page 498](#)
- [Virtual domain policies on page 499](#)
- [Disabling a virtual domain on page 501](#)

Virtual domain basics

A virtual domain (VDOM) is a complete FortiADC instance that runs on the FortiADC platform. The VDOM feature supports multi-tenant deployments. To do this, you create a virtual domain configuration object that contains all of the system and feature configuration options of a full FortiADC instance, and you provision an administrator account with privileges to access and manage only that VDOM.

Note: The super user **admin** can access all VDOMs that have been created on the system, but the administrator accounts that are provisioned for a VDOM can access only that particular VDOM.

To use the VDOM feature, complete the following steps:

1. Enable the virtual domain feature.
2. Create a virtual domain configuration object.
3. Assign network interfaces and administrators to the virtual domain.

Enabling the virtual domain feature

You can use the web UI to enable the virtual domain feature. By default, the virtual domain feature is not enabled, and the GUI for virtual domain configuration is hidden.

Before you begin:

- You must have super user permission (user **admin**) to enable the virtual domain feature.

To enable the virtual domain feature:

1. Go to System > Settings.
The configuration page displays the **Basic** tab.
2. Click edit, on the top right.
3. Enable **Virtual Domain**, a button on the far lower left.
4. Save the configuration.

Note: You can also enable the virtual domain feature from under the Dashboard > Status tab.

[Super admin login with virtual domain on page 498](#) shows the landing page after the **admin** administrator logs into the system when the virtual domain feature is enabled. From here, the **admin** administrator can create virtual domains, assign network interfaces to virtual domains, create admin users for virtual domains, and navigate to the system and feature configuration pages for the virtual domains, including the root (default) domain.

When a non-admin user with a delegated administrator account logs in, the landing page is the standard landing page. Such users cannot perform the tasks related to virtual domain administration that the **admin** administrator performs.

Super admin login with virtual domain



Creating a virtual domain

By default, FortiADC has a predefined virtual domain named root that you cannot delete or modify. The **admin** user can add, delete, enable, and disable virtual domains.

Before you begin:

- You must have super user permission (user **admin**) to create virtual domains.
- You must have super user permission (user **admin**) to assign network interfaces to virtual domains.

To create a virtual domain:

1. Go to Virtual Domain.
2. Click **Create New**, enter a unique name for the virtual domain.
3. Save the configuration.

Assigning network interfaces and admin users to VDOMs

By default, all network interfaces are assigned to the root virtual domain. After you have created the virtual domain, you can assign network interfaces to it.

To assign a network interface to a virtual domain:

1. Go to Networking > Interface.
2. Double-click an interface configuration or click **Create New** to create one.
3. Configure interface settings and select the virtual domain.
4. Save the configuration.

When virtual domain administrators log into the FortiADC system, they only see configuration settings and data for the virtual domain that you assigned them to. They do not see the Virtual Domains menu in the navigation pane.

To create an administrator for a virtual domain:

1. Go to System > Administrator.
2. Click **Create New** to create an administrator.
3. Configure administrator settings and select the virtual domain.
4. Save the configuration.

Virtual domain policies

FortiADC allows you to create and impose custom policies or restrictions on each virtual domain you have added. For each virtual domain, you can configure the maximum range for its Dynamic Resources and Static Resources. Dynamic Resources are related to a virtual domain's performance, while Static Resources are related to its configuration. The Vdom configuration dialog ([Vdom configuration on page 499](#)) also shows a virtual domain's current configuration and workload settings, which serve as good reference points for you to fine-tune the virtual domain.

Vdom configuration

Vdom

Name: root

Dynamic Resources

	Current	Max	Range
L4 CPS	0	0	Range: 0-1000000
L7 CPS	0	0	Range: 0-1000000 per second
L7 RPS	0	0	Range: 0-1000000 per second
SSL CPS	0	0	Range: 0-1000000
SSL Throughput	0	0	Range: 0-10000000 Kbps
Concurrent Session	0	0	Range: 0-1000000
Inbound	0	0	Range: 0-40000000 Kbps
Outbound	0	0	Range: 0-40000000 Kbps

Static Resources

	Current	Max	Range
Virtual Server	8	0	Range: 0-1024
Real Server	7	0	Range: 0-1024
Health Check	5	0	Range: 0-1024
Source Pool	2	0	Range: 0-1024
Error Page	0	0	Range: 0-1024
Local User	0	0	Range: 0-1024
User Group	0	0	Range: 0-1024

Save Cancel Refresh

VDOM configuration parameters

Parameter	Description
Dynamic Resources	
L4 CPS	Shows the L4 CPS data transfer rate in kilobyte per second (kB/s) at the last page refresh. Note: You can set the VDOM's maximum L4 CPS data transfer rate by specifying a desired value in the box. Valid values range from 0 to 1,000,000.
L7 CPS	Shows the L7 CPS data transfer rate in kilobyte per second (kB/s) at the last page refresh. Note: You can set the VDOM's maximum L7 CPS data transfer rate by specifying a desired value in the box. Valid values range from 0 to 1,000,000.
L7 RPS	Shows the L7 RPS data transfer rate in kilobyte per second (kB/s) at the last page refresh. Note: You can set the VDOM's maximum L7 RPS data transfer rate by specifying a desired value in the box. Valid values range from 0 to 1,000,000.
SSL CPS	Shows the SSL CPS data transfer rate in kilobyte per second (kB/s) at the last page refresh. Note: You can set the VDOM's maximum SSL CPS data transfer rate by specifying a desired value in the box. Valid values range from 0 to 1,000,000.
SSL Throughput	Shows the SSL throughput rate in kilobyte per second (kB/s) at the last page refresh. Note: You can set the VDOM's maximum SSL throughput rate by specifying a desired value in the box. Valid values range from 0 to 1,000,000.
Concurrent Session	Shows the number of concurrent sessions at the last page refresh. Note: You can set the VDOM's maximum number of concurrent sessions by specifying a desired value in the box. Valid values range from 0 to 1,000,000.
Inbound	Shows the inbound TCP data transfer rate in kilobyte per second (kB/s) at the last page refresh. Note: You can set the VDOM's maximum inbound <i>TCP</i> data transfer rate by specifying a desired value in the box. Valid values range from 0 to 4,000,000.
Outbound	Shows the outbound TCP data transfer rate in kilobyte per second (kB/s) at the last page refresh. Note: You can set the VDOM's maximum outbound <i>TCP</i> data transfer rate by specifying a desired value in the box. Valid values range from 0 to 4,000,000.
Static Resources	
Virtual Server	Shows the number of virtual servers at the last page refresh. Note: You can set the maximum number of virtual servers that can be configured on this VDOM by specifying a desired value in the box. Valid values range from 0 to 1024.
Real Server	Shows the number of real servers at the last page refresh.

Parameter	Description
	Note: You can set the maximum number of real servers that can be configured on this VDOM by specifying a desired value in the box. Valid values range from 0 to 1024.
Health Check	Shows the number of health check objects at the last page refresh. Note: You can set the maximum number of health check objects that can be configured on this VDOM by specifying a desired value in the box. Valid values range from 0 to 1024.
Source Pool	Shows the number of source pools at the last page refresh. Note: You can set the maximum number of source pools that can be configured on this VDOM by specifying a desired value in the box. Valid values range from 0 to 1024.
Error Page	Shows the number of error pages at the last page refresh. Note: You can set the maximum number of error pages that can be configured on this VDOM by specifying a desired value in the box. Valid values range from 0 to 1024.
Local User	Shows the number of local users at the last page refresh. Note: You can set the maximum number of local users that can be configured on this VDOM by specifying a desired value in the box. Valid values range from 0 to 1024.
User Group	Shows the number of user groups at the last page refresh. Note: You can set the maximum number of user groups that can be configured on this VDOM by specifying a desired value in the box. Valid values range from 0 to 1024.

Disabling a virtual domain

To disable the virtual domain feature:

1. Go to System > Settings > **Basic**.
2. Click the edit function on the far top right, opening up the Basic dialogue.
3. On the far left bottom, you can disable the Virtual Domain.
4. The ADC should refresh on its own.

Chapter 17: SSL Transactions

This chapter includes the following topics:


- [SSL offloading on page 502](#)
- [SSL decryption by forward proxy on page 504](#)
- [SSL profile configurations on page 508](#)
- [Certificate guidelines on page 513](#)
- [SSL/TLS versions and cipher suites on page 513](#)
- [Exceptions list on page 517](#)
- [SSL traffic mirroring on page 517](#)

SSL offloading

You can use FortiADC in a Layer-7 load-balancing topology to offload SSL decryption from the real server farm, as illustrated in [SSL offloading on page 503](#). In such a deployment, the FortiADC unit uses a copy of the real server certificate and its private key to negotiate the SSL connection. It acts as an SSL proxy for the servers, using the certificates and their private keys to:

- authenticate itself to clients
- decrypt requests
- encrypt responses

When session data has been decrypted, you can use the FortiADC content rewriting, content routing, and web application firewall features.

SSL offloading

FortiADC forwards data unencrypted to the servers, and the servers can maximize performance because they are processing HTTP and not HTTPS transactions.

To realize the benefits of SSL offloading and maintain security, you must deploy the FortiADC appliance in a trusted network with a direct path to the real servers so that the connection between the FortiADC and the real server does not have to be re-encrypted. For example, you connect FortiADC and the real servers through the same switch, and all are physically located on the same locked rack.

In cases where traffic is forwarded along untrusted paths toward the real servers, you can use a real server SSL profile to re-encrypt the data before forwarding it to the real servers.

Basic steps:

1. Import the X.509 v3 server certificates and their private keys that ordinarily belong to the backend servers, as well as any certificate authority (CA) or intermediate CA certificates that are used to complete the chain of trust between your clients and servers.
2. Configure a local certificate group that includes the server's local certificate and the Intermediate CA group that contains the Intermediate CAs.
3. Configure an application profile and a client SSL profile (if needed) that reference the local certificate group and specify the allowed SSL/TLS versions and list of SSL ciphers that can be used for the SSL connection between the client and the FortiADC unit. Select this profile when you configure the virtual server.
4. Configure a real server SSL profile that enables or disables SSL for the connection between the FortiADC unit and the real server. If enabled, specify the SSL/TLS versions and the list of SSL ciphers that can be used. Select this profile when you configure the real server pool.

SSL decryption by forward proxy

You can use SSL decryption by forward proxy in cases where you cannot copy the server certificate and its private key to the FortiADC unit because it is either impractical or impossible (in the case of outbound traffic to unknown Internet servers).

When SSL forward proxy is enabled, FortiADC becomes a proxy to both sides of the connection. The server certificate and its private key used to negotiate the SSL connection with the client are dynamically derived from the certificate presented by the real server and optionally chained with an Intermediate CA trusted by the client.


Basic steps:

1. Import a special Intermediate CA and its private key to the local certificate store that you have provisioned for SSL forward proxy operations.
2. Configure an Intermediate CA group. (Optional)
3. Configure a certificate caching object (or use the pre-defined one).
4. Configure a client SSL profile that enables SSL proxy, references the local certificate, and specifies the allowed SSL/TLS versions and list of SSL ciphers that can be used for the SSL connection between the client and the FortiADC unit. Select this profile when you configure the virtual server.
5. Configure all settings required for backend SSL.

Layer-7 deployments

[Layer 7 SSL decryption by forward proxy on page 505](#) illustrates a Layer 7 SSL forward proxy deployment similar to the SSL offloading example—inbound traffic to your server farm. When the FortiADC virtual server receives the ClientHello message, it selects a real server and sends its own ClientHello to the server to set up its own SSL session with it (represented by the dashed line in the figure). FortiADC uses the certificate presented by the server to derive the certificate to present to the client. This derived certificate is signed by an Intermediate CA that is trusted by the client, so the client completes its handshake with the FortiADC, and FortiADC can decrypt the traffic.

Layer 7 SSL decryption by forward proxy



[Layer 7 SSL decryption methods on page 505](#) summarizes the pros and cons of Layer 7 SSL decryption methods.

Layer 7 SSL decryption methods

Method	Pros	Cons
SSL offloading	Better performance. No feature limitations. In most cases, you do not need to maintain SSL functionality (certificates and keys, SSL ports) on the real servers.	You must be able to copy the local certificates and private keys from the real servers.

Method	Pros	Cons
SSL forward proxy	You do not need to copy the local certificates and keys from the real servers. Instead, you add only one Intermediate CA and private key to be used for all the HTTPS servers.	Performance cost associate with SSL proxy operations and certificate re-signing. You need to maintain SSL functionality on the real servers. Incompatible with some features because the server must be selected before the client request is decrypted: Incompatible features include: • Some load balancing methods (only Round Robin and Least Connection are supported) • Some persistence methods (only Source Address, Source Address Hash, Source Address-Port Hash, and SSL Session ID are supported) • Client SNI Required option • Content routing

Layer 2 deployments


You can use FortiADC in a Layer 2 sandwich topology to offload SSL decryption tasks from FortiGate.

[Layer 2 SSL decryption by forward proxy on page 506](#) shows the topology. To decrypt traffic to and from external HTTPS destinations, you must use SSL forward proxy.

When the FortiADC virtual server receives the ClientHello message, it sends its own ClientHello to the destination server in order to fetch the server certificate so that it can be manipulated. The FortiGate and second FortiADC in the network path must be configured to pass-through this HTTPS traffic. FortiADC uses the server certificate to derive a certificate to present to the client. This derived certificate is signed by an Intermediate CA that is trusted by the client, so the client completes its handshake with the first FortiADC, and FortiADC decrypts the traffic.

In a sandwich deployment like this one, you do not want to re-encrypt the traffic until it egresses the second FortiADC. You control server-side SSL with the real server SSL profile configuration, discussed next.

Layer 2 SSL decryption by forward proxy




SSL profile configurations

The application profile and client SSL profile determine the settings` for the client-FortiADC connection; the real server SSL profile determines settings for the FortiADC-real server connection. This granularity gives you flexibility in how you leverage FortiADC's SSL transaction capabilities. For example, in the case of SSL offloading, your goal is to eliminate SSL transactions on the real servers so that you can configure a server-side SSL profile that does not use SSL. Or it could be the case that the back-end real servers support only SSLv2, but you want to use the more secure TLSv1.2 for the client-FortiADC segment.

[SSL profiles on page 508](#) illustrates the basic idea of client-side and server-side profiles.

SSL profiles




The call-outs in [Layer 2 sandwich profiles on page 511](#) have guidance for the two types of profiles used in a Layer 2 sandwich deployment.

In this deployment, the FortiADC 1 virtual server is of a Layer-2 HTTPS virtual server configuration. Its client SSL profile supports SSL forward proxy, including the special local signing CA. For Layer-2 virtual servers, the "real server" target is the next hop. In this case, the real server target is the FortiGate pool. Because SSL is not enabled in the real server SSL profile, FortiADC 1 does not re-encrypt the SSL connection. (However, you can configure allowed SSL versions and

ciphers in the client SSL profile, and you can also configure an SSL certificate verification policy to enforce rules and checks on the destination server certificate.) The client SSL profile settings are used when re-encrypting the server response traffic in the return segment to the client.

The FortiADC 2 virtual server is a Layer 2 HTTP virtual server configuration. It receives unencrypted traffic from FortiGate. Its server pool is the next hop gateway. On its server side, FortiADC uses the real server SSL profile settings when it encrypts the outbound SSL connection and decrypts the inbound response traffic.

Layer 2 sandwich profiles



For information on virtual server profile configuration objects, see [Configuring Application profiles](#).

For information on real server SSL configuration objects, see [Configuring real server SSL profiles](#).

Certificate guidelines

When a client browser requests an HTTPS connection to a web server, the server presents a server certificate to the client for verification. The client checks the content of the certificate against a local browser database of Certificate Authorities, and if it finds a match, the connection is made. If no match is found, the browser displays a warning that asks if you want to continue with the connection.

To avoid this warning, you must upload an Intermediate CA signed by one of the CA vendors that has its root certificates preinstalled in the web browsers. When the vendor issues you a local server certificate for your website, it typically includes the Intermediate CAs in your package.

For SSL offloading deployments, you create a local certificate group that references the local certificate for the server and its Intermediate CA group (a group that references all Intermediate CAs the vendor provided with your certificate package).

For SSL decryption by forward proxy deployments, you create a local certificate group that references any local certificate and an Intermediate CA group that includes the Intermediate CA and private key configuration you have provisioned for the SSL forward proxy operations.



You are not required to obtain SSL certificates from SSL vendors. You can use an enterprise certificate server (like Microsoft CertSrv) or open-source tools like OpenSSL or to generate them. Note, however, that a web browser will not trust the certificate unless it is associated with a certificate installed in the browser. If you use your own tools to generate the Intermediate CA, you must distribute that certificate to client browsers in whatever manner you typically do that—automatic update package from IT, manual distribution, and so on.

For information on importing certificates and configuring certificate configuration objects, see [Manage and validate certificates](#).

SSL/TLS versions and cipher suites

An SSL cipher is an algorithm that performs encryption and decryption. It transforms plain text into a coded set of data (cipher text) that is not reversible without a key. During the SSL handshake phase of the connection, the client sends a list of the ciphers it supports. FortiADC examines the client cipher list in the order it is specified, chooses the first cipher that matches a cipher specified in the virtual server configuration, and responds to the client. If none of the ciphers offered by the client are in the cipher suite list for the virtual server, the SSL handshake fails.

To see the list of ciphers supported by the browser you are using, go to a link maintained by the Leibniz University of Hannover Distributed Computing & Security (DCSec) Research Group:

<https://cc.dcsec.uni-hannover.de/>

FortiADC SLB profiles support a specific list of [RSA ciphers](#), [PFS ciphers](#), [ECDHE ciphers](#), [ECDSA ciphers](#), and [eNull ciphers](#).

[Cipher suites with RSA key exchange on page 514](#) lists supported RSA ciphers.

Cipher suites with RSA key exchange

Abbreviation	Cipher Suite	Protocol	Kx	Au	Enc	MAC
AES256-GCM-SHA384	TLS_RSA_WITH_AES_256_GCM_SHA384	TLS 1.2	RSA	RSA	AESGCM (256)	AEAD
AES256-SHA256	TLS_RSA_WITH_AES_256_CBC_SHA256	TLS 1.2	RSA	RSA	AES(256)	SHA
AES256-SHA	TLS_RSA_WITH_AES_256_CBC_SHA	SSL 3.0 TLS 1.2, 1.1, 1.0	RSA	RSA	AES(256)	SHA
AES128-GCM-SHA256	TLS_RSA_WITH_AES_128_GCM_SHA256	TLS 1.2	RSA	RSA	AESGCM (128)	AEAD
AES128-SHA256	TLS_RSA_WITH_AES_128_CBC_SHA256	TLS 1.2	RSA	RSA	AES(128)	SHA
AES128-SHA	TLS_RSA_WITH_AES_128_CBC_SHA	SSL 3.0 TLS 1.2, 1.1, 1.0	RSA	RSA	AES(128)	SHA
RC4-SHA	SSL_RSA_WITH_RC4_128_SHA	SSL 3.0	RSA	RSA	RC4	SHA
	TLS_RSA_WITH_RC4_128_SHA	TLS 1.2, 1.1, 1.0	RSA	RSA	RC4	SHA
RC4-MD5	SSL_RSA_WITH_RC4_128_MD5	SSL 3.0	RSA	RSA	RC4	MD5
	TLS_RSA_WITH_RC4_128_MD5	TLS 1.2, 1.1, 1.0	RSA	RSA	RC4	MD5
DES-CBC3-SHA	SSL_RSA_WITH_3DES_EDE_CBC_SHA	SSL 3.0	RSA	RSA	DES-CBC3	SHA
	TLS_RSA_WITH_3DES_EDE_CBC_SHA	TLS 1.2, 1.1, 1.0	RSA	RSA	DES-CBC3	SHA

With RSA ciphers, the server's public RSA key is part of the server certificate and is typically very long lived. It is not uncommon for the same public key to be used for months or years. This creates a potential problem: if an SSL server's private key were to be leaked or stolen, all connections made in the past using that key would be vulnerable. If someone has recorded your SSL connections, they can use the stolen private key to decrypt them.

[Cipher suites with DHE/EDH key exchange on page 515](#) lists supported Perfect Forward Secrecy (PFS) ciphers with DHE/EDH key exchange. With PFS, a fresh public key is created for every single connection. That means that an adversary would need to break the key for each connection individually to read the communication.

Cipher suites with DHE/EDH key exchange

Abbreviation	Cipher Suite	Protocol	Kx	Au	Enc	MAC
DHE-RSA-AES256-GCM-SHA384	TLS_DHE_RSA_WITH_AES_256_GCM_SHA384	TLS 1.2	DH	RSA	AES256	SHA384
DHE-RSA-AES256-SHA256	TLS_DHE_RSA_WITH_AES_256_CBC_SHA256	TLS 1.2	DH	RSA	AES256	SHA256
DHE-RSA-AES256-SHA	TLS_DHE_RSA_WITH_AES_256_CBC_SHA	SSL 3.0 TLS 1.2, 1.1, 1.0	DH	RSA	AES256	SHA256
DHE-RSA-AES128-GCM-SHA256	TLS_DHE_RSA_WITH_AES_128_GCM_SHA256	TLS 1.2	DH	RSA	AES128	SHA256
DHE-RSA-AES128-SHA256	TLS_DHE_RSA_WITH_AES_128_CBC_SHA256	TLS 1.2	DH	RSA	AES128	SHA256
DHE-RSA-AES128-SHA	TLS_DHE_RSA_WITH_AES_128_CBC_SHA	SSL 3.0 TLS 1.2, 1.1, 1.0	DH	RSA	AES128	SHA
EDH-RSA-DES-CBC3-SHA	TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA	SSL 3.0 TLS 1.2, 1.1, 1.0	DH	RSA	3DES	SHA

Cipher suites with ECDHE key exchange on page 515 lists supported PFS ciphers with Elliptic curve Diffie–Hellman Ephemeral key (ECDHE) key exchange. ECDHE is significantly faster than DHE. The supported suites include both the Elliptic Curve Digital Signature Algorithm (ECDSA) and RSA key authentication (Au) algorithms.

Cipher suites with ECDHE key exchange

Abbreviation	Cipher Suite	Protocol	Kx	Au	Enc	MAC
ECDHE-ECDSA-AES256-GCM-SHA384	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384	TLS 1.2	ECDH	ECDSA	AESGCM256	AEAD
ECDHE-ECDSA-AES256-SHA384	TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384	TLSv1.2	ECDH	ECDSA	AES256	SHA384
ECDHE-ECDSA-AES256-SHA	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	SSL 3.0 TLS 1.2, 1.1, 1.0	ECDH	ECDSA	AES256	SHA
ECDHE-ECDSA-AES128-GCM-SHA256	TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	TLSv1.2	ECDH	ECDSA	AESGCM128	AEAD
ECDHE-ECDSA-AES128-SHA256	TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256	TLSv1.2	ECDH	ECDSA	AES128	SHA256

Abbreviation	Cipher Suite	Protocol	Kx	Au	Enc	MAC
ECDHE-ECDSA-AES128-SHA	TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA	SSL 3.0 TLS 1.2, 1.1, 1.0	ECDH	ECDSA	AES128	SHA
ECDHE-ECDSA-RC4-SHA	TLS_ECDHE_ECDSA_WITH_RC4_128_SHA	SSL 3.0 TLS 1.2, 1.1, 1.0	ECDH	ECDSA	RC4	SHA
ECDHE-ECDSA-DES-CBC3-SHA	TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA	SSL 3.0 TLS 1.2, 1.1, 1.0	ECDH	ECDSA	3DES	SHA
ECDHE-RSA-AES256-GCM-SHA384	TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	TLS 1.2	ECDH	RSA	AESGCM256	AEAD
ECDHE-RSA-AES256-SHA384	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384	TLS 1.2	ECDH	RSA	AES256	SHA384
ECDHE-RSA-AES256-SHA	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	TLS 1.2	ECDH	RSA	AES256	SHA
ECDHE-RSA-AES128-GCM-SHA256	TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	TLS 1.2	ECDH	RSA	AESGCM128	AEAD
ECDHE-RSA-AES128-SHA256	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256	TLS 1.2	ECDH	RSA	AES128	SHA256
ECDHE-RSA-AES128-SHA	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA	SSL 3.0	ECDH	RSA	AES128	SHA
ECDHE-RSA-RC4-SHA	TLS_ECDHE_RSA_WITH_RC4_128_SHA	SSL 3.0	ECDH	RSA	RC4	SHA
ECDHE-RSA-DES-CBC3-SHA	TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA	SSL 3.0	ECDH	RSA	3DES	SHA



Profiles support TLS_AES_128_GCM_SHA256 and TLS_AES_256_GCM_SHA384 for TLSv1.3. They will be set automatically when TLSv1.3 is selected in ssl version.

In addition, profiles support an eNull cipher option. This option represents all cipher suites that do not apply encryption to the application data (integrity check is still applied). The exact cipher suite used depends on the SSL/TLS version used. As an example, in SSL v3.0, eNULL includes NULL-MD5, NULL-SHA, ECDH-RSA-NULL-SHA, ECDH-ECDSA-NULL-SHA, and some other non-encryption cipher suites.

Finally, profiles support a user-specified cipher list. You can specify a colon-separated list of OpenSSL cipher suite short names. The names are validated against the form of the cipher suite short names published on the OpenSSL website:

<https://www.openssl.org/docs/manmaster/apps/ciphers.html>

Exceptions list

In some jurisdictions, SSL interception and decryption by forward proxy is disfavored for some types of websites or disallowed entirely. If necessary, you can use the L2 Exception List configuration to define destinations that should not have its sessions decrypted. You can leverage FortiGuard web filter categories, and you can configure a list of additional destinations.

You associate the L2 Exception List configuration with virtual servers that are in the path of outbound traffic. The virtual server evaluates whether an exception applies before processing the initial SSL client hello. If an exception applies, that connection is passed through, and it is not decrypted.

For information on creating the configuration, see [Configuring an L2 exception list](#).

SSL traffic mirroring

FortiADC supports mirroring packets (HTTPS/TCPs) to specified network interfaces. When the feature is enabled, SSL traffic will be mirrored to the specified ports by the virtual server after it has been decrypted. See the following figures.

