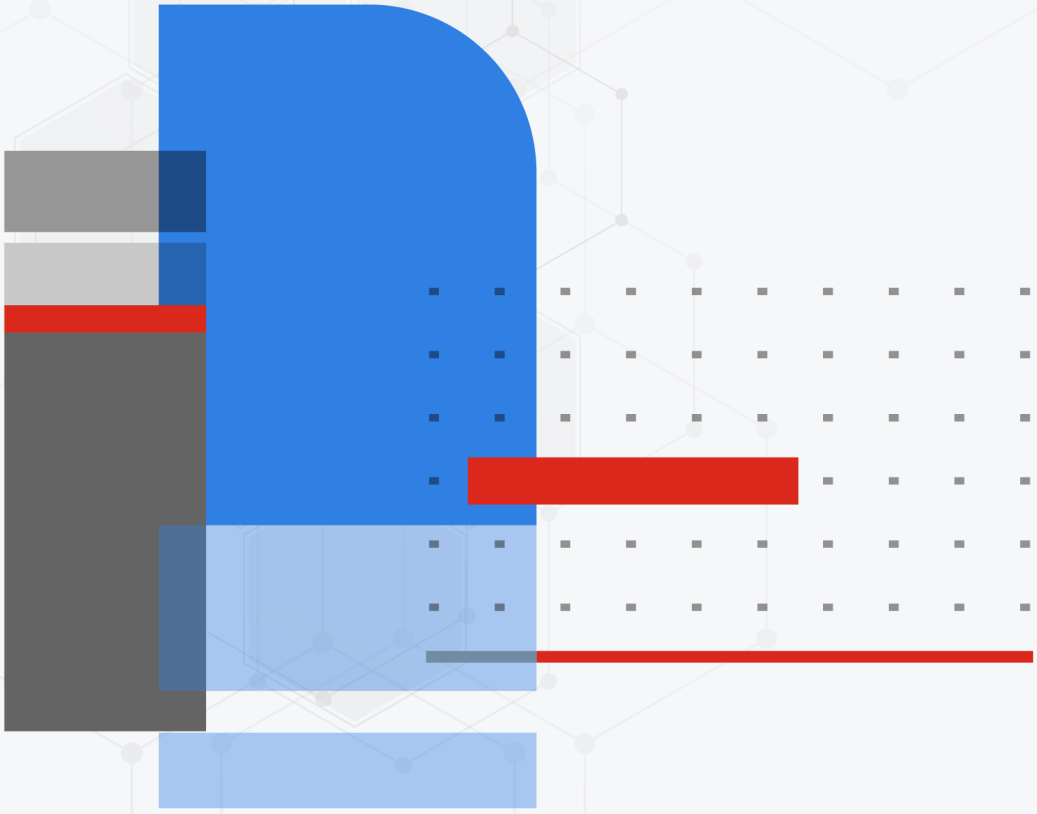
Administration Guide

FortiADC 7.6.1



FORTINET DOCUMENT LIBRARY
https://docs.fortinet.com

FORTINET VIDEO LIBRARY
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/training-certification

FORTINET TRAINING INSTITUTE
https://training.fortinet.com

FORTIGUARD LABS
https://www.fortiguard.com

END USER LICENSE AGREEMENT
https://www.fortinet.com/doc/legal/EULA.pdf

FEEDBACK
Email: techdoc@fortinet.com



November 25, 2024
FortiADC 7.6.1 Administration Guide

TABLE OF CONTENTS

Change Log **13**
Introduction **14**
Features 14
Basic network topology 14
Scope 16
Key Concepts and Features **17**
Server load balancing 17
Feature summary 17
Authentication 18
Caching 18
Compression 19
Decompression 19
Content rewriting 19
Content routing 19
Scripting 19
SSL transactions 19
Link load balancing 20
Global load balancing 20
Security 20
High availability 21
Virtual Domain (VDOM) and Administrative Domain (ADOM) 21
Getting Started **22**
Step 1: Install the appliance 22
Step 2: Configure the management interface 23
Step 3: Configure basic network settings 25
Step 4: Test connectivity to destination servers 29
Step 5: Complete product registration, licensing, and upgrades 30
Validating a VM license with no internet connection 31
Step 6: Configure a basic server load balancing policy 31
Step 7: Test the deployment 34
Step 8: Back up the configuration 36
Dashboard **37**
Widgets 37
Dashboard management tools 38
Adding a dashboard 39
Editing a dashboard 39
Deleting a dashboard 40
Adding Features 40
Security Fabric **42**
Automation 42
Creating automation stitches 42
Configuring Automation Triggers 54
Configuring Automation Actions 67

Diagnose commands ... 77
Fabric connectors ... 77
FortiSIEM Connector ... 77
FortiAnalyzer Connector ... 78
FortiSandbox Connector ... 82
FortiADC Manager Connector ... 83
FortiGSLB Connector ... 85
FortiClient EMS Connector ... 86
FortiGate Security Fabric Connector ... 89
External connectors ... 94
Amazon Web Services (AWS) Connector ... 95
Oracle Cloud Infrastructure (OCI) Connector ... 98
Kubernetes Connector ... 100
Splunk Connector ... 104
SAP Connector ... 105
IP Address Connector ... 106
FortiView ... 109
Logical Topology ... 109
Server Load Balance Logical Topology ... 109
Link Load Balance Logical Topology ... 110
Global Load Balance Logical Topology ... 111
Server Load Balance ... 112
Virtual Servers ... 112
Data Analytics ... 119
All Sessions ... 121
Security ... 121
OWASP Top 10 ... 121
OWASP Top 10 Compliance ... 124
Threat Map ... 129
Data Analytics ... 130
Security Logs ... 133
Blocked IP ... 138
System ... 139
Event Logs ... 139
Automation ... 140
Global Load Balance ... 141
Host ... 141
Data Analytics ... 142
Link Load Balance ... 142
Gateway ... 142
ZTNA FortiClient endpoint ... 143
System ... 145
Settings ... 145
Configuring basic system settings ... 145
Sending FortiADC Threat Telemetry to FortiGuard ... 148
Configuring system time ... 150
Updating firmware ... 152

Configuring an SMTP mail server 155
Pushing/pulling configurations 156
Backing up and restoring configuration 158
Configuring DNS settings for non-root VDOM 162
Cloud Auto Scaling 163
Azure LB Backend 166
Virtual Domain 171
Virtual Domain (VDOM) and Administrative Domain (ADOM) overview 171
Enabling the Virtual Domain feature and selecting the Virtual Domain Mode 174
Creating a virtual domain 175
Assigning administrator users and network interfaces to VDOMs 175
Virtual domain policies 176
Disabling a virtual domain 179
High Availability 180
Configuring HA settings 183
HA configuration synchronization 192
Monitoring an HA cluster 194
Updating firmware for an HA cluster 200
HA Troubleshooting 202
Deploying an Active-Passive cluster 207
Deploying an Active-Active cluster 218
Deploying an Active-Active-VRRP cluster 236
Traffic Group 246
Administrator 247
Managing administrator users 248
Creating administrator users 249
Creating REST API administrator users 251
Configuring access profiles 253
Enabling password policies 255
Global Resources 256
WCCP 258
SNMP 259
Configuring SNMP 260
Downloading SNMP MIBs 261
Configuring SNMP v1/v2 262
Configuring SNMP v3 263
Replacement Messages 266
Enable the pre-login banner 266
FortiGuard 266
Connecting to FortiGuard services 267
Configuring FortiGuard service settings 270
Debug 274
Certificate 275
Manage Certificates 275
Verify 292
OCSP 297

Network ... **304**
Interface ... 304
Configuring network interfaces ... 307
Configuring the management interface ... 313
Configuring virtual overlay networks ... 315
Linking VDOMs for inter-VDOM routing ... 317
Transparent mode ... 322
Routing ... 322
Configuring static routes ... 323
Configuring policy routes ... 324
Configuring OSPF routes ... 325
Configuring OSPFv3 routes ... 329
Configuring ISP routes ... 333
Configuring BGP routes ... 335
Configuring Bidirectional Forwarding Detection (BFD) ... 339
Access list vs. prefix list ... 340
Configuring an IPv4 access list ... 341
Configuring an IPv6 access list ... 342
Configuring an IPv4 prefix list ... 343
Configuring an IPv6 prefix list ... 344
NAT ... 345
Configuring source NAT ... 345
Configuring 1-to-1 NAT ... 348
QoS ... 350
Configuring the QoS filter ... 351
Configuring the QoS IPv6 filter ... 352
Configuring a QoS queue ... 353
Packet Capture ... 353
Shared Resources ... **355**
Health Check ... 355
Configuring health checks ... 355
Monitoring health check status ... 365
Configuring health check scripts ... 366
Schedule Group ... 368
Address ... 368
Configuring IPv4 address groups ... 369
Creating IPv4 address objects ... 370
Configuring IPv6 address groups ... 370
Creating IPv6 address objects ... 371
Managing ISP address books ... 372
Service ... 375
Creating service groups ... 375
Creating service objects ... 376
Server Load Balance ... **378**
Server load balancing configuration overview ... 382
Virtual Server ... 383
Configuring virtual servers ... 384

Using content rewriting rules 397
Configuring content routes 412
Using source pools 416
Using schedule pools 429
Using clone pools 430
Load Balancing for Layer 4 Virtual Servers 434
Application Resources 450
Configuring Application profiles 451
Configuring Client SSL profiles 512
Configuring HTTP2 profiles 518
Configuring HTTP3 profiles 521
Configuring load-balancing (LB) methods 526
Configuring persistence rules 527
Configuring error pages 535
Configuring decompression rules 536
Configuring Captcha 541
Application Optimization 542
Creating a Page Speed configuration 543
Creating Page Speed profiles 545
Configuring compression rules 546
Using caching features 548
Real Server Pool 551
Using real server pools 552
Configuring real servers 559
Configuring real server SSL profiles 562
Scripting 569
Using HTTP scripting 569
Predefined HTTP scripts 572
Multi-script support 580
Using Stream scripting 588
SSL-FP Resources 594
Configuring an L2 exception list 594
Creating a Web Filter Profile configuration 595
Viewing Web Filter categories 596
Configuring certificate caching 596
Link Load Balance 598
Link load balancing basics 598
Link load balancing configuration overview 601
Link Policy 603
Link Group 604
Configuring a link group 605
Configuring gateway links 607
Configuring LLB persistence rules 608
Configuring proximity route settings 610
Virtual Tunnel 611
Global Load Balance 613
Global load balancing basics 613
Global load balancing configuration overview 615

GLB Wizard 617
FQDN 619
Configuring hosts 619
Configuring virtual server pools 621
Configuring location lists 623
Configuring GLB settings 623
Zone Tools 626
Configuring a Global DNS policy 626
Configuring DNS zones 627
Configuring general settings 634
Configuring the trust anchor key 648
Configuring DNS64 649
Configuring the DSSET list 650
Configuring an address group 651
Configuring remote DNS servers 652
Configuring the response rate limit 652
Global Object 653
Configuring servers 653
Configuring a global load balance link 658
Configuring data centers 658
Web Application Firewall 660
Web application firewall basics 660
Web application firewall configuration overview 663
OWASP TOP10 665
WAF Profile 667
Configuring a WAF Profile 667
Configuring WAF Action objects 671
Configuring WAF Exception objects 673
Known Web Attacks 680
Configuring a Web Attack Signature policy 680
Using the Signature Creation Wizard 685
Common Attacks Detection 688
Configuring an Advanced Protection policy 688
Configuring an HTTP Protocol Constraint policy 690
Configuring an SQL/XSS Injection Detection policy 693
Configuring a CSRF protection policy 696
Sensitive Data Protection 697
Configuring a Cookie Security policy 697
Configuring an HTTP Header Security policy 700
Data Loss Prevention 704
FortiGuard DLP service 705
Configuring a DLP Policy 707
Configuring a Sensitive Data Type object 710
Configuring a DLP Sensor object 713
Configuring a DLP Dictionary object 715
Input Validation 718
Configuring an Input Validation policy 718
Configuring a Parameter Validation rule 719

Configuring a Hidden Field rule ... 721
Configuring a File Restriction rule ... 722
Access Protection ... 729
Configuring a URL Protection policy ... 730
Configuring a Brute Force Attack Detection policy ... 731
Configuring a Credential Stuffing Defense policy ... 732
CORS Protection ... 733
Configuring a CORS Protection Rule ... 733
Configuring an Allowed Origin List ... 736
Configuring a CORS Headers List ... 738
API Protection ... 739
JSON Protection ... 739
XML Protection ... 742
OpenAPI Validation ... 746
API Gateway ... 748
API Discovery ... 751
Bot Mitigation ... 756
Configuring a Bot Detection policy ... 756
Configuring a Threshold Based Detection policy ... 762
Configuring a Biometrics Based Detection policy ... 767
Configuring a Fingerprint Based Detection policy ... 769
Web Vulnerability Scanner ... 771
Scan Profile ... 774
Scan History ... 776
Scan Integration ... 778
Web Anti-Defacement ... 782
WAF Adaptive Learning ... 784
Configuring an Adaptive Learning policy ... 787
Viewing Adaptive Learning Analysis and Recommendation ... 793
Practical Applications of Adaptive Learning ... 801
Adaptive Learning troubleshooting and debugging ... 811
Advanced Bot Protection (ABP) ... 812
Prerequisites ... 813
Basic configuration ... 814
Enabling the Advanced Bot Protection connector ... 815
Obtaining the Application ID from the FortiGuard ABP User Portal ... 816
Configuring an Advanced Bot Protection policy ... 821
Advanced Bot Protection troubleshooting and debugging ... 826
Network Security ... 829
Security features basics ... 829
Firewall ... 829
Configuring an IPv4 firewall policy ... 830
Configuring an IPv6 firewall policy ... 832
Configuring an IPv4 connection limit policy ... 834
Configuring an IPv6 connection limit policy ... 835
Intrusion Prevention ... 836

AntiVirus 841
Creating an AV profile 842
Setting AV quarantine policies 844
Setting the AV service level 846
IP Reputation 847
Managing IP Reputation policy settings 848
Configuring an IP reputation exception 849
Configuring an IP reputation block list 849
Geo IP Protection 850
Using the Geo IP block list 850
Using the Geo IP allowlist 852
Special Geo codes 853
Zero Trust Network Access (ZTNA) 854
ZTNA Overview 854
ZTNA telemetry, tags, and policy enforcement 854
ZTNA in FortiADC server load balancing 855
Prerequisites 856
Basic ZTNA configuration 857
How device identity and trust context is established with FortiClient EMS 858
Configuring FortiClient EMS Connector for ZTNA 859
Verifying EMS CA certificate, ZTNA tag, and FortiClient endpoint synchronized from FortiClient EMS 863
Configuring a ZTNA Profile 865
ZTNA troubleshooting and debugging 866
DoS Protection 867
DoS Protection Profile 867
Application 869
Configuring an HTTP Access Limit policy 870
Configuring an HTTP Connection Flood policy 871
Configuring an HTTP Request Flood policy 873
Configuring a DNS Query Flood Protection policy 874
Configuring a DNS Reverse Flood Protection policy 877
Networking 879
Configuring an IP fragmentation policy 879
Configuring a TCP SYN flood protection policy 880
Configuring a TCP slow data flood protection policy 881
Configuring a TCP connection access flood protection policy 882
User Authentication 883
Authentication Policy 883
User Group 885
Configuring user groups 886
Configuring customized authentication forms 888
Local User 892
Remote User 893
Using an LDAP authentication server 893
Using a RADIUS authentication server 903
Using an NTLM authentication server 905

Using a TACACS+ authentication server 906
Authentication Relay 908
Using HTTP Basic SSO 912
SAML 913
Configure an SAML service provider 914
Import IDP Metadata 917
AD FS Proxy 917
Adding an AD FS Publish 918
Adding an AD FS Proxy 919
OAuth Proxy 920
Deploying OAuth 2.0 authentication 921
Log & Report 924
Using the traffic log 925
Using the security log 932
Using the script log 937
Log Setting 937
Configuring local log settings 937
Configuring syslog settings 939
Configuring OFTP settings for FortiAnalyzer logs 943
Configuring fast stats log settings 945
Report Setting 946
Configuring reports 946
Configuring report emails 947
Configuring report queries 948
Configuring fast reports 950
AI Threat Analytics 953
Overview 953
AI Threat Analytics troubleshooting and debugging 958
SSL Advanced Services 962
SSL offloading 962
SSL decryption by forward proxy 964
SSL profile configurations 967
Certificate guidelines 970
SSL/TLS versions and cipher suites 970
Exceptions list 975
SSL traffic mirroring 976
HSM Integration 979
Best Practices and Fine Tuning 986
Regular backups 986
Rebooting, resetting, and shutting down the system 986
SCP support for configuration backup 987
Security 987
Performance tips 989
High availability 989

FortiADC Use Cases **991**
DNS Load Balancing and DDoS Protection 992
User Scenario: DDoS Mitigation and DNS Reliability for Global Online Services 992
Global Server Load Balancing in Hybrid Environments 999
User Scenario: Implementing GSLB with FortiADC for a Hybrid E-Commerce Environment 999
Selecting the Optimal Mode: VDOM vs. ADOM 1019
Use-case for VDOM 1019
Use-case for ADOM 1029
SAP Application Integration 1037
User Scenario: Enhancing SAP Application Delivery with the FortiADC SAP Connector 1037
Multi-Cloud Application Load Balancing 1052
User Scenario: Multi-Cloud Traffic Distribution with FortiADC 1052
Leveraging DevOps Tools (Ansible/Terraform) to Manage FortiADC 1066
User Scenario: Automating FortiADC Management with Ansible and Terraform 1066
Troubleshooting **1072**
Logs 1072
Tools 1072
execute commands 1072
diagnose commands 1073
System dump 1074
Diff 1075
Solutions by issue type 1076
Resetting the configuration 1083
Restoring firmware (“clean install”) 1083
Additional resources 1086
Appendices **1087**
Maximum Configuration Values 1088
Supported image versions for EOS models 1096

Change Log

Date	Change Description
November 25, 2024	FortiADC 7.6.1 Administration Guide initial release.