The feature supports both IPv4 and IPv6. FortiADC can send traffic to up to four outgoing interfaces, including aggregated and VLAN interfaces. Mirrored traffic is transmitted as a single packet stream, using the original client-side source and destination IP address and port numbers. The source and destination MAC addresses are 0 (zero) in mirrored traffic. The feature requires a virtual server set to Layer 7 or Layer 2, with a profile configured for HTTPS or TCPS. It is supported on all FortiADC platforms.

To configure SSL traffic mirroring

1. Go to Virtual Server. Go to the far right and click **Create New**. You have to click Advanced Mode if you want traffic mirroring.
2. In the **Basic** tab, go to **Type**, and set it to Layer 7.
3. Then go to the **General** tab. Go under **Resources** to Profile.
4. Select either LB_PROF_HTTPS (**not** just HTTP, without the 's') or LB_PROF_TCPS
5. When you do this, **SSL Traffic Mirror** will appear as a tab to the right of General.
6. Go to SSL Traffic Mirror and **enable** it.
7. Click **Save**.
8. Click **Create New**. Two options will drop down: Basic and Advanced.
9. Select **Advanced**.

Virtual Server								Content Rewriting	Content Routing	NAT Source Pool	Schedule Pool	Clone Pool	
Add Filter								Create New					
	Name	Type	Address	Port	Profile	Status	Availability	Basic Mode					
	DNS	Layer 7	10.125.2.16	53	DNS	Enable		Advanced Mode					
	Mysql	Layer 7	10.125.2.16	3306	LB_PROF_HTTP	Enable							
	Radius_1812	Layer 7	10.125.2.16	1812	radius	Enable							
	Radius_1813	Layer 7	10.125.2.16	1813	radius	Enable							
	SMTP	Layer 7	10.125.2.16	25	SMTP	Enable							
	Web-VIP1	Layer 7	10.125.2.15	80	test	Enable							
	Web_L4	Layer 4	10.125.2.16	80	TCP	Enable							

10. Set the type to Layer 7.

Virtual Server

Basic

General

Security

Application Optimization

Monitoring

Name

Required config name. No spaces.

Status

Disable

Enable

Maintain

Traffic Group

default

Type

Layer 7

Layer 4

Layer 2

Address Type

IPv4

IPv6

Specifics

Schedule Pool

Create New

NSP1

SNP2

Transaction Rate Limit


0

Default: 0 (disabled) Range: 0-1048567 transactions per second


Save

Cancel

11. Click on the **Profile** tab. It will drop down to reveal a list of options. Choose only between LB_PROF_TCPS and LB_PROF_HTTPS.



12. The SSL Mirror tab appears.



Go into it and enable traffic mirroring.

To enable this feature in a policy, execute the following command:

```
config load-balance virtual-server
edit vs-name
set ssl-mirror enable
set ssl-mirror-intf port1 port2
next
end
```

Chapter 18: Advanced Networking

This chapter includes the following topics:

- [Configuring static routes on page 366](#)
- [Configuring policy routes on page 368](#)
- [OSPF on page 529](#)
- [ISP routes on page 533](#)
- [BGP on page 536](#)
- [Configuring an Access List on page 541](#)
- [Configuring an Access IPv6 List on page 542](#)
- [Configuring a Prefix List on page 542](#)
- [Configuring a Prefix IPv6 List on page 543](#)
- [NAT on page 520](#)
- [Configure source NAT on page 521](#)
- [Configure 1-to-1 NAT on page 524](#)
- [QoS on page 526](#)
- [Configuring the QoS filter on page 528](#)
- [Configuring the QoS IPv6 filter on page 527](#)
- [Configuring a QoS queue on page 527](#)
- [Packet capture on page 551](#)
- [TCP multiplexing on page 195](#)
- [Reverse path route caching on page 533](#)
- [Transparent mode on page 543](#)

NAT

A number of network address translation (NAT) methods map packet IP address information for the packets that are received at the ingress network interface into the IP address space you configure. Packets with the new IP address are forwarded through the egress interface.

You can configure NAT per virtual server within the virtual server configuration.

This section describes the system-wide, policy-based NAT feature. The system-wide feature supports:

- **SNAT**—Translates the packet header source IP address to the configured address. See [Configure source NAT](#).
- **1-to-1 NAT**—Maps the public IP address for an interface to an IP address on a private network. See [Configure 1-to-1 NAT](#).
- **Port forwarding**—Maps an external published protocol port to the actual port. Configuration for port forwarding is included in the configuration for 1-to-1 NAT.

Configure source NAT


You use source NAT (SNAT) when clients have IP addresses from private networks. This ensures you do not have multiple sessions from different clients with source IP 192.168.1.1, for example. Or, you can map all client traffic to a single source IP address because a source address from a private network is not meaningful to the FortiADC system or backend servers.

[SNAT on page 521](#) illustrates SNAT. The SNAT rule matches the source and destination IP addresses in incoming traffic to the ranges specified in the policy. If the client request matches, the system translates the source IP address to an address from the SNAT pool. In this example, a client with private address 192.168.1.1 requests a resource from the virtual server address at 192.0.2.1 (not the real server address 10.0.0.1; the real server address is not published). The two rule conditions match, so the system translates the source IP to the next address in the SNAT pool—10.1.0.1. SNAT rules do not affect destination addresses, so the destination address in the request packet is preserved.

The system maintains this NAT table and performs the inverse translation when it receives the server-to-client traffic. Be sure to configure the backend servers to use the FortiADC address as the default gateway so that server responses are also rewritten by the NAT module.

Note: This SNAT feature is not supported for traffic to virtual servers. Use the virtual server SNAT feature instead.

SNAT



Before you begin:

- You must know the IP addresses your organization has provisioned for your NAT design.
- You must have Read-Write permission for System settings.

To configure source NAT:

1. Go to Networking > NAT.
The configuration page displays the Source tab.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in [Source NAT configuration on page 523](#).
4. Save the configuration.
5. Reorder rules, as necessary.

Source NAT configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Source	Address/mask notation to match the source IP address in the packet header. For example, 192.0.2.0/24.
Destination	Address/mask notation to match the destination IP address in the packet header. For example, 10.0.2.0/24.
Egress Interface	Interface that forwards traffic.
Translation Type	• IP Address—Select to translate the source IP to a single specified address. • Pool—Select to translate the source IP to the next address in a pool. • No NAT—Select to avoid translating the source IP.
Translation to IP Address	Note: This option applies only when the Translation Type is set to IP address. Specify an IPv4 address. The source IP address in the packet header will be translated to this address.
Pool Address Range	Note: This option applies only when Translation Type is set to Pool. Specify the first IP address in the SNAT pool.
No NAT	Note: This option applies only when Translation Type is set to No-NAT
To	Specify the last IP address in the SNAT pool.
Traffic Group	Select a traffic group. Otherwise, the system will use the default traffic group.
Reordering	
	After you have saved a rule, reorder rules as necessary. The rules table is consulted from top to bottom. The first rule that matches is applied and subsequent rules are not evaluated.

Configure 1-to-1 NAT


You can use 1-to-1 NAT when you want to publish public or “external” IP addresses for FortiADC resources but want the communication among servers on the internal network to be on a private or “internal” IP address range.

[One-to-One NAT on page 524](#) illustrates 1-to-1 NAT. The NAT configuration assigns both external and internal (or “mapped”) IP addresses to Interface 1. Traffic from the external side of the connection (such as client traffic) uses the external IP address and port. Traffic on the internal side (such as the virtual server communication with real servers) uses the mapped IP address and port.

1-to-1 NAT is supported for traffic to virtual servers. The address translation occurs before the ADC has processed its rules, so FortiADC server load balancing policies that match source address (such as content routing and content rewriting rules) should be based on the mapped address space.

The system maintains this NAT table and performs the inverse mapping when it sends traffic from the internal side to the external side.

One-to-One NAT



Before you begin:

- You must know the IP addresses your organization has provisioned for your NAT design.
- You must have Read-Write permission for System settings.

To configure one-to-one NAT:

1. Go to Networking > NAT.
2. Click the **1-to-1 NAT** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [1-to-1 NAT configuration on page 526](#).
5. Save the configuration.
6. Reorder rules, as necessary.

1-to-1 NAT configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
External Interface	Interface that receives traffic.
External Address Range	Specify the first address in the range. The last address is calculated after you enter the mapped IP range.
Mapped Address Range	Specify the first and last addresses in the range.
Port Forwarding	
Port Forwarding	Select to enable.
Protocol	• TCP • UDP
External Port Range	Specify the first port number in the range. The last port number is calculated after you enter the mapped port range.
Mapped Port Range	Specify the first and last port numbers in the range.
Traffic Group	Select a traffic group. Otherwise, the system will use the default.
Reordering	
	After you have saved a rule, reorder rules as necessary. The rules table is consulted from top to bottom. The first rule that matches is applied and subsequent rules are not evaluated.

QoS

You can use quality-of-service (QoS) policies to provision bandwidth for any traffic that matches the rule. You might consider QoS policies for latency- or bandwidth-sensitive services, such as VoIP and ICMP.

The FortiADC system does not provision bandwidth based on the TOS bits (also called differentiated services) in the IP header to control packet queueing. Instead, the system provisions bandwidth based on a source/destination/service matching tuple that you specify.

Note: The QoS policy feature is not supported for traffic to virtual servers.

Basic steps

1. Configure a [QoS queue](#).
2. Configure a [QoS filter](#) or [QoS IPv6 filter](#).

Configuring a QoS queue

You must configure a queue before you configure a filter.

Before you begin:

- You must have Read-Write permission for System settings.

To configure a QoS queue:

1. Go to Networking > QoS.
2. Click the **QoS Queue** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [QoS queue configuration on page 527](#)
5. Save the configuration.

QoS queue configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Bandwidth	Maximum bandwidth rate. Specify a number and a unit abbreviation. For example, specify 100K for 100 Kbps, 10M for 10 Mbps, and 1G for 1Gbps.

Configuring the QoS IPv6 filter

A QoS filter is the policy that assigns traffic to the QoS queue.

Before you begin:

- You must have a good understanding and knowledge of traffic in your network that requires QoS provisioning.
- You must have created the address configuration objects and service configuration objects that define the matching tuple for QoS rules. Use the Shared Resources menu firewall address and service object configuration editor.
- You must have created a QoS queue configuration object.
- You must have Read-Write permission for System settings.

To configure QoS filter:

1. Go to Networking > QoS.
2. Click the **QoS IPv6 Filter** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [QoS IPv6 filter configuration on page 528](#).
5. Save the configuration.
6. Reorder rules, as necessary.

QoS IPv6 filter configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Status	Enable/disable the filter.
Queue	Select the queue that will be used for packets that match the filter criteria.
Service	Select a service object to use to form the matching tuple.
Source	Select a source address object to use to form the matching tuple.
Destination	Select a destination address object to use to form the matching tuple.
Ingress Interface	Select the interface that receives traffic.
Egress Interface	Select the interface that forwards traffic.
	After you have saved a rule, reorder rules as necessary. The rules table is consulted from top to bottom. The first rule that matches is applied and subsequent rules are not evaluated.

Configuring the QoS filter

A QoS filter is the policy that assigns traffic to the QoS queue.

Before you begin:

- You must have a good understanding and knowledge of traffic in your network that requires QoS provisioning.
- You must have created the address configuration objects and service configuration objects that define the matching tuple for QoS rules. Use the Shared Resources menu firewall address and service object configuration editor.
- You must have created a QoS queue configuration object.
- You must have Read-Write permission for System settings.

To configure QoS filter:

1. Go to Networking > QoS.
2. Click the QoS Filter tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [QoS filter configuration on page 529](#).

5. Save the configuration.
6. Reorder rules, as necessary.

QoS filter configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Status	Enable/disable the filter.
Queue	Select the queue that will be used for packets that match the filter criteria.
Service	Select a service object to use to form the matching tuple.
Source	Select a source address object to use to form the matching tuple.
Destination	Select a destination address object to use to form the matching tuple.
Ingress Interface	Select the interface that receives traffic.
Egress Interface	Select the interface that forwards traffic.
	After you have saved a rule, reorder rules as necessary. The rules table is consulted from top to bottom. The first rule that matches is applied and subsequent rules are not evaluated.

OSPF

OSPF (Open Shortest Path First) is described in RFC2328, OSPF Version 2. It is a link-state interior routing protocol. Compared with RIP, OSPF can provide scalable network support and faster convergence times. OSPF is widely used in large networks such as ISP backbone and enterprise networks. FortiADC supports OSPF version 2.

By the support HA for OSPF route injection feature, the virtual server IP/IPv6 address can be injected into the OSPF domain, and can be advertised or withdrawn according to the health state of the real server.

Before you begin:

- You must know how OSPF has been implemented in your network, and you must know the configuration details of the implementation.
- You must have Read-Write permission for System settings.

To configure OSPF:

1. Go to Networking > Routing.
2. Click the **OSPF** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [OSPF configuration on page 530](#).
5. Save the configuration.

OSPF configuration

Settings	Guidelines
Router	32-bit number that sets the router-ID of the OSPF process. The router ID uses dotted decimal notation. The router-ID must be an IP address of the router, and it must be unique within the entire OSPF domain to the OSPF speaker.
Default Metric	The default is 10.
Distance	The default is 110.
Default Information Originate	• Disable—Default. • Enable—Originate an AS-External (type-5) LSA describing a default route into all external routing capable areas of the specified metric and metric type. • Always—The default is always advertised even when there is no default route present in the routing table.
Default Information Metric	The default is -1, which equals to the Default Metric.
Default Information Metric Type	Select either of the following: • 1—If selected, the metric equals to the Default Information Metric, plus the Default Metric. • 2—(Default) If selected, the metric equals to the Default Information Metric.
Redistribute Connected	Enable/disable to redistribute connected routes to OSPF, with the metric type and metric set if specified. Redistributed routes are distributed into OSPF as Type-5 External LSAs into links to areas.
Redistribute Connected Metric	The default is -1, which equals to the Default Metric.
Redistribute Connected Metric Type	Select either of the following: • 1—If selected, the metric equals to the Redistribute Connected Metric, plus the Default Metric. • 2—(Default) If selected, the metric equals to the Redistribute Connected Metric.
Redistribute Static	Enable/disable to redistribute static routes to OSPF, with the metric type and metric set if specified. Redistributed routes are distributed to OSPF as Type-5 External LSAs into links to areas.
Redistribute Static Metric	The default is -1, which equals to the Default Metric.
Redistribute Static Metric Type	• 1—If selected, the metric equals to the Redistribute Static Metric, plus the Default Metric. • 2—(Default) If selected, the metric equals to the Redistribute Static Metric.
Area Authentication	
Area	32-bit number that identifies the OSPF area. An OSPF area is a smaller part of the larger OSPF network. 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.
Authentication	Specify an authentication type:

Settings	Guidelines
	• None—Also called null authentication. No authentication is 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. • Text—A simple password is used. The password is a 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. • MD5—Use OSPF cryptographic authentication. A shared secret key is used 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.
Network	
Prefix	Address/mask notation to specify the subnet.
Area	Select an area configuration.
Interface	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Interface	Select the interface to enable OSPF for it.
Ignore MTU	Enable/disable to ignore the interface MTU. Disabled by default.
Network Type	• Broadcast • Point to Point • Point to Multipoint
Retransmit Interval	Interval for retransmitting Database Description and Link State Request packets. The default is 5 seconds.
Transmit Delay	Increment LSA age by this value when transmitting. The default is 1 second.
Cost	Set link cost for the specified interface. The cost value is set to router-LSA's metric field and used for SPF calculation. The default is 0.
Priority	The router with the highest priority will be more eligible to become Designated Router. Setting the value to 0 makes the router ineligible to become Designated Router. The default is 1.
Dead Interval	Number of seconds for RouterDeadInterval timer value used for Wait Timer and Inactivity Timer. This value must be the same for all routers attached to a common network. The default is 40 seconds.
Hello Interval	Number of seconds between hello packets sent on the configured interface. This value must be the same for all routers attached to a common network. The default is 10 seconds.

Settings	Guidelines
Authentication	Specify an authentication type. All OSPF interfaces that want to learn routes from each other must be configured with the same authentication type and password or MD5 key (one match is enough). Options are: • None—Also called null authentication. No authentication is 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. • Text—A simple password is used. The password is a 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. • MD5—Use OSPF cryptographic authentication. A shared secret key is used 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.
Text	If using text authentication, specify a password string. Passwords are limited to 8 characters.
MD5	If using MD5 authentication, select an MD5 configuration name.
HA Router	
Router	You use the HA Router list configuration in an HA active-active deployment. On each HA cluster node, add an HA Router configuration that includes an entry for each cluster node. When the appliance is in standalone mode, it uses the primary OSPF Router ID; when it is in HA mode, it uses the HA Router list ID. Specify a 32-bit number that sets the router-ID of the OSPF process. The router ID uses dotted decimal notation. The router-ID must be an IP address of the router, and it must be unique within the entire OSPF domain to the OSPF speaker.
Node	HA Node ID (0-7).
MD5 Key List	
Name	Configuration name. You select this name in the OSPF Interface configuration. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Member	
Key ID	A number 1-255. Each member key ID must be unique to its member list.
Key	A string of up to 16 characters to be hashed with the cryptographic MD5 hash function.

ISP routes

ISP routes can be used for outbound traffic and link load balancing traffic.

Routes for outbound traffic are chosen according to the following priorities:

1. Link local routes—Self-traffic uses link local routes.
2. LLB Link Policy route—Configured policy routes have priority over default routes.
3. Policy route—Configured policy routes have priority over default routes.
4. Static route / ISP route / OSPF route—Priority is based on the distance metric. By default, distance for static routes is 10, for ISP routes is 20, and for OSPF routes is 110. The distance metric is configurable for static routes and OSPF routes, but not ISP routes.
5. Default LLB Link Policy route—Default routes have lower priority than configured routes.
6. Default static route / OSPF route—Default routes have lower priority than configured routes.

Before you begin:

- You must have read-write permission for system settings.

Note: Adding a new ISP route does not affect existing sessions. Deleting or editing an ISP route causes the related sessions to be re-created.

To configure ISP Routes:

1. Go to Networking > Routing.
2. Click the **ISP** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in [ISP Route configuration on page 533](#).
5. Save the configuration.

ISP Route configuration


Settings	Guidelines
Destination	Select an ISP address book configuration object. Note: Two ISP routes cannot reference the same ISP address book. The ISP routing feature does not support multipath routing.
Gateway	IP address of the gateway router that can route packets to the destination IP address that you have specified.

Reverse path route caching

By default, reverse path route caching is enabled. FortiADC caches a reverse path route for inbound traffic so it can forward reply packets to the ISP link that forwarded the corresponding request packet. This is useful when your site receives traffic from multiple ISP links. For example, in [Reverse path route caching enabled on page 534](#), the reverse path pointer ensures that client traffic received from ISP1 is returned through ISP1.

Note: FortiADC does not support IPv6 traffic reverse path route caching.

Reverse path route caching enabled



When reverse path caching is not enabled, the system forwards reply packets based on the results of routing lookup.

To enable/disable reverse path route caching, use the `config router setting` CLI command:

```
FortiADC-VM # config router setting
FortiADC-VM (setting) # get
rt-cache-strict : disable
rt-cache-reverse : enable
ip-forward : enable
ip6-forward : enable
FortiADC-VM (setting) # set rt-cache-reverse disable
FortiADC-VM (setting) # end
FortiADC-VM # get router setting
rt-cache-strict : disable
rt-cache-reverse : disable
ip-forward : enable
ip6-forward : enable
```

The `rt-cache-strict` option is disabled by default. Enable it when you want to send reply packets only via the same interface that received the request packets. When enabled, source interface becomes part of the matching tuple that FortiADC uses to identify sessions, so reply traffic is forwarded from the same interface that received the traffic. (Normally each session is identified by a 5-tuple: source IP, destination IP, protocol, source port, and destination port.)

If the `rt-cache-reverse` option is enabled, you can use the `config rt-cache-reverse-exception` command to maintain an exceptions list for source IP addresses that should be handled differently. For example, if you configure an exception for 192.168.1.0/24, FortiADC will not maintain a pointer to the ISP for traffic from source 192.168.1.18. Reply packets will be forwarded based on the results of routing lookup.

```
FortiADC-docs # config router setting
FortiADC-docs (setting) # get
rt-cache-strict : disable
rt-cache-reverse : enable
ip-forward : enable
ip6-forward : enable
icmp-redirect-send : disable
FortiADC-docs (setting) # config rt-cache-reverse-exception
FortiADC-docs (rt-cache-rever~e) # edit 1
Add new entry '1' for node 3740
FortiADC-docs (1) # set ip-netmask 192.168.1.0/24
FortiADC-docs (1) # end
FortiADC-docs (setting) # end
```

BGP

BGP stands for Border Gateway Protocol, which 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. Often classified as a path-vector protocol and sometimes as a distance-vector routing protocol, BGP exchanges routing and reachability information among autonomous systems over the Internet.

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.

How BGP works

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.

By the support HA for BGP route injection feature, the virtual server IP/IPv6 address can be injected into the BGP domain, and can be advertised or withdrawn according to the health state of the real server.

IBGP vs. 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.

For more information on BGP routing, see "Chapter 3 - Advanced Routing" of the *FortiOS Handbook for FortiOS 5.4.1*.

Before you begin, you must:

- Know how BGP has been implemented in your network, i.e., the configuration details of the implementation..
- Have Read-Write permission for System settings.
- Have configured all the needed access (IPv6) lists and prefix (IPv6) lists. See [Access list vs. prefix list](#).

To configure BGP:

1. Click Networking > Routing.
2. Click the BGP tab.
3. Make the desired entries and/or seldctions as described in [BGP configuration on page 538](#).
4. Click Save when done.

BGP configuration

Settings	Guidelines
AS	Enter the AS (Autonomous System) number of the BGP router. Valid values are from 0 to 4294967295. Note: Per RFC 6996, the first and last ASNs of the original 16-bit integers, namely 0 and 65535, and the last ASN of the 32-bit numbers, namely 4,294,967,295, are reserved and should not be used by operators; ASNs 64,512 to 65,534 of the original 16-bit AS range, and 4,200,000,000 to 4,294,967,294 of the 32-bit range are reserved for private use, which means that they can be used internally but should not be announced to the global Internet.
Router ID	Enter the 32-bit number that sets the router-ID of the BGP process. The router ID uses dotted decimal notation. The router-ID must be the IP address of the router, and it must be unique within the entire BGP domain to the BGP speaker.
Redistribute OSPF	Enable/Disable (default) the redistribution of OSPF routes to the BGP process.
Redistribute Connected	Enable/Disable (default) the redistribution of connected routes to the BGP process.
Redistribute Static	Enable/Disable (default) the redistribution of static routes to the BGP process.
Redistribute IPv6 Connected	Enable/Disable (default) the redistribution of connected IPv6 routes to the BGP process.
Redistribute IPv6 Static	Enable/Disable (default) the redistribution of static IPv6 routes to the BGP process.
Always Compare MED	Enable/Disable (default) the comparison of Multi-Exit Discriminator (MED) for paths from neighbors in different ASs (Autonomous Systems).
Deterministic MED	Enable/Disable (default) the deterministic comparison of Multi-Exit Discriminator (MED) values among all paths received from the same AS (Autonomous System).
Bestpath Compare Router ID	Enable/Disable (default) the BGP routing process to compare identical routes received from different external peers during the best-path selection process and to select the route with the lowest router ID as the best path.
Network	
Type	Select either of the following (IP address) types: • IPv4 • IPv6
IPv4 Prefix	If IPv4 is selected (above), specify the IPv4 prefix in the format of 0.0.0.0/0.
IPv6 Prefix	If IPv6 is selected (above), specify the IPv6 prefix in the format of ::/0.
Save	Be sure to click Save after you are done with configuring the network.
Neighbor	
Remote AS	Specify the remote AS (Autonomous System) number of the BGP neighbor you are creating. Valid values are from 1 to 4294967295.
Type	Select either of the following: • IPv4 • IPv6

Settings	Guidelines
IP/IPv6	Specify the IPv4 address or IPv6 address for the BGP neighbor.
Interface	Click to select the interface for the BGP neighbor.
Port	Specify the port of the BGP neighbor.
Keep Alive	Specify the frequency (in seconds) at which the BGP neighbor sends out <i>keepalive</i> message to its peer. Valid values are from 0 to 65535, with 60 seconds being the default.
Hold Time	Specify the "wait time" or pause (in seconds) the BGP neighbor declares a peer dead after failing to receive a <i>keepalive</i> message from it. Valid values are from 0 to 65535, with 180 (seconds) being the default. When the minimum acceptable hold time is configured on a BGP router, a remote BGP peer session can be established only when the latter is advertising a hold time equal to, or greater than, the minimum acceptable hold time configured on the former. If the minimum acceptable hold time is greater than the configured hold time, then the next time the remote BGP peer tries to establish a session with the local BGP router, it will fail and the local BGP router will notify the remote BGP peer saying "unacceptable hold time".
Distribute List In/Distribute IPv6 List In	Click to select an Access List or Access IPv6 List. The BGP router will apply the selected access list to inbound advertisements to the BGP neighbor when distributing BGP neighbor information. Note: It is highly recommended that you have the Prefix List or the IPv6 Prefix List configured before configuring BGP Routing.
Distribute List Out/Distribute IPv6 List Out	Click to select an Access List or Access IPv6 List. The BGP router will apply the selected access list to outbound advertisements to the neighbor when distributing BGP neighbor information. Note: It is highly recommended that you have the Access List or the Access IPv6 List configured before configuring BGP Routing.
Prefix List In/Prefix IPv6 List In	Click to select an Prefix List or Prefix IPv6 List. The BGP router will apply the selected Prefix (IPv6) List to inbound advertisements to the neighbor when distributing BGP neighbor information. Note: It is highly recommended that you have the Prefix List or the Prefix IPv6 List configured before configuring BGP Routing.
Prefix List Out/Prefix IPv6 List Out	Click to select an Prefix List or Prefix IPv6 List. The BGP router will apply the selected Prefix (IPv6) List to outbound advertisements to the neighbor when distributing BGP neighbor information. Note: It is highly recommended that you have the Prefix List or the Prefix IPv6 List configured before configuring BGP Routing.
Weight	Assign a weight to a neighbor connection. Valid values are from 0 to 65535. By default, routes learned through another BGP peer carries a weight value of 0, whereas routes sourced by the local router carry a default weight value of 32768. Initially, all routes learned from a neighbor will have an assigned weight. The route with the greatest weight is chosen as the preferred route when multiple routes are available to a network.