Introduction

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

The FortiADC 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

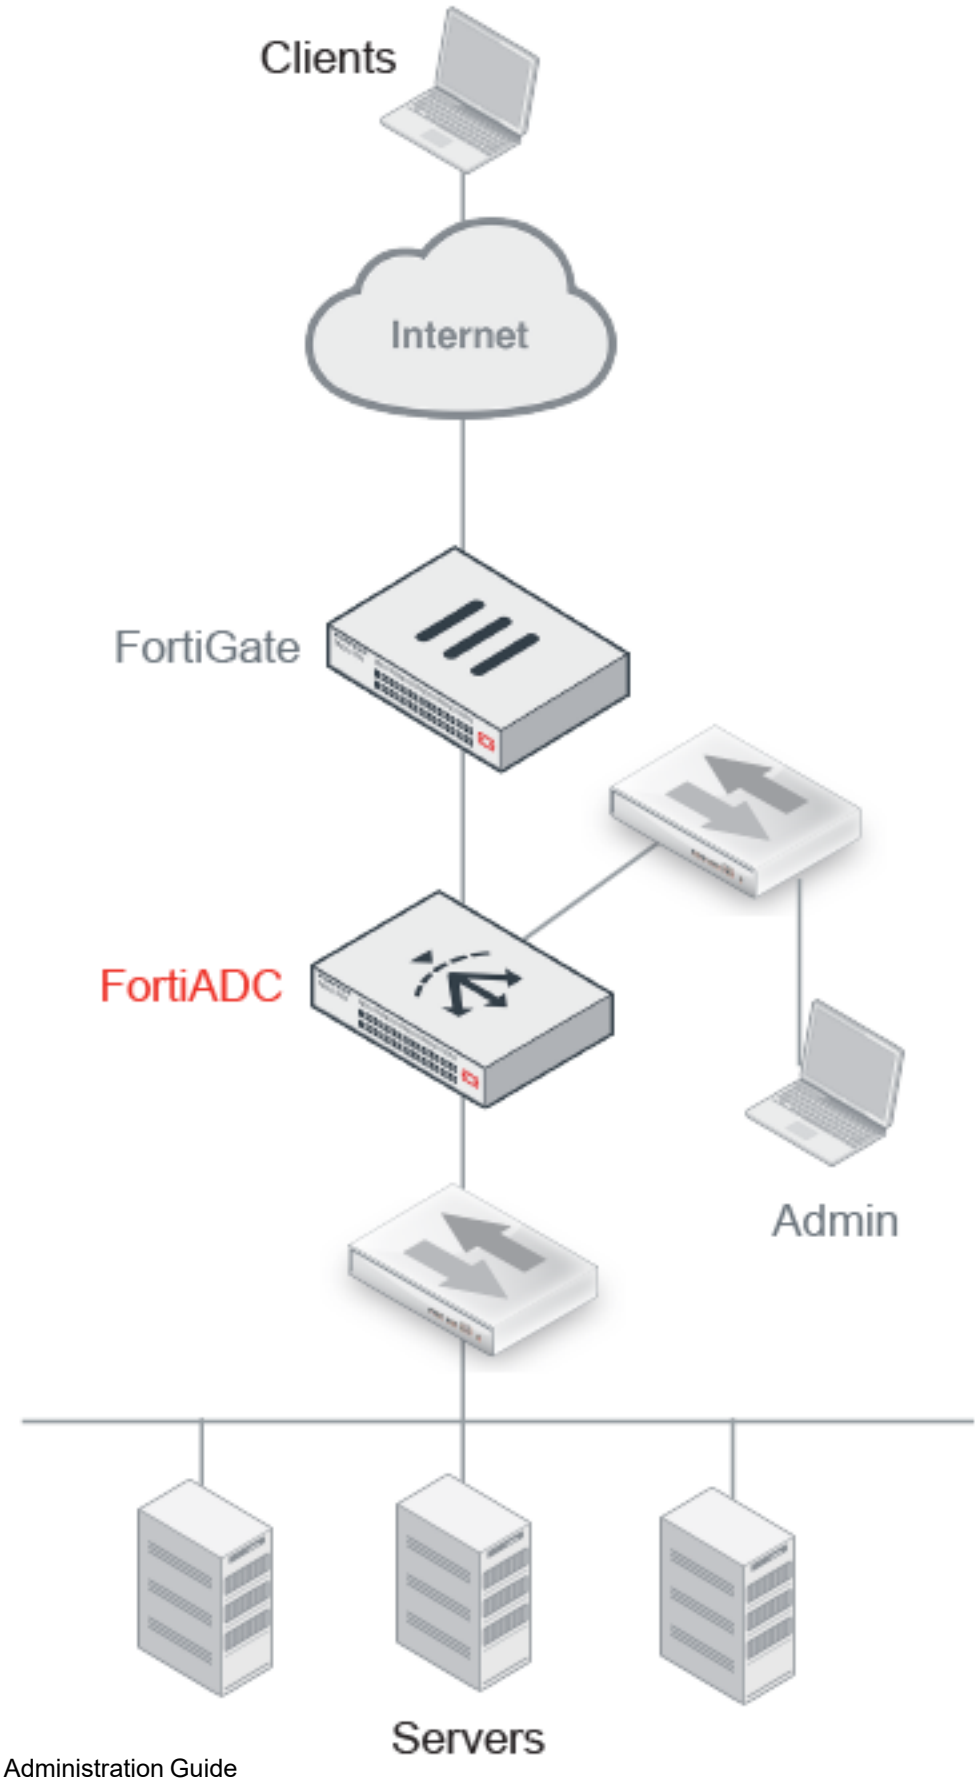
The FortiADC family includes physical appliances and virtual appliances.

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 14 shows a basic Router Mode deployment.

Basic network topology


Clients
Internet
FortiGate
FortiADC
Admin
Servers

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.

Key Concepts and Features

This section includes the following topics:

- Server load balancing on page 17
- Link load balancing on page 20
- Global load balancing on page 20
- Security on page 20
- High availability on page 21
- Virtual Domain (VDOM) and Administrative Domain (ADOM) on page 21

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

The table below 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.
Server management	• Warm up • Rate limiting

Features	Summary
	• 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, SMTP, FTP, Diameter, RTSP, RTMP, MySQL, MSSQL 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 Server Load Balance.

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 Authentication Policy.

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 HTTP scripting.

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 SSL Advanced Services.

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 Link Load Balance.

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 Global Load Balance.

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.
- AntiVirus—Provide protection against a variety of threats, including both known and unknown malicious codes (malware) and Advanced Target Attacks (ATA).
- Denial of service protection—Drop half-open connections to protect the system from a SYN flood attack.

For detailed information, see 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 High Availability.

Virtual Domain (VDOM) and Administrative Domain (ADOM)

A Virtual Domain (VDOM) is a complete FortiADC instance that runs on the FortiADC platform. VDOM configuration objects contain all of the system and feature configuration options of a full FortiADC instance and can be used to divide a FortiADC into two or more virtual units that function independently, allowing it to support multi-tennant deployments.

The VDOM feature supports two Virtual Domain Modes that allow the VDOMs to function independently with its own networking or as Administrative Domains (ADOMs) with shared networking between all ADOMs. When the VDOM is in the Independent Network mode, you can provision an administrator account with privileges to access and manage only their assigned VDOM. The VDOM user can then configure their VDOM as desired untethered to other VDOMs. Alternatively, when the VDOM is in Share Network mode, it functions as an ADOM that shares the same networking interfaces and routing between all the ADOMs. The ADOM functionality enables the administrator to constrain access privileges to a subset of server load-balancing servers by defaulting all interface settings to the root ADOM. For detailed information, see Virtual Domain.

Getting Started

This section 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:
   - New administrator password (required)
   - 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:

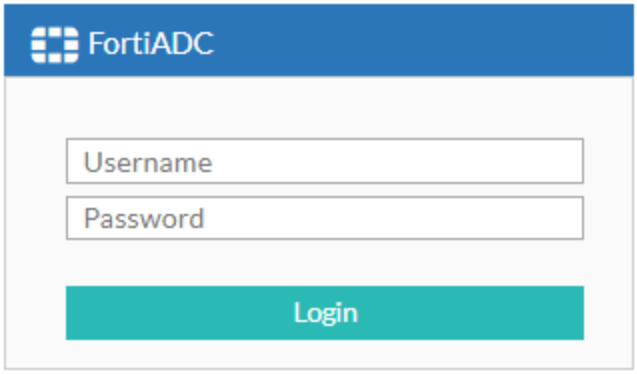
- You must know the IP address for the default gateway of the management subnet and the IP address that you plan to assign the management interface.
- You need access to the machine room in which a physical appliance has been installed. With physical appliances, you must connect a cable to the management 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 23.

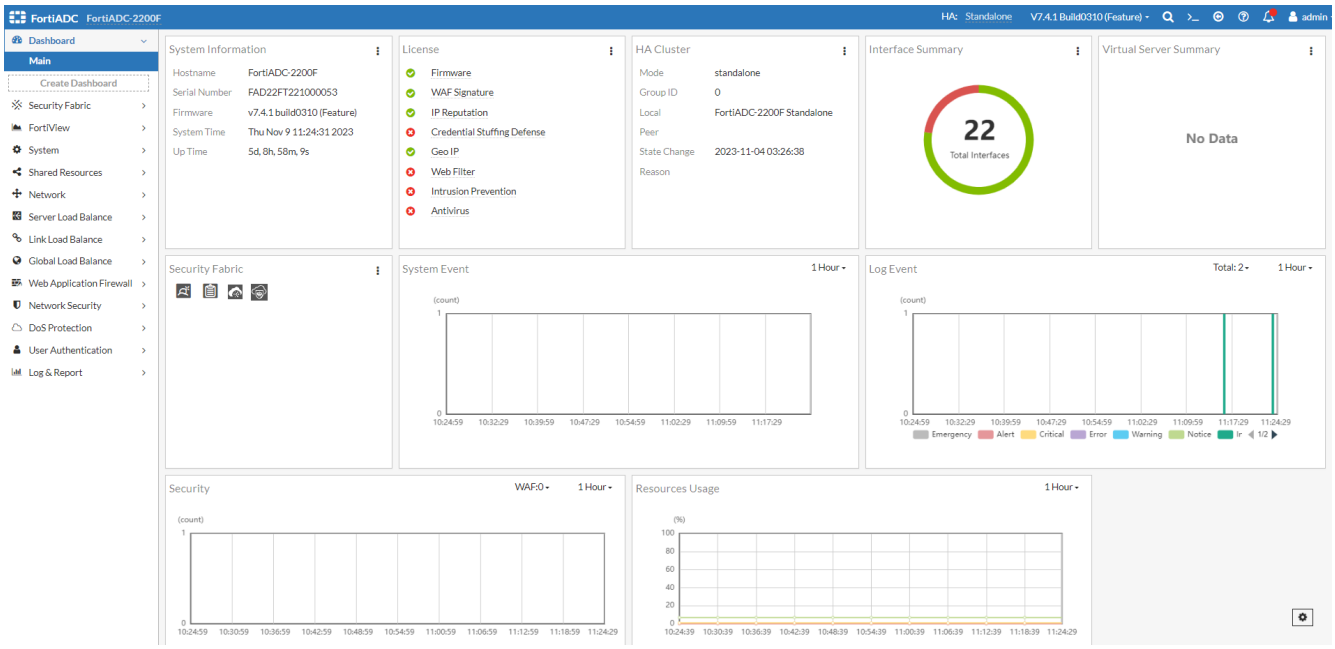
   Login page



4. Enter the username **admin** and set up a new password.

The system displays the dashboard. See Dashboard after initial login on page 24.

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 create a new 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:

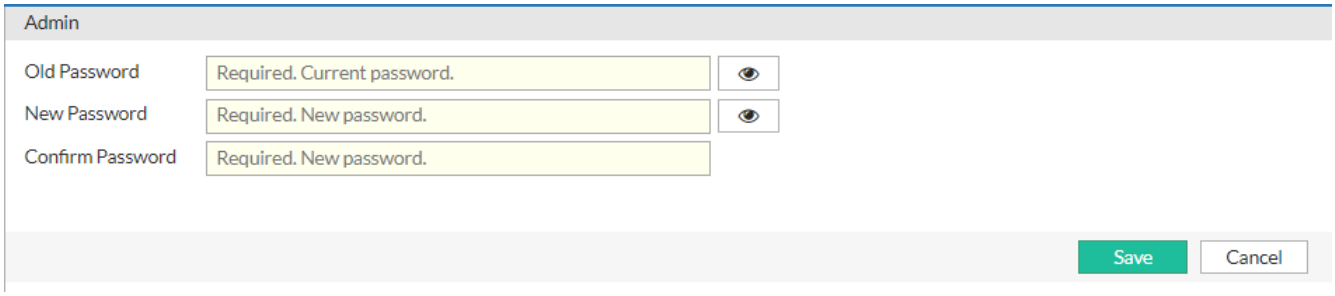
- You must know the IP address for the NTP servers your network uses to maintain system time.
- You must know the IP addresses that have been provisioned for the traffic interfaces for your FortiADC deployment.
- You must know the IP address for the primary and secondary DNS servers your network uses 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 26.
3. Change the password and save the configuration.

For detailed information on configuring administrator accounts, see Managing administrator users on page 248.

System administrator change password editor



CLI commands:

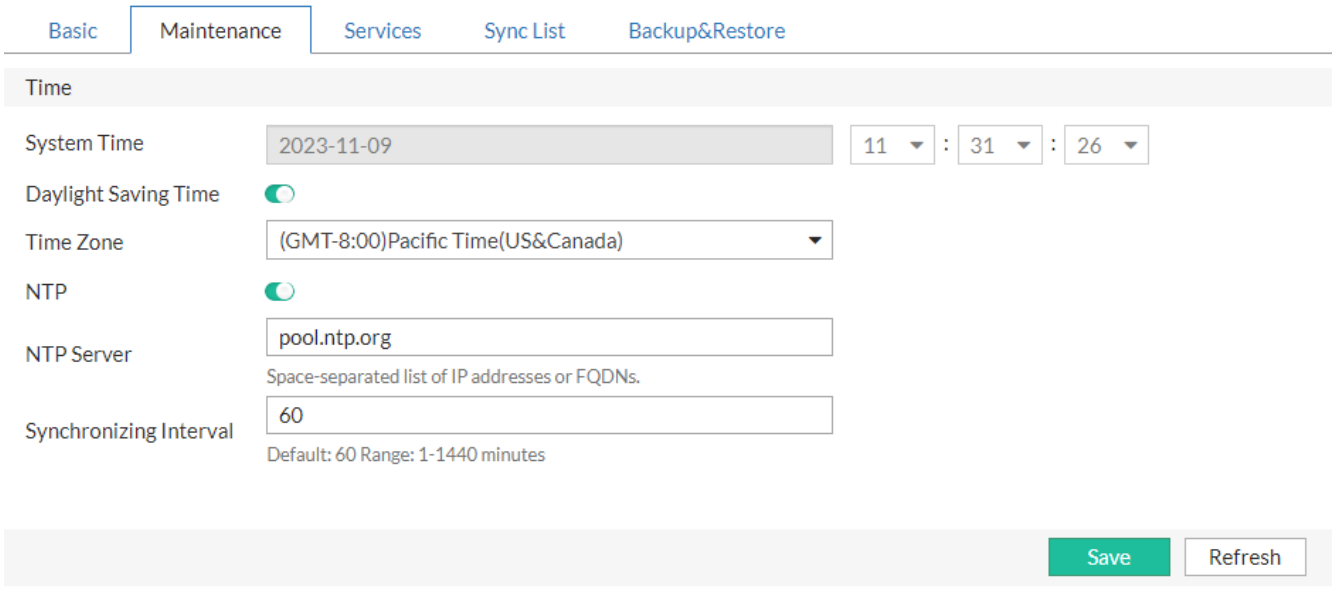
```
FortiADC-VM # config system admin
FortiADC-VM (admin) # edit admin
FortiADC-VM (admin) # set password <string>
Current password for 'admin':
FortiADC-VM (admin) # end
```

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 26.
3. Enter NTP settings and save the configuration.

For detailed information, see Configuring system time.

System time configuration page



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:

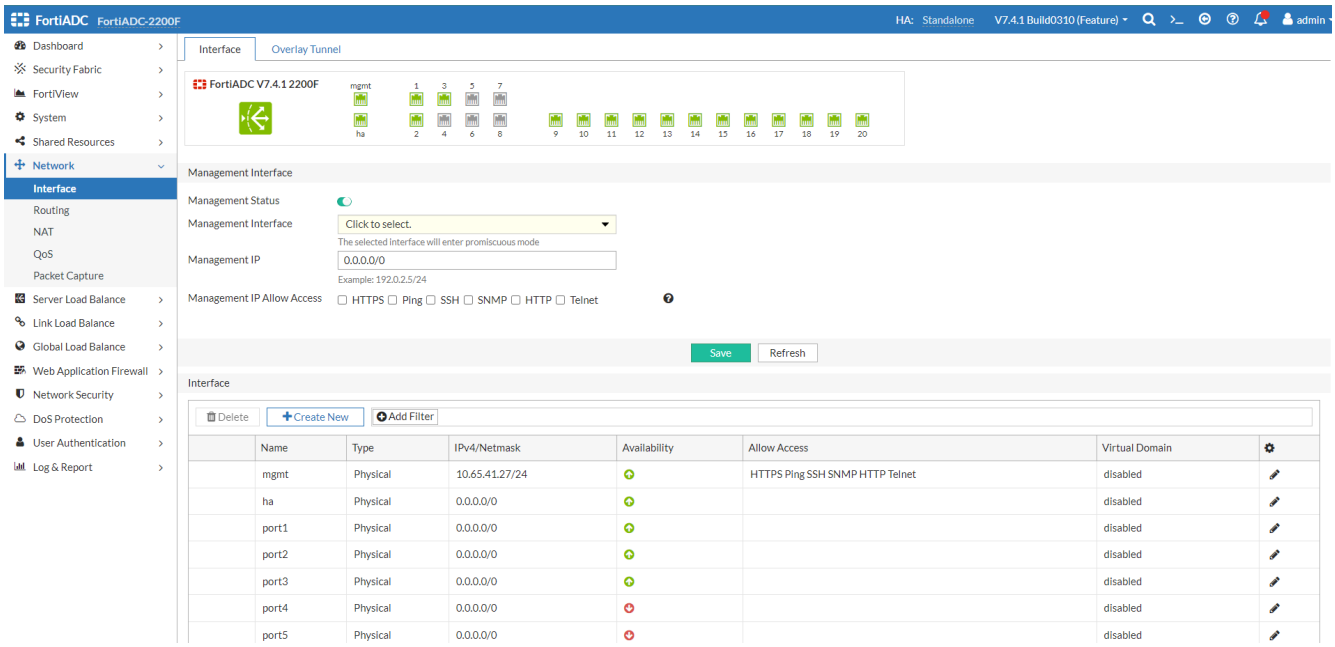
```
config system time manual
set zone <timezone_index>
set daylight-saving-time {enable | disable}
end
execute date <MM/DD/YY> <HH:MM:SS>
```

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 27.
3. Enter the IP address and other interface settings and save the configuration.

For detailed information, see Configuring network interfaces.

Network interface configuration page



CLI commands:

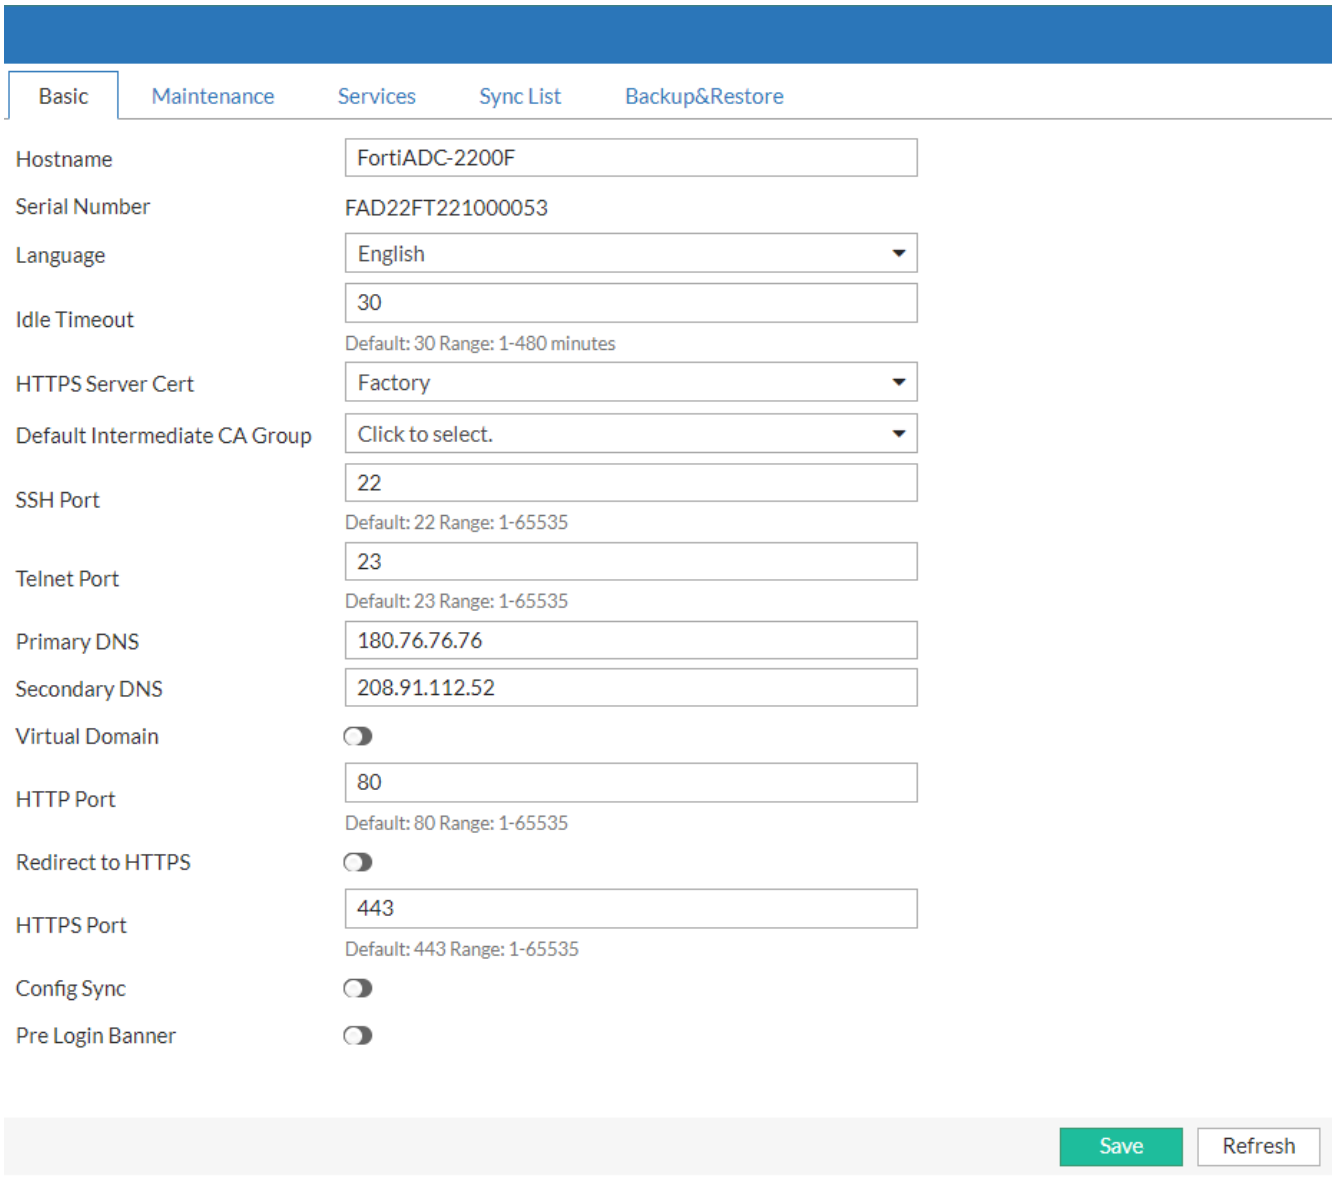
```
config system interface
edit <interface_name>
set ip <ip&netmask>
set allowaccess {http https ping snmp ssh telnet}
end
end
```

To configure DNS:

1. Go to **System > Settings** to display the **Basic** configuration page. See DNS configuration page on page 28.
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



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 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 Step 5: Complete product registration, licensing, and upgrades on page 30.
- 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 > FortiGuard**.
2. Under Status, click **Upgrade License** to upload or locate the license file.

To upgrade your firmware:

1. Go to the **System > FortiGuard**.
2. On the **Maintenance** tab, navigate to the **Firmware** section.
3. 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 > FortiGuard**.
2. Under the **Update Schedule** pane, find the edit function (pencil icon) on the top right.
3. Go into the window enable **Tunneling status**.
4. Complete the configuration as described in the table below.

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

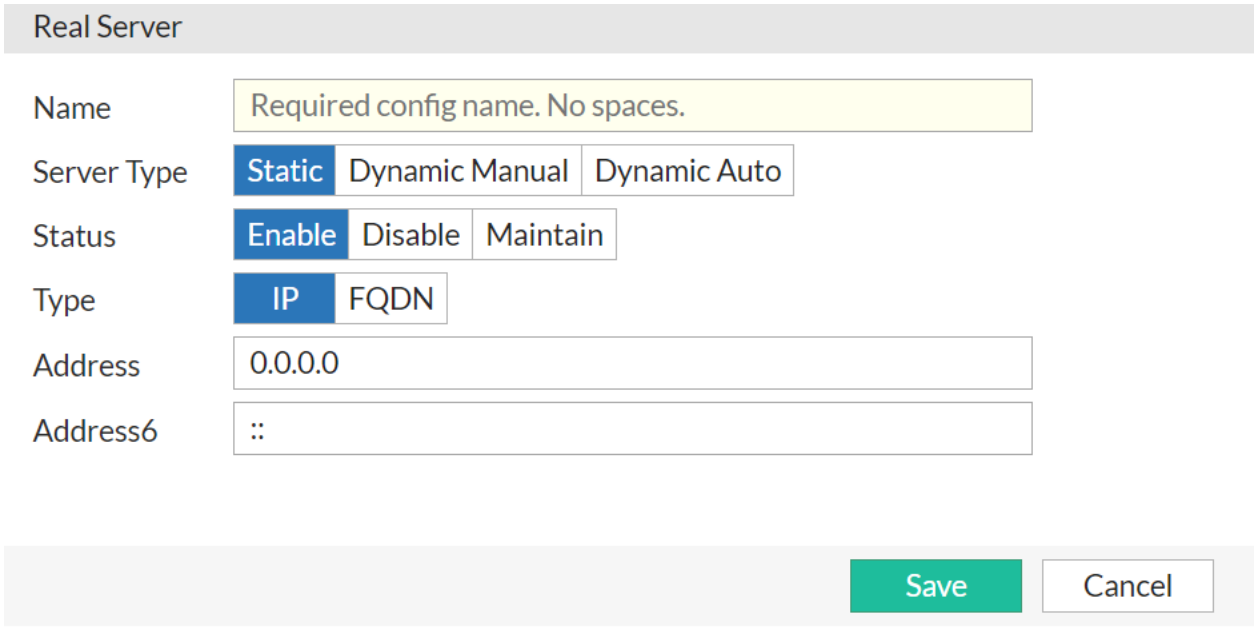
For complete information on server load balancing features, start with Server load balancing basics.

To configure the real server and server pool:

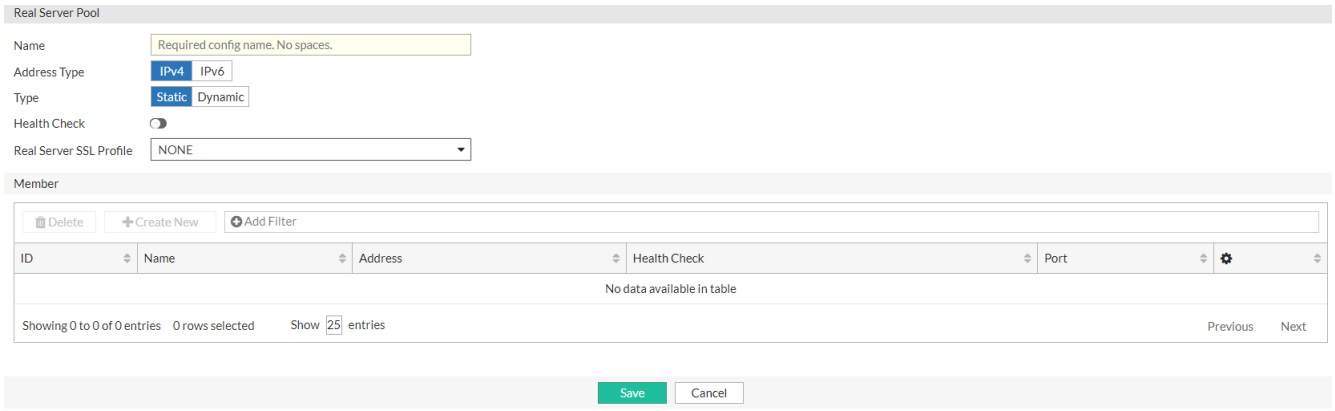
1. Go to **Server Load Balance > Real Server Pool** and click the **Real Server** tab to display the configuration page.
2. Click **Create New** to display the configuration editor. See Real server configuration page on page 32
3. Complete the configuration and click **Save**. Repeat the steps to create more real servers as needed.
4. Click the **Real Server Pool** tab to display the configuration page.
5. Click **Create New** to display the configuration editor. See Step 6: Configure a basic server load balancing policy on page 31Real server configuration page on page 32
6. Complete the configuration and click **Save**.
7. Under Member, click **Create New** to display the Edit Member configuration editor. See Real server pool member configuration page on page 33.
8. Complete the member configuration and click **Save**.