Settings	Guidelines
Save	Be sure to click Save after you are done with configuring the Neighbor.
HA Router ID List	
Router ID	Use the HA Router list configuration in an HA active-active deployment. On each HA cluster node, add an HA Router configuration that includes an entry for each cluster node. When the appliance is in standalone mode, it uses the primary BGP Router ID; when it is in HA mode, it uses the HA Router list ID. Specify a 32-bit number that sets the router-ID of the BGP process. The router ID uses dotted decimal notation. The router-ID must be an IP address of the router, and it must be unique within the entire BGP domain to the BGP speaker.
Node	Specify the HA Node ID (0-7).
Save	Be sure to click Save after you are done with configuring the HA Router ID List.

Note: The Access List and Prefix List features are mutually exclusive. Therefore, do NOT apply both to any neighbor in any direction (inbound or outbound) when configuring BGP routing.

Route health injection (RHI)

Route health injection (RHI) allows for advertising routes to virtual server IP addresses based on the health status of the corresponding service. For FortiADC deployment, routes to virtual server IP addresses can be injected into the dynamic routing protocol like BGP, OSPF, etc. and spread through the network. The status of a virtual server depends on factors such as the status of its real servers, the scheduled if the schedule pool is enabled. For example, if there is at least one available real server (virtual server is healthy), the route to the virtual server IP address will be injected and spread to the neighbors as long as the virtual server IP is added into the BGP network. Conversely, the route to the virtual server IP will not be injected if no real server is available (virtual server is unhealthy).

Access list vs. prefix list

Access lists and prefix lists are different mechanisms that you can use to control traffic into and out of a network.

Access lists

Access lists allow you to filter packets so that you can permit or deny them from crossing specified network interfaces. You can control whether packets are forwarded or blocked at the routers' interfaces based on the criteria set in the access lists.

Access lists fall into two categories: standard and extended. A standard access list (1-99) only checks the source addresses of all IP packets, whereas an extended access list (100-199) checks both source and destination addresses, specific UDP/TCP/IP protocols, and destination ports.

[Range comparison between standard access list and extended access list on page 540](#) below provides a comparison between standard access lists and extended access lists in terms of range.

Range comparison between standard access list and extended access list

Access List Type	Range
Standard	1-99, 1300-1999
Extended	100-199, 2000-2699

Note: For this release, FortiADC only supports user-defined access lists. It does NOT support either standard or extended access lists. Access lists are NOT required for BGP routing configuration. However, if you want to include access lists in BGP routing configuration, we highly recommend that you have them configured ahead of time.

Prefix list

Prefix lists are used to configure filter IP routes. They are configured with the permit or deny keywords to either allow or block the prefix based on the matching conditions. A prefix list is made up of an IP address and a bit mask. The IP address can be a classful network, a subnet, or a single host route, whereas the bit mask can be a numeric value ranging from 1 to 32. An implicit deny is applied to the route that matches any entry in the prefix list.

A prefix list contains one or multiple sequential entries which are evaluated sequentially, starting with the entry with the lowest sequence number. Evaluation of a prefix against a prefix list comes to an end when a match is found and the permit or deny statement is applied to that network.

Although extended access lists, and, to some extent, standard access lists, can be utilized to match prefix announcements, prefix lists are considered more graceful.

Note: Prefix lists are NOT required for BGP routing configuration. However, if you want to include prefix lists in BGP routing configuration, we highly recommend that you have them configured ahead of time.

Configuring an Access List

FortiADC D-Series units support IPv4 access lists over BGP routing. If you are configuring BGP routing using IPv4, you must configure access lists using the IPv4 protocol.

To configure an access list:

1. Click Networking > Routing.
2. Click the Access List tab.
3. Click **Create New**.
4. Enter a unique name for the new access list. **Note:** The name can be up to 35 alphanumeric characters long, including . (period), : (colon), _ (underscore), and - (hyphen). No space is allowed.
5. Enter a brief description of the access list. **Note:** The description can be up to 1023 alphanumeric characters long, with no restriction on use of special characters. Space between characters is allowed.
6. Click **Save**. The newly created access list entry appears in the access list table.
7. Click the Edit button to open the Access List dialog.
8. In the Rule pane, click **Create New**. The Access List > Edit Rule tab opens.
9. For Action, select the **Permit** or **Deny** radio button.
10. For IPv4 Prefix, enter the IPv4 address/subnet mask in the format of 0.0.0.0./0.
11. Click **Save** when done.
12. Repeat Steps 8 through 11 above to add as many rules to the access list as needed.
13. Click **X** to close the Access List dialog when done.

Configuring an Access IPv6 List

FortiADC D-Series units support IPv6 access lists over BGP routing. If you are configuring BGP routing using IPv6, you must configure access lists using the IPv6 protocol.

To configure an Access IPv6 List:

1. Go to Network > Routing.
2. Click the Access IPv6 List tab.
3. Click **Add**.
4. Enter a unique name for the new access list. **Note:** The name can be up to 35 alphanumeric characters long, including . (period) , : (colon), _ (underscore), and - (hyphen). No space is allowed.
5. Enter a brief description of the access list. **Note:** The description can be up to 1023 alphanumeric characters long, with no restriction on use of special characters. Space between characters is allowed.
6. Click **Save**. The newly created access list entry appears in the access list table.
7. Click the **Edit** button to open the Access IPv6 List dialog.
8. In the Rule pane, click **Add**. The Access IPv6 List > Edit Rule tab pens.
9. For Action, select the **Permit** or **Deny** radio button.
10. For IPv6 Prefix, enter the IPv6 address/subnet mask in the format of ::/0.
11. Click **Save** when done.
12. Repeat Steps 8 through 11 above to add as many rules to the access list as needed.
13. Click **X** to close the Access IPv6 List dialog when done.

Configuring a Prefix List

FortiADC D-Series units support IPv4 prefix lists over BGP routing. If you are configuring BGP routing using IPv4, you must configure access lists using the IPv4 protocol.

To configure a Prefix list:

1. Go to Network > Routing.
2. Click the Prefix List tab.
3. Click **Create New**.
4. Enter a unique name for the new access list. **Note:** The name can be up to 35 alphanumeric characters long, including . (period) , : (colon), _ (underscore), and - (hyphen). No space is allowed.
5. Enter a brief description of the access list. **Note:** The description can be up to 1023 alphanumeric characters long, with no restriction on use of special characters. Space between characters is allowed.
6. Click **Save**. The newly created access list entry appears in the access list table.
7. Click the **Edit** button to open the Prefix List dialog.
8. In the Rule pane, click **Create New**. The Prefix List > Edit Rule tab pens.
9. For Action, select the Permit or Deny radio button.
10. For IPv4 Prefix, enter the IPv4 address/subnet mask in the format of 0.0.0.0/0.
11. For GE, set the GE (greater than and equal to) values.
12. For LE, set the LE (less than and equal to) values
13. Click **Save** when done.

14. Repeat Steps 8 through 13 above to add as many rules to the access list as needed.
15. Click **X** to close the Prefix List dialog when done.

Configuring a Prefix IPv6 List

FortiADC D-Series units support IPv6 prefix lists over BGP routing. If you are configuring BGP routing using IPv6, you must configure access lists using the IPv6 protocol.

To configure a Prefix IPv6 List:

1. Go to Network > Routing.
2. Click the Prefix IPv6 List tab.
3. Click **Create New**.
4. Enter a unique name for the new access list. **Note:** The name can be up to 35 alphanumeric characters long, including . (period) , : (colon), _ (underscore), and - (hyphen). No space is allowed.
5. Enter a brief description of the access list. **Note:** The description can be up to 1023 alphanumeric characters long, with no restriction on use of special characters. Space between characters is allowed.
6. Click **Save**. The newly created access list entry appears in the access list table.
7. Click the **Edit** button to open the Prefix IPv6 List dialog.
8. In the Rule pane, click **Create New**. The Prefix IPv6 List > Edit Rule tab pens.
9. For Action, select the Permit or Deny radio button.
10. For IPv6 Prefix, enter the IPv6 address/subnet mask in the format of ::/0.
11. For GE, set the GE (greater than and equal to) values.
12. For LE, set the LE (less than and equal to) values
13. Click **Save** when done.
14. Repeat Steps 8 through 13 above to add as many rules to the access list as needed.
15. Click **X** to close the Prefix IPv6 List dialog when done.

Transparent mode

In transparent mode, the FortiADC appliance (the load balancer) splits a subnet into two VLANs and bridges them together. This allows you to insert the appliance into an existing network without modifying the IP addressing.

To support deploy FortiADC in transparent mode, you must first create a softswitch interface on the appliance. All traffic that FortiADC does not supported can directly pass through this soft-switch interface without interruption, and FortiADC-supported traffic, such as LLDB and DHCP, needs to be terminated.

Keep in mind that the FortiADC soft-switch does not participate in the STP node, and all STP BPDU will be forwarded by this soft-switch interface directly.

For more information, see [FortiADC Transparent Configuration Guide](#).

Chapter 19: Best Practices and Fine Tuning

This chapter is a collection of best practice tips and fine-tuning guidelines. It includes the following topics:

- [Regular backups on page 544](#)
- [Security on page 544](#)
- [Performance tips on page 546](#)
- [High availability on page 547](#)

Regular backups

Make a backup before executing disruptive operations, such as:

- Upgrading the firmware
- Running the CLI commands `execute factoryreset` or `execute restore`
- Clicking the **Reset** button in the System Information widget on the dashboard

Always password-encrypt your backups.

Security

This section lists tips to further enhance security.

Topology

- Virtual servers can be on the same subnet as physical servers. This configuration creates a one-arm load balancer.

For example, the virtual server 10.0.0.2/24 could forward to the physical server 10.0.0.3-200.

If you are deploying gradually, you might want to initially install your FortiADC in a one-arm topology during the transition phase, and route traffic to it only after you have configured FortiADC to handle it.

Long term, this is *not* recommended. Unless your network's routing configuration prevents it, it could allow clients that are aware of the physical server's IP address to bypass the FortiADC appliance by accessing the physical server directly.

- Make sure web traffic cannot bypass the FortiADC appliance in a complex network environment.
- FortiADC appliances are *not* general-purpose firewalls. While they are security-hardened network appliances, security is not their primary purpose, and you should not allow traffic to pass through without inspection. FortiADC and FortiGate complement each other to improve security, availability, and performance. To protect your servers, install the FortiADC appliance or appliances between the servers and a general purpose firewall such as a FortiGate. *FortiADC complements, and does not replace, general purpose firewalls.*
- Disable all network interfaces that should not receive any traffic.

For example, if administrative access is typically through port1, the Internet is connected to port2, and servers are connected to port3, you would disable ("bring down") port4. This would prevent an attacker with physical access from connecting a cable to port4 and thereby gaining access if the configuration inadvertently allows it.

Administrator access

- As soon as possible during initial setup, give the default administrator, `admin`, a password. This super-administrator account has the highest level of permissions possible, and access to it should be limited to as few people as possible.
- Change all administrator passwords regularly. Set a policy—such as every 60 days—and follow it. (Mark the **Change Password** check box to reveal the password dialog.)
- Instead of allowing administrative access from any source, restrict it to trusted internal hosts. On those computers that you have designated for management, apply strict patch and security policies. Always password-encrypt any configuration backup that you download to those computers to mitigate the information that attackers can gain from any potential compromise.
- Do not use the default administrator access profile for all new administrators. Create one or more access profiles with limited permissions tailored to the responsibilities of the new administrator accounts.
- By default, an administrator login that is idle for more than 30 minutes times out. You can change this to a longer period in Timeout, but Fortinet does not recommend it. Left unattended, a web UI or CLI session could allow anyone with physical access to your computer to change system settings. Small idle timeouts mitigate this risk.
- Administrator passwords should be at least 8 characters long and include both numbers and letters.
- Restrict administrative access to a single network interface (usually port1), and allow only the management access protocols needed.
- Use only the most secure protocols. Disable ping, except during troubleshooting. Disable HTTP, SNMP, and Telnet unless the network interface only connects to a trusted, private administrative network.
- Disable all network interfaces that should not receive any traffic.

- For example, if administrative access is typically through port1, the Internet is connected to port2, and servers are connected to port3, you would disable (“bring down”) port4. This would prevent an attacker with physical access from connecting a cable to port4 and thereby gaining access if the configuration inadvertently allows it.
- Immediately revoke certificates that have been compromised. If possible, automate the distribution of certificate revocation lists.

Performance tips

When configuring the system and its features, there are many settings and practices that can yield better performance.

System performance

- Delete or disable unused policies. The system allocates memory with each server policy, regardless of whether it is actually in active use. Configuring extra policies will unnecessarily consume memory and decrease performance.
- To reduce latency associated with DNS queries, use a DNS server on your local network as your primary DNS.
- If your network’s devices support them, you can create one or more VLAN interfaces. VLANs reduce the size of a broadcast domain and the amount of broadcast traffic received by network hosts, thus improving network performance.
- If you have enabled the server health check feature and one of the servers is down for an extended period, you can improve system performance by disabling group membership for the physical server, rather than allowing the server health check to continue checking for the server’s responsiveness.

Reducing the impact of logging on performance

- If you have a FortiAnalyzer, store FortiADC logs on the FortiAnalyzer to avoid resource usage associated with writing logs to the local hard disk.
- If you do not need a traffic log, disable it to reduce the use of system resources.
- Reduce repetitive log messages. Use the alert email settings to define the interval that emails are sent if the same condition persists following the initial occurrence.
- Avoid recording log messages using low severity thresholds, such as information or notification, to the local hard disk for an extended period of time. Excessive logging frequency saps system resources and can cause undue wear on the hard disk and may cause premature failure.

Reducing the impact of reports on system performance

Generating reports can be resource intensive. To avoid performance impacts, consider scheduling report generation during times with low traffic volume, such as at night and on weekends.

Keep in mind that most reports are based upon log messages. All caveats regarding log performance also apply.

Reducing the impact of packet capture on system performance

Packet capture can be useful for troubleshooting but can be resource intensive. To minimize the impact on system performance, use packet capture only during periods of minimal traffic. Use a local console CLI connection rather than a

Telnet or SSH CLI connection, and be sure to stop the command when you are finished.

High availability

We recommend that you deploy high availability (HA). Keep these points in mind when setting up a cluster:

- Isolate HA interface connections from your overall network.
Heartbeat and synchronization packets contain sensitive configuration information and can consume considerable network bandwidth. For best results, directly connect the two HA interfaces using a crossover cable. If your system uses switches instead of crossover cables to connect the HA heartbeat interfaces, those interfaces must be reachable by Layer 2 multicast.
- When configuring an HA pair, pay close attention to the options ARP Packet Numbers and ARP Packet Interval.
The FortiADC appliance broadcasts ARP packets to the network to ensure timely failover. Delayed broadcast intervals can slow performance. Set the value of ARP Packet Numbers no higher than needed.
When the FortiADC appliance broadcasts ARP packets, it does so at regular intervals. For performance reasons, set the value for ARP Packet Interval no greater than required.
Some experimentation might be needed to set these options at their optimum value.

We recommend that you configure an SNMP community and enable the **HA heartbeat failed** option to generate a message if the HA heartbeat fails.

Chapter 20: Troubleshooting

This chapter includes the following topics:

- [Logs on page 548](#)
- [Tools on page 548](#)
- [Save debug file on page 553](#)
- [Solutions by issue type on page 554](#)
- [Resetting the configuration on page 561](#)
- [Restoring firmware \(“clean install”\) on page 561](#)
- [Additional resources on page 564](#)

Logs

Log messages often contain clues that can aid you in determining the cause of a problem.

Depending on the type, log messages may appear in either the event, attack, or traffic logs. The FortiADC appliance must be enabled to record event, attack, and traffic log messages; otherwise, you cannot analyze the log messages for events of that type. To enable logging of different types of events, go to Log & Report > Log Settings.

During troubleshooting, you may find it useful to lower the logging severity threshold for more verbose logs, to include more information on less severe events. To configure the severity threshold, go to Log & Report > Log Settings.

Tools

This section gives an overview of the following troubleshooting tools:

- [execute commands](#)
- [diagnose commands](#)
- [System dump](#)
- [Packet capture](#)
- [Diff](#)

execute commands

execute commands

You can use the command-line interface (CLI) execute commands to run system management utilities, such as backups, upgrades and reboots; and network diagnostic utilities, such as nslookup, ping, traceroute, and tcpdump.

The following example shows the list of execute commands:

```
FortiADC-VM # execute ?
```



```
backup backup
caching caching management
certificate certificate
checklogdisk find and auto correct errors on the log disk
clean clean
config-sync config sync
date set/get date and time
discovery-glb-virtual-server Sync virtual servers from glb server, add them to the virtual
    server list
dumpsystem dump system information for debugging purpose
dumpsystem-file manipulate the dumped debugging information
factoryreset reset to factory default
fixlogdisk correct errors on the log disk
formatlogdisk format log disk to enhance performance
geolookup lookup geography information for IP address
glb-dprox-lookup lookup GLB dynamic proximity information
glb-persistence-lookup lookup GLB persistence information
ha ha
isplookup lookup ISP name and isp-address for IP address
log log management
nslookup nslookup
packet-capture packet-capture <Port Number> [filter] (Only IPv4)
packet-capture-file packet-capture-file
packet-capture6 packet-capture6 <Port Number> [filter] (Include IPv6)
ping ping <host name | host ip>
ping-option ping option settings
ping6 ping <host name | host ipv6>
ping6-option ping6 option settings
reboot reboot the system
reload reload appliance
restore restore
shutdown shutdown appliance
ssh Simple SSH client.
statistics-db statistics db management
telnet Simple telnet client.
traceroute traceroute
vm vm
web-category-test Test a url find its web-category
```

For details, see the [FortiADC CLI Reference](#).

diagnose commands

You can use the CLI diagnose commands to gather diagnostic information that can be useful to Fortinet Customer Care when diagnosing any issues with your system. The commands are similar to the Linux commands used for debugging hardware, system, and IP networking issues.

The most important command for customers to know is `diagnose debug report`. This prepares a report you can give to your Fortinet support contact to assist in debugging an issue.

The following examples show the lists of diagnose commands:

```
FortiADC-VM # diagnose ?
debug debug
hardware hardware
```

```
llb llb
netlink netlink
server-load-balance server-load-balance
sniffer sniffer
system system
```

```
FortiADC-VM # diagnose debug ?
application set/get debug level for daemons
cli set/get debug level for CLI and CMDB
config-error-log read/clear config error information
crashlog crashlog
disable disable debug output
enable enable debug output
flow flow
info show debug info
kernel set/get debug level for kernel
report Report for tech support.
timestamp timestamp
```

```
FortiADC-VM # diagnose hardware get ?
deviceinfo list device status and information
ioport read data from an I/O port
pciconfig list information on PCI buses and connected devices
sysinfo list system hardware information
```

```
FortiADC-VM # diagnose netlink ?
backlog set netlink backlog length
device display network devices statistic information
interface netlink interface
ip ip
ipv6 ipv6
neighbor netlink neighbor
neighbor6 netlink neighbor for ipv6
route netlink routing table
route6 netlink routing table
tcp display tcp statistic information
udp display udp statistic information
```

```
FortiADC-VM # diagnose system ?
top show top process
vm check vm state
```

For details, see the [FortiADC CLI Reference](#).

System dump

The system includes utilities for generating system dump files that can help Fortinet support engineers analyze an issue for you. The CLI and Web UI versions have different usage:

- CLI—Used to dump kernel and user space information when the system is still responsive.
- Web UI—Used to dump kernel information when the system is deeply frozen.

The following is an example of CLI command usage:

```
FortiADC-VM # execute dumphsystem
This operation will reboot the system!
Do you want to continue? (y/n)y
Begins to dump userspace information
Begins to dump kernel information

FortiADC-VM # execute dumphsystem-file list
-rw----- 1 0 0 96719189 Mar 15 13:35 coredump-2016-03-15-13_35
-rw-r--r-- 1 0 0 16654391 Mar 15 13:34 user_coredump_2016_03_15_13_34_46.tar.bz2

FortiADC-VM # execute dumphsystem-file upload tftp coredump-2016-03-15-13_35 172.30.184.77
coredump-2016-03-15- 7% |** | 7152k 0:09:58 ETA
```

To use the web UI system dump utility:

1. Go to System > Debug.
2. Click **System Dump** to generate the file.
After the file has been generated, you are logged out. When you log back in and revisit the page, the system dump file appears in the file list.
3. Select the file and click **Export** to download the file.

Packet capture

The tcpdump utility is supported through the CLI and web UI.

See the [FortiADC CLI Reference](#) for information on using the CLI command.

Use the following procedure to use the web UI version.

Before you begin:

- You must have a good understanding of tcpdump and filter expressions. See <http://www.tcpdump.org/manpages/pcap-filter.7.html>.
- You must have Read-Write permission for System settings.

To use the web UI version of tcpdump:

1. Go to Networking > Packet Capture.
2. Click **Create New** to open the Packet Capture editor, and specify your packet capture settings as shown in the figure below.
3. Use the controls to start, stop, and download the packet capture. See [Packet capture toolbar on page 552](#).

Packet capture configuration page

Packet capture toolbar

ID	Interface	Filter Criteria	Maximum Packet Count
1	port1	host=192.0.2.0/24 & port=80 & protocol=tcp	10

Diff


You can compare backups of the core configuration file with your current configuration. This can be useful if, for example:

- A previously configured feature is no longer functioning, and you are not sure what in the configuration has changed.

- You want to recreate something configured previously, but do not remember what the settings were.

Difference-finding programs, such as [WinMerge](#) and the original [diff](#) can help you to quickly find all changes. They can compare your configurations, line by line, and highlight parts that are new, modified, or deleted.

Configuration differences highlighted in WinMerge




For instructions, see the documentation for your diff program.

Save debug file

A Save System State feature allows you to create an archive of your various configuration files, logs and other details used to help in diagnosing any issues that may arise. The file can be saved locally or uploaded to an FTP server.

1. Go to Global > System > Debug. Click on the **Save Debug File**.
2. The save debug will run. Only when the Running Status becomes "standby" can you save another debug file. Note: You can have at most three debug files.
3. When the file is ready (standby in running status), download it via GUI or upload it to your FTP server.



Save debug file

Settings	Guidelines
Name	system_debug_file_<date>_<time>
Running Status	Running status indicates the final download status. Running —the system is still collecting and compressing the debug file to generate a download file. Standby —file ready to download.
Upload Status	The upload FTP status. Running —it's uploading. Standby —upload completed.

Solutions by issue type

Recommended solutions vary by the type of issue.:

- [Login issues](#)
- [Connectivity issues](#)
- [Resource issues](#)

Login issues

If an administrator is entering his or her correct account name and password, but cannot log in from some or all computers, examine that account's trusted host definitions. It should include all locations where that person is allowed to log in, such as your office, but should *not* be too broad.

Connectivity issues

One of your first tests when configuring a new policy should be to determine whether allowed traffic is flowing to your servers. Investigate the following connectivity issues if traffic does not reach the destination servers:

- Is there a FortiADC policy for the destination servers? By default, FortiADC allows traffic to reach a backend server. However, the virtual servers must also be configured before traffic can pass through.
- If your network utilizes secure connections (HTTPS) and there is no traffic flow, is there a problem with your certificate?

Checking hardware connections

If there is no traffic flowing from the FortiADC appliance, you want to rule out hardware problems.

To check hardware connections:

- Ensure the network cables are properly plugged in to the interfaces on the FortiADC appliance.
- Ensure there are connection lights for the network cables on the appliance.
- Change the cable if the cable or its connector are damaged or you are unsure about the cable's type or quality.
- Connect the FortiADC appliance to different hardware to see if that makes a difference.
- In the web UI, go to System > Networking > Interface and ensure the link status is up for the interface. If the status is down (down arrow on red circle), edit the configuration to change its status to Up.

You can also enable an interface in CLI, for example:

```
config system interface
edit port2
set status up
end
```

If any of these checks solve the problem, it was a hardware connection issue. You should still perform some basic software tests to ensure complete connectivity.

If the hardware connections are correct and the appliance is powered on but you cannot connect using the CLI or web UI, you may be experiencing bootup problems. See [Restoring firmware \("clean install"\)](#).

Checking routing

The `ping` and `tracert` utilities are useful for investigating issues with network connectivity and routing.

Since you typically use these tools to troubleshoot, you can allow ICMP, the protocol used by these tools, in firewall policies and on interfaces only when you need them. Otherwise, disable ICMP for improved security and performance.

By default, FortiADC appliances do not respond to `ping` and `tracert`. However, if the appliance does not respond, and there are no firewall policies that block it, ICMP type 0 (`ECHO_RESPONSE`) might be effectively disabled.

To enable ping and tracert responses:

1. Go to Networking > Interface.
2. Select the row for the network interface and click the edit icon.

3. Under Allow Access, enable ping.
4. Save the update.

The appliance should now respond when another device such as your management computer sends a `ping` or `tracert` to that network interface.



Note: Disabling ping only prevents the system from *receiving* ICMP type 8 (ECHO_REQUEST) and traceroute-related UDP. It does not disable CLI commands such as `execute ping` or `execute traceroute` that send such traffic.

To verify routes between clients and your servers:

1. Attempt to connect *through* the FortiADC appliance, from a client to a backend server, via HTTP and/or HTTPS. If the connectivity test fails, continue to the next step.
2. Use the `ping` command on both the client and the server to verify that a route exists between the two. Test traffic movement in both directions: from the client to the server, and the server to the client. Servers do not need to be able to initiate a connection, but must be able to send reply traffic along a return path. If the routing test *succeeds*, continue with Step 4. If the routing test *fails*, continue to the next step.
3. Use the `tracert` or `traceroute` command on both the client and the server (depending on their operating systems) to locate the point of failure along the route. If the route is broken when it reaches the FortiADC appliance, first examine its network interfaces and routes. To display network interface addresses and subnets, enter the CLI command:

```
show system interface
```

 To display all recently-used routes with their priorities, enter the CLI command:

```
diagnose netlink route list
```

 You may need to verify that the physical cabling is reliable and not loose or broken, that there are no IP address or MAC address conflicts or blacklisting, misconfigured DNS records, and otherwise rule out problems at the physical, network, and transport layer. If these tests *succeed*, a route exists, but you cannot connect using HTTP or HTTPS, an application-layer problem is preventing connectivity.
4. For application-layer problems, on the FortiADC, examine the:
 - virtual server policy and all components it references
 - certificates (if connecting via HTTPS)
 - server service/daemon

On routers and firewalls between the host and the FortiADC appliance, verify that they permit HTTP and/or HTTPS connectivity between them.