For detailed information, refer to the online help or see Using real server pools on page 552 and Configuring real servers on page 559

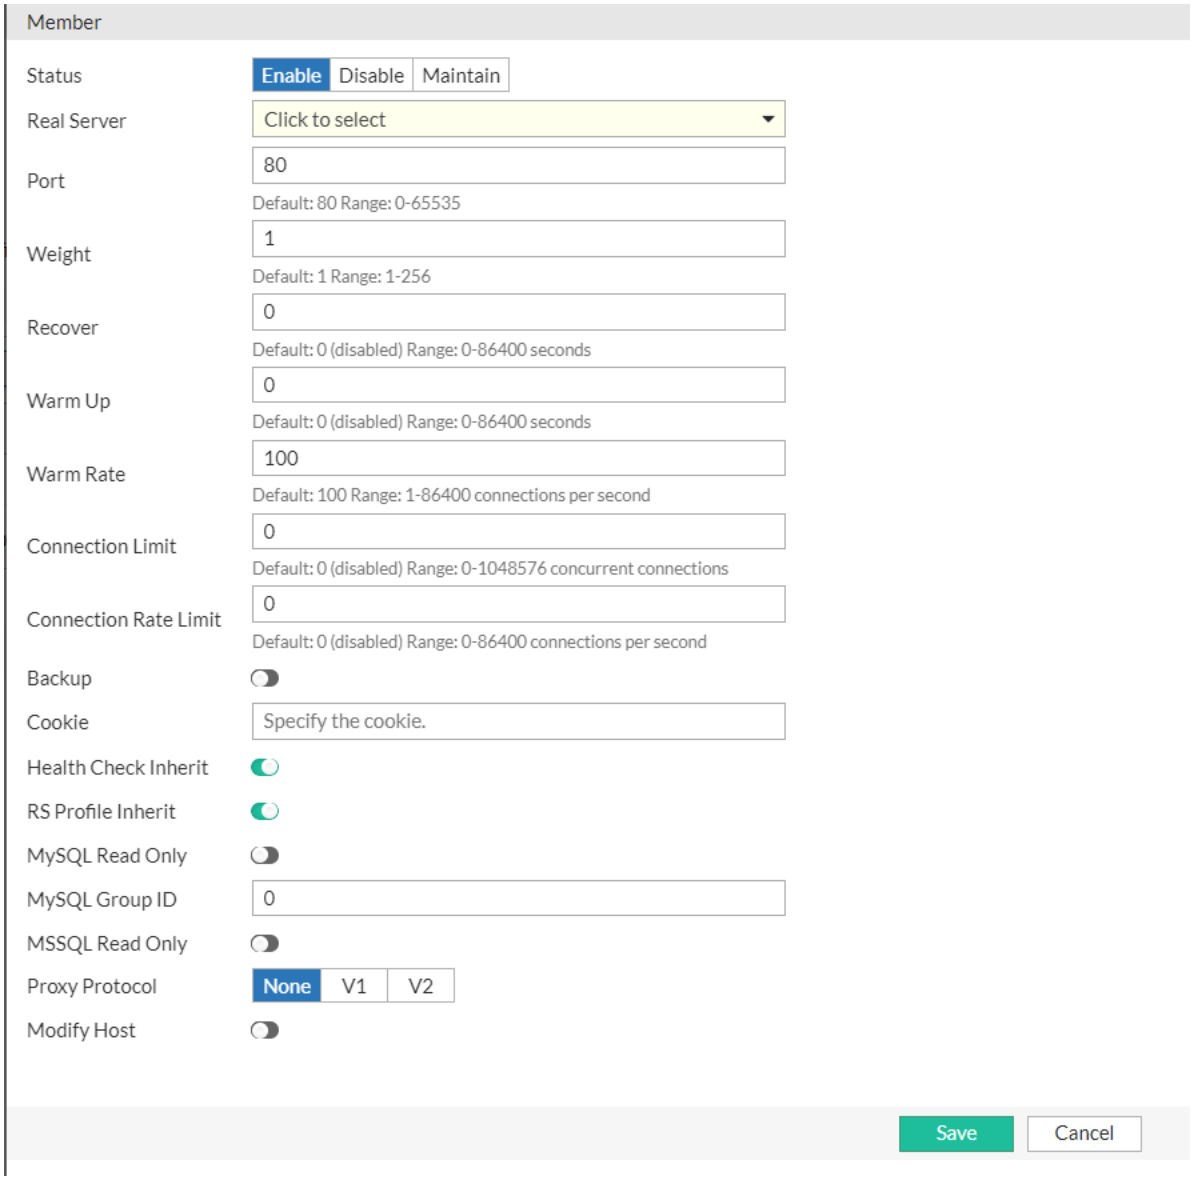
Real server configuration page



Real server pool basic configuration page



Real server pool member configuration page

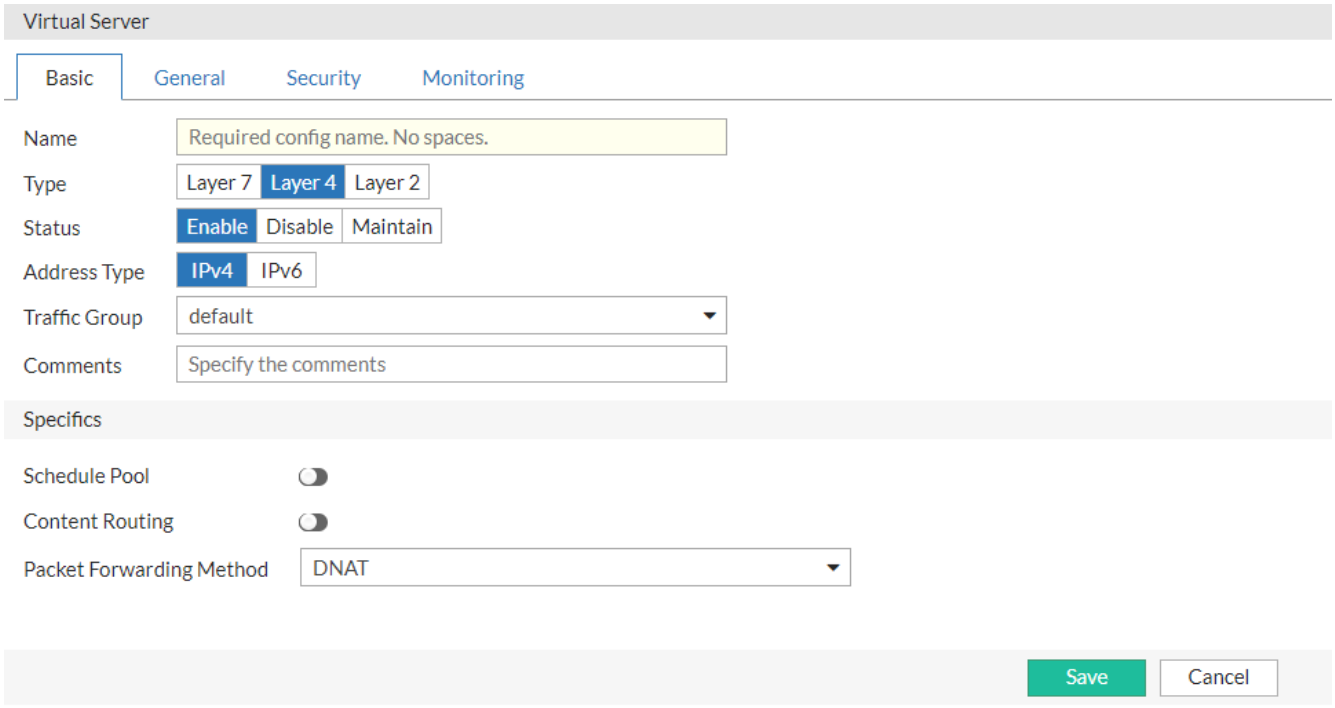


To configure the virtual server:

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 34.
3. Complete the configuration and click **Save**.

For detailed information, refer to the online help or see Configuring virtual servers.

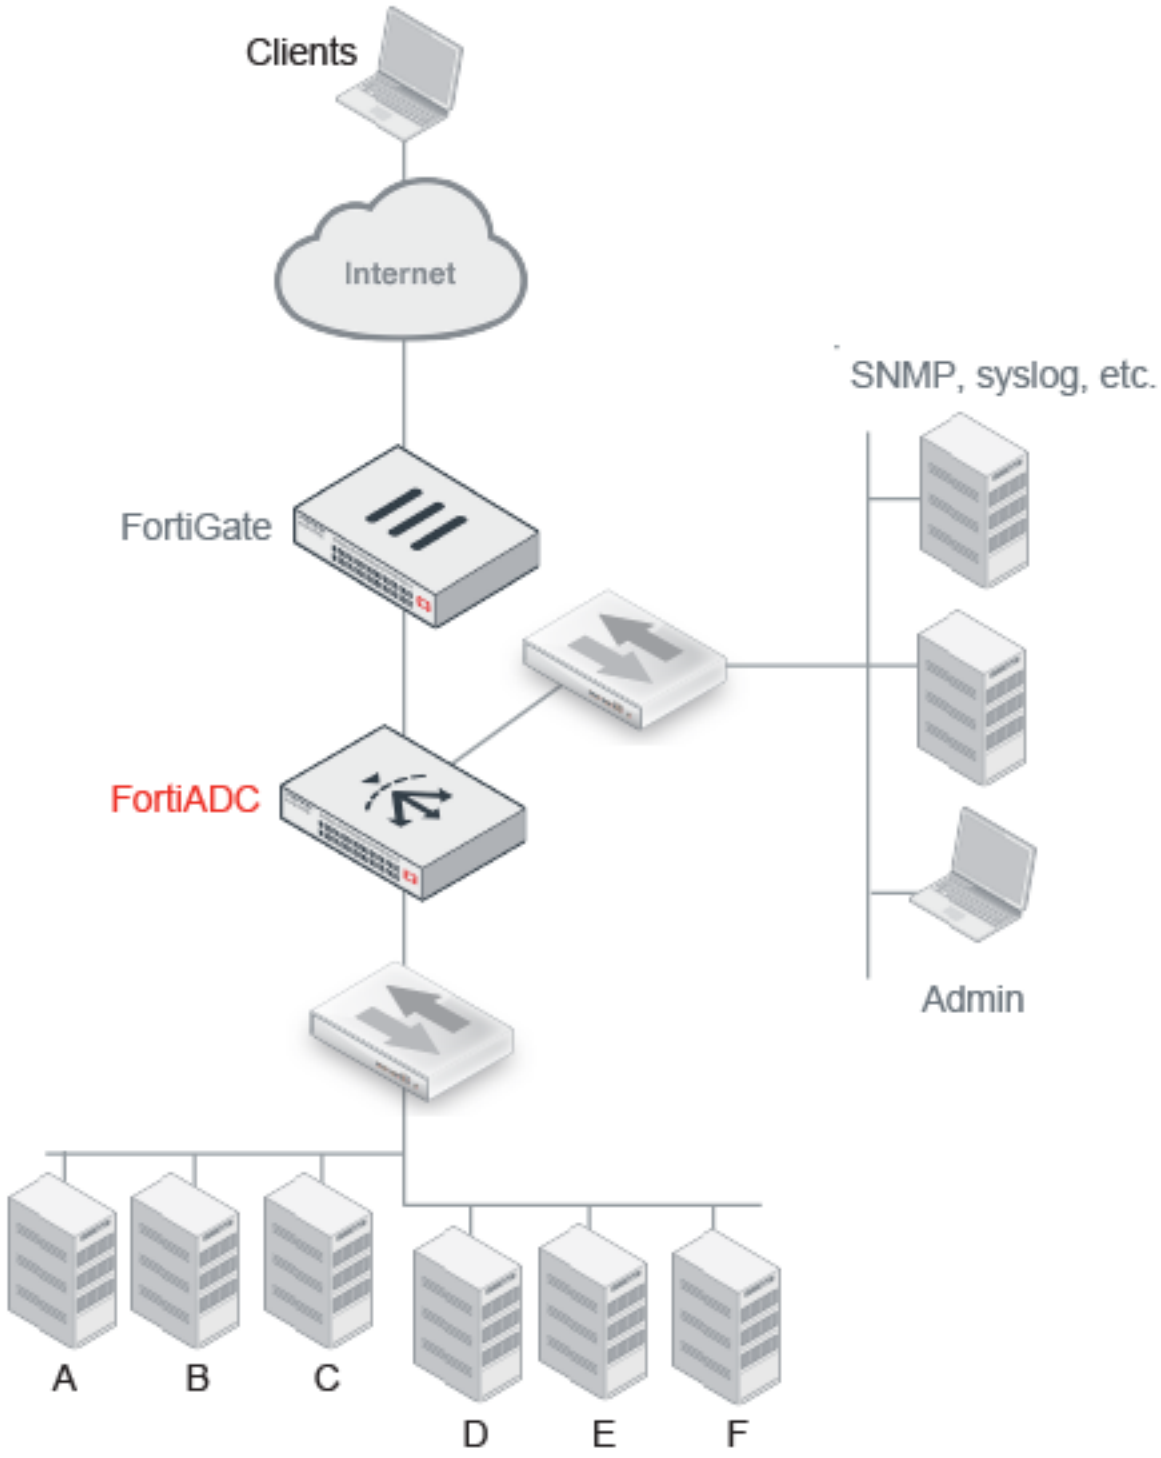
Virtual server configuration page



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 34 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**.

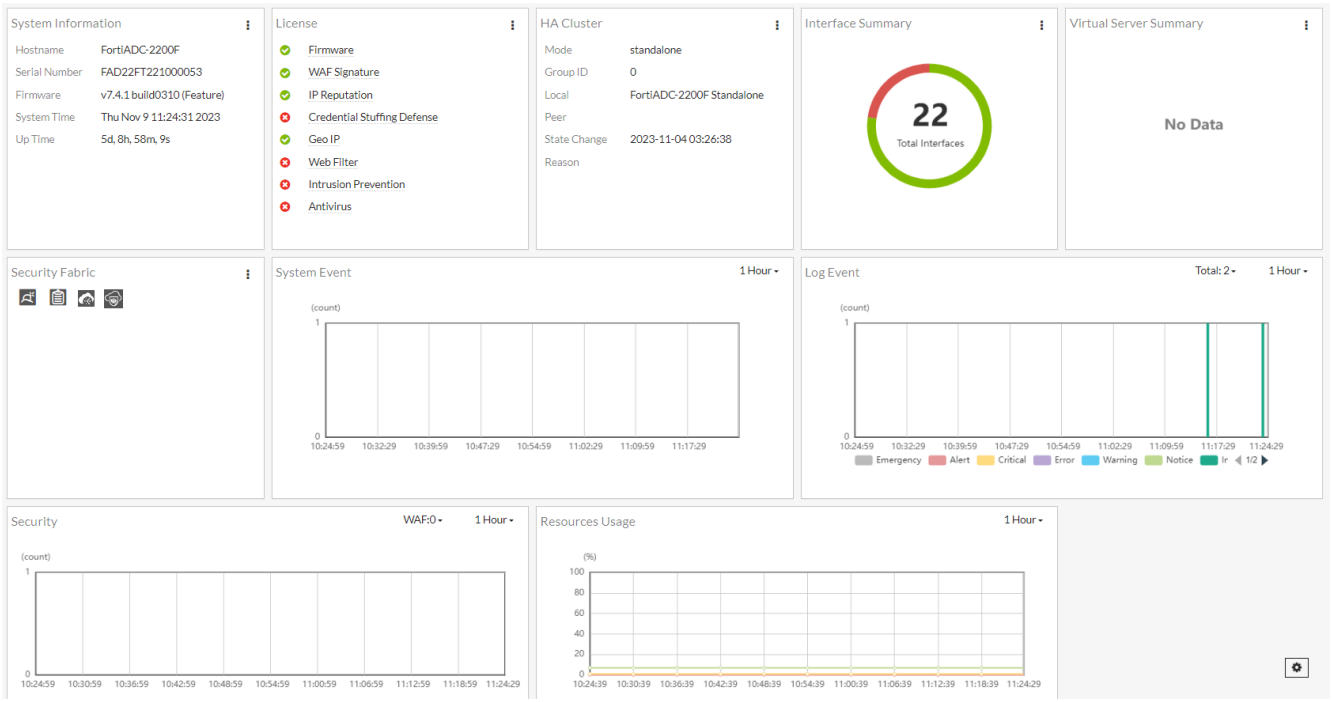
For detailed information, refer to the online help or see Backing up and restoring configuration.

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 section covers the following topics:

- Widgets on page 37
- Dashboard management tools on page 38



Widgets

The default Dashboard page displays a collection of 13 widgets, which fall into three categories: Information, Configuration, and Statistics.

The following table highlights the information contained in each of the widgets.

Widget	Category	Description and Utilities
System Info	Information	Lists the following information about the system: • Hostname • Serial Number • Firmware version, build number, and version

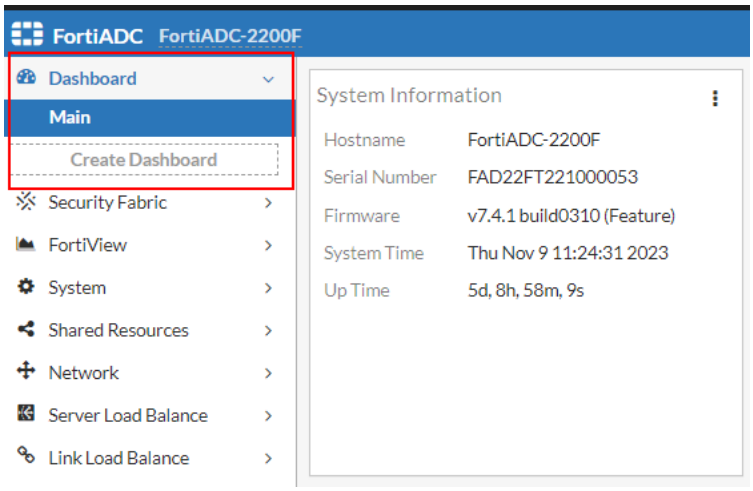
Widget	Category	Description and Utilities
		maturity (Feature or Mature). Feature — New features are included in this version. Mature — This version primarily includes bug fixes and/or vulnerability patches, and no new or major features. • System Time • Up Time
License Info	Information	Lists the FortiGuard Service Licenses and their status.
HA Info	Information	Lists the information about the HA Cluster.
Interface Summary	Configuration	Shows type, name, mode, IP address, status and allowed forms of access for each port, as well as its virtual domain.
Virtual Server Summary	Configuration	Shows the status of your virtual servers, including health, throughput, and pool number.
Security Fabric	Configuration	Shows the Security Fabric connectors and their connection status.
System Event	Statistics	Shows the number of system events that has occurred within a specified interval.
Log Event	Statistics	Tells you how many emergencies, alerts, and warnings, etc., that are going on.
Resources Usage	Statistics	Shows the resource usage in percentage for CPU, RAM, and Disk.
Server Load Balance	Statistics	Shows your throughput, concurrent, and connections, while displaying the inbound/outbound throughput as a graph.
Security Logs	Statistics	Shows the number of security events that has triggered the WAF within a specified interval.
OWASP Top 10 Threats	Statistics	Shows the OWASP Top 10 Threats that has been detected within a specified interval.
Global Load Balance	Statistics	Shows the host and response per second.

Click **See Detail** to see more, it should show up when you **hover** over the widget.

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:

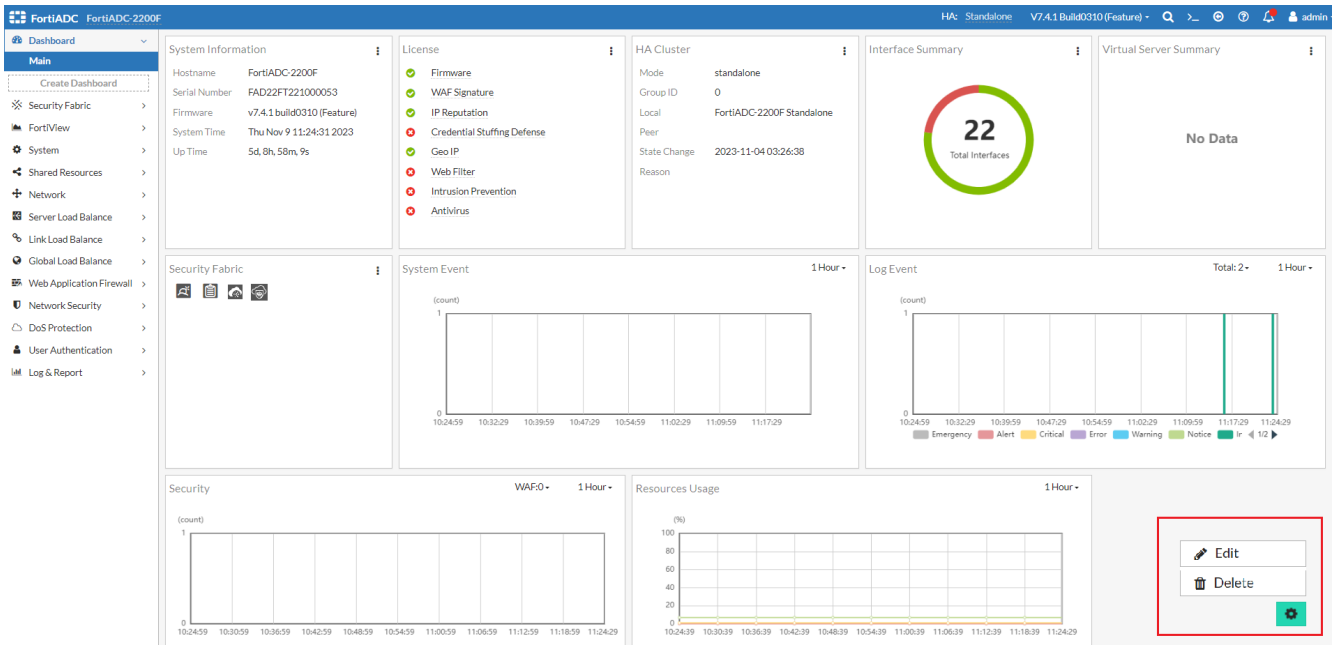
1. Click **Create Dashboard.**
2. Specify a unique name for the dashboard.
3. Click **save**. The name of the dashboard appears under Dashboard>Main on the navigation bar.

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. Hover over the dashboard until you see the gear icon appear in the bottom right corner. Click to view the dashboard modification options.
3. Click **Edit** to display the dashboard editor.
4. 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. Hover over the dashboard until you see the gear icon appear in the bottom right corner. Click to view the dashboard modification options.
3. Click **Delete** as shown in the illustration above.
4. Read the warning message onscreen.
5. Click Delete if you've decided to remove the selected dashboard.

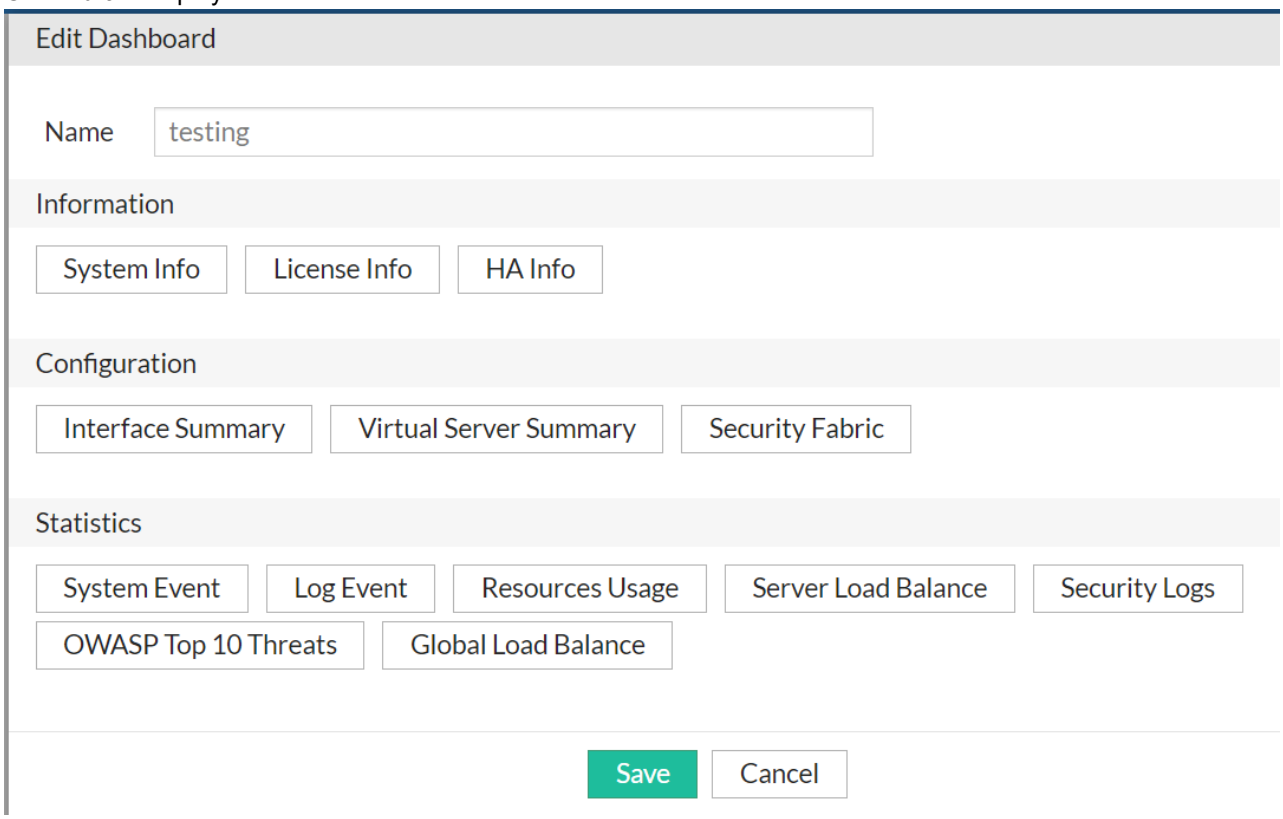
Adding Features

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 navigation bar, select the name of the dashboard.
2. Hover over the dashboard until you see the gear icon appear in the bottom right corner. Click to view the dashboard modification options.

3. Click **Edit** to display the dashboard editor.



4. Decide which widget to display on the dashboard by selecting or deselecting the widget name.
5. Click **Save** when done.

Security Fabric

The Fortinet Security Fabric provides an intelligent architecture that interconnects discrete security solutions into an integrated whole to detect, monitor, block, and remediate attacks across the entire attack surface. All elements in the Security Fabric work together as a team to share policy, threat intelligence, and application flow information that expand network visibility to provide fast threat detection in real time and the ability to initiate and synchronize a coordinated response, no matter which part of the network is being compromised.

This section describes the following topics:

- Automation on page 42 — Automation Stitches pair an event trigger with one or more actions to monitor the network and take the designated actions automatically when the Security Fabric detects a threat.
- Fabric connectors on page 77 — Fabric connectors provide integration with Fortinet products to automate the process of managing dynamic security updates without manual intervention.
- External connectors on page 94 — External connectors provide integration with public and private cloud solutions to ensure changes to cloud environment attributes are automatically updated in the Security Fabric.

Automation

Automation Stitches automate the activities between the different components in the Security Fabric, which decreases the response times to security events. Events from any source in the Security Fabric can be monitored, and action responses can be set up to any destination.

An automation stitch consists of two parts: the trigger and the actions. The trigger is the condition or event on the FortiADC that activates the action, for example, a specific log or a failed login attempt. The action is what the FortiADC does in response to the trigger. This allows you to monitor your network and take appropriate action when a threat is detected by automating certain actions in response to certain triggers. For example, you can automate alert emails to be sent in response to specific security events, which allows for far more granular log-based alerting than Alert Emails configured under the Log & Report settings.

This section includes:

- Creating automation stitches
- Configuring Automation Triggers on page 54
- Configuring Automation Actions on page 67
- Diagnose commands

Creating automation stitches

Automation stitches pair a trigger with one or more response actions to allow FortiADC to automatically respond with the action(s) once the trigger condition is met.

From the GUI, **Security Fabric > Automation** page, you can create an automation stitch by selecting a **Trigger** event type and the corresponding **Action** that you would like to automate from the same configuration editor.

FortiADC supports eight trigger event types and six response actions for automation.

- Triggers: Security Events, SLB Metrics, Period Block IP, HA Failover, System Metrics, System Events, Interface Metrics, Schedule and FortiADC Log.
- Actions: CLI Script, Email, Syslog, SNMP Trap, Webhook, and FortiGate IP Ban.

However, some response actions are only supported for certain trigger types. The table below lists each trigger type and their available response actions.

	Security Events	SLB Metrics	Period Block IP	HA Failover	System Metrics	System Events	Interface Metrics	Schedule	FortiADC Log
CLI Script	✓	✓	✗	✓	✓	✓	✓	✓	✓
Email	✓	✓	✓	✓	✓	✓	✓	✓	✓
Syslog	✓	✓	✓	✓	✓	✓	✓	✓	✓
SNMP Trap	✓	✓	✓	✓	✓	✓	✓	✗	✓
Webhook	✓	✓	✓	✓	✓	✓	✓	✓	✓
FortiGate IP Ban	✗	✗	✓	✗	✗	✗	✗	✗	✗

FortiADC offers Predefined Automation Stitch configurations on page 50 you can use to get started.