Testing for connectivity with ping

The `ping` command sends a small data packet to the destination and waits for a response. The response has a timer that may expire, indicating that the destination is unreachable via ICMP.

ICMP is part of Layer 3 on the OSI Networking Model. `ping` sends Internet Control Message Protocol (ICMP) ECHO_REQUEST ("ping") packets to the destination, and listens for ECHO_RESPONSE ("pong") packets in reply.

Some networks block ICMP packets because they can be used in a ping flood or denial of service (DoS) attack if the network does not have anti-DoS capabilities, or because `ping` can be used by an attacker to find potential targets on the network.

Beyond basic existence of a possible route between the source and destination, `ping` tells you the amount of packet loss (if any), how long it takes the packet to make the round trip (latency), and the variation in that time from packet to packet (jitter).

If `ping` shows *some* packet loss, investigate:

- cabling to eliminate loose connections
- ECMP, split horizon, or network loops
- all equipment between the ICMP source and destination to minimize hops

If `ping` shows *total* packet loss, investigate:

- cabling to eliminate incorrect connections
- all firewalls, routers, and other devices between the two locations to verify correct IP addresses, routes, MAC lists, and policy configurations

If `ping` finds an outage between two points, use `traceroute` to locate exactly where the problem is.

To use ping:

Log into the CLI via either SSH, Telnet, or the CLI Console widget of the web UI.

1. If you want to adjust the behavior of `execute ping`, first use the `execute ping-options` command.
2. Enter the command:

```
execute ping <destination_ipv4>
```

where `<destination_ipv4>` is the IP address of the device that you want to verify that the appliance can connect to, such as `192.168.1.1`.

3. If the appliance can reach the host via ICMP, output similar to the following appears:

```
PING 192.168.1.1 (192.168.1.1): 56 data bytes
64 bytes from 192.168.1.1: icmp_seq=0 ttl=253 time=6.5 ms
64 bytes from 192.168.1.1: icmp_seq=1 ttl=253 time=7.4 ms
64 bytes from 192.168.1.1: icmp_seq=2 ttl=253 time=6.0 ms
64 bytes from 192.168.1.1: icmp_seq=3 ttl=253 time=5.5 ms
64 bytes from 192.168.1.1: icmp_seq=4 ttl=253 time=7.3 ms

--- 192.168.1.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 5.5/6.5/7.4 ms
```

If the appliance cannot reach the host via ICMP, output similar to the following appears:

```
PING 10.0.0.1 (10.0.0.1): 56 data bytes
Timeout ...
Timeout ...
Timeout ...
Timeout ...
Timeout ...

--- 10.0.0.1 ping statistics ---
5 packets transmitted, 0 packets received, 100% packet loss
```

“100% packet loss” and “Timeout” indicates that the host is not reachable.



To verify that routing is bidirectionally symmetric, you should also ping the appliance.

Testing routes and latency with traceroute

The traceroute utility sends ICMP packets to test each hop along the route. It sends three packets to the destination, and then increases the time to live (TTL) setting by one, and sends another three packets to the destination. As the TTL increases, packets go one hop farther along the route until they reach the destination.

Most traceroute commands display their maximum hop count—that is, the maximum number of steps it will take before declaring the destination unreachable—before they start tracing the route. The TTL setting may result in routers or firewalls along the route timing out due to high latency.

Where ping only tells you if the signal reached its destination and returned successfully, traceroute shows each step of its journey to its destination and how long each step takes. If you specify the destination using a domain name, the traceroute output can also indicate DNS problems, such as an inability to connect to a DNS server.

By default, the traceroute utility uses UDP with destination ports numbered from 33434 to 33534. The traceroute utility usually has an option to specify use of ICMP ECHO_REQUEST (type 8) instead, as used by the Windows tracert utility. If you have a firewall and you want traceroute to work from both machines (Unix-like systems and Windows) you will need to allow both protocols inbound through your firewall (UDP ports 33434 - 33534 and ICMP type 8).

To use traceroute:

1. Log into the CLI via either SSH, Telnet, or the CLI Console widget of the web UI.
2. Enter the command:

```
execute traceroute {<destination_ipv4> | <destination_fqdn>}
```

where {<destination_ipv4> | <destination_fqdn>} is a choice of either the device's IP address or its fully qualified domain name (FQDN).

For example, you might enter:

```
execute traceroute www.example.com
```

If the appliance *has* a complete route to the destination, output similar to the following appears:

```
traceroute to www.fortinet.com (66.171.121.34), 32 hops max, 84 byte packets
 1 172.16.1.2 0 ms 0 ms 0 ms
 2 209.87.254.221 <static-209-87-254-221.storm.ca> 2 ms 2 ms 2 ms
 3 209.87.239.129 <core-2-g0-1-1104.storm.ca> 2 ms 1 ms 2 ms
 4 67.69.228.161 2 ms 2 ms 3 ms
 5 64.230.164.17 <core2-ottawa23_POS13-1-0.net.bell.ca> 3 ms 3 ms 2 ms
 6 64.230.132.234 <core2-ottawatic_POS5-0-0.net.bell.ca> 20 ms 20 ms 20 ms
 7 64.230.132.58 <core4-toronto21_POS0-12-4-0.net.bell.ca> 24 ms 21 ms 24 ms
 8 64.230.138.154 <bx4-toronto63_so-2-0-0-0.net.bell.ca> 8 ms 9 ms 8 ms
 9 64.230.185.145 <bx2-ashburn_so2-0-0.net.bell.ca> 23 ms 23 ms 23 ms
10 12.89.71.9 23 ms 22 ms 22 ms
11 12.122.134.238 <cr2.wswdc.ip.att.net> 100 ms 12.123.10.130 <cr2.wswdc.ip.att.net> 101 ms
    102 ms
12 12.122.18.21 <cr1.cgil.ip.att.net> 101 ms 100 ms 99 ms
13 12.122.4.121 <cr1.sffca.ip.att.net> 100 ms 98 ms 100 ms
14 12.122.1.118 <cr81.sj2ca.ip.att.net> 98 ms 98 ms 100 ms
15 12.122.110.105 <gar2.sj2ca.ip.att.net> 96 ms 96 ms 96 ms
16 12.116.52.42 94 ms 94 ms 94 ms
```

```
17 203.78.181.10 88 ms 87 ms 87 ms
18 203.78.181.130 90 ms 89 ms 90 ms
19 66.171.121.34 <fortinet.com> 91 ms 89 ms 91 ms
20 66.171.121.34 <fortinet.com> 91 ms 91 ms 89 ms
```

Each line lists the routing hop number, the IP address and FQDN (if any) of that hop, and the 3 response times from that hop. Typically a value of <1ms indicates a local router.

If the appliance does not have a complete route to the destination, output similar to the following appears:

```
traceroute to 10.0.0.1 (10.0.0.1), 32 hops max, 84 byte packets
1 172.16.1.2 0 ms 0 ms 0 ms
2 172.16.1.10 0 ms 0 ms 0 ms
3 * * *
4 * * *
```

The asterisks (*) indicate no response from that hop in the network routing.

Examining the routing table

When a route does not exist, or when hops have high latency, examine the routing table. The routing table is where the FortiADC appliance caches recently used routes.

If a route is cached in the routing table, it saves time and resources that would otherwise be required for a route lookup. If the routing table is full and a new route must be added, the oldest, least-used route is deleted to make room.

To check the routing table in the CLI, enter:

```
diagnose netlink route list
```

Examining server daemons

If a route exists, but you cannot connect to the web UI using HTTP or HTTPS, an application-layer problem is preventing connectivity.

Verify that you have enabled HTTPS and/or HTTP on the network interface. Also examine routers and firewalls between the host and the FortiADC appliance to verify that they permit HTTP and/or HTTPS connectivity between them. Finally, you can also use the CLI command to verify that the daemons for the web UI and CLI, such as `sshd`, `cli`, `nginx`, and `php-fpm` are running and not overburdened:

```
diagnose system top delay 10
```

Checking port assignments

If you are attempting to connect to FortiADC on a given network port, and the connection is expected to occur on a different port number, the attempt will fail. For a list of ports used by FortiADC, see [Appendix B: Port Numbers](#).

Performing a packet trace

When troubleshooting malformed packet or protocol errors, it helps to look inside the protocol headers of packets to determine if they are traveling along the route you expect, and with the flags and other options you expect.



If you configure virtual servers on your FortiADC appliance, packets' destination IP addresses will be those IP addresses, not the physical IP addresses (i.e., the IP address of port1, etc.). An ARP update is sent out when a virtual IP address is configured.

If the packet trace shows that packets *are* arriving at your FortiADC appliance's interfaces but no HTTP/HTTPS packets egress, check that:

- Physical links are firmly connected, with no loose wires
- Network interfaces are brought up
- Link aggregation peers, if any, are up
- VLAN IDs, if any, match
- Virtual servers exist, and are enabled
- Matching policies exist, and are enabled
- If using HTTPS, valid server/CA certificates exist
- IP-layer and HTTP-layer routes, if necessary, match
- Servers are responsive, if server health checks are configured and enabled

Checking the SSL/TLS handshake & encryption

If the client is attempting to make an HTTPS connection, but the attempt fails after the connection has been initiated, during negotiation, the problem may be with SSL/TLS. Symptoms may include error messages such as:

- `ssl_error_no_cypher_overlap`
(Mozilla Firefox 9.0.1)
- `Error 113 (net::ERROR_SSL_VERSION_OR_CIPHER_MISMATCH): Unknown error.`
(Google Chrome 16.0.912.75 m)

The handshake is between the client and FortiADC. If the connection cannot be established, verify that the browser supports one of the key exchanges, encryption algorithms, and authentication (hashes) offered by FortiADC.

If you are not sure which cipher suites are currently supported, you can use SSL tools such as [OpenSSL](#) to discover support. For example, you could use this client-side command to know whether the server or FortiADC supports strong (HIGH) encryption:

```
openssl s_client -connect example.com:443 -cipher HIGH
```

or supports deprecated or old versions such as SSL 2.0:

```
openssl s_client -ssl2 -connect example.com:443
```

Resource issues

This section includes troubleshooting questions related to sluggish or stalled performance.

Monitoring traffic load

Heavy traffic loads can cause sustained high CPU or RAM usage. If this is unusual, no action is required. However, sustained heavy traffic load might indicate that you need a more powerful FortiADC model.

In the web UI, you can view traffic load two ways:

- Monitor current HTTP traffic on the dashboard. Go to System > Dashboard > Virtual Server and examine the throughput graphs.
- Examine traffic history in the traffic log. Go to Logs & Report > Log Browsing > Traffic Log.

DoS attacks

A prolonged denial of service (DoS) can bring your servers down if your FortiADC appliance and your network devices are not configured to prevent it. To prevent DoS attacks, enable the DoS and connection limit features. Also, configure protections on your FortiGate and other network devices. DoS attacks can use a variety of mechanisms. For in-depth protection against a wide variety of DoS attacks, you can use a specialized appliance such as FortiDDoS.

In the web UI, you can watch for attacks in two ways:

- Monitor current traffic on the dashboard. Go to System > Dashboard and examine the system-wide throughput.
- Examine attack history in the traffic log. Go to Logs & Report > Log Browsing > Security Log.

Resetting the configuration

If you will be selling your FortiADC appliance, or if you are not sure what part of your configuration is causing a problem, you can reset it to its default settings and erase data. (If you have not updated the firmware, this is the same as resetting to the factory default settings.)



Important: Back up the configuration before performing a factory reset.

To delete your data from the system, connect to the CLI and enter this command:

```
execute formatlogdisk
```

To reset the configuration, connect to the CLI and enter this command:

```
execute factoryreset
```

Restoring firmware (“clean install”)

Restoring (also called re-imaging) the firmware can be useful if:

- you are unable to connect to the FortiADC appliance using the web UI or the CLI
- you want to install firmware *without* preserving any existing configuration (i.e. a “*clean install*”)
- a firmware version that you want to install requires a different size of system partition (see the Release Notes accompanying the firmware)
- a firmware version that you want to install requires that you format the boot device (see the Release Notes accompanying the firmware)

The procedure in this section applies to physical appliances. Restoring firmware re-images the boot device. Also, restoring firmware can only be done during a boot interrupt, before network connectivity is available, and therefore requires a local console connection to the CLI. It cannot be done through an SSH or Telnet connection.



Alternatively, if you cannot physically access the appliance's local console connection, connect the appliance's local console port to a terminal server to which you have network access. Once you have used a client to connect to the terminal server over the network, you will be able to use the appliance's local console through it. However, be aware that from a remote location, you may not be able to power cycle the appliance if abnormalities occur.

For virtual appliances, you can use VMware to backup and restore virtual appliance images.



Important: Back up the configuration before performing a clean install.

To restore the firmware:

1. Download the firmware file from the Fortinet Customer Service & Support website:
<https://support.fortinet.com/>
2. Connect your management computer to the FortiADC console port using a RJ-45-to-DB-9 serial cable or a null-modem cable.
3. Initiate a local console connection from your management computer to the CLI of the FortiADC appliance, and log in as the `admin` administrator, or an administrator account whose access profile contains Read-Write permissions in the Maintenance category.
4. Connect port1 of the FortiADC appliance directly or to the same subnet as a TFTP server.
5. Copy the new firmware image file to the root directory of the TFTP server.
6. If necessary, start your TFTP server. (If you do not have one, you can temporarily install and run one such as `tftpd` ([Windows](#), [Mac OS X](#), or [Linux](#)) on your management computer.)



TFTP is not secure, and it does not support authentication. You should run it only on trusted administrator-only networks, and never on computers directly connected to the Internet. Turn off `tftpd` off immediately after completing this procedure.

7. Verify that the TFTP server is currently running, and that the FortiADC appliance can reach the TFTP server. To use the FortiADC CLI to verify connectivity, enter the following command:
`execute ping 192.168.1.168`
where `192.168.1.168` is the IP address of the TFTP server.
8. Enter the following command to restart the FortiADC appliance:
`execute reboot`
As the FortiADC appliances starts, a series of system startup messages appear.
Press any key to display configuration menu.....
9. Immediately press a key to interrupt the system startup.



You have only 3 seconds to press a key. If you do not press a key soon enough, the FortiADC appliance 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.
[Q]: Quit menu and continue to boot with default firmware.
[H]: Display this list of options.
```

Enter G,F,B,Q, or H:

Please connect TFTP server to Ethernet port "1".

10. If the firmware version requires that you first format the boot device before installing firmware, type `F`. Format the boot disk before continuing.

11. Type `G` to get the firmware image from the TFTP server.

The following message appears:

```
Enter TFTP server address [192.168.1.168]:
```

12. Type the IP address of the TFTP server and press Enter.

The following message appears:

```
Enter local address [192.168.1.188]:
```

13. Type a temporary IP address that can be used by the FortiADC appliance to connect to the TFTP server.

The following message appears:

```
Enter firmware image file name [image.out]:
```

14. Type the file name of the firmware image and press Enter.

The FortiADC appliance downloads the firmware image file from the TFTP server and displays a message similar to the following:

```
MAC:00219B8F0D94
```

```
#####
```

```
Total 28385179 bytes data downloaded.
```

```
Verifying the integrity of the firmware image..
```

```
Save as Default firmware/Backup firmware/Run image without saving:[D/B/R]?
```



If the download fails after the integrity check with the error message:

invalid compressed format (err=1)

but the firmware matches the integrity checksum on the Fortinet Customer Service & Support website, try a different TFTP server.

15. Type `D`.

The FortiADC appliance downloads the firmware image file from the TFTP server. The FortiADC appliance installs the firmware and restarts. The time required varies by the size of the file and the speed of your network connection.

The FortiADC appliance reverts the configuration to default values for that version of the firmware.

16. To verify that the firmware was successfully installed, log in to the CLI and type:

```
get system status
```

The firmware version number is displayed.

17. Either reconfigure the FortiADC appliance or restore the configuration file.

Additional resources

Fortinet also provides these resources:

- The Release Notes provided with your firmware
- [Technical documentation](#) (reference guides, installation guides, and other documents)
- [Knowledge base](#) (technical support articles)
- [Forums](#)
- [Online campus](#) (tutorials and training materials)

If you have problem using FortiADC, check within your organization first. You can save time and effort during the troubleshooting process by checking if other FortiADC administrators have experienced a similar problem before.


If you cannot resolve the issue on your own, contact [Fortinet Customer Service & Support](#).

Chapter 21: System Dashboard

The default Dashboard page opens when you log into the system root (or a virtual domain). You can also navigate to the Dashboard from any other pages of the GUI by selecting Dashboard>Main on the navigation bar.

This chapter covers the following topics:

- [Widgets on page 566](#)
- [Dashboard management tools on page 566](#)



Widgets

The default Dashboard page displays a collection of 10 widgets, which fall into following categories:

- Summary
- Interface Summary
- Virtual Server Summary
- Host Summary
- Link Group Summary
- Log Event
- Server Load Balance
- Security
- Global Load Balance

Click **See Detail** to see more, it should show up when you **hover** over the widget.

[Dashboard widgets on page 566](#) highlights the information contained in each of the widgets.


Dashboard widgets

Widget	Description and Utilities
Interface Summary	Shows type, name, mode, IP address, status and allowed forms of access for each port, as well as its virtual domain.
Virtual Server Summary	Shows the status of your virtual servers, including health, throughput, and pool number.
Host Summary	Shows the health and response of your host.
Link Group Summary	Gives the status of your link group.
Log Event	Tells you how many emergencies, alerts, and warnings, etc., that are going on.
Server Load Balance	Shows your throughput, concurrent, and connections, while displaying the inbound/outbound throughput as a graph.
Security	Shows you the security status.
Global Load Balance	Shows the host and response per second

Dashboard management tools

When you click the Dashboard bar on the very left column, it will drop down and likely show the default dashboard, **Main**.

Dashboard pop-up list menu



Adding a dashboard

This option allows you to create your own dashboard with widgets of your choice.

To add a custom dashboard:


1. Click **Create Dashboard**.
2. Specify a unique name for the dashboard.
3. Click **save**. The name of the dashboard appears under Dashboard>Main on the navigation bar.

Editing a dashboard

Note: This option only applies to custom dashboards that you have created. The Main (default) dashboard cannot be edited.

To edit a custom dashboard:

1. On the navigation bar, select the name of the dashboard.
2. Click **Edit** on the far right top corner.
3. Rename the dashboard, if you like. (Note: If you change the name, be sure to click the Save button.)



Deleting a dashboard

Note: This option only applies to custom dashboards that you have created. The Main (default) dashboard cannot be deleted.

To delete a custom dashboard:

1. On the navigation bar, select the name of the dashboard.
2. Click **Delete** on the far right, as shown in the illustration above.
3. Read the warning message onscreen.
4. Click Delete if you've decided to remove the selected dashboard.

Adding Features

Note: You can add onto the dashboard various summaries and reports, say, on Server Load Balance, thus allowing you to have 'one look' at many of the FortiADC's features.

To add or remove features:

1. On the top right, select **Edit**.
2. Open up the **Edit Dashboard** window.
3. Decide which features you want to see, by shifting the button to **On**.
4. Click Save when done.

Chapter 22: FortiView

The FortiView pages display important information about your FortiADC appliance, which includes the logical topology of real-server pools and their members within each virtual server, server load-balancing information, security, and some other system events and alerts.

The information is organized by topic as follows:

- [Physical Topology on page 569](#)
- [HA Status on page 570](#)
- [Server Load Balance on page 570](#)
- [Logical Topology on page 570](#)
- [Virtual Servers on page 576](#)
- [Data Analytics on page 581](#)
- [Traffic Logs on page 585](#)
- [Link Load Balance on page 587](#)
- [Logical Topology on page 587](#)
- [Link Group on page 588](#)
- [Global Load Balance on page 589](#)
- [Logical Topology on page 589](#)
- [Host on page 590](#)
- [Data Analytics](#)
- [Security on page 590](#)
- [Threat Map on page 590](#)
- [Data Analytics on page 591](#)
- [Security Logs on page 593](#)
- [All Segments on page 594](#)
- [Event Logs on page 594](#)
- [Alerts on page 595](#)
- [All Sessions on page 596](#)

Physical Topology

This page displays the physical topology of your FortiADC network structure. It shows your FortiADC appliance or appliances identified by serial number and the real servers connected to it

Note: This page is read-only.

HA Status

The HA Status page shows the information about FortiADC's HA configurations and performance, as shown in [HA Status on page 570](#). It has the following sections:

- HA Cluster—Shows the serial number, node ID, IP address, and source configuration of the each device in HA mode.
- Link—Shows the link status: up or down.
- System—Shows the system status: pass or fail.
- Remote IP—Shows the remote IP addresses and their status: up or down.
- Sync Statistics—Shows the number of sent and received sync packets.
- Device Management Errors—Shows the number of device management errors by duplicate node ID and by version mismatch.
- Traffic Status—Shows traffic group name, current device node, next device node, preempt, and floating IP addresses.

HA status page

1. Click System > **High Availability**.
2. It will reveal the first diagram. There is, however, an extra step you must take.
3. Hover your mouse over the gray area, here highlighted in red for your convenience.
4. It will display the **HA Status** page, as a pop-up.

Server Load Balance


The FortiView>Server Load Balance menu shows server load-balancing configurations on your FortiADC. It has the following sub-menus:

- [Logical Topology](#)
- [Virtual Server](#)
- [Data Analytics](#)
- [Traffic Logs](#)

Logical Topology

The Server Load Balance>Logical Topology page uses the tree-view format to show the internal configuration of each virtual server on your FortiADC appliance. Depending on the actual configuration, the diagram may show content routing, schedule pools, real-server pools, and real-server pool members configured on a virtual server, as illustrated in [Logical topology on page 570](#).

Logical topology



The image above is a partial screen capture of the FortiView > Logical Topology page. It shows the internal configuration of a virtual server named "L7_HTTP", which has the following configurations on it:

- A real-server pool named "HTTPServicePool" which contains 9 members (real servers) in it.
- It is using Port 7, which is up (working).

Apart from viewing the internal configurations of virtual servers, you can also drill down into the components (except for content routing and schedule group) for details by clicking their corresponding icons. Below highlights what you will see when you click any of the following icons:

- Virtual server—Opens the page with details of that virtual server. See [Virtual server details on page 577](#)
- Real-server pool—Opens the page with details of the real-server pool. See [Real server pool details on page 580](#)
- Real server —Opens the page showing details of the real server. See [Real-server pool member details on page 575](#)

Virtual server details

This page shows detailed information about the virtual server you select.

Go to FortiView > Server Load Balance > Virtual Server.


Select the virtual server you want by clicking on its name, on the left side; it will lead you into the page illustrated below.

Below the virtual server name are four tabs, which allow you to display the data about the virtual server by

- [Analytics](#)
- [Health](#)
- [Client](#)
- [Session](#)
- [Persistence](#)
- [Statistics](#)

Analytics

The Analytics page provides real-time analysis of data about the virtual server using colored icons, charts, and diagrams, etc. See the following figure:



In the upper-right corner of the page is a drop-down box. Click the down arrow to pull down the drop-down menu which contains for setting the time frame for the graph the bottom of the page. The options are:

- 1 Hour
- 6 Hour
- 1 Day
- 1 Week
- 1 Month
- 1 Year


In the lower-right corner of the page is another drop-down box which contains data options you can choose to show in the graph. The options are:

- End to End Timing (default)
- Throughput
- Concurrent Connections
- Connections per Second
- Request

Health

This page uses a bar graph to show the virtual server's health status in a specific time frame, as shown in the following figure:


Health



In the upper-right corner of the page is a drop-down menu, which provides the time frames that you can choose from for the graph. The options are the same as those described in the section above.

Client

This page depicts the clients of the virtual server across the globe, as illustrated in the following figure:



The Client page has the following sections:

- **Location**—This part of the page shows the top five countries in the world where most of the client traffic is coming from. The dots on the map show the locations of those countries. Mouse over a dot to see the name of that country in the tool tip. The + (plus) and – (minus) signs allow you to zoom in or out on the map. The table below the map shows percentage of client traffic from each of those countries: the green up arrows indicate that traffic is increasing; the percentage in green indicates the percentage increase in client traffic since the last data was sampled; and the percentage in black indicates the percentage of traffic each of the counties accounts for in total client traffic.
- **Device**—This part of the page shows the types of devices that the clients are using, the percentage increase in the use of each of the devices since the last data was sampled, and the percentage of a type of device among all devices that are used.
- **Browser**—This part of the page shows the web browsers that the clients are using, the percentage increase in the use of each of the browsers, and the percentage of each of the browsers among all browsers that are used.

- **Operating System**—This part of the page shows the operating systems that the clients are using, the percentage increase in the use of each of the operating systems since the last data was sampled, and the percentage that each operating system accounts for among all the operating systems that are used.
- **Top URLs**—This part of the page shows the top five web browsers that the clients are using, and the percentage that each of them accounts for among all the browsers that are used.

Session

This page shows all the active sessions that the virtual server currently maintains. The table provides the same information and tools as described in [All Sessions on page 596](#)

Persistence

This page shows all the active persistence sessions that the virtual server currently maintains. The table provides the same information and tools as described in [All Sessions on page 596](#)

Statistics

This page shows the statistics for this particular VS, its counters. After reboot, all the counters will restart on their own. It is limited to HTTP/HTTPS virtual servers.

Real server pool details

The real server pool details page ([on page 574](#)) shows detailed information about the real server pool you select (click) on the FortiView > Logical Topology page. See [Logical Topology on page 570](#)

The top of the page shows the name of the real server pool and the virtual server to which it is assigned. Below the real server pool name are two tabs—Members and Health. The former shows information about the members (real servers) in the real server pool, whereas the latter shows the health state of the real server pool in general.

Member

The Member pages (see the image above) shows key information about the real servers in a real server pool, as described in [Real server pool member information on page 574](#).

Real server pool member information

Column title	Description
Name	The name of a real server pool member (real server). Note: Clicking the name of a real server opens the page with detailed information about the real server.
Status	Shows the status of a real server pool member, which can be either of the following: • Enable • Disable
Address	The IP address of a real server pool member (real server).

Column title	Description
Port	The port used by a real server pool member.
Weight	The weight assigned to a real server pool member.
Throughput (bits/sec)	The graph shows the change in a real server's throughput in bits per second over the specified period of time. Note: If you mouse over a specific point in the graph, a tool tip will pop up showing the number of bits per second that a real server pool member transmits at that time point.
Concurrent	The graph shows the change in the number of concurrent connections with the real server pool member over the specified period of time. Note: If you mouse over a specific point in the graph, a tool tip will pop up showing the number of concurrent connections at that time point.
Health	The color of the heart icon indicates the health state of a real server pool member, which can be either of the following: - Green = healthy - Red = Unhealthy


Health

This graph shows the overall health status of the real server pool.

Real-server pool member details

This page shows detailed information about the real server pool member selected on the FortiView > Logical Topology page. See the following figure:

Real server pool member details



Across the top of the page is the name of the real server pool member, preceded by the name of the virtual server and the name of the real server pool. The page has two display options—Analytics and Health, as represented by the two tabs below the name of the real server pool member.

Analytics


The Analytics page uses charts and diagrams to help you analyze data related to the real server pool member. The diagram and the pie chart in the upper part of the page show the dynamic changes in server round-trip time and application response time.

The page has two drop-down menus which allow you to set the time frame and data type displayed in the line chart at the bottom of the page.

Virtual Servers

The FortiView>Server Load Balance>Virtual Server page ([Virtual server on page 576](#)) is a table that shows some key configuration and traffic information about the virtual servers that have the FortiView feature enabled on them. You can enable FortiView on a virtual server using Server Load Balance>Virtual Server>Add>Advanced Mode>Traffic Log>FortiView>ON. You can also show or hide all the virtual servers on or from this page using the Enable All or Disable All button across the top of the table, regardless whether you have FortiView enabled or not when configuring the virtual servers.

Virtual server



[Virtual Server table on page 576](#) describes the information on the FortiView > Server Load Balance > Virtual Server page.

Virtual Server table

Column title	Description
Name	The name of a virtual server Note: Clicking the name of a virtual server opens the page with detailed information about the virtual server.

Column title	Description
Type	The type of virtual servers, which can be one of the following: • I2 = Layer 2 • I4 = Layer 4 • I7 = Layer 7
Address	The IP address of a virtual server. Note: For Layer-2 virtual servers, this field shows 0.0.0.0.
Port	The port used by a virtual server, which depends on the type of traffic the port is handling.
Pool	The name of a real-server pool configured on a virtual server. Note: Clicking the name of a real-server pool opens the page with details of that real-server pool.
Throughput (bits/sec)	The graph shows the change in a virtual server's throughput in terms of bits per second over the past 24 hours. Note: The data was sampled at 60 different time points over the last 24 hours (i.e., once every 24 minute). If you mouse over a specific point in the graph, a tool tip will pop up showing the throughput for that time point.
Concurrent	The graph shows the change in the number of concurrent connections with the virtual server over the last 24 hours. Note: The data was sampled at 60 different time points over the last 24 hours (i.e., once every 24 minute). If you mouse over a specific point in the graph, a tool tip will pop up showing the number of concurrent connections at that time point.
Connections (counts/sec)	The graph shows the change in the number of connections with the virtual server over the last 24 hours. Note: The data was sampled at 60 different time points over the last 24 hours (i.e., once every 24 minute). If you mouse over a specific point in the graph, a tool tip will pop up showing the number of connections for that time point.
Health	The color of the heart icon indicates the health state of a virtual server, which can be either of the following: • Green = healthy • Red = Unhealthy

Virtual server details

This page shows detailed information about the virtual server you select.

Go to FortiView > Server Load Balance > Virtual Server.

Select the virtual server you want by clicking on its name, on the left side; it will lead you into the page illustrated below.


Below the virtual server name are four tabs, which allow you to display the data about the virtual server by

- [Analytics](#)
- [Health](#)
- [Client](#)

- [Session](#)
- [Persistence](#)
- [Statistics](#)

Analytics

The Analytics page provides real-time analysis of data about the virtual server using colored icons, charts, and diagrams, etc. See the following figure:



In the upper-right corner of the page is a drop-down box. Click the down arrow to pull down the drop-down menu which contains for setting the time frame for the graph the bottom of the page. The options are:

- 1 Hour
- 6 Hour
- 1 Day
- 1 Week
- 1 Month
- 1 Year


In the lower-right corner of the page is another drop-down box which contains data options you can choose to show in the graph. The options are:

- End to End Timing (default)
- Throughput
- Concurrent Connections
- Connections per Second
- Request

Health

This page uses a bar graph to show the virtual server's health status in a specific time frame, as shown in the following figure:


Health



In the upper-right corner of the page is a drop-down menu, which provides the time frames that you can choose from for the graph. The options are the same as those described in the section above.

Client

This page depicts the clients of the virtual server across the globe, as illustrated in the following figure:



The Client page has the following sections:

- **Location**—This part of the page shows the top five countries in the world where most of the client traffic is coming from. The dots on the map shows the locations of those countries. Mouse over a dot to see the name of that country in the tool tip. The + (plus) and – (minus) signs allow you to zoom in or out on the map. The table below the map shows percentage of client traffic from each of those countries: the green up arrows indicate that traffic is increasing; the percentage in green indicates the percentage increase in client traffic since the last data was sampled; and the percentage in black indicates the percentage of traffic each of the counties accounts for in total

client traffic.

- **Device**—This part of the page shows the types of devices that the clients are using, the percentage increase in the use of each of the devices since the last data was sampled, and the percentage of a type of device among all devices that are used.
- **Browser**—This part of the page shows the web browsers that the clients are using, the percentage increase in the use of each of the browsers, and the percentage of each of the browsers among all browsers that are used.
- **Operating System**—This part of the page shows the operating systems that the clients are using, the percentage increase in the use of each of the operating systems since the last data was sampled, and the percentage that each operating system accounts for among all the operating systems that are used.
- **Top URLs**—This part of the page shows the top five web browsers that the clients are using, and the percentage that each of them accounts for among all the browsers that are used.

Session

This page shows all the active sessions that the virtual server currently maintains. The table provides the same information and tools as described in [All Sessions on page 596](#)

Persistence

This page shows all the active persistence sessions that the virtual server currently maintains. The table provides the same information and tools as described in [All Sessions on page 596](#)

Statistics

This page shows the statistics for this particular VS, its counters. After reboot, all the counters will restart on their own. It is limited to HTTP/HTTPS virtual servers.

Real server pool details

The real server pool details page ([on page 580](#)) shows detailed information about the real server pool you select (click) on the FortiView > Logical Topology page. See [Logical Topology on page 570](#)

The top of the page shows the name of the real server pool and the virtual server to which it is assigned. Below the real server pool name are two tabs—Members and Health. The former shows information about the members (real servers) in the real server pool, whereas the latter shows the health state of the real server pool in general.

Member

The Member pages (see the image above) shows key information about the real servers in a real server pool, as described in [Real server pool member information on page 580](#).

Real server pool member information

Column title	Description
Name	The name of a real server pool member (real server). Note: Clicking the name of a real server opens the page with detailed information about the real server.

Column title	Description
Status	Shows the status of a real server pool member, which can be either of the following: • Enable • Disable
Address	The IP address of a real server pool member (real server).
Port	The port used by a real server pool member.
Weight	The weight assigned to a real server pool member.
Throughput (bits/sec)	The graph shows the change in a real server's throughput in bits per second over the specified period of time. Note: If you mouse over a specific point in the graph, a tool tip will pop up showing the number of bits per second that a real server pool member transmits at that time point.
Concurrent	The graph shows the change in the number of concurrent connections with the real server pool member over the specified period of time. Note: If you mouse over a specific point in the graph, a tool tip will pop up showing the number of concurrent connections at that time point.
Health	The color of the heart icon indicates the health state of a real server pool member, which can be either of the following: • Green = healthy • Red = Unhealthy

Health

This graph shows the overall health status of the real server pool.

Data Analytics

The FortiView>Server Load Balance>Data Analytics page shows, initially, the **Dynamic Charts** page.

This is among three tabs:

- [Dynamic Charts on page 581](#)
- [Static Charts on page 585](#)
- [Statistics on page 585](#)

First we will speak of the **Dynamic Charts** page.

Dynamic Charts

In this tab you can customize your data analytics chart by using the **Add Widget** button, to create charts of your own. See the table below for details.

Note: Normally, the Data Analytics page automatically refreshes itself every few seconds so that new data can be added to the charts. You can stop the page from refreshing by clicking the **Enabled** button across the top of the page. The charts stop refreshing. as soon as the button turns to **Disabled**.

To add a widget (chart):

1. Click FortiView > Server Load Balance > Data Analytics.
2. Click the Add Widget button to open the Fast Report dialog.
3. Make the entries and selections as described in [Data Analytics Widget on page 582](#).
4. Click Save when done.

Data Analytics Widget

Chart/Graph	Description
Name	Enter a unique name for a chart.
SLB Subtype	Click the down arrow and select a server load-balancing data you want to show in the chart. • Top Source IP—Most used source IP addresses • Top Destination IP— Most used destination IP addresses • Top Browser— Most used web browsers • Top OS— Most used operating systems • Top Device—The type of device (PC vs. Mobile) with the most traffic • Top Domain— Most used domains • Top URL—

Chart/Graph	Description
	Most used URLs. • Top Referrer—Referrers which forwarded most traffic • Top Source Country—The countries where most of the traffic originated • Top Session—Sessions with the most traffic
History Chart	A "history" chart shows historical data that the system captured over a specific time period in the past. The option is turned OFF (disabled) by default, but you can click the button to turn it ON (enable it).

Chart/Graph	Description
	Note: If this option is turned off, the chart will be a pie chart. If it is turned on, then you will see a bar chart for most of the data types except for Session Total and Throughput Total which use line charts instead. Both bar charts and line charts have a time-range selector in their upper-right corner which allows you to select one of the following: • 10 Minutes • 1 Hour • 1 Day • 1 Week • 1 Month
Time Range	Click the down arrow to select one of the following time ranges: • 10 Minutes • 1 Hour • 1 Day • 1 Week • 1 Month Note: This option becomes unavailable if History Chart is enabled.
Data Type	Select either of the following: • Bandwidth

Chart/Graph	Description
	(default) Session
Top X	Specify a maximum value for the X axis. Note: The default is 5, but the valid values are from 3 to 7.
Top Y	Specify a maximum value for the Y axis. Note: The default is 5, but the valid values are from 3 to 7.

Static Charts

Here, by default, there will show up two static charts. The first will measure **SSL TPS**, i.e. transactions per second. The second measures **Compression Throughput**.

You have a time range of:

- 1 Hour
- 6 Hours
- 1 Day
- 1 Week
- 1 Month
- 1 Year

Statistics

Statistics shows you the status of the whole VDOM.

Traffic Logs

The FortiView>Server Load Balance>Traffic Logs page shows server load-balancing traffic logs that the system has generated.

Selecting log categories

The logs are organized into 10 categories, as indicated by the radio buttons across the top of the page. They are:

- SLB Layer 4
- SLB HTTP
- SLB TCPS
- SLB RADIUS
- GLB
- SLB SIP
- SLB RDP
- SLB DNS
- SLB RTSP
- SLB SMTP
- SLB DIAMETER
- SLB MySQL
- LLB


You can view any of these types of logs by clicking the corresponding radio button, and the page will be populated with logs that are available in that category. If no logs are available in that category, the page will come up blank (with no logs).

Viewing SLB traffic log details

All logs are presented in a tabular format, with each row being a log entry. The log table shows some key information contained in the logs, which may vary slightly depending on the log category you select.

You can view details of a log by clicking the corresponding Preview button, as illustrated below.

SLB traffic log details



The screenshot shows the FortiView interface with the 'Traffic Log' tab selected. Below the tab are radio buttons for various log types: SLB Layer 4, SLB HTTP, SLB TCPS, SLB RADIUS, GLB, SLB SIP, SLB RDP, SLB DNS, SLB RTSP, SLB SMTP, SLB RTMP, SLB DIAMETER, SLB MySQL, and LLB. The 'SLB HTTP' radio button is selected. Below the radio buttons are buttons for 'Filter Setting', 'Download', and 'Refresh'. The main area contains a table with the following columns: Date, Time, Source, Received Bytes, Destination, Sent Bytes, Service, Method, URL, Return Code, Virtual Server, Real Server Name, and a gear icon for settings. The table displays six log entries for HTTP traffic on 2018-09-19.

Date	Time	Source	Received Bytes	Destination	Sent Bytes	Service	Method	URL	Return Code	Virtual Server	Real Server Name	
2018-09-19	16:42:25	6.1.1.1	121	6.1.1.102	519652	https	get	/jde.html	200	1_JDE-Test	rs1	
2018-09-19	16:42:25	6.1.1.1	121	6.1.1.102	519652	https	get	/jde.html	200	1_JDE-Test	rs1	
2018-09-19	16:42:25	6.1.1.1	121	6.1.1.101	519652	https	get	/jde.html	200	1_JDE-Production	rs1	
2018-09-19	16:42:24	6.1.1.1	121	6.1.1.101	519652	https	get	/jde.html	200	1_JDE-Production	rs1	
2018-09-19	16:42:24	6.1.1.1	121	6.1.1.101	519652	https	get	/jde.html	200	1_JDE-Production	rs1	
2018-09-19	16:42:24	6.1.1.1	121	6.1.1.102	519652	https	get	/jde.html	200	1_JDE-Test	rs1	

Downloading SLB traffic logs

In the upper-right corner of the FortiView > Server Load Balance > Virtual Server page is a Download button. It enables you to download logs and save them in a .tar file. It comes in handy when you want to back up the logs for further analysis.

You can view the downloaded logs using a text-editing application. Below are some the most popular text editors you can use:

- WordPad (built-in in Microsoft Windows)
- Notepad ++
- EditPlus,
- Sublime

[View log messages in a text editor on page 587](#) shows the first three log entries when viewed in a text editor.

View log messages in a text editor