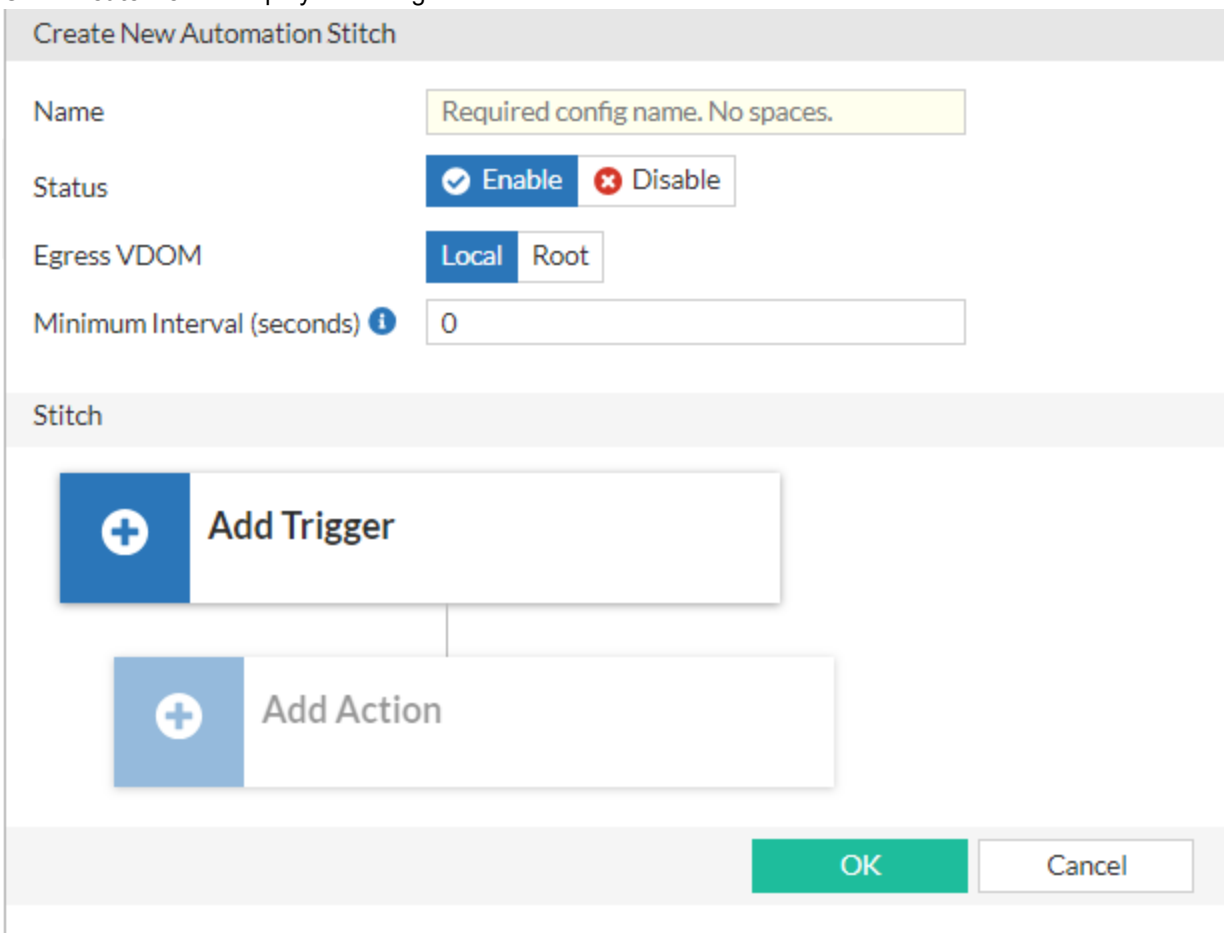
Before you begin:

- You must have Global Administrator access. Ensure that your admin account settings has **Global Admin** set to **Yes**. For more information, see Creating administrator users on page 249.

To configure an automation stitch:

1. Go to **Security Fabric > Automation**.
   The configuration page displays the **Stitch** tab.

2. Click **Create New** to display the configuration editor.

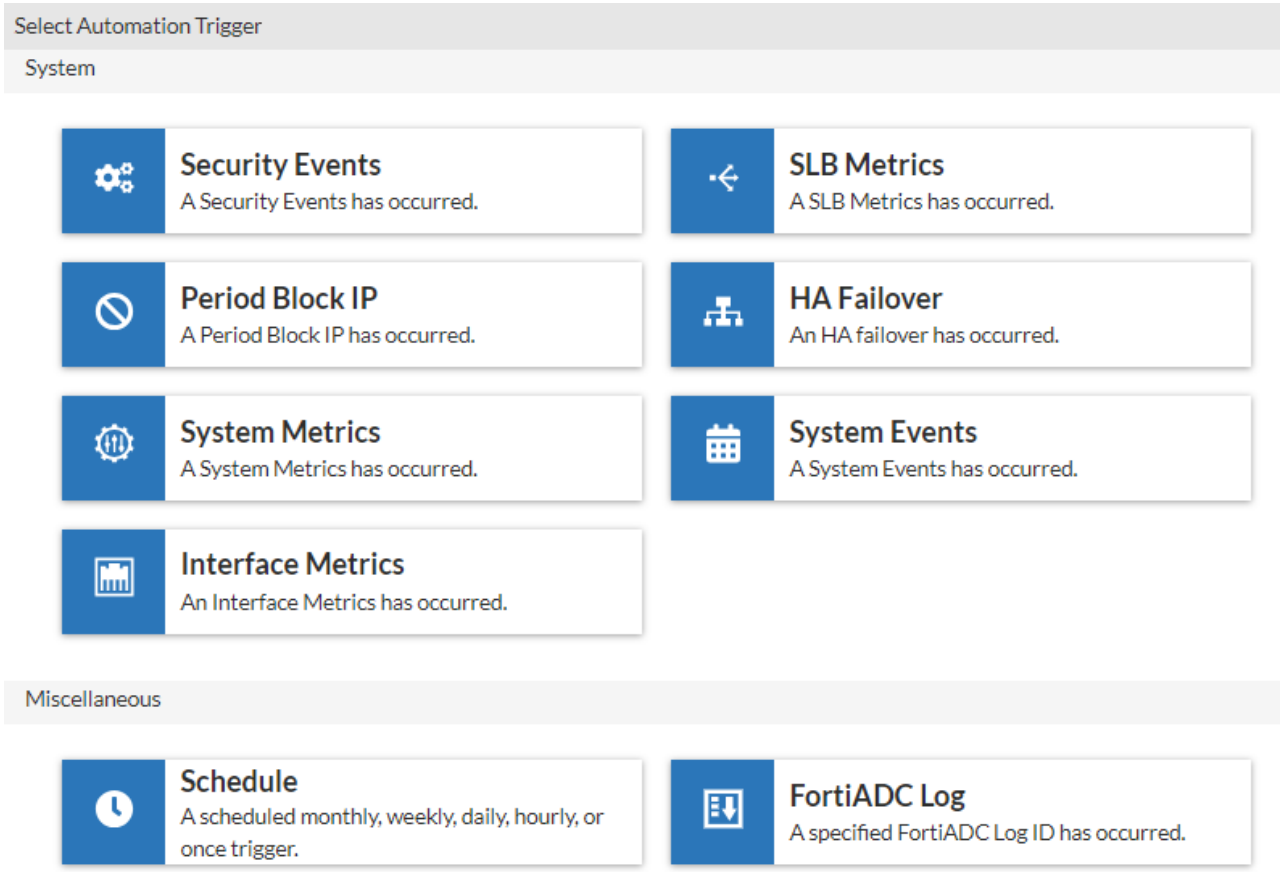


3. Configure the following settings for the Automation Stitch:

Setting	Description
Name	Enter a name for the new automation stitch. The configuration name cannot be edited once it has been saved.
Status	Enable/disable the automation stitch.
Egress VDOM	The Egress VDOM determines the VDOM from which the alert packets will be sent, regardless of the local VDOM from which the automation is configured. This affects automation actions that require alert packets to be sent, which include Syslog, SNMP Trap, Webhook, and Email. Actions such as Syslog, SNMP Trap, and Webhook can egress from either the local or root VDOM. However, for Email actions, the Egress VDOM **must** be Root to correspond with the SMTP server configured in Global Settings. Select the Egress VDOM from which the alert packets will be sent: • Local — Alert packets will be sent from the local VDOM from which the automation is configured. • Root — Alert packets will be sent from the Root VDOM.

4. Under the **Stitch** section, click **Add Trigger** to select an Automation Trigger event and configure the settings specific to each trigger event type. Some trigger events are predefined while some trigger events are user-defined.

For example, the System Event trigger provides a list of predefined system events for selection, whereas the SLB Metrics trigger requires users to define the alert metrics. For details about each trigger event type, see Configuring Automation Triggers on page 54.



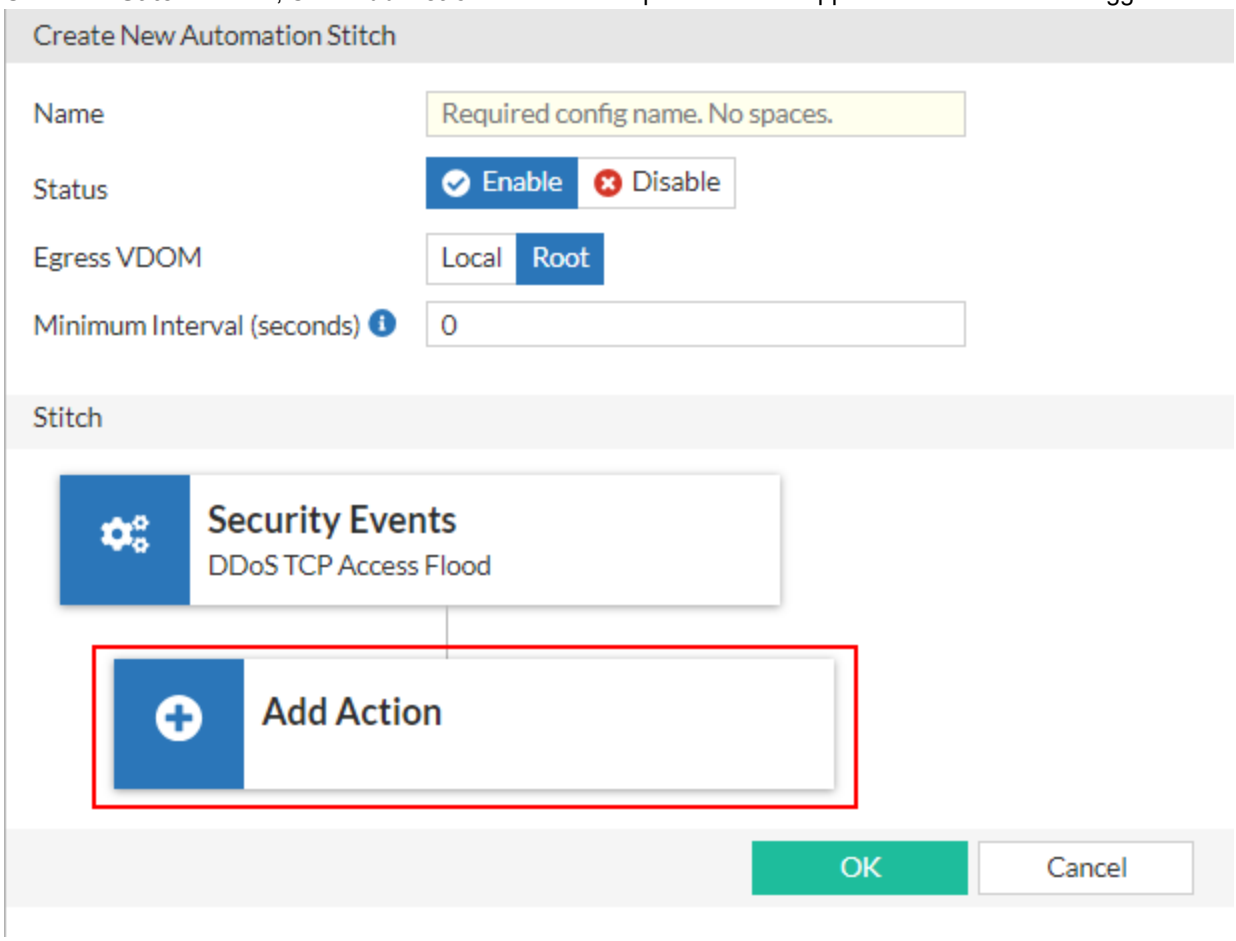
Trigger	Description
Security Events	
Apply to	Select whether to apply the security events automation stitch to All or VS: • All — All related events will trigger the Alert action. • VS — Only specified Virtual Server related events will trigger the Alert action.
Virtual Server	The **Virtual Server** option appears if **Apply to** is **VS**. Specify the virtual server. This is required.
Event	Select the security events (such as DDoS SYNFLOOD attack start, bot detected, etc.) that will trigger the action. The list of available security events is predefined. For details, see Configuring Automation Triggers on page 54.
Advanced Settings	Click **Advanced Settings** to display additional settings for **Rolling Window**.
Rolling Window	Enable to define a **Rolling Window Time** and **Number of Occurence**. The **Rolling Window Time** sets a period of time in which a number of events must take place for an action to be triggered. The number of events that must take place within this period of time is set in the **Number of Occurrences** option.

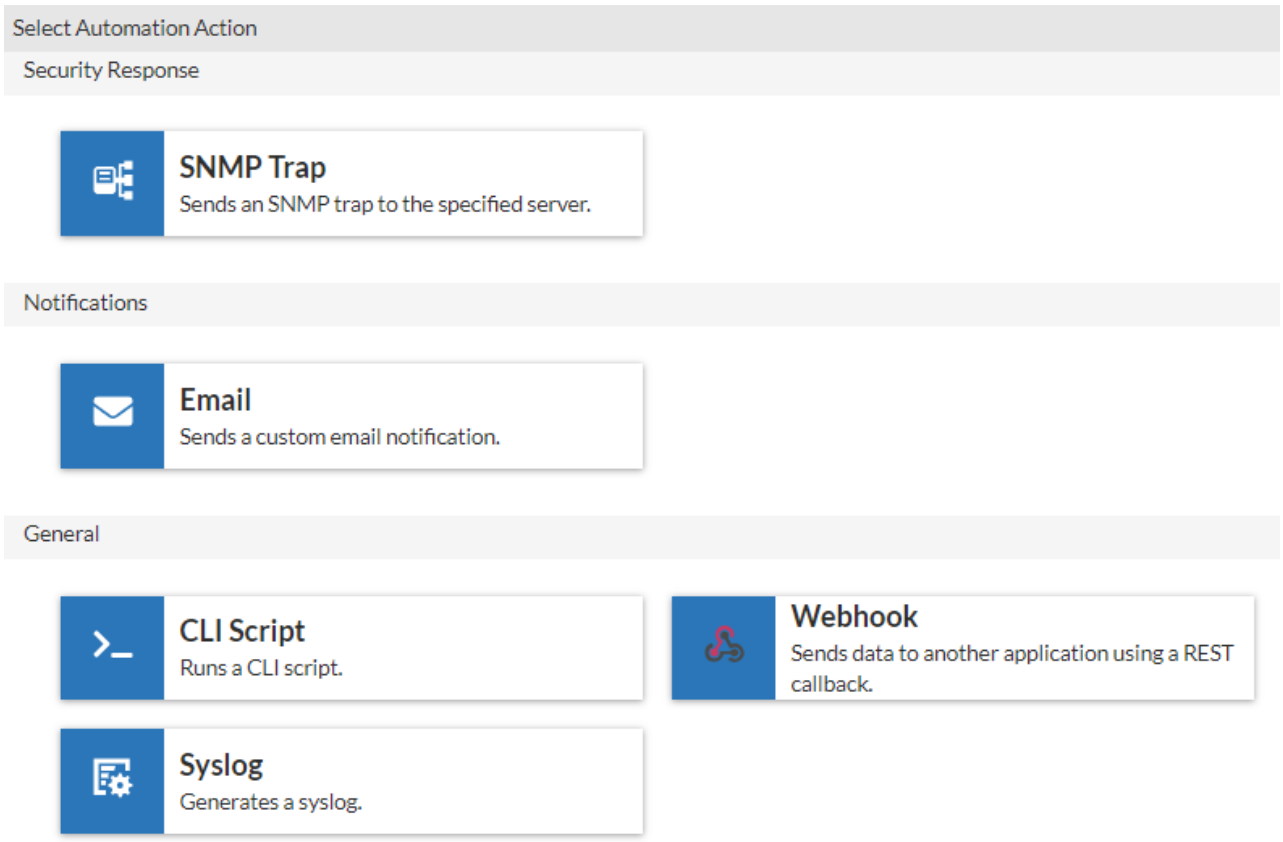
Trigger	Description
Rolling Window Time	The **Rolling Window Time** option appears if **Rolling Window** is enabled. Specify the range of time (in seconds) for the rolling window.
Number of Occurrences	The **Number of Occurrences** option appears if **Rolling Window** is enabled. Specify the number of events that must take place before FortiADC will trigger the action.
SLB Metric	
Alert	Select a user-defined Alert trigger or create a new alert trigger for SLB Metrics. For details, see Configuring Automation Triggers on page 54.
Period Block IP	
Period Block IP	Select this trigger to retrieve the Source IP addresses from the Period Block list.
HA Failover	
Event	Select the HA failover events (such as HA peer lost) that will trigger the action. The list of available HA failover events is predefined. For details, see Configuring Automation Triggers on page 54.
Advanced Settings	Click **Advanced Settings** to display additional settings for **Rolling Window**.
Rolling Window	Enable to define a **Rolling Window Time** and **Number of Occurence**. The **Rolling Window Time** sets a period of time in which a number of events must take place for an action to be triggered. The number of events that must take place within this period of time is set in the **Number of Occurrences** option.
Rolling Window Time	The **Rolling Window Time** option appears if **Rolling Window** is enabled. Specify the range of time (in seconds) for the rolling window.
Number of Occurrences	The **Number of Occurrences** option appears if **Rolling Window** is enabled. Specify the number of events that must take place before FortiADC will trigger the action.
System Metrics	
Alert	Select a user-defined Alert trigger or create a new alert trigger for System Metrics. For details, see Configuring Automation Triggers on page 54.
System Events	
Apply to	Select whether to apply the system events automation stitch to All, VS or Real Server: • All — All related events will trigger the Alert action. • VS — Only specified Virtual Server related events will trigger the Alert action. • Real Server — Only the specified Virtual Server, Pool, and Real Server related events will trigger the Alert action.
Virtual Server	The **Virtual Server** option appears if **Apply to** is **VS** or **Real Server**. Specify the virtual server. This is required if **Apply** to is **VS**.
Pool	The **Real Server** option appears if **Apply to** is **Real Server**.