```
date=2017-10-17 time=01:57:10 log_id=100 type=event subtype=config
pri=information vd=root msg_id=434747638 user=admin ui=SSH
(172.30.176.110) action=add cfgpath=router static cfgobj=<No.>
cfgattr=2 logdesc=Change the configuration msg=added a new entry '2'
for "router static" on domain "root"
date=2017-10-17 time=01:57:10 log_id=100 type=event subtype=config
pri=information vd=root msg_id=434747637 user=admin ui=SSH
(172.30.176.110) action=edit cfgpath=system interface
cfgobj=allowaccess cfgattr=->ping logdesc=Change the configuration
msg=changed settings 'port12' for "system interface"
date=2017-10-17 time=01:57:10 log_id=100 type=event subtype=config
pri=information vd=root msg_id=434747636 user=admin ui=SSH
(172.30.176.110) action=edit cfgpath=system interface cfgobj=ip
cfgattr=0.0.0.0/0->4.1.255.254/16 logdesc=Change the configuration
msg=changed settings 'port12' for "system interface"
```

Link Load Balance

The FortiView>Link Load Balance menu shows link load-balancing configurations on your FortiADC. It has two sub-menus:

- [Logical Topology](#)
- [Link Group](#)

Logical Topology

The **Link Load Balance>Logical Topology** page shows the logical topology of link groups that have been configured.

Adding link groups

To add a link group:

1. Click the Add Link Group button.
2. Make desired entries or selections as described in [Configuring a link group on page 204](#)
3. Click Save when done.

Note: While in Editor View, you can click any component in the logical topology to edit or delete it.

Filtering link groups

The Add Filters button on top of the page allows you to customize the logical topology by:

- Availability
- Gateway Status
- Link Group Name
- Gateway Name
- Gateway IP

To add a filter:

1. Click the Add Filters button.
2. Select the filter.

You can use the same steps to apply multiple filters. Applied filters appear in front of the Add Filters button in the order they are added. You can remove a filter by clicking the x sign on it.

Link Group

The Link Load Balance>Link Group page shows link group configurations in a tabular format. It provides the following information about each gateway:

- Name
- IP Address
- Availability (Up or Down)
- Inbound Bandwidth
- Outbound Bandwidth
- Health Check

Monitoring traffic

You can display traffic going through a gateway using charts by selecting the corresponding check box in the Monitor column.

Monitoring traffic on a link

Editing gateway configuration

You can edit the configuration of a gateway for a link group by clicking the corresponding Edit button. For instructions on how to edit a gateway configuration, see [Configuring gateway links on page 206](#)

Global Load Balance

The FortiView>Global Load Balance menu shows global load-balancing configurations on your FortiADC. It has two sub-menus:

- [Logical Topology on page 589](#)
- [Host on page 590](#)
- [Data Analytics](#)

Logical Topology

The FortiView>**Global Load Balance**>**Logical Topology** page shows the logical topology of your global load balance configurations.

Adding hosts

To add a host:

1. Click Add Host.
2. Make desired entries or selections as described in [Configuring hosts on page 221](#)
3. Click Save when done.

Note: While in Editor View, you click any component in the logical topology to edit or delete it.

Filtering hosts

The Add Filters button on top of the page allows you to customize the logical topology by:

- Availability
- Host
- Domain Name
- VS Pool
- Server
- Server Member
- Data Center

To add a filter:

1. Click the Add Filters button.
2. Select the desired filter from the drop-down list menu.

Note: You can use the same steps to apply multiple filters. Applied filters appear in front of the Add Filters button in the order they are added. You can remove a filter by clicking the x sign on it.

Host

The FortiView>**Global Load Balance**>**Host** page shows global load-balancing host configurations in a tabular format. The first thing you see is the **Summary**. It shows you the health of all the hosts. If you want to see it in more detail, you can **Enable All Analytics**.

You will see the following information about each host:

- Name
- Host Domain
- Total Response
- Current
- Virtual Server Pool
- Response
- Health

To no longer view the **Analytics**, click **Disable All Analytics**.

Editing a host

Click the name of the host to edit it. For instructions on how to edit a global load balance host, see [Configuring hosts on page 221](#)

Click the number next to arrow icon in Virtual Server Pool to show the virtual servers inside the pool.

Security

The FortiView>Security menu shows network security information captured by FortiADC. The page has three sub-menus:

- [Threat Map](#)
- [Data Analytics](#)
- [WAF Security Logs](#)

Threat Map

The FortiView>Security>Threat Map page depicts the security threats to your FortiADC devices in real time. The darker part of the world map represents the part of the world at night, whereas the lighter areas are parts of the world in daylight. The device icons represent your FortiADC appliances deployed at various locations in the world. The shooting stars represent the live attacks on your FortiADC appliances as they occur.

The table at the bottom of the map lists the live threats as they occur, with the following information about each threat:

- Location—The country and the IP address where an attack comes from.
- Threat—The name or brief description of a threat
- Severity (score)—The level of severity of a threat, which can be high, medium, or low.
- Time—The date and time when an attack occurs.


The severity of threats are color-coded:

- High — Red
- Medium — Yellow
- Low — White.

The map and the table complement each other, showing you when the attacks occur, pinpointing where they come from, and telling you the nature and severity of the attacks so that you can make well-informed decision as to how to react to those threats.

You can open the Threat Map page by clicking FortiView > Security > Threat Map. [Threat map on page 591](#) shows the Threat Map with only one FortiADC appliance.

Threat map



Data Analytics

The FortiView>Security>Data Analytics page shows Web application firewall information in charts called "widgets". By default, the page is empty. You must create charts of your own using the Add Widget button.

Note: Normally, the Data Analytics page automatically refreshes itself every a few seconds so that new data can be added to the charts. You can stop the page from refreshing by clicking the Enabled button across the top of the page. The charts stop refreshing, as soon as the button turns to Disabled.

To add a widget (chart):

1. Click FortiView>Security>Data Analytics.
2. Click the Add Widget button to open the Fast Report dialog.

3. Make the entries and selections as described in [Data Analytics widget on page 592](#).
4. Click Save when done,

Data Analytics widget

Chart/Graph	Description
Name	Enter a unique name for a chart.
Attack Subtype	Click the down arrow and select a server load-balancing data you want to show in the chart. • Top Attack Type for All • Top Attack Type by VS for All • Top VS for DDoS • Top Destination Country for DDoS • Top VS for GEO • Top Source for GEO • Top Destination for GEO • Top Source Country for GEO • Top Destination Country for GEO • Top Action by Source for GEO • Top Action by Source Country for GEO • Top Category by VS for IP Reputation • Top Source for IP Reputation • Top Destination for IP Reputation • Top Source Country for IP Reputation • Top Destination Country for IP Reputation • Top Attack Type by VS for WAF • Top Attack Type by Source Country for WAF • Top Attack Type by Source for WA • Top Attack Type by Destination Country for WAF • Top Attack Type by Destination for WAF • Top Platform Name by Destination for AV • Top Platform Name by Destination Country for AV • Top Platform Name by Source for AV • Top Platform Name by VS for AV • Top Reference by Destination for AV • Top Reference by Destination Country for AV • Top Reference by Source for AV • Top Reference by Source Country for AV • Top Reference by VS for AV
History Chart	A "history" chart shows historical data that the system captured over a specific time period in the past. The option is turned OFF (disabled) by default, but you can click the button to turn it ON (enable it). Note: If this option is turned off, you will get a pie chart when you save the widget. If it is turned on, then you will see a bar chart. Both bar charts and line charts have a time-range selector in their upper-right corner which allows you to select one of the following: • 10 Minutes

Chart/Graph	Description
	• 1 Hour • 1 Day • 1 Week • 1 Month
Time Range	Click the down arrow to select one of the following time ranges: • 10 Minutes • 1 Hour • 1 Day • 1 Week • 1 Month Note: This option becomes unavailable if History Chart is enabled.
Data Type	Note: For this 4.8.1 release, Count is the only option and is selected by default. No action is needed.
Top X	Specify a maximum value for the X axis. Note: The default is 5, but the valid values are from 3 to 7.
Top Y	Specify a maximum value for the Y axis. Note: The default is 5, but the valid values are from 3 to 7.

Security Logs

The FortiView>Security>Security Logs page displays Web application firewall logs that the system has generated, from **Log & Report > Log Browsing**. It has two types of logs: **Security** and **Aggregate**.

Security Log

The security log can show you two logs, the AV Log or the WAF Log. You can choose between them by clicking on the upper right. There, you will also have the option to select the timespan for the logs generated.

Click on the graph to see information.

Note: The information parameters for the WAF and AV Logs are identical. Also, these logs will also appear in the Aggregate tab.

AV/WAF Log	Description
Date	Log date
Time	Log time
Severity	Rule severity
Source	Source IP address
Destination	Destination IP address.

Aggregate Log

The Aggregate Log provides an aggregated view of security logs within a selected time frame.

There are five types of aggregated security logs:

- Synflood—Traffic logged by the SYN Flood feature
- Geo—Traffic logged by the Geo IP block list feature
- IP Reputation—Traffic logged by the IP Reputation feature
- WAF—Traffic logged by the WAF feature
- AV—Traffic logged by the Anti virus module

To view an aggregate log:

1. Click log type.
2. Select a time frame.
3. Click **Refresh** to apply the filter and redisplay the log.

The following table shows the detailed information of an aggregated GEO log. The other aggregated logs show the same details.

Details of an aggregated GEO log

Column	Example	Description
Date	2016-12-02	Log date
Time	10:27:01	Log time
Count	1	For DoS, number of timeouts sent per destination
Severity	high	Always “high” for DoS
Source	173.177.99.94	Source IP address
Destination	10.61.2.100	Destination IP address
Action	deny	Policy action

All Segments

The FortiView>All Segments menu shows the logs, alerts, and session information. It has following sub-menus:


- [System Events](#)
- [Alerts](#)
- [All Sessions](#)

Event Logs

The FortiView>All Segments>Event Logspage shows all system event logs that FortiADC generated.

Setting log filters

1. Go to the far right and locate the System and Time figures, as highlighted in red.
2. Choose the filters.



- The logs are presented in a tabular format, with each row being a log entry. The log table shows some key information contained in the logs.
- You can drag a blue rectangle over the graph to show the logs for a certain span of time.

Alerts

The FortiView>All Segments>Alerts page shows the alert messages that the system has generated.

Setting alert filters

You can use the Filter Setting button in the upper-left corner of the page to filter logs displayed on the page.

To set your filter:

1. Click **Filter Setting**.
2. The following diagram will appear below, with an **Apply** button.
3. Select between the following filters: Trigger Time, Alert, Priority, Message.
4. Click OK when done.

Trigger Time	Alert
2018-09-28 14:43:40	SYS_CRL_expires
2018-09-28 14:43:25	SEC_Generic_Attack_Detected

You can apply multiple filters to the page. All filters you have configured will appear under the Filter Setting button in the order they are created. To remove a filter, click the x sign on it; to clear all filters, click Remove All Filters.

Viewing alerts

The alert messages are presented in a tabular format, with each row being an alert entry. The alert table shows some basic information about each alert. You can view details of an alert by clicking the log, which will drop down with the following information: Timestamp, Resource Name, level, Summary.

You can also remove alerts from the page by clicking the corresponding x button.

All Sessions

The FortiView>All Segments>All Sessions page has two tabs, which open the Session Table and Persist Table, respectively.

Viewing the Session or Persist Table

The Session Table shows information about the sessions that FortiADC has established. The page shows the live sessions only. Expired sessions are removed from the table when the page refreshes.

To view the Session or Persistence Table:

1. Click FortiView>All Segments>All Sessions.
2. Select the Session Table or Persist Table tab.

You can use the Filter Setting button (located in the upper-left corner of the page) to filter the sessions displayed on the page.

To set the filter:

1. Click the Filter Setting button.
2. Select between the following filters: Trigger Time, Alert, Priority, Message.
3. Click OK when done.

You can apply multiple filters. All filters you have configured will appear under the Filter Setting button in the order they are created. To remove a filter, click the x sign on it.

Note: The Clear button (next to Filter Setting), if clicked, clears all sessions in the table. If you click the button by mistake, you can always re-populate the page with session data by clicking the Refresh button.

Appendix A: Fortinet MIBs

FortiADC MIBs on page 597 lists the management information bases (MIBs) used with FortiADC.

FortiADC MIBs

MIB or RFC	Description
Fortinet Core MIB	This Fortinet-proprietary MIB enables your SNMP manager to query for system information and to receive traps that are common to multiple Fortinet devices.
FortiADC MIB	This Fortinet-proprietary MIB enables your SNMP manager to query for FortiADC-specific information and to receive FortiADC-specific traps.
RFC 1213 (MIB II)	The FortiADC SNMP agent supports MIB II groups, except: There is no support for the EGP group from MIB II (RFC 1213, section 3.11 and 6.10). Protocol statistics returned for MIB II groups (IP, ICMP, TCP, UDP, and so on) do not accurately capture all FortiADC traffic activity. More accurate information can be obtained from the information reported by the FortiADC MIB.
RFC 3635 (Ethernet-like MIB)	The FortiADC SNMP agent uses any of the objects in the Ethernet-like interface types specification (dot3StatsIndex).

You can download the Fortinet MIB files from the Fortinet Customer Service & Support website, <https://support.fortinet.com/>. See [FortiADC MIB download on page 598](#).

To view a trap or query's name, object identifier (OID), and description, open its MIB file in a plain text editor.

To communicate with the FortiADC SNMP agent, you must first compile these MIBs into your SNMP manager. If the standard MIBs used by the SNMP agent are already compiled into your SNMP manager, you do not have to compile them again. The FortiADC SNMP implementation is read-only.

All traps sent include the message, the FortiADC appliance's serial number, and hostname.

FortiADC MIB download

[Home](#)[Asset](#)[Assistance](#)[Download](#)[Feedback](#)Firmware
Images

Fortinet Firmware Images And Software Releases

Welcome to the Firmware Images download center for Fortinet's extensive line of security solutions.

Select Product



FortiADC

[Release Notes](#)[Download](#)

Image File Path

[/ FortiADC/ v4.00/ MIB/](#)

Image Folders/Files

[Up to higher level directory](#)

	Name	Size (KB)	Date Created	Date Modified		
	FORTINET-CORE-MIB.mib	13	2015-09-08 08:09:17	2015-09-08 08:09:17	HTTPS	Checksum
	FORTINET-FORTIADC-MIB.mib	15	2015-09-08 08:09:16	2015-09-08 08:09:16	HTTPS	Checksum

Appendix B: Port Numbers

Communications between the FortiADC system, clients, servers, and FortiGuard Distribution Network (FDN) require that any routers and firewalls between them permit specific protocols and port numbers.

[Default ports used by FortiADC for outgoing traffic on page 599](#) and [Default ports used by FortiADC for incoming traffic \(listening\) on page 599](#) list the default port assignments that FortiADC uses for outgoing and incoming traffic, respectively.

Default ports used by FortiADC for outgoing traffic

Port Number	Protocol	Purpose
N/A	ARP	HA failover of network interfaces.
N/A	ICMP	- Server health checks. <code>execute ping</code> and <code>execute traceroute</code>.
25	TCP	SMTP for alert email.
53	UDP	DNS queries.
69	UDP	TFTP for backups, restoration, and firmware updates. See commands such as <code>execute backup</code> or <code>execute restore</code> .
80	TCP	Server health checks.
123	UDP	NTP synchronization.
162	UDP	SNMP traps.
389	TCP	LDAP authentication queries.
443	TCP	- FortiGuard polling. - Server health checks.
514	UDP	Syslog.
6055	UDP	HA heartbeat. Layer 2 multicast.
6056	UDP	HA configuration synchronization. Layer 2 multicast.

Default ports used by FortiADC for incoming traffic (listening)

Port Number	Protocol	Purpose
N/A	ICMP	<code>ping</code> and <code>traceroute</code> responses.
22	TCP	SSH administrative CLI access.
23	TCP	Telnet administrative CLI access.
53	UDP	DNS queries from clients for global load balancing and inbound link load balancing.

Port Number	Protocol	Purpose
80	TCP	• HTTP administrative web UI access. • Predefined HTTP service. Only occurs if the service is used by a virtual server.
161	UDP	SNMP queries.
443	TCP	• HTTPS administrative web UI access. Only occurs if the destination address is a network interface's IP address. • Predefined HTTPS service. Only occurs if the service is used by a virtual server, and if the destination address is a virtual server.
6055	UDP	HA heartbeat. Layer 2 multicast.
6056	UDP	HA configuration synchronization. Layer 2 multicast.

Appendix C: Scripts

You can embed Lua scripts to perform tasks that are not supported by the built-in feature set.

This appendix provides guidance for getting started. It includes the following topics:

- [Events and actions on page 601](#)
- [Predefined scripts on page 602](#)
- [Predefined commands on page 605](#)
- [Control structures on page 624](#)
- [Operators on page 625](#)
- [String library on page 626](#)
- [Special characters on page 627](#)
- [Examples on page 629](#)

For general information about Lua, visit <http://www.lua.org/docs.html>.

Events and actions

Scripts are associated with a particular virtual server, and they are event-driven. A script is triggered when the associated virtual server receives an HTTP request or response. Then, it does the programmed action.



You can set different script priorities when you run multiple scripts at once. See [Prioritize scripts on page 637](#) for more information.

[Script events and actions on page 601](#) provides the syntax, usage, and examples of the predefined commands that are useful for writing scripts.

Script events and actions

Event/Action	Description
Event	
HTTP_REQUEST	The virtual server receives a complete HTTP request header.
HTTP_RESPONSE	The virtual server receives a complete HTTP response header.
RULE_INIT	The event is used to initialize global or static variables used within a script. It is triggered when a script is added or modified, or when the device starts up, or when the software is restarted.
VS_LISTENER_BIND	The virtual server tries to bind.
SERVER_BEFORE_	The virtual server is going to connect to the backend real server.

Event/Action	Description
CONNECT	
SERVER_CONNECTED	The HTTP proxy deems that the backend real server is connected.
AUTH_RESULT	The authentication (HTML Form / HTTP-basic) is done.
HTTP_RESPONSE_CONTINUE	Triggered immediately when the system receives a 100 continue response from the server.
HTTP_DATA_FETCH_SET_DEMO	FortiADC reads the body of every HTTP request, and can manipulate the data depending on settings.
HTTP_DATA_REQUEST	Triggered whenever an HTTP:collect command finishes processing, after collecting the requested amount of data.
HTTP_REQUEST_SEND	Triggered immediately before a request is sent to a server.
HTTP_RESPONSE_CONTINUE	Triggered immediately when the system receives a 100 Continue response from the server.
HTTP_DATA_RESPONSE	Triggered when an HTTP:collect command finishes processing on the server side of a connection.
CLIENTSSL_HANDSHAKE	The virtual server receives a complete HTTPS handshake on the client side.
SERVERSSL_HANDSHAKE	FortiADC receives a complete HTTPS handshake on the server side.
CLIENTSSL_RENEGOTIATE	The virtual server receives a re-connection request from a peer.
SERVERSSL_RENEGOTIATE	FortiADC sends a re-connection request to a peer.
TCP_ACCEPTED	The virtual server receives a complete TCP connection.
TCP_CLOSED	The virtual server close a TCP connection.
Action	
in Lua mode	An action defined by a Lua script that uses predefined commands and variables to manipulate the HTTP request/response or select a content route.

Predefined scripts

[Predefined scripts on page 602](#) provides the syntax, usage, and examples of the predefined commands that are useful for writing scripts.

Predefined scripts

Predefined script	Description

Predefined script	Description
INSERT_RANDOM_MESSAGE_ID_DEMO	Inserts a 32-bit hex string into the HTTP header with a parameter "Message-ID". Note: You can use the script directly, without making any changes.
GENERAL_REDIRECT_DEMO	FortiADC redirects HTTP requests to a set location.
USE_REQUEST_HEADERS_in_OTHER_EVENTS	FortiADC uses a session ID to obtain data from that session.
COMPARE_IP_ADDR_2_ADDR_GROUP_DEMO	FortiADC tries to find the client IP address in an internal list and returns the result.
HTTP_2_HTTPS_REDIRECTION_FULL_URL	FortiADC redirects an HTTP request.
REWRITE_HTTP_2_HTTPS_in_LOCATION	FortiADC changes an HTTP location given in an HTTP response with an HTTPS location.
REWRITE_HTTPS_2_HTTP_in_LOCATION	FortiADC changes an HTTPS location given in an HTTP response with an HTTP location.
REWRITE_HTTP_2_HTTPS_in_REFERER	FortiADC changes a HTTP referer given in an HTTP response with an HTTPS referer.
REWRITE_HTTPS_2_HTTP_in_REFERER	FortiADC changes a HTTPS referer given in an HTTP response with an HTTP referer.
HTTP_DATA_FETCH_SET_DEMO	FortiADC reads the body of every HTTP request, and can manipulate the data depending on settings.
HTTP_DATA_FIND_REMOVE_REPLACE_DEMO	FortiADC reads the body of every HTTP request and will find and replace data in the body.
MULTIPLE_SCRIPT_CONTROL_DEMO_1	When multiple scripts are running, this will determine the priority of each script.
MULTIPLE_SCRIPT_CONTROL_DEMO_2	When multiple scripts are running, this will determine the priority of each script.
HTTP_REQUEST_SEND	Triggered immediately before a request is sent to a server.
AES_DIGEST_SIGN_2F_COMMANDS	Demonstrates how to use AES to encryption/decryption data and some tools to generate the digest.
AUTH_COOKIE_BAKE	Allows you to retrieve the baked cookie and edit the cookie content.
AUTH_EVENTS_n_COMMANDS	Used to get the information from authentication process.
CLASS_SEARCH_n_MATCH	Demonstrates how to use the class_match and class_search utility function.

Predefined script	Description
CONTENT_ROUTING_by_URI	Routes to a pool member based on URI string matches. You should not use this script as is. Instead, copy it and customize the URI string matches and pool member names.
CONTENT_ROUTING_by_X_FORWARDED_FOR	Routes to a pool member based on IP address in the X-Forwarded-For header. You should not use this script as is. Instead, copy it and customize the X-Forwarded-For header values and pool member names.
COOKIE_COMMANDS	Demonstrate the cookie command to get the whole cookie in a table and how to remove/insert/set the cookie attribute.
COOKIE_COMMANDS_USAGE	Demonstrate the sub-function to handle the cookie attribute "SameSite" and others.
COOKIE_CRYPTO_COMMANDS	Used to perform cookie encryption/decryption on behalf of the real server.
CUSTOMIZE_AUTH_KEY	Demonstrate how to customize the crypto key for authentication cookie.
GEOIP_UTILITY	Used to fetch the GEO information country and possible province name of an IP address.
HTTP_2_HTTPS_REDIRECTION	Redirects requests to the HTTPS site. You can use this script without changes
HTTP_DATA_FETCH_SET_DEMO	"Collects data in HTTP request body or HTTP response body. In HTTP_REQUEST or HTTP_RESPONSE, you could collect specified size data with "size" in collect(). In HTTP_DATA_REQUEST or HTTP_DATA_RESPONSE. You could print the data use "content", calculate data length with "size", and rewrite the data with "set". Note: Do NOT use this script ""as is"". Instead, copy it and manipulate the collected data."
IP_COMMANDS	Used to get various types IP Address and port number between client and server side.
MANAGEMENT_COMMANDS	Allow you to disable/enable rest of the events from executing.
OPTIONAL_CLIENT_AUTHENTICATION	Performs optional client authentication. Note: Before using this script, you must have the following four parameters configured in the client-ssl-profile: • client-certificate-verify—Set to the verify you'd like to use to verify the client certificate. • client-certificate-verify-option—Set to optional • ssl-session-cache-flag—Disable. • use-tls-tickets—Disable. "
REDIRECTION_by_STATUS_CODE	Redirects requests based on the status code of server HTTP response (for example, a redirect to the mobile version of a site). Do NOT use this script "as is". Instead, copy it and customize the condition in the server HTTP response status code and the URL values.
REDIRECTION_by_USER_AGENT	Redirects requests based on User Agent (for example, a redirect to the mobile version of a site). You should not use this script as is. Instead, copy it and customize the User

Predefined script	Description
	Agent and URL values
REWRITE_HOST_n_PATH	Rewrites the host and path in the HTTP request, for example, if the site is reorganized. You should not use this script as is. Instead, copy
SNAT_COMMANDS	Allows you to overwrite client source address to a specific IP for certain clients, also support IPv4toIPv6 or IPv6toIPv4 type. Note: Make sure the flag SOURCE ADDRESS is selected in the HTTP or HTTPS type of profile.
SOCKOPT_COMMAND_USAGE	Allows user to customize the TCP_send buffer and TCP_receive buffer size.
SPECIAL_CHARACTERS_HANDLING_DEMO	Shows how to use those "magic characters" which have special meanings when used in a certain pattern. The magic characters are () . % + - * ? [] ^ \$
SSL_EVENTS_n_COMMANDS	Demonstrate how to fetch the SSL certificate information and some of the SSL connection parameters between server and client side.
TCP_EVENTS_n_COMMANDS	Demonstrate how to reject a TCP connection from a client in TCP_ACCEPTED event.
URL_UTILITY_COMMANDS	Demonstrate how to use those url tools to encode/decode/parser/compare .
UTILITY_FUNCTIONS_DEMO	Demonstrates how to use the basic string operations and random number/alphabet, time, MD5, SHA1, SHA2, BASE64, BASE32, table to string conversion, network to host conversion utility function.

Predefined commands

[Predefined commands on page 605](#) provides the syntax, usage, and examples of the predefined commands that are useful for writing scripts.

Predefined commands

Syntax	Usage and Example
Global	
debug("msg", ...)	Write the message to the debug buffer. For example: debug("HTTP Request method is %s.\n", HTTP:method_get()) Debug strings can be written to the console when the event is triggered. This is helpful when you are testing your scripts.

Syntax	Usage and Example
	To enable debug strings to be written to the console, use the following CLI commands: <pre>diagnose debug enable diagnose debug module httpproxy scripting</pre>
cmp_addr(addr, addr_group)	Used to match one IP address against a group of IP addresses. It can automatically detect IPv4 and IPv6 and can be used to compare IPv4 addresses with IPv6 addresses. For example: <pre>cmp_addr("192.3.2.1/24", "192.3.2.0/32") cmp_addr("::ffff:192.3.2.1/120", "::ffff:192.3.2.0/128") cmp_addr("192.3.2.1/24", "::ffff:192.3.2.0/128")</pre> Input format: For an IPv4 ip_addr/[mask], the mask can be a number between 0 and 32 or a dotted format like 255.255.255.0 For an IPv6 ip_addr/[mask], the mask can be a number between 0 and 128. FortiADC supports address group for the second argument. <pre>when RULE_INIT{ --initialize the address group here addr_group = "192.168.1.0/24" --first network address addr_group = addr_group..",::ffff:172.30.1.0/120" --second network address --so on and so forth } when HTTP_REQUEST{ client_ip=HTTP:client_addr() match_ip=cmp_addr(client_ip, addr_group) }</pre>
log("fmt", ...)	Writes log messages into the SLB log category in the script log part. You must enable Script log and SLB sub-category under the Script log on the log setting page. For example: <pre>log("This HTTP Request method is %s.\n", HTTP:method_get())</pre> Note: \ and % are handled in a unique way. Special characters that the log supports are :~!@#\$%^&*()_+{}[]. If you want to print out % in the log, you must use %%; if you want to print out \, you must use \\.
rand()	Generates a random number. For example: <pre>a = rand() debug("a=%d\n",a)</pre>
time()	Returns the current time as an integer. For example: The following code will return the current time, in Unix time format, as an integer and store it in variable "t".

Syntax	Usage and Example
	<code>t=time()</code>
<code>ctime()</code>	Returns the current time as a string. For example: The following code will return the current time as a string and store it in variable "ct". <code>ct=ctime()</code>
<code>md5()</code>	Calculates the MD5 of a string input and stores the results in an intermediate variable. For example: The following code will calculate the MD5 of the string provided and store it in variable "Md". <code>Str="test string\1\2"</code> <code>Md=md5(str)</code>
<code>md5_hex()</code>	Calculates the MD5 of a string input of a string input and outputs the results in HEX format. The following code will calculate the MD5 of the string provided and store it, in HEX format, in variable "re_hex". <code>Str="abc"</code> <code>re_hex=md5_hex(str)</code>
<code>sha1()</code>	Calculates the SHA1 of a string input of a string input and stores the results in an intermediate variable. The following code will calculate the SHA1 of the string provided and store it in variable "sha". <code>Str="abc"</code> <code>sha=sha1(str)</code>
<code>sha1_hex()</code>	Calculates the SHA1 of a string input of a string input and outputs the results in HEX format. The following code will calculate the SHA1 of the string provided and store it, in HEX format, in variable "sha". <code>Str="abc"</code> <code>sha=sha1_hex(str)</code>
<code>b64_enc()</code>	Encodes a string input in base64 and outputs the results in string format. The following code will encode the string provided and store it in the variable "en". <code>Str="abc"</code> <code>en=b64_enc(str)</code>
<code>b64_dec()</code>	Decodes a base64 encoded string input and outputs the results in string format. The following code will encode the string provided and store it in the variable "en". <code>Str="abc"</code> <code>en=b64_dec(str)</code>
<code>htonl()</code>	Converts a long integer input into network byte order and outputs the results in string format.

Syntax	Usage and Example
	The following code will convert the integer provided and store it, as a string, in the variable "b". <pre>a=32 b=htonl(a)</pre>
ntohl()	Converts a long integer input into host byte order and outputs the results in string format. The following code will convert the integer provided and store it, as a string, in the variable "b". <pre>a=32 b=ntohl(a)</pre>
htons()	Converts a short integer input into network byte order and outputs the results in string format. The following code will convert the integer provided and store it, as a string, in the variable "b". <pre>a=32 b=htons(a)</pre>
ntohs()	Converts a short integer input into host byte order and outputs the results in string format. The following code will convert the integer provided and store it, as a string, in the variable "b". <pre>a=32 b=ntohs(a)</pre>
string.format()	Converts an integer to string format. The following code will convert the integer provided and store it, as a string, in the variable "b". <pre>a=32 b=string.format(a)</pre> You may also use the function as shown in the code below. The string "12,pi=3.14" will be stored in variable "b". <pre>a=12 b=string.format("%s,pi=%.4f",a,3.14);</pre>
string.char()	Converts a number in string format to its corresponding ASCII char. The following code will convert the string provided and store it in the variable "test". In this case, string.char() will return "a". <pre>str=97 test=string.char(str)</pre>
{<variable>:byte(1,-1)}	Creates a table with the codes of all characters in the variable. This table can be used to recreate the original string using the table_to_string() command. The following code will create a table, then store the variable 'str' in the table. In this case, variable "t" is the table, and t[1] is 97, t[2] is 98, t[3] is 99, t[4] is 1, t[5] is 2, t[6] is 0. <pre>str="abc\1\2\0" t={str:byte(1,-1)}</pre>
{<variable>:sub(i,j)}	Returns a sub-string of the variable indexed from i to j. The following code will return the string "abc" and store it into variable "t".

Syntax	Usage and Example
	<pre>str="abc\1\2\0" t={str:sub(1,3)}</pre>
table_to_string ()	Converts a table to string format. The following code will convert the table "t" and store it, as a string, in the variable "str". The string stored in "str" at the end is "abc\1". <pre>t={}; t[1]=97; t[2]=98; t[3]=99; t[4]=1; str=table_to_string(t);</pre>
to_HEX	Converts a string to HEX format. The following code will convert the string "str" and store it to "hex" in HEX format. <pre>str="\0\123\3" hex=to_HEX(str);</pre>
crc32(str);	Returns the crc32 check value of the string, or 0 if it is an empty string, For example: <pre>when HTTP_REQUEST { str = "any string for crc32 calculation" crc = crc32(str); debug("rc is %d\n", crc); }</pre>
new_key = key_ gen(str_pass, str_salt, iter_ num, len_num); "	Creates an AES key to encrypt/decrypt data, either generated by password or user specific defined. For example: <pre>when HTTP_REQUEST { new_key = key_gen("pass", "salt", 32, 32); debug("new key in hex is %s\n", to_HEX(new_key)); }</pre>
aes_enc(t)	Encrypts a string using AES algorithm, For example: <pre>when HTTP_REQUEST { t={}; t["message"] = "value"; t["key"] = "aaaaaaaaaabbbbb"; t["size"]=128 enc = aes_enc(t) debug("encrypted in hex is %s, after b64 encoding %s\n", to_HEX (enc), b64_enc_str(enc)); }</pre>