Trigger	Description
	Specify the pool. This is optional.
Real Server	The **Real Server** option appears if **Apply to** is **Real Server**. Specify the real server. This is required.
Event	Select the system events (such as bad PSU fan, good device fan, etc.) that will trigger the action. The list of available System events is predefined. For details, see Configuring Automation Triggers on page 54.
Advanced Settings	Click **Advanced Settings** to display additional settings for **Rolling Window**.
Rolling Window	Enable to define a **Rolling Window Time** and **Number of Occurence**. The **Rolling Window Time** sets a period of time in which a number of events must take place for an action to be triggered. The number of events that must take place within this period of time is set in the **Number of Occurrences** option.
Rolling Window Time	The **Rolling Window Time** option appears if **Rolling Window** is enabled. Specify the range of time (in seconds) for the rolling window.
Number of Occurrences	The **Number of Occurrences** option appears if **Rolling Window** is enabled. Specify the number of events that must take place before FortiADC will trigger the action.
Interface Metric	
Alert	Select a user-defined Alert trigger or create a new alert trigger for Interface Metrics. For details, see Configuring Automation Triggers on page 54.
Schedule	
Schedule	Select a user-defined Alert trigger or create a new alert trigger for Schedule. For details, see Configuring Automation Triggers on page 54.
FortiADC Log	
Event	Select a user-defined Event trigger or create a new alert trigger for FortiADC Log. For details, see Configuring Automation Triggers on page 54.

5. Click **OK** to commit the Automation Trigger selection and exit from the dialog.
Once an Automation Trigger has been selected, you can now add an automation response action.

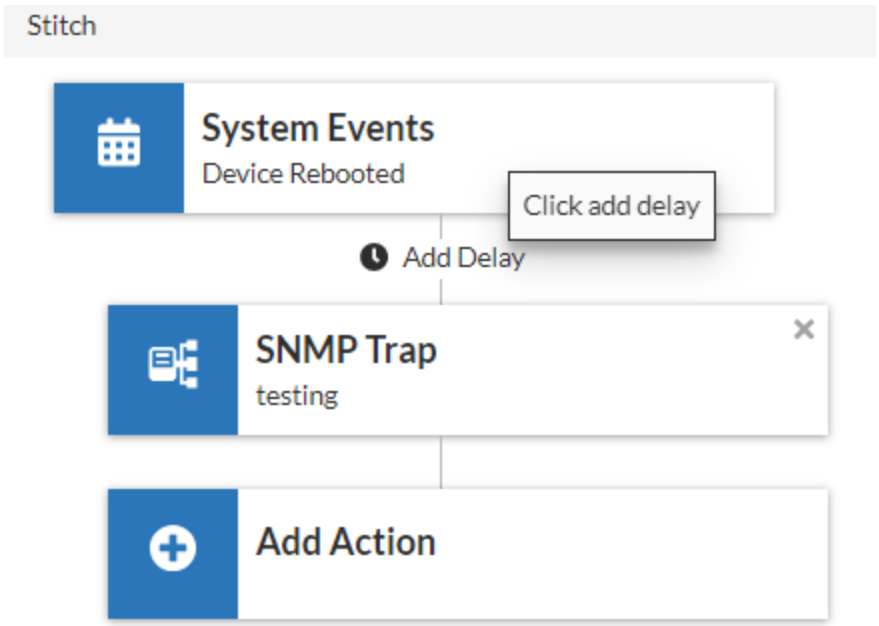
6. Under the **Stitch** section, Click **Add Action** to select a response action supported for the selected trigger event.





Configure the settings specific to each response action. Each **Action** is user-defined. For details about each response action, see Configuring Automation Actions on page 67.

7. Click **OK** to commit the Automation Action selection and exit from the dialog.
   Once the first Automation Action has been added, you have the option to stitch more actions to the trigger.
8. Optionally, you can add a delay to trigger the response action, or delay between multiple response actions. After adding an Automation Action to your Trigger or a second Action, the **Add Delay** option appears.



Click **Add Delay** to specify the delay (in seconds) to trigger the response action or between multiple response actions. The default is 0 second delay, and any configuration changes to the Trigger or Action in the stitch will set the delay back to default.

9. Click **OK** to save the Automation Stitch.
   The newly created automation stitch appears on the **Security Fabric > Stitch** page, under its trigger event type.

After configuring the automation stitch, you may test it through CLI command `diagnose debug module alertd`.

Predefined Automation Stitch configurations

The following Automation Stitch configurations have predefined trigger events but no response actions selected. You may clone these predefined configurations and use them as a template.

Name	Type	Trigger events
HA_Template	HA Failover	HA Peer Lost HA Master Failover
Admin_Template	System Events	User Login User Logout
Configuration_Template	System Events	Config Create Config Delete Config Update
System_basic_Template	System Events	Lost Log Disk High CPU Usage High Disk Usage High Memory Usage SSD MWI Near Threshold SSD MWI Reached Threshold

Name	Type	Trigger events
Health_check_Template	System Events	Real Server HC Down Real Server HC Up Virtual Server Down Virtual Server Up Gateway HC Down Link Group HC Down Gateway HC Up Link Group HC Up GLB Real Server Not Available GLB Real Server Available GLB Virtual Server Not Available GLB Virtual Server Available GLB GW Not Available GLB GW Available
Certificate_Template	System Events	Certificate Expire
SNMP_sys_event_Template	System Events	High CPU Temp Normal CPU Temp High Device Temp Normal Device Temp High PSU Temp Normal PSU Temp Slow PSU Fan Slow Device Fan Bad PSU Fan Good PSU Fan Bad Device Fan Good Device Fan High Voltage Low Voltage High Power Supply Low Power Supply High PSU Voltage Low PSU Voltage PSU Failure Lost Log Disk High CPU Usage High Disk Usage High Memory Usage SSD MWI Near Threshold

Name	Type	Trigger events
		SSD MWI Reached Threshold Device Rebooted Device Upgrade Completed User Login User Logout ARP Conflict Logical Interface Up Logical Interface Down Logical Interface Disabled Log Full FW SNAT Port Exhausted Real Server HC Down Real Server HC Up Real Server Enabled Real Server Disabled Real Server Maintain Mode Real Server Connection Rate Start Real Server Connection Rate Stop Real Server Connection Limit Start Real Server Connection Limit Stop Virtual Server Down Virtual Server Up Virtual Server Enabled Virtual Server Disabled Virtual Server Maintain Mode Virtual Server Connection Rate Start Virtual Server Connection Rate Stop Virtual Server Connection Limit Start Virtual Server Connection Limit Stop Virtual Server Transaction Rate Start Virtual Server Transaction Rate Stop Virtual Server IP Pool Limit Certification Expire Gateway HC Down Link Group HC Down Gateway HC Up Link Group HC Up Gateway Inbound Bandwidth Gateway Outbound Bandwidth Gateway Inbound Spillover

Name	Type	Trigger events
		Gateway Outbound Spillover Gateway Total Spillover GLB Real Server Not Available GLB Real Server Available GLB Virtual Server Not Available GLB Virtual Server Available GLB GW Not Available GLB GW Available Config Create Config Delete Config Update OCSP Response Expires SSL Certificate Revoked CRL Expires
SNMP_sec_event_Template	Security Events	DDoS SYNFLOOD attack start DDoS SYNFLOOD attack stop Request Blocked XSS Attack Detected SQL Injection Attack Detected Generic Attack Detected URL Pattern Violate Detected Protocol Constraint Detected Bot Detected Geo Violate Detected Reputation Violate Detected Virtual Server Authentication Failed JSON Violate Detected XML Violate Detected SOAP Violate Detected Web Anti Defacement Detected CSRF Violate Detected Brute Force Detected Data Leak Violate Detected HTML Validation Detected DDoS IP Fragmentation DDoS TCP Slow Data Attack DDoS TCP Access Flood DDoS HTTP Connection Flood DDoS HTTP Request Flood

Name	Type	Trigger events
		DDoS HTTP Access Limit OPENAPI Violate Detected CORS Violate Detected SEC Threshold Violate Detected SEC Biometrics Base Detected
SNMP_HA_event_Template	HA Failover	HA Peer Lost HA Master Failover

Configuring Automation Triggers

On the **Security Fabric > Automation > Trigger** tab, you can view the list of available automation trigger events that are predefined or user-defined. After defining your automation triggers, you can combine them with response actions to create an automation stitch. For details, see Creating automation stitches on page 42

FortiADC supports eight trigger event types, wherein some events are predefined and some must be user-defined.

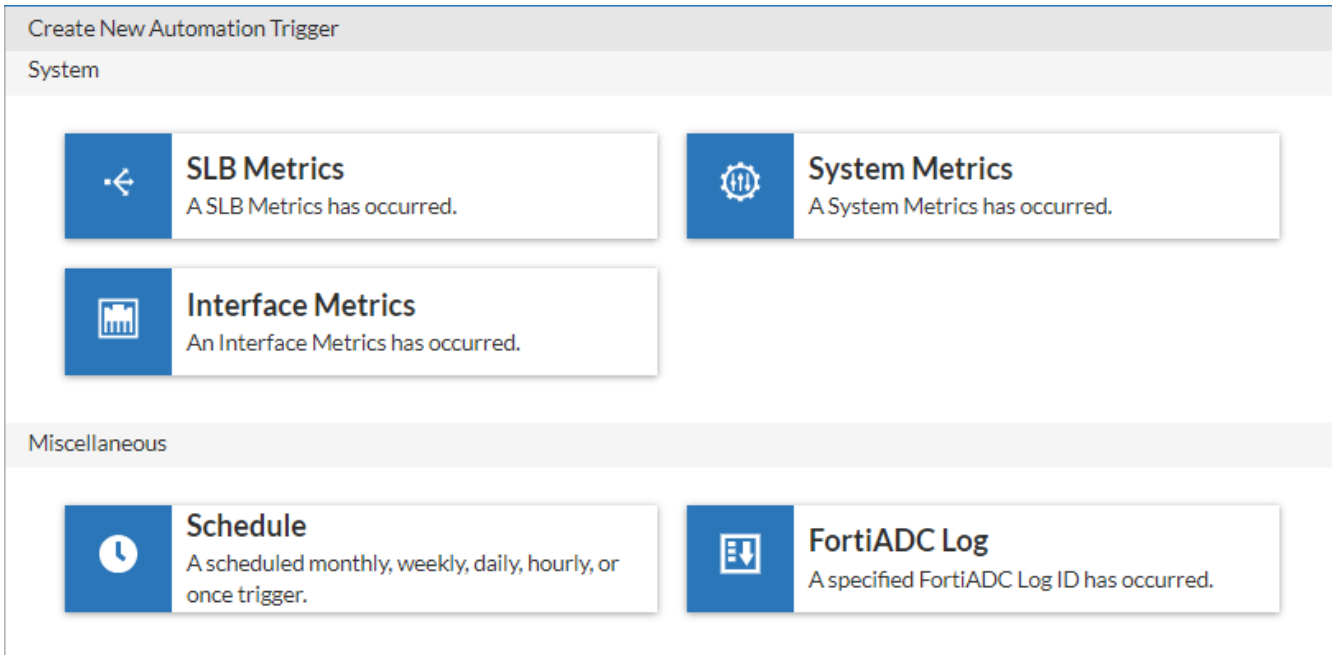
Predefined Triggers:

- Security Events — Uses security events such as "DDoS SYNFLOOD attack start" or "bot detected" as the alert trigger.
- HA Failover — Uses HA failover events such as "HA peer lost" as the alert trigger.
- System Events — Uses system events such as "bad PSU fan" or "good device fan" as the alert trigger.

See Predefined automation trigger events on page 64 for the full list of predefined events available for each trigger type.

User-defined Triggers:

- SLB Metrics on page 55 — Uses server load balance performance metrics as the alert trigger.
- Period Block IP — Uses the FortiADC Source IP addresses that have been blocked by WAF as trigger events for the automated response actions. To view or release the blocked IPs, see Blocked IP on page 138.
- System Metrics on page 57 — Uses system metrics such as "average CPU usage" or "average memory usage" as the alert trigger.
- Interface Metrics on page 59 — Uses network interface events as the alert trigger.
- Schedule on page 61 — Uses user-defined schedules as the alert trigger.
- FortiADC Log on page 62 — Uses a specific FortiADC logged event occurrence (identified through log ID) as the alert trigger.