aes_dec(t)	Decrypts a string using AES algorithm. For example: <pre>when HTTP_REQUEST { t={}; t["message"] = enc; t["key"] = "aaaaaaaaaabbbbb"; t["size"]=128 dec = aes_dec(t); debug("decrypted in hex is %s\n", to_HEX(dec));</pre>

Syntax	Usage and Example
	} }
EVP_Digest (alg, str)	EVP_Digest for oneshot digest calculation. For example: when HTTP_REQUEST { alg = "MD5"; data = "your data" re = EVP_Digest(alg, data); debug("the digest in hex is %s\n", to_HEX(re)); }
HMAC(alg, str, key)	HMAC message authentication code. For example: when HTTP_REQUEST { alg = "MD5"; --must be "MD5", "SHA1", "SHA256", "SHA384", "SHA512" data = "your data" key = "11234567890ab"; re = HMAC(alg, data, key); debug("the HMAC in hex is %s\n", to_HEX(re)); }
HMAC_verify (alg, data, key, verify)	Check if the signature is same as the current digest. when HTTP_REQUEST { alg = "MD5"; data = "your data" verify = "your result to compare" key = "11234567890ab"; re = HMAC_verify(alg, data, key, verify); if re then debug("verified\n") else debug("not verified\n") end }
G2F(alg, key)	Returns a G2F random value . For example: when HTTP_REQUEST { alg = "MD5"; key = "11234567890ab"; re = G2F(alg, key); debug("the G2F value is %d\n", re); }
class_match (str, method, list);	Used to match the string against an element in list: when HTTP_REQUEST { url = HTTP:uri_get() status, count, t = class_match(url, "starts_with", url_ list); debug("status %s, count %s\n", status, count); for k,v in pairs(t) do debug("index %s, value %s\n", k,v); end }
class_search (list, method, str);	Used to search the an element in the list against a string: when HTTP_REQUEST { status, count, t = class_search(url_list, "starts_with", url); --or "ends_with", "equals", "contains" for k,v in pairs(t) do debug("index %s, value %s\n", k,v); end }
ip2country_ name(ip)	Return the GEO information (country name) of an IP address. when HTTP_REQUEST { cip = IP:client_addr(); cnm = ip2country_name (cip); debug("cname %s\n", cnm); }
ip2countryPro v_name(ip)	Return the the GEO information (country name + possible province name) of an IP address. when HTTP_REQUEST { cip = IP:client_addr(); cnm = ip2countryProv_ name(cip); debug("cname %s\n", cnm); }
url_enc(str)	Converted the url into a valid ASCII format. when HTTP_REQUEST { url = "http://foor bar/@"; enc = url_enc(url); debug("encoded url is %s\n", enc); }
url_dec(str)	converted the encoding-url into a original url.

Syntax	Usage and Example
	<pre>when HTTP_REQUEST { url = "http://foor.bar/test/"; enc = url_enc (url); debug("encoded url is %s\n", enc); }</pre>
<code>url_parser(str)</code>	Extracte the url and host are converted to lower case letters. <pre>when HTTP_REQUEST { url = "http://foo:bar@w1.superman.com/very/long/path.html?p1=v1&p2=v2#more- details" purl = url_parser(url); if purl then debug("parsed url scheme %s, host %s, port %s, path %s, query %s, fragment %s, username %s, passowrd %s\n", purl["scheme"], purl["host"], purl ["port"],purl["path"], purl["query"], purl["fragment"], purl ["username"], purl["password"]); end }</pre>
<code>url_compare(url1, url2)</code>	Compare two url string, return true if it's the same. <pre>when HTTP_REQUEST { url1 = "http://www.example.com/url/path/data" url2 = "http://WWW.example.com:80/url/path/data" if url_compare (url1, url2) then debug("url match\n"); else debug("url not match\n"); end }</pre>
<code>rand_hex(int)</code>	Generate a random number in HEX: <pre>str = rand_hex(16);</pre>
<code>rand_alphanum(int)</code>	Generate a random alphabet+number sequence: <pre>str = rand_alphanum(16);</pre>
<code>rand_seq(int)</code>	Generate a random in sequence: <pre>str = rand_seq(16)</pre>
<code>md5_str(str)</code>	Calculate the MD5 of a string input and stores the results in an intermediate variable, In some cases you need a this version to deal with it. For example: <pre>Md=md5_str(input); --input can be a cert in DER format</pre>
<code>md5_hex_str(str)</code>	Calculates the MD5 of a string input of a string input and outputs the results in HEX format, In some cases you need a this version to deal with it. For example: <pre>Md=md5_hex_str(input); --input can be a cert in DER format</pre>
<code>sha1_str()</code>	Calculates the SHA1 of a string input of a string input and stores the results in an intermediate variable, In some cases you need a this version to deal with it. For example: <pre>result=sha1_str(input); --input can be a cert in DER format</pre>
<code>sha1_hex_str()</code>	Calculates the SHA1 of a string input of a string input and outputs the results in HEX format, In some cases you need a this version to deal with it. For example: <pre>result=sha1_hex_str(input); --input can be a cert in DER format</pre>
<code>sha256()</code>	Calculates the SHA256 of a string input of a string input and stores the results in an intermediate variable. The following code will calculate the SHA256 of the string provided and store it in variable "sha256". <pre>Str="abc" sha256=sha256(str)</pre>

Syntax	Usage and Example
sha256_hex()	Calculates the SHA256 of a string input of a string input and outputs the results in HEX format. The following code will calculate the SHA256 of the string provided and store it, in HEX format, in variable "sha256". Str="abc" sha256=sha256_hex(str)
sha256_str()	Calculates the SHA256 of a string input of a string input and stores the results in an intermediate variable. In some cases you need a this version to deal with it. For example: result=sha256_str(input); --input can be a cert in DER format
sha256_hex_str() ()	Calculates the SHA256 of a string input of a string input and outputs the results in HEX format. In some cases you need a this version to deal with it. For example: result=sha256_hex_str(input); --input can be a cert in DER format
sha384()	Calculates the SHA384 of a string input of a string input and stores the results in an intermediate variable. The following code will calculate the SHA384 of the string provided and store it in variable "sha384". Str="abc" sha384=sha384(str)
sha384_hex()	Calculates the SHA384 of a string input of a string input and outputs the results in HEX format. The following code will calculate the SHA384 of the string provided and store it, in HEX format, in variable "sha384". Str="abc" sha384=sha384_hex(str)
sha384_str()	Calculates the SHA384 of a string input of a string input and stores the results in an intermediate variable. In some cases you need a this version to deal with it. For example: result=sha384_str(input); --input can be a cert in DER format
sha384_hex_str() ()	Calculates the SHA384 of a string input of a string input and outputs the results in HEX format. In some cases you need a this version to deal with it. For example: result=sha384_hex_str(input); --input can be a cert in DER format
sha512()	Calculates the SHA512 of a string input of a string input and stores the results in an intermediate variable. The following code will calculate the SHA512 of the string provided and store it in variable "sha512". Str="abc" sha512=sha512(str)
sha512_hex()	Calculates the SHA512 of a string input of a string input and outputs the results in HEX format. The following code will calculate the SHA512 of the string provided and store it, in HEX format, in variable "sha512". Str="abc" sha512=sha512_hex(str)
sha512_str()	Calculates the SHA512 of a string input of a string input and stores the results in an intermediate variable. In some cases you need a this version to deal with it. For example: result=sha512_str(input); --input can be a cert in DER format
sha512_hex_str() ()	Calculates the SHA512 of a string input of a string input and outputs the results in HEX format. In some cases you need a this version to deal with it. For example: result=sha512_hex_str(input); --input can be a cert in DER format

Syntax	Usage and Example
b32_enc()	Encodes a string input in base32 and outputs the results in string format. The following code will encode the string provided and store it in the variable "en". Str="abc" en=b32_enc(str)
b32_enc_str(str)	Encodes a string input in base32 and outputs the results in string format. In some cases you need a this version to deal with it. For example: result=b32_enc_str(input); --input can be a cert in DER format
b32_dec()	Decodes a base32 encoded string input and outputs the results in string format. The following code will encode the string provided and store it in the variable "dec". Str="abc" dec=b32_dec(str)
b32_dec_str()	Decodes a base32 encoded string input and outputs the results in string format. In some cases you need a this version to deal with it. For example: result=b32_dec_str(input); --input can be a cert in DER format
get_pid()	Return the PID value of the VS process. For exmaple: debug("VS PID is : %d\n", get_pid());
HTTP	
cookie_list	Returns a list of cookies: their names and values. For example: ret=HTTP:cookie_list() for k,v in pairs(ret) do debug("cookie name %s, value %s\n", k,v); end
cookie	Allows you to GET/SET its value and its attribute, REMOVE a whole cookie, GET the whole cookie in HTTP RESPONSE, and INSERT a new cookie. For example: t={}; t["name"]="test" t["parameter"]="value";--value, cookie, path, domain, expires, secure, maxage, max-age, httponly, version, port t["action"]="get"--get, set, remove, insert ret = HTTP:cookie(t) if ret then debug("get cookie value succeed %s\n",ret); else debug("get cookie value failed\n"); end
cookie_crypto	The provided function response_encrypt_cookie can be used to perform cookie encryption in HTTP RESPONSE and request_decrypt_cookie can be used to perform cookie decryption in HTTP REQUEST. For example: --Decrypt cookie "test" in HTTP REQUEST before forwarding to real servers local t={}; t["name"]="cookieName" t["action"]="encrypt"--encrypt, or decrypt t["key"]="0123456789ABCDEF"; t["prefix"]="XXXX";

Syntax	Usage and Example
	<pre> t["size"]=size-- 128, 192, or 256, the corresponding key length is 16, 24, and 32 if HTTP: cookie_crypto(t) then debug("Encrypt cookie succeed\n"); else debug("Encrypt cookie failed\n"); end </pre>
respond	Allows you to return a customized page, For example: <pre> when HTTP_REQUEST{ tt={} tt["code"] = 200; tt["content"] = "XXXXXX Test Page XXXXXXXX\n\n\n\n"; status = HTTP:respond(tt); debug("HTTP_respond() status: %s\n", status); } </pre>
header_get_names()	Returns a list of all the headers present in the request or response. For example: <pre> --use header and value headers = HTTP:header_get_names() for k, v in pairs(headers) do debug("The value of header %s is %s.\n", k, v) end --only use the header name for name in pairs(headers) do debug("The request/response includes header %s.\n", name) end </pre>
header_get_values(header_name)	Returns a list of value(s) of the HTTP header named <header_name>, with a count for each value. Note that the command returns all the values in the headers as a list if there are multiple headers with the same name. For example: <pre> cookies=HTTP:header_get_values("Cookie") for k, cnt in pairs(cookies) do debug("initially include cookie %s cnt %d\n", k, v) end </pre>
header_get_value(header_name)	Returns the value of the HTTP header named <header_name>. Returns false if the HTTP header named <header_name> does not exist. Note: The command operates on the value of the last header if there are multiple headers with the same name. For example: <pre> host = HTTP:header_get_value("Host") </pre>
header_remove(header_name)	Removes all headers names with the name <header_name>. For example: <pre> HTTP:header_remove("Cookie") </pre>

Syntax	Usage and Example
header_remove2 (header_name, countid)	header_get_values() returns a count ID for each item. This count ID can be used in both header_remove2() and header_replace2() to remove and replace a certain header of a given name referenced by the count ID. For example: <pre>cookies=HTTP:header_get_values("Set-Cookie") for k, v in pairs(cookies) do debug("include cookie %s cnt %d\n", k, v) end if HTTP:header_remove2("Set-Cookie", 1) then debug("remove 1st cookie\n") end</pre>
header_insert (header_name, value)	Inserts the named HTTP header(s) and value(s) into the end of the HTTP request or response. For example: <pre>HTTP:header_insert("Cookie", "cookie=server1")</pre>
header_replace (header_name, value)	Replaces the value of the last occurrence of the header named <header_name> with the string <value>. Performs a header insertion if the header is not present. For example: <pre>HTTP:header_replace("Host", "www.fortinet.com")</pre>
header_replace2 (header_name, value, countid)	header_get_values() returns a count ID for each item. This count ID can be used in both header_remove2() and header_replace2() to remove and replace a certain header of a given name referenced by the count ID. For example: <pre>cookies=HTTP:header_get_values("Set-Cookie") for k, v in pairs(cookies) do debug("include cookie %s cnt %d\n", k, v) end if HTTP:header_replace2("Set-Cookie", "new2=value2", 2) then debug("replace 2nd cookie by new2=value2\n") end</pre>
header_exists (header_name)	Returns true if the named header is present and not empty on the request or response. For example: <pre>if HTTP:header_exists("Cookie") then ... end</pre>
header_count (header_name)	Returns the number of HTTP headers present in the request or response. For example: <pre>count = HTTP:header_count("Cookie")</pre>
method_get()	Return the string of the HTTP request method. For example:

Syntax	Usage and Example
	<code>method = HTTP:method_get()</code>
<code>method_set</code> (string)	Set the HTTP request method to the string " <i>value</i> ". For example: <code>HTTP:method_set("POST")</code>
<code>path_get()</code>	Returns the path part of the HTTP request. For example: <code>path = HTTP:path_get()</code>
<code>path_set(string)</code>	Sets the path part of the HTTP request. The client will not see the update unless the web application uses the requested path to generate response headers and/or content. If you want the client to see the update to the path in the browser's address bar, you can send an HTTP redirect using <code>HTTP:redirect</code> or <code>HTTP:respond</code> . For example: <code>HTTP:path_set("/other.html")</code>
<code>uri_get()</code>	Returns the URI given in the request. For example: <code>uri = HTTP:uri_get()</code>
<code>uri_set(string)</code>	Changes the URI passed to the server. It should always start with a slash. For example: <code>HTTP:uri_set("/index.html?value=xxx")</code>
<code>query_get()</code>	Returns the query part of the HTTP request. For example: <code>query = HTTP:query_get()</code>
<code>query_set</code> (string)	Sets the query part of the HTTP request. For example: <code>HTTP:query_set("value=xxx")</code>
<code>redirect("URL", ...)</code>	Redirects an HTTP request or response to the specified URL. For example: <code>Host = HTTP:header_get_value("host")</code> <code>Path = HTTP:path_get()</code> <code>HTTP:redirect("https://%s%s", Host, Path)</code>
<code>redirect_with_cookie(URL, cookie)</code>	Redirects an HTTP request or response to the specified URL with Cookie. For example: <code>HTTP:redirect_with_cookie("www.example.com", "server=nginx")</code>
<code>redirect_t</code>	Redirects an HTTP request or response to the URL specified in the table. For example: <code>a={}</code> <code>a["url"]="http://192.168.1.7"</code> <code>a["code"]="303"</code> <code>a["cookie"]="test=server"</code> <code>HTTP:redirect_t(a)</code>

Syntax	Usage and Example
version_get()	Returns the HTTP version of the request or response. For example: vers = HTTP:version_get()
version_set (string)	Sets the HTTP version of the request or response. For example: HTTP:version_set("1.0")
status_code_ get()	Returns the response status code output as string. For example: responsestatus=HTTP:status_code_get()
status_code_ set(string)	Sets the response status code. For example: HTTP:status_code_set("301")
code_get()	Returns the response status code,output as integer. For example: responsestatus=HTTP:code_get()
code_set (integer)	Sets the response status code. For example: HTTP:code_set(301)
reason_get()	Returns the response reason. For example: HTTP:reason_get()
reason_set (string)	Sets the response reason. For example: HTTP:reason_set(string)
rand_id()	Returns a random string of 32-long in hex format, which can be inserted directly as an HTTP header. For example: ID=HTTP:rand_id() HTTP:header_insert("Message-ID", ID)
client_addr()	Returns the client IP address of a connection for an HTTP_REQUEST packet, which is the source address for the HTTP_REQUEST packet. It's a destination address. For example: CIP=HTTP:client_addr()
local_addr()	For HTTP_REQUEST, returns the IP address of the virtual server the client is connected to; for HTTP_RESPONSE, returns the incoming interface IP address of the return packet. For example: LIP=HTTP:local_addr()
remote_addr()	Returns the IP address of the host on the far end of the connection. For example: RIP=HTTP:remote_addr()

Syntax	Usage and Example
<code>server_addr()</code>	Returns the IP address of the server in HTTP_RESPONSE. SIP=HTTP:server_addr()
<code>close()</code>	Closes an HTTP connection using code 503. For example: HTTP:close()
<code>client_port()</code>	Returns the client port number in a string format. For example: HTTP:client_port()
<code>local_port()</code>	Returns the local port number in a string format. For example: HTTP:local_port()
<code>remote_port()</code>	Returns the remote port number in a string format. For example: HTTP:local_port()
<code>server_port()</code>	Returns the server port number in a string format. For example: HTTP:server_port()
<code>client_ip_ver()</code>	Returns the client IP version number. For example: HTTP:client_ip_ver()
<code>server_ip_ver()</code>	Returns the server IP version number. For example: HTTP:server_ip_ver()
<code>collect</code>	Collects data. You may specify a specific amount using the length argument. Used in HTTP_REQUEST or HTTP_RESPONSE. For example: t={}; t["size"]=1000; --optional HTTP:collect(t);
<code>payload (size)</code>	Returns the size of the buffered content. Used in HTTP_DATA_REQUEST or HTTP_DATA_RESPONSE. For example: t={}; t["operation"]="size" sz=HTTP:payload(t); --return value is an int
<code>payload (content)</code>	Returns the buffered content in a string. Used in HTTP_DATA_REQUEST or HTTP_DATA_RESPONSE. For example: t={}; t["operation"]="content" t["offset"]=12; --optional t["size"]=20; --optional

Syntax	Usage and Example
	<pre>ct = HTTP:payload(t); --return value is a string</pre>
payload (set)	Replaces the buffered data with new data. Used in HTTP_DATA_REQUEST or HTTP_DATA_RESPONSE. For example: <pre>t={}; t["operation"]="set" t["offset"]=12; --optional t["size"]=20; --optional t["data"]="new data to insert"; ret = HTTP:payload(t); --returns true if operation succeeds</pre>
payload (find)	Searches for a particular string in the buffered data. Used in HTTP_DATA_REQUEST or HTTP_DATA_RESPONSE. For example: <pre>t={}; t["operation"]="find" t["data"]="sth"; -- can also be a regular expression, like (s.h) t["offset"]=12; --optional t["size"]=20; --optional t["scope"]="first" -- the scope field can be either "first" or "all" ct = HTTP:payload(t); --returns the number of occurrences found</pre>
payload (remove)	Removes a particular string from the buffered data. Used in HTTP_DATA_REQUEST or HTTP_DATA_RESPONSE. For example: <pre>t={}; t["operation"]="remove" t["data"]="sth"; -- can also be a regular expression, like (s.h) t["offset"]=12; t["size"]=20; t["scope"]="first" -- or "all" ct = HTTP:payload(t); --returns number of occurrences removed</pre>
payload (replace)	Replaces a particular string or regular expression with a new string. Used in HTTP_DATA_REQUEST or HTTP_DATA_RESPONSE. For example: <pre>t={}; t["operation"]="replace" t["data"]="sth"; -- can be a regular expression, like (s.h) t["new_data"]="sth new"; t["offset"]=12; --optional t["size"]=20; --optional t["scope"]="first" -- or "all" ct = HTTP:payload(t); --returns number of occurrences replaced</pre>
set_event	Sets a request or response event. For example: <pre>t={};</pre>

Syntax	Usage and Example
	<pre>t["event"] = "data_res"; --can be req,res,data_req, or data_res t["operation"] = "disable"; HTTP:set_event(t)</pre>
set_auto	Sets an automatic request or response event. For example: <pre>t={}; t["event"] = "data_res"; --can be req, res, data_req, or data_res t["operation"] = "disable"; HTTP:set_auto(t)</pre>
Load Balance	
routing (content_route)	Selects a content route. For example: <pre>LB:routing("content2")</pre>
TCP	
reject()	Allow you to reject a TCP connection from a client. Can be used in TCP_ACCEPTED event. For example: <pre>when TCP_ACCEPTED { --Check if the st is true or false if st then TCP:reject(); end }</pre>
set_snat_ip(str)	Allows user to set the backend TCP connection's source address and port. For example: <pre>when TCP_ACCEPTED { addr_group = "172.24.172.60/32" client_ip = IP:client_addr() matched = cmp_addr(client_ip, addr_group) if matched then if TCP:set_snat_ip("10.106.3.124") then debug("set SNAT ip to 10.106.3.124\n"); end end }</pre>
clear_snat_ip()	Allows you to clear whatever customized ip you ever set using set_snat_ip(). For example: <pre>when TCP_ACCEPTED { if TCP:clear_snat_ip() then debug("Clear SNAT IP !\n"); end }</pre>
sockopt(t)	Allows user to customize the send buffer and receive buffer size. For example: <pre>when VS_LISTENER_BIND { local t = {};</pre>

Syntax	Usage and Example
	<pre> t["op"] = "get"; t["message"] = "snd_buf" --"snd_buf" or "rcv_buf" if TCP:sockopt(t) then debug("tcp send buffer is %d\n", tcp_snd_buf); else debug("get tcp send buffer failed\n"); end } </pre>
SSL	
version()	Allows you to GET the SSL version, can be used in CLIENTSSL_HANDSHAKE / SERVERSSL_HANDSHAKE / CLIENTSSL_RENEGOTIATE / SERVERSSL_RENEGOTIATE event. For example: <pre>ver=SSL:version();</pre>
cipher()	Allows you to GET the SSL cipher, can be used in CLIENTSSL_HANDSHAKE / SERVERSSL_HANDSHAKE / CLIENTSSL_RENEGOTIATE / SERVERSSL_RENEGOTIATE event. For example: <pre>ci=SSL:cipher();</pre>
alg_keysize()	Allows you to GET the SSL key size, can be used in CLIENTSSL_HANDSHAKE / SERVERSSL_HANDSHAKE / CLIENTSSL_RENEGOTIATE / SERVERSSL_RENEGOTIATE event. For example: <pre>alg_keysize=SSL:alg_keysize()</pre>
npn()	Allows you to GET the SSL NPN extension, can be used in CLIENTSSL_HANDSHAKE / SERVERSSL_HANDSHAKE / CLIENTSSL_RENEGOTIATE / SERVERSSL_RENEGOTIATE event. For example: <pre>npn=SSL:npn();</pre>
alpn	Allows you to GET the SSL ALPN extension, can be used in CLIENTSSL_HANDSHAKE / SERVERSSL_HANDSHAKE / CLIENTSSL_RENEGOTIATE / SERVERSSL_RENEGOTIATE event. For example: <pre>alpn=SSL:alpn();</pre>
sni()	Allows you to GET the SSL SNI, can be used in CLIENTSSL_HANDSHAKE / SERVERSSL_HANDSHAKE / CLIENTSSL_RENEGOTIATE / SERVERSSL_RENEGOTIATE event. For example: <pre>sni=SSL:sni();</pre>
client_cert()	Returns the client certificate status, can be used in CLIENTSSL_HANDSHAKE / SERVERSSL_HANDSHAKE / CLIENTSSL_RENEGOTIATE / SERVERSSL_RENEGOTIATE event. For example: <pre>client_cert=SSL:client_cert()</pre>
session(t)	Allows you to GET SSL session id / Reused / Remove from cache, can be used in CLIENTSSL_HANDSHAKE / SERVERSSL_HANDSHAKE / CLIENTSSL_RENEGOTIATE / SERVERSSL_RENEGOTIATE event. For example: <pre> t={} t["operation"]="get_id"--or "remove" "reused" sess_id=SSL:session(t); if sess_id then </pre>

Syntax	Usage and Example
	<pre> sess_id=to_HEX(sess_id) debug("client sess id %s\n", sess_id); else sess_id="FALSE" end </pre>
cert(t)	Allows you to GET the cert info between local or remote, can be used in CLIENTSSL_HANDSHAKE / SERVERSSL_HANDSHAKE / CLIENTSSL_RENEGOTIATE / SERVERSSL_RENEGOTIATE event. For example: <pre> t={} t["direction"]="remote";--or "local" t["operation"]="count";-- or "index", or "issuer" cert=SSL:cert(t) if cert then debug("has %s certs\n", cert) else debug("no cert\n") end </pre>
peer_cert(str)	Returns the peer certificate, can be used in CLIENTSSL_HANDSHAKE / SERVERSSL_HANDSHAKE / CLIENTSSL_RENEGOTIATE / SERVERSSL_RENEGOTIATE event. For example <pre> cder=SSL:peer_cert("der");--for remote leaf certificate, the input parameter can be "info" or "der" or "pem". if cder then hash=sha1_hex_str(cder); debug("whole cert sha1 hash is %s\n", hash); end </pre>
IP	
client_addr()	Returns the client IP address, can be used in all events except VS_LISTENER_BIND. For example: <pre> cip=IP:client_addr() </pre>
local_addr()	Returns the local IP address, can be used in all events except VS_LISTENER_BIND / SERVER_BEFORE_CONNECT. For example: <pre> lip=IP:local_addr() </pre>
remote_addr()	Returns the remote IP address, can be used in all events except VS_LISTENER_BIND / SERVER_BEFORE_CONNECT. For example: <pre> rip=IP:remote_addr() </pre>
client_port()	Returns the client IP port number, can be used in all events except VS_LISTENER_BIND. For example: <pre> cp=IP:client_port() </pre>
local_port()	Returns the local port number, can be used in all events except VS_LISTENER_BIND / SERVER_BEFORE_CONNECT. For example:

Syntax	Usage and Example
	<pre> lp=IP:local_port() </pre>
remote_port()	Returns the remote port number can be used in all events except VS_LISTENER_BIND / SERVER_BEFORE_CONNECT. For example: <pre> rp=IP:remote_port() </pre>
client_ip_ver()	Returns the client IP version, can be used in all events except VS_LISTENER_BIND. For example: <pre> cipv=IP:client_ip_ver(); </pre>
server_addr()	Returns the server IP address, can be used in server-side events. For example: <pre> sip=IP:server_addr() </pre>
server_port()	Returns the server port number, can be used in server-side events. For example: <pre> sp=IP:server_port() </pre>
server_ip_ver()	Returns the server IP version, can be used in server-side events. For example: <pre> sipv=IP:server_ip_ver(); </pre>
Management	
get_session_id()	Returns the session id, can be used in all events except VS_LISTENER_BIND. For example: <pre> sid=MGM:get_session_id() debug("sess id %s\n", sid); </pre>
rand_id()	Returns the rand id, can be used in all events except VS_LISTENER_BIND. For example: <pre> sid=MGM:get_session_id() debug("rand id %s\n", rid); </pre>
set_event(t)	Allow user to disable/enable rest of the events from executing by disabling this event. For example: <pre> t={}; t["event"]="req"; -- can be "req", "res", "data_req", "data_res", "ssl_client", "ssl_server", "tcp_accept", "tcp_close", "ssl_renego_client", "ssl_renego_server", "server_connected", "server_close", "server_ before_connect", "vs_listener_bind", "auth_result", "cookie_bake" t["operation"]="disable"; -- can be "enable", and "disable" MGM:set_event(t); debug("disable rest of the HTTP_REQUEST events\n"); </pre>
set_auto(t)	Allow user to enable/disable automatic re-enabling. For example: <pre> t={}; t["event"]="req"; -- can be "req", "res", "data_req", "data_res", "ssl_server", "ssl_renego_server", "server_connected", "server_close", "server_before_connect" t["operation"]="disable"; -- can be "enable", and "disable" MGM:set_auto(t); debug("disable automatic re-enabling of the HTTP_REQUEST events\n"); </pre>
Auth	
get_baked_cookie()	Allows you to retrieve the baked cookie, For example: <pre> when COOKIE_BAKE { cookie = AUTH:get_baked_cookie() debug("baked cookie %s\n", cookie); } </pre>

Syntax	Usage and Example
set_baked_cookie(cookie)	Allows you to customize cookie attributes the baked cookie, For example: when COOKIE_BAKE { cookie = AUTH:get_baked_cookie() --add new attribute HttpOnly new_cookie = cookie.."; HttpOnly"; AUTH:set_baked_cookie(new_cookie); }
on_off()	Returns the authentication is required or not, For example: on_off = AUTH:on_off();
success()	Returns the authentication is successful or not, For example: succ = AUTH:success();
form_based()	Returns the authentication is HTTP form based or not , For example: fm = AUTH:form_based()
user()	Returns the authentication of the user name , For example:user = AUTH:user()
pass()	Returns the authentication of the password , For example:pass = AUTH:pass()
usergroup()	Returns the authentication of the user group , For example:userg = AUTH:usergroup()
realm()	Returns the authentication of the realm , For example:userg = AUTH:usergroup()
host()	Returns the authentication of the host , For example:host = AUTH:host()
Proxy	
set_auth_key(str)	Allows user to customize the crypto key FADC used for encrypt/decrypt authentication cookie, For example: when VS_LISTENER_BIND { AUTH_KEY = ""0123456789ABCDEF0123456789ABCDEF"" if PROXY:set_auth_key(AUTH_KEY) then debug("set auth key succeed\n"); end }

Control structures

[Lua control structures on page 624](#) lists the Lua control structures.

Lua control structures

Type	Structure
if then else	if condition1 then

Type	Structure
	<pre> ... elseif condition2 then ... else ... end </pre>
for	<pre> --fetch all values of table 't' for k, v in pairs(t) do ... end </pre>

Operators

[Lua operators on page 625](#) lists the FortiADC operators.

Lua operators

FortiADC Operator	Operator sub-type	Description
- +	Arithmetic	Unary minus, unary plus.
~	Bitwise	Bitwise NOT.
not	Logical	Performs a logical "not" on a value.
* / %	Arithmetic	Multiple, divide, remainder.
//		Floor division.
^		Exponentiation.
+ -	Arithmetic	Add and subtract.
<< >>	Bitwise	Left and right shift.
< > <= >=	Relational	Boolean less, greater, less than or equal, and greater than or equal.
== !=	Relational	Boolean equal and not equal.
&	Bitwise	Bitwise AND.
~	Bitwise	Bitwise exclusive OR.
	Bitwise	Bitwise OR.

FortiADC Operator	Operator sub-type	Description
and	Logical	Performs a logical "and" comparison between two values.
or	Logical	Performs a logical "or" comparison between two values.
starts_with (a,b)	String	Tests to see if String a starts with String b. Returns true or false.
ends_with (a,b)	String	Tests to see if String a ends with String b. Returns true or false.
contains	String	Checks to see whether String a contains String b. Returns true or false.
match	String	Searches for a specified string.
..		The string concatenation operator in Lua is denoted by two dots ('..'). If both operands are strings or numbers, then they are converted to a string. It's the same as <code>__concat</code> .

String library

The FortiADC OS supports only the [Lua string library](#). All other libraries are disabled. The string library includes the following string-manipulation functions:

- `string.byte(s, i)`
- `string.char(i1,i2...)`
- `string.dump(function)`
- `string.find(s, pattern)`
- `string.format`
- `string.gmatch`
- `string.gsub`
- `string.len`
- `string.lower`
- `string.match`

- `string.rep`
- `string.reverse`
- `string.sub`
- `string.upper`
- `string.starts_with`
- `string.ends_with`

For example: `uri:starts_with (b)`, `uri:ends_with (b)`

Note:

- If you want to do regular expression match, you can use `string.match` with Lua patterns.
- All relational operators `>`, `<`, `>=`, `<=`, `~=`, `==` apply to strings. Especially, `==` can be used to test if one string equals to another string.
- `string.find` can be used to test whether one string contains another string.

For a tutorial on scripting with the Lua string library, see <http://lua-users.org/wiki/StringLibraryTutorial>.

Special characters

This section discusses the use of special characters in FortiADC scripting.

Log and debug

FortiADC supports the special characters as listed in [Special characters and ways to handle them on page 627](#) in log and debug scripts.

Special characters and ways to handle them

Character	Name
~	Tilde
!	Exclamation
@	At sign
#	Number sign (hash)
\$	Dollar sign
^	Caret
&	Ampersand
*	Asterisk
(	Left parenthesis
)	Right parenthesis
_	Underscore
+	Plus

Character	Name
{	Left brace
}	Right brace
[	Left bracket
]	Right bracket
.	Full stop
?	Question mark

When written in a string, these characters look like this (between double quotes: "~!@#\$%^&*()_+{}[] . ?")

Note: The back slash (\) and the percent (%) signs are handled in a unique way in log and debug scripts. To print out %, you must use %%; to print out \, you must use \\.

HTTP data body commands

HTTP data body commands, such as `find`, `remove`, and `replace` support regular expression, which treats special characters such as (between double quote) "\$^?*. | () [] {} \" in a special way. You MUST escape these characters to demolish their special meaning. [Special characters in HTTP data body commands on page 628](#) shows how to escape these special characters.

Special characters in HTTP data body commands

To print out ...	You MUST use ...
\$	\\\$
^	\\^
?	\\?
*	*
+	\\+
.	\\.
	\\
\\	\\\\
(and)	\\(and \\)
{and}	\\{and \\}
[and]	\\[and \\]

Note:

- { and } are special because the script syntax looks for the matching { and }. So be sure to use them in pairs.
- The `find`, `remove`, and `replace` commands use special expression. Particularly, `p.ge` will match the whole word page and remove and replace the whole word page. However, `p*ge` will remove and replace only the `ge` part.
- The HTTP data body `set` command does not support regular expression. Only `\\` is special in the `set` command, and you must use `\\` for it.

Examples

This section provides example scripts for common use cases. It includes the following examples:

- Select content routes based on URI string matches. See [Select content routes based on URI string matches on page 629](#)
- Rewrite the HTTP request host header and path. See [Rewrite the HTTP request host header and path on page 630](#)
- Rewrite the HTTP response Location header. See [Rewrite the HTTP response Location header on page 631](#)
- Redirect HTTP to HTTPS using Lua string substitution. See [Redirect HTTP to HTTPS using Lua string substitution on page 631](#)
- Redirect mobile users to the mobile version of a website. See [Redirect mobile users to the mobile version of a website on page 631](#)
- Insert random message ID into a header. See [Insert random message ID into a header on page 632](#)
- General HTTP redirect. See [General HTTP redirect on page 632](#)
- Use request headers in other events. See [Use request headers in other events on page 632](#)
- Compare IP address to address group. See [Compare IP address to address group on page 633](#)
- Redirect HTTP to HTTPS. See [Redirect HTTP to HTTPS on page 633](#)
- Rewrite HTTP to HTTPS in location. See [Rewrite HTTP to HTTPS in location on page 633](#)
- Rewrite HTTP to HTTPS in referer. See [Rewrite HTTP to HTTPS in referer on page 634](#)
- Rewrite HTTPS to HTTP in location. See [Rewrite HTTPS to HTTP in location on page 634](#)
- Rewrite HTTPS to HTTP in referer. See [Rewrite HTTPS to HTTP in referer on page 634](#)
- Fetch data from HTTP events. See [Fetch data from HTTP events on page 634](#)
- Replace HTTP body data. See [Replace HTTP body data on page 635](#)
- Run multiple scripts. See [Run multiple scripts on page 636](#)
- Prioritize scripts. See [Prioritize scripts on page 637](#)



Tip: The examples show debug strings. Debug strings can be written to the console when the event is triggered. This is helpful when you are testing your scripts.

To enable debug strings to be written to the console, use the following CLI commands:

```
diagnose debug enable
diagnose debug module httpproxy scripting
```

Select content routes based on URI string matches

The content routing feature has rules that match HTTP requests to content routes based on a Boolean AND combination of match conditions. If you want to select routes based on a Boolean OR, you can configure multiple rules. The content routing rules table is consulted from top to bottom until one matches.

In some cases, it might be simpler to get the results you want using a script. In the following example, each rule selects content routes based on OR match conditions.

Content routing example

```
when RULE_INIT {
  debug("get header init 1\n")
}
```

```
when HTTP_REQUEST{
  uri = HTTP:uri_get()
  if uri:find("sports") or uri:find("news") or uri:find("government") then
    LB:routing("sp2")
    debug("uri %s matches sports|news|government\n", uri);
  elseif uri:find("finance") or uri:find("technology") or uri:find("shopping") then
    LB:routing("sp3")
    debug("uri %s matches finance|technology|shopping\n", uri);
  elseif uri:find("game") or uri:find("bbs") or uri:find("testing") then
    LB:routing("sp4")
    debug("uri %s matches game|bbs|testing\n", uri);
  elseif uri:find("billing") or uri:find("travel") or uri:find("weibo") then
    LB:routing("sp5")
    debug("uri %s matches billing|travel|weibo\n", uri);
  else
    debug("no matches for uri: %s \n", uri);
  end
}
```

To use a script for content routing:

1. Create the content route configuration objects. In the example above, sp2, sp3, sp4, and sp5 are the names of the content route configuration objects. You do not need to configure matching conditions for the content routes, however, because the script does the content matching.
2. Create a script that matches content to the content route configuration objects, as shown above. Create a configuration object for the script.
3. In the virtual server configuration:
 - a. Enable content routing and select the content route configuration objects.
 - b. Select the script.

Rewrite the HTTP request host header and path

You can use the content rewriting feature to rewrite the HTTP request Host header or the HTTP request URL. If you need more granular capabilities, you can use scripts. The following example rewrites the HTTP Host header and path.

Rewrite the HTTP Host header and path in a HTTP request

```
when RULE_INIT {
  debug("rewrite the HTTP Host header and path in a HTTP request \n")
}

when HTTP_REQUEST{
  host = HTTP:header_get_value("Host")
  path = HTTP:path_get()
  if host:lower():find("myold.hostname.com") then
    debug("found myold.hostname.com in Host %s \n", host)
    HTTP:header_replace("Host", "mynew.hostname.com")
    HTTP:path_set("/other.html")
  end
}
```

Note: You might find it useful to use a combination of string manipulation functions. For example, this script uses `lower()` to convert the Host strings to lowercase in combination with `find()`, which searches for the Host header for a match:

```
host:lower():find("myold.hostname.com").
```

Rewrite the HTTP response Location header

You can use the content rewriting feature to rewrite the HTTP response Location header. If you are more comfortable using Lua string substitution, you can write a script to get the results you want. The following example rewrites the HTTP response Location header.

Rewrite the HTTP body in the response

```
when RULE_INIT {
  debug("rewrite the HTTP response replacing myold.hostname.com with mynew.hostname.com \n")
}

when HTTP_RESPONSE{
  location = HTTP:header_get_value("Location")
  if location:lower():find("myold.hostname.com") then
    debug("found myold.hostname.com in Location %s \n", location)
    HTTP:header_replace("Location", "mynew.hostname.com")
  end
}
```

Redirect HTTP to HTTPS using Lua string substitution

You can use the content rewriting feature to redirect an HTTP request to an HTTPS URL that has the same host and request URL using a PCRE regular expression. If you are more comfortable using Lua string substitution, you can write a script to get the results you want. The following example redirects users to the HTTPS location.

Redirect HTTP to HTTPS

```
when RULE_INIT {
  debug("http to https redirect\n")
}

when HTTP_REQUEST{
  host = HTTP:header_get_value("Host")
  path = HTTP:path_get()
  HTTP:redirect("https://%s%s",host,path);
}
```

Redirect mobile users to the mobile version of a website

The content rewriting feature does not support matching the User-Agent header. You can write a script that detects User-Agent headers that identify mobile device users and redirect them to the mobile version of a website.

Redirect mobile users to the mobile version of a website by parsing the User-Agent header

```
when RULE_INIT {
  debug("detect User-Agent and go to mobile site\n")
}

when HTTP_REQUEST{
```

```
path = HTTP:path_get()
debug("path=%s\n",path)
agent = HTTP:header_get_value("User-Agent")
if agent:lower():find("iphone") or agent:lower():find("ipad") then
debug("found iphone or ipad in User-Agent %s \n", agent)
HTTP:redirect("https://m.mymobilesite.com%s",path)
end
}
```

Insert random message ID into a header

FortiADC offers the feature to insert messages and message IDs into HTTP request headers.

```
when HTTP_REQUEST{
ID=HTTP:rand_id()-- a 32-long string of HEX symbols
HTTP:header_insert("Message-ID",ID)
}
```

General HTTP redirect

You can redirect both HTTP requests and HTTP responses to a given location.

```
GENERAL_REDIRECT_DEMO:
when HTTP_REQUEST{
--can be used in both HTTP_REQUEST and HTTP_RESPONSE
--code and cookie are optional, code can be 301, 302, 303, 307, 308, if missed, 302 is used
t={}
t["code"] = 302;
t["url"] = "www.example.com"
t["cookie"] = "name=value; Expires=Wed, 09 Jun 2021 10:18:14 GMT"
HTTP:redirect_t(t);
}
```

Use request headers in other events

You can get stored request headers by using the session ID.

```
when RULE_INIT{
--initialize the global so-called "environment" variable
env={}
}
when HTTP_REQUEST{
sess_id = HTTP:get_session_id()
req={}
--store whatever you want to req, take url as an example
req["url"] = HTTP_uri_get()
env[id] = req
}
when HTTP_RESPONSE{
sess_id = HTTP:get_session_id()
req = env[id]
--now you can access the stored request headers
}
```

```
debug("my stored request url is %s\n", req["url"]);
}
when HTTP_DATA_REQUEST{
  sess_id = HTTP:get_session_id()
  req = env[id]
  --now you can access the stored request headers
  debug("my stored request url is %s\n", req["url"]);
}
when HTTP_DATA_RESPONSE{
  sess_id = HTTP:get_session_id()
  req = env[id]
  --now you can access the stored request headers
  debug("my stored request url is %s\n", req["url"]);
}
```

Compare IP address to address group

You can compare IP addresses to an internal list of IP addresses. The script will return different results signifying whether the IP is in the list.

```
when RULE_INIT{
  --initialize the address group here
  --for IPv4 address, mask can be a number between 0 to 32 or a dotted format
  --support both IPv4 and IPv6, for IPv6, the mask is a number between 0 and 128
  addr_group = "192.168.1.0/24"
  addr_group = addr_group..",172.30.1.0/255.255.0.0"
  addr_group = addr_group..",::ffff:172.40.1.0/120"
}
when HTTP_REQUEST{
  client_ip = HTTP:client_addr()
  matched = cmp_addr(client_ip, addr_group)
  if matched then
    debug("client ip found in address group\n");
  else
    debug("client ip not in address group\n");
  end
}
```

Redirect HTTP to HTTPS

You can redirect an HTTP request from an HTTP location to an HTTPS location.

```
when HTTP_REQUEST{
  Host = HTTP:header_get_value("host")
  Url = HTTP:uri_get()
  HTTP:redirect("https://%s%s", Host, Url)
}
```

Rewrite HTTP to HTTPS in location

You can rewrite HTTP request headers to replace all HTTP addresses with HTTPS addresses in the redirect location.

```
when HTTP_RESPONSE{
  loc = HTTP:header_get_value("Location")
```

```
if loc then
newloc = string.gsub(loc, "http", "https") --replace all http by https in the redirect
    location
HTTP:header_replace("Location", newloc);
end
}
```

Rewrite HTTP to HTTPS in referer

You can rewrite HTTP request headers to replace all HTTP addresses with HTTPS addresses in the redirect referer.

```
when HTTP_RESPONSE{
ref = HTTP:header_get_value("Referer")
if ref then
newref = string.gsub(ref, "http", "https") --replace all http by https in the referer header
HTTP:header_replace("Referer", newref);
end
}
```

Rewrite HTTPS to HTTP in location

You can rewrite HTTP request headers to replace all HTTPS addresses with HTTP addresses in the redirect location.

```
when HTTP_RESPONSE{
loc = HTTP:header_get_value("Location")
if loc then
newloc = string.gsub(loc, "https", "http") --replace all https by http in the redirect
    location
HTTP:header_replace("Location", newloc);
end
}
```

Rewrite HTTPS to HTTP in referer

You can rewrite HTTP request headers to replace all HTTPS addresses with HTTP addresses in the redirect referer.

```
when HTTP_RESPONSE{
ref = HTTP:header_get_value("Referer")
if ref then
newref = string.gsub(ref, "https", "http") --replace all https by http in the referer header
HTTP:header_replace("Referer", newref);
end
}
```

Fetch data from HTTP events

You can collect data from both HTTP request and HTTP response events. You can then manipulate this data.

```
when HTTP_REQUEST{
--HTTP:collect command can be used in both HTTP_REQUEST and HTTP_RESPONSE events
--size is optional, otherwise, it will collect up to the full length or when 1.25M is reached
```

```
t={}
t["size"] = 100;
HTTP:collect(t)
}
when HTTP_DATA_REQUEST{
--check the size of the content
t={};
t["operation"]="size";
sz=HTTP:payload(t);
debug("content size: %s\n", sz);
--fetch the collected content
--offset and size are optional
t={};
t["operation"]="content";
t["offset"] = 0;
t["size"] = sz;
ct=HTTP:payload(t);
debug("content: %s\n", ct);
--do your own manipulation on the collected content
--replace the collected content by your new data
--offset and size are optional
t={};
t["operation"]="set";
t["offset"] = 0;
t["size"] = sz;
t["data"]="NEW DATA to SEND";
ret = HTTP:payload(t);
debug("set ret %s\n", ret);
}
```

Replace HTTP body data

You can find, remove, and replace data in the body of an HTTP request.

```
when HTTP_REQUEST{
--HTTP:collect command can be used in both HTTP_REQUEST and HTTP_RESPONSE events
--size is optional, otherwise, it will collect up to the full length or when 1.25M is
  reached
t={}
t["size"] = 100;
HTTP:collect(t)
}
when HTTP_DATA_REQUEST{
--check the size of the content
t={};
t["operation"]="size";
sz=HTTP:payload(t);
debug("content size: %s\n", sz);
--find a string or a regular expression in the buffered data
--offset, size and scope are optional, if scope is missing, "all" is assumed
t={};
t["operation"]="find";
t["offset"] = 0;
t["size"] = sz;
t["scope"] = "all";-- or "first"
t["data"] = "your string or a regular expression to find";
```

```
if HTTP:payload(t) then
debug("found %d occurrences\n", ret);
else
debug("not found\n");
end
--remove a string or a regular expression in the buffered data
--offset, size and scope are optional, if scope is missing, "all" is assumed
t={};
t["operation"]="remove";
t["offset"] = 0;
t["size"] = sz;
t["scope"] = "all";-- or "first"
t["data"] = "your string or a regular expression to find";
if HTTP:payload(t) then
debug("removed %d occurrences\n", ret);
else
debug("not found\n");
end
--replace a string or a regular expression in the buffered data by a new string
--offset, size and scope are optional, if scope is missing, "all" is assumed
t={};
t["operation"]="replace";
t["offset"] = 0;
t["size"] = sz;
t["scope"] = "all";-- or "first"
t["data"] = "your string or a regular expression to find";
t["new_data"] = "your new data";
if HTTP:payload(t) then
debug("replaced %d occurrences\n", ret);
else
debug("not found\n");
end
end
}
```

Run multiple scripts

You can run multiple scripts in FortiADC. When running multiple scripts, you may set a priority number for each script. FortiADC will run them in order from lowest priority to highest priority. The default priority is 500. If two scripts have the same priority number, they will be executed in the order in which they were added.

--script 1:

```
when HTTP_REQUEST priority 500 {
LB:routing("cr1")
}
```

--script 2:

```
when HTTP_RESPONSE priority 500 {
HTTP:close()
}
```

--script 3:

```
when HTTP_REQUEST priority 400 {
LB:routing("cr2")
}
```


--script 4:

```
when HTTP_RESPONSE priority 600 {
  HTTP:close()
}
```

Prioritize scripts

While running multiple scripts, you can prioritize scripts. Add a priority number to each script when you create it, and FortiADC will run them in order from lowest priority to highest priority. The default priority is 500. If two scripts have the same priority number, they will be executed in the order in which they were added.

```
when RULE_INIT priority 14 {
```

--This is one of the script to demo the control of multiple scripts

--please change the priority of each event according to your need

```
  debug("INIT in script 1\n");
}
when HTTP_REQUEST priority 12 {
  debug("HTTP_REQUEST in script 1\n");
```

--add your own manipulation here

--you can disable rest of the HTTP_REQUEST events from executing by disabling this event

```
  t={};
  t["event"]="req"; -- can be "req", "res", "data_req", and "data_res"
  t["operation"]="disable"; -- can be "enable", and "disable"
  HTTP:set_event(t);
  debug("disable rest of the HTTP_REQUEST events in script 1\n");
```

--you can also disable other events, say HTTP_RESPONSE, DATA events

--in the case of keep-alive, all events will be re-enabled automatically even though they are disabled in previous TRANSACTION using the HTTP:set_event(t) command. To disable this automatic re-enabling behavior, you can call HTTP:set_auto(t) as below

```
  t={};
  t["event"]="req"; -- can be "req", "res", "data_req", and "data_res"
  t["operation"]="disable"; -- can be "enable", and "disable"
  HTTP:set_auto(t);
  debug("disable automatic re-enabling of the HTTP_REQUEST events in script 1\n");
  --you can also disable automatic re-enabling for other events, say HTTP_RESPONSE, DATA events
}
```

or

```
when RULE_INIT priority 24 {
  --This is one of the script to demo the control of multiple scripts
  --please change the priority of each event according to your need
  debug("INIT in script 2\n");
}
when HTTP_REQUEST priority 24 {
  debug("HTTP_REQUEST in script 2\n");
  --add your own manipulation here
  --you can disable rest of the HTTP_REQUEST events from executing by disabling this event
  t={};
  t["event"]="req"; -- can be "req", "res", "data_req", and "data_res"
```

```
t["operation"]="disable"; -- can be "enable", and "disable"
HTTP:set_event(t);
debug("disable rest of the HTTP_REQUEST events in script 2\n");
--you can also disable other events, say HTTP_RESPONSE, DATA events
--in the case of keep-alive, all events will be re-enabled automatically even though they are
    disabled in previous TRANSACTION using the HTTP:set_event(t) command. To disable this
        automatic re-enabling behavior, you can call HTTP:set_auto(t) as below
t={};
t["event"]="req"; -- can be "req", "res", "data_req", and "data_res"
t["operation"]="disable"; -- can be "enable", and "disable"
HTTP:set_auto(t);
debug("disable automatic re-enabling of the HTTP_REQUEST events in script 2\n");
--you can also disable automatic re-enabling for other events, say HTTP_RESPONSE, DATA events
}
```

Appendix D: Maximum Configuration Values

[Maximum configuration objects - hardware models on page 639](#) and [Maximum configuration objects - virtual appliances on page 641](#) show the maximum number of configuration objects by hardware or VM model. For more information specific to your FortiADC appliance, refer to your model's *QuickStart Guide* or *Datasheet*.

Note: The maximum number of Layer-7 virtual servers that each model supports varies, depending on the available system memory and the number of features enabled on the unit.

Maximum configuration objects - hardware models

Parameters		60F	100F	200F	300F	400F	1000F	2000F	4000F	200D	300D	400D	700D	1500D	2000D	4000D
System																
Administration	Administrative users	300	300	300	300	300	300	300	300	300	300	300	300	300	300	300
	Access profiles	16	16	16	64	64	64	64	64	16	64	64	64	64	64	64
	Virtual domains (VDOMs)	2	10	10	10	10	45	60	90	10	10	10	30	45	60	90
Certificates	Any configuration object	256	256	256	256	256	256	256	256	256	256	256	256	256	256	256
Shared Resources	Address	1024	1024	1024	2048	2048	2048	2048	4096	1024	2048	2048	2048	2048	2048	4096
	Address group	256	256	256	256	256	256	256	256	256	256	256	256	256	256	256
	Health checks	128	128	128	256	256	256	256	512	128	256	256	256	256	256	512
	ISP address book	32	32	32	32	32	32	32	32	32	32	32	32	32	32	32
	Schedule	256	256	256	256	256	256	256	256	256	256	256	256	256	256	256
	Schedule group	64	64	64	64	64	64	64	64	64	64	64	64	64	64	64
	Service	1024	1024	1024	2048	2048	2048	2048	4096	1024	2048	2048	2048	2048	2048	4096
	Service group	256	256	256	256	256	256	256	256	256	256	256	256	256	256	256
SNMP	SNMP community	16	16	16	16	16	16	16	16	16	16	16	16	16	16	16
	SNMP community Host	16	16	16	16	16	16	16	16	16	16	16	16	16	16	16
	SNMP user	16	16	16	16	16	16	16	16	16	16	16	16	16	16	16
Networking																
Interface	Physical network interfaces	5	4	6	8	10	21	25	15	4	4	4	12	12	20	24

Appendix D: Maximum Configuration Values

Parameters		60F	100F	200F	300F	400F	1000F	2000F	4000F	200D	300D	400D	700D	1500D	2000D	4000D
Routing	VLAN interfaces	256	256	256	512	512	512	512	1024	256	512	512	512	512	512	1024
	ARP table entries (per VDOM)	4096	4096	4096	4096	4096	4096	4096	4096	4096	4096	4096	4096	4096	4096	4096
	Static routes	2048	2048	2048	4096	4096	4096	4096	4096	2048	4096	4096	4096	4096	4096	4096
	Policy routes	64	64	64	128	128	128	128	256	64	128	128	128	128	128	256
	ISP routes	32	32	32	32	32	32	32	32	32	32	32	32	32	32	32
NAT	Any configuration object	256	256	256	256	256	256	256	256	256	256	256	256	256	256	256
QoS	Any configuration object	256	256	256	256	256	256	256	256	256	256	256	256	256	256	256
Packet capture	Table	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5
User																
	Any configuration object	256	256	256	256	256	256	256	256	256	256	256	256	256	256	256
Server Load Balancing																
Virtual Servers	L4	1024	1024	1024	2048	2048	2048	2048	4096	1024	2048	2048	2048	2048	2048	4096
	L7	128	256	256	512	512	1024	2048	4096	256	512	512	1024	1024	1024	2048
	L7 HTTPs	64	128	128	256	256	512	512	2048	128	256	256	512	512	512	1024
Real Server Pool	Pools	1024	1024	1024	2048	2048	2048	2048	4096	1024	2048	2048	2048	2048	2048	4096
	Pool members	1024	1024	1024	2048	2048	2048	2048	4096	1024	2048	2048	2048	2048	2048	4096
	Real server SSL profiles	256	256	256	256	256	256	256	256	256	256	256	256	256	256	256
Resources	Profiles	256	256	256	256	256	256	256	256	256	256	256	256	256	256	256
	Cache policies	256	256	256	256	256	256	256	256	256	256	256	256	256	256	256
	Compression policies	256	256	256	256	256	256	256	256	256	256	256	256	256	256	256
	Persistence policies	128	128	128	256	256	256	256	512	128	256	256	256	256	256	512
	Method policies	64	64	64	128	128	128	128	256	64	128	128	128	128	128	256
	Authentication policies	256	256	256	256	256	256	256	256	256	256	256	256	256	256	256
	Scripts	256	256	256	256	256	256	256	256	256	256	256	256	256	256	256

Appendix D: Maximum Configuration Values

Parameters		60F	100F	200F	300F	400F	1000F	2000F	4000F	200D	300D	400D	700D	1500D	2000D	4000D
Content Rules	Content routing rules	256	256	256	512	512	512	512	1024	256	512	512	512	512	512	1024
	Content rewriting rules	256	256	256	512	512	512	512	1024	256	512	512	512	512	512	1024
Link Load Balancing																
Link Group	Gateway	1024	1024	1024	2048	2048	2048	2048	4096	1024	2048	2048	2048	2048	2048	4096
	Link group	512	512	512	1024	1024	1024	1024	2048	512	1024	1024	1024	1024	1024	2048
	Link group member	1024	1024	1024	2048	2048	2048	2048	4096	1024	2048	2048	2048	2048	2048	4096
Virtual Tunnel Group	Virtual tunnel group	512	512	512	1024	1024	1024	1024	2048	512	1024	1024	1024	1024	1024	2048
	Virtual tunnel member	256	256	256	256	256	256	256	256	256	256	256	256	256	256	256
Policy	LLB policy rule	512	512	512	1024	1024	1024	1024	1024	512	1024	1024	1024	1024	1024	2048
Global Load Balancing																
Any configuration object		256	256	256	256	256	256	256	256	256	256	256	256	256	256	256
Security																
Any configuration object		256	256	256	256	256	256	256	256	256	256	256	256	256	256	256
Log & Report																
Remote Syslog Servers		3	3	3	3	3	3	3	3	3	3	3	3	3	3	3

Maximum configuration objects - virtual appliances


Parameters		VM01	VM02	VM04	VM08
System					
Administration	Administrative users	300	300	300	300
	Access profiles	8	16	64	64
	Virtual domains (VDOMs)	0	0	5	10
Certificate	Any configuration object	256	256	256	256

Parameters		VM01	VM02	VM04	VM08
Shared Resources	Address	512	1024	2048	4096
	Address group	256	256	256	256
	Health checks	64	128	256	512
	ISP address book	32	32	32	32
	Schedule	256	256	256	256
	Schedule group	64	64	64	64
	Service	512	1024	2048	4096
	Service group	256	256	256	256
SNMP	SNMP community	16	16	16	16
	SNMP community host	16	16	16	16
	SNMP user	16	16	16	16
Networking					
Interfaces	Physical network interfaces	10	10	10	10
	VLAN interfaces	128	256	512	1024
Routing	ARP table entries (per VDOM)	4096	4096	4096	4096
	Static routes	1024	2048	4096	4096
	Policy routes	32	64	128	256
	ISP routes	32	32	32	32
NAT	Any configuration object	256	256	256	256
QoS	Any configuration object	256	256	256	256
Packet Capture	Table	5	5	5	5
User					
	Any configuration object	256	256	256	256
Server Load Balancing					
Virtual Servers	L4	512	1024	2048	4096
	L7	128	256	512	1024
	L7 HTTPs	64	128	256	512
Real Server Pool	Pools	512	1024	2048	4096
	Pool members	512	1024	2048	4096
	Real server SSL profile	256	256	256	256

Parameters		VM01	VM02	VM04	VM08
Resources	Profiles	256	256	256	256
	Cache policies	256	256	256	256
	Compression policies	256	256	256	256
	Persistence policies	128	128	128	256
	Method policies	32	64	128	256
	Authentication policies	256	256	256	256
	Scripts	256	256	256	256
Content Rules	Content routing rules	128	256	512	1024
	Content rewriting rules	128	256	512	1024
Link Load Balancing					
Link Group	Gateway	512	1024	2048	4096
	Link group	256	512	1024	2048
	Link group member	512	1024	2048	4096
Virtual Tunnel	Virtual tunnel	256	512	1024	2048
	Virtual tunnel member	256	256	256	256
Policy	LLB policy rule	256	512	1024	2048
Global Load Balancing					
Any configuration object		256	256	256	256
Security					
Any configuration object		256	256	256	256
Log & Report					
Remote Syslog Servers		3	3	3	3



FORTINET®



Copyright© 2020 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.