Before you begin:

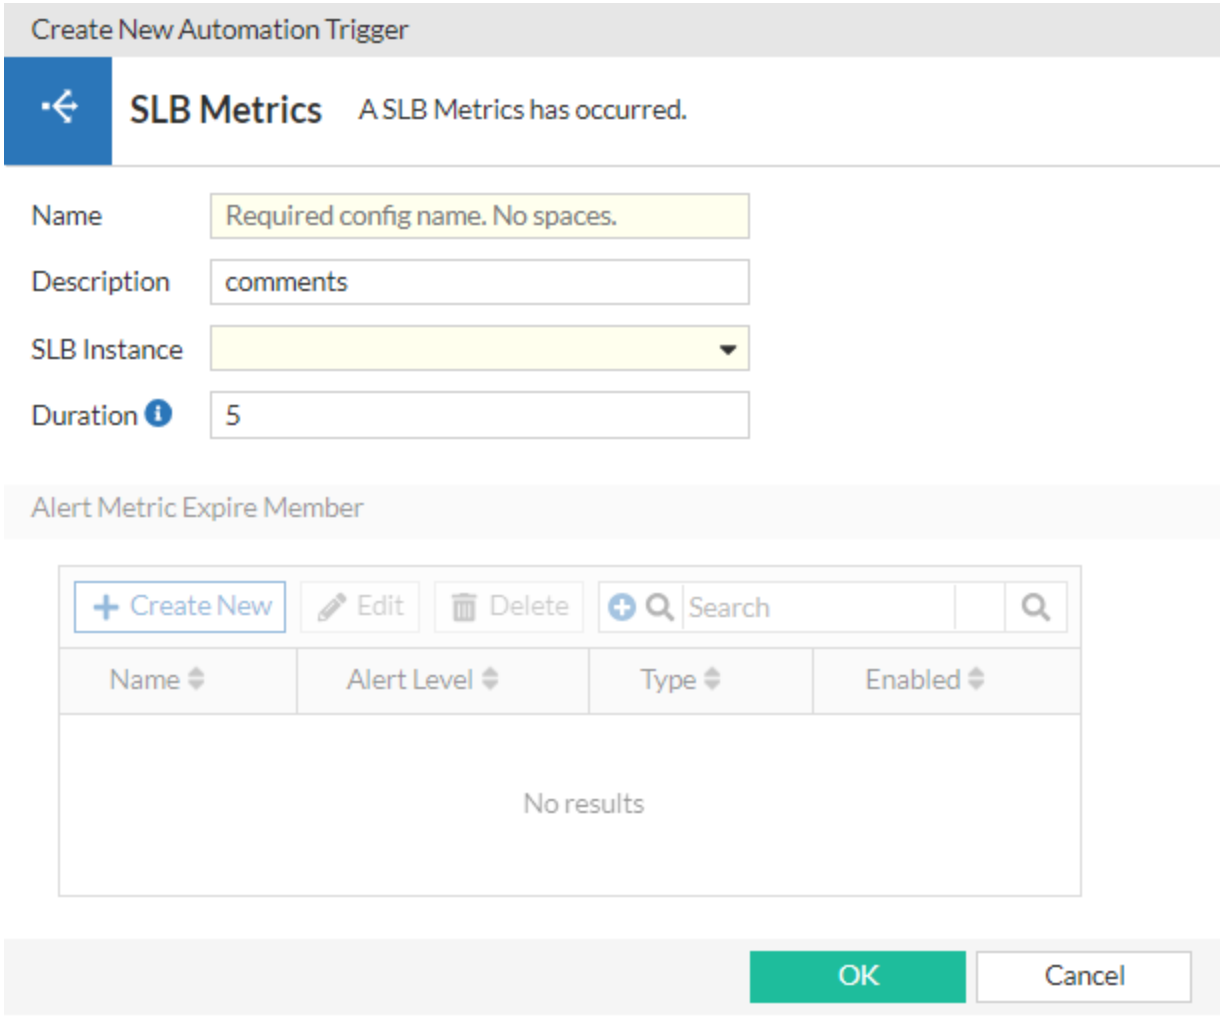
- You must have Global Administrator access. Ensure that your admin account settings has **Global Admin** set to **Yes**. For more information, see Creating administrator users on page 249.

SLB Metrics

To configure an SLB Metrics trigger alert:

1. Go to **Security Fabric > Automation**.
2. Click the **Trigger** tab.
3. Click **Create New** to display the Create New Automation Trigger configuration page.

4. Under the **System** section, click **SLB Metrics** to display the configuration editor.

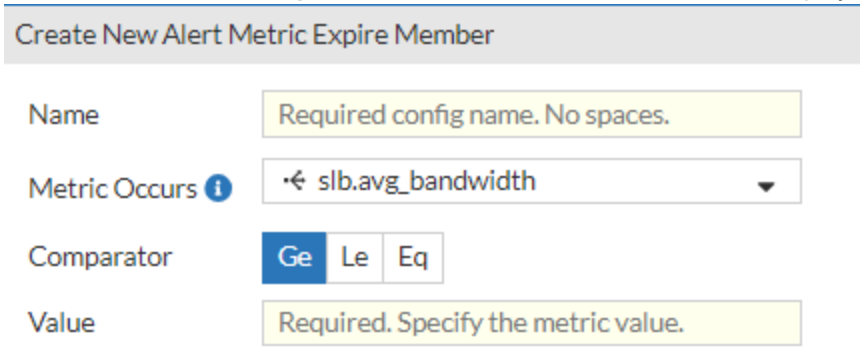


5. Configure the following trigger alert settings:

Setting	Description
Name	Enter a name for the new SLB Metrics trigger alert. The configuration name cannot be edited once it has been saved.
Description	Optionally, you can add a description about this trigger alert configuration.
SLB Instance	Select the virtual server on which the SLB Metrics trigger applies.
Duration	Specify the metric duration in seconds. Range: 5-3600 seconds.

6. Click **OK** to commit the SLB Metrics trigger settings.
Once the SLB Metrics trigger alert configuration has been saved, you can then add the alert member configurations under the **Alert Metric Expire Member** section.

7. Under the **Alert Metric Expire Member** section, click **Create New** to display the configuration editor.



8. Configure the following trigger alert member settings:

Setting	Description
Name	Enter a name for the new SLB Metrics trigger alert member. The configuration name cannot be edited once it has been saved.
Metric Occurs	Select the server load balance performance metric events that will trigger the action.
Comparator	The metric is compared to the Value field according to the selected option: • Ge—greater than • Le—less than • Eq—equal to The action will be triggered if the specified value satisfies the selected option.
Value	Specify the metric value that the **Comparator** uses to determine if the metric triggers an action (for example, 2 milliseconds).

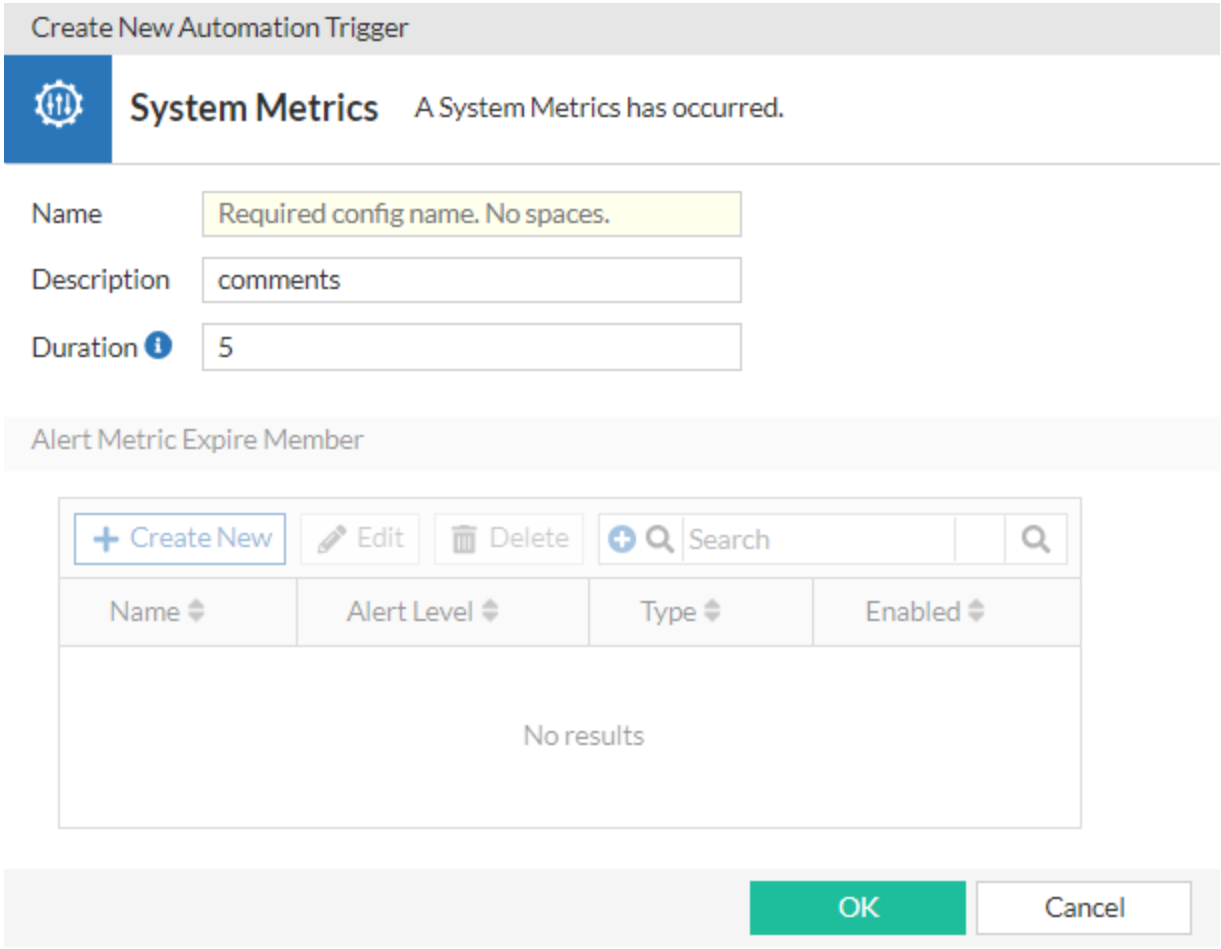
9. Click **OK** to commit the trigger alert member settings.
   The newly created trigger alert member is added under the **Alert Metric Expire Member** section.
10. Click **OK** to save the SLB Metrics trigger alert configuration.

System Metrics

To configure a System Metrics trigger alert:

1. Go to **Security Fabric > Automation**.
2. Click the **Trigger** tab.
3. Click **Create New** to display the Create New Automation Trigger configuration page.

4. Under the **System** section, click **System Metrics** to display the configuration editor.

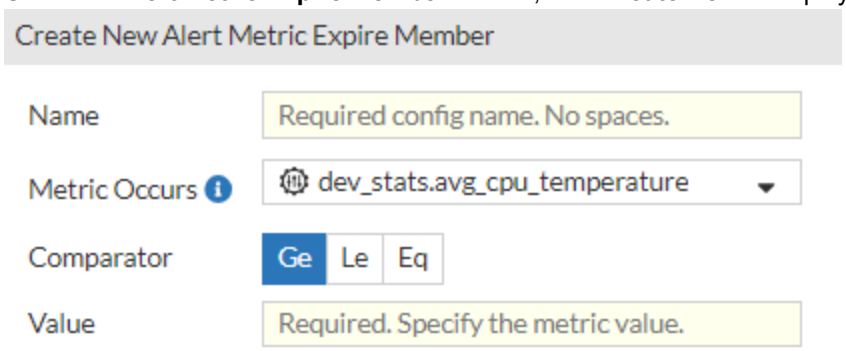


5. Configure the following trigger alert settings:

Setting	Description
Name	Enter a name for the new System Metrics trigger alert. The configuration name cannot be edited once it has been saved.
Description	Optionally, you can add a description about this trigger alert configuration.
Duration	Specify the metric duration in seconds. Range: 5-3600 seconds.

6. Click **OK** to commit the System Metrics trigger settings.
Once the System Metrics trigger alert configuration has been saved, you can then add the alert member configurations under the **Alert Metric Expire Member** section.

7. Under the **Alert Metric Expire Member** section, click **Create New** to display the configuration editor.



8. Configure the following trigger alert member settings:

Setting	Description
Name	Enter a name for the new System Metrics trigger alert member. The configuration name cannot be edited once it has been saved.
Metric Occurs	Select the system metrics events (average CPU usage, average memory usage, etc.) that will trigger the action.
Comparator	The metric is compared to the Value field according to the selected option: • Ge—greater than • Le—less than • Eq—equal to The action will be triggered if the specified value satisfies the selected option.
Value	Specify the metric value that the **Comparator** uses to determine if the metric triggers an action (for example, 2 milliseconds).

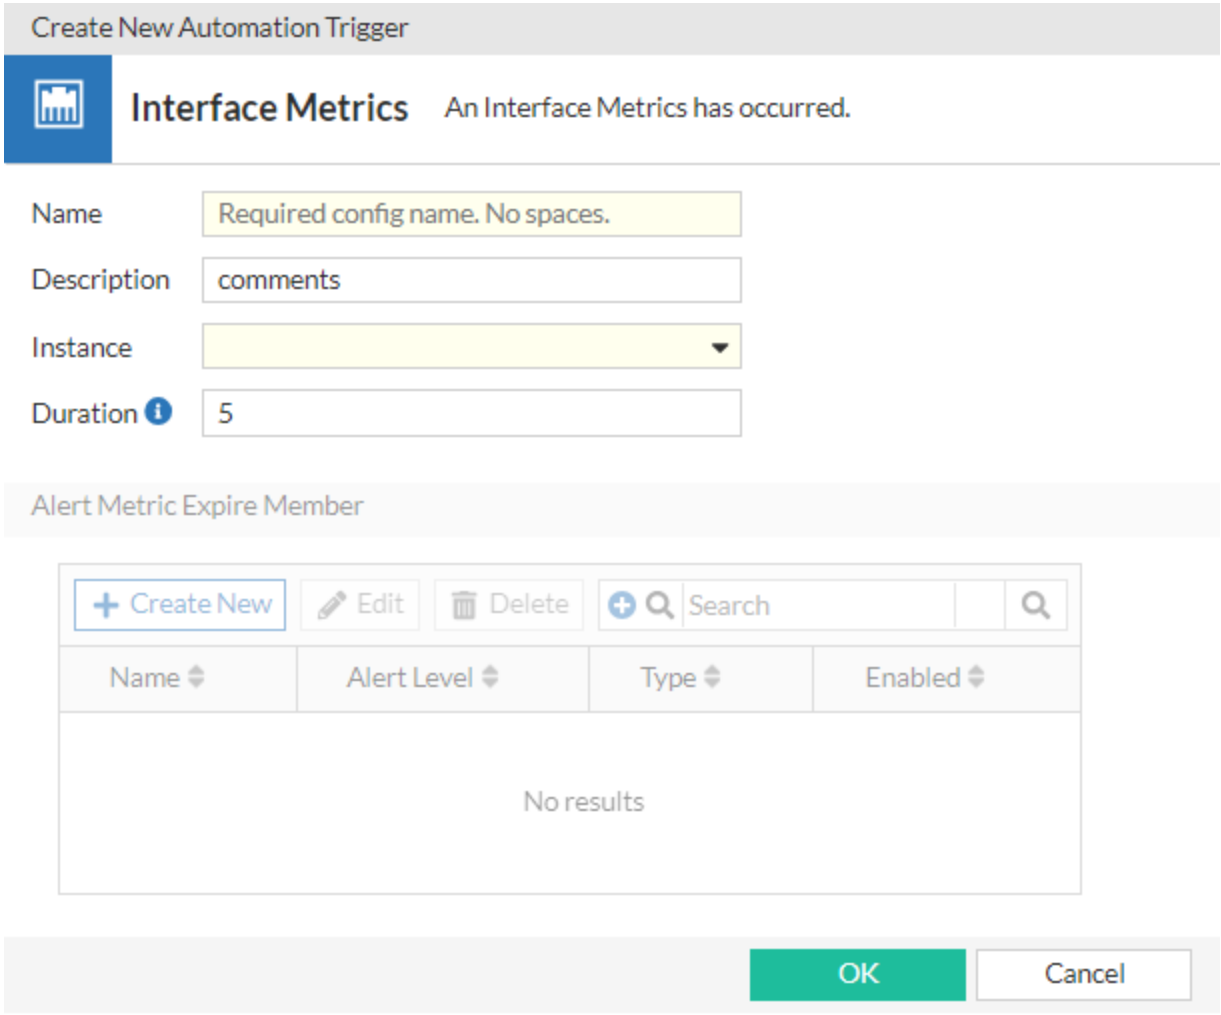
9. Click **OK** to commit the trigger alert member settings.
   The newly created trigger alert member is added under the **Alert Metric Expire Member** section.
10. Click **OK** to save the System Metrics trigger alert configuration.

Interface Metrics

To configure an Interface Metrics trigger alert:

1. Go to **Security Fabric > Automation**.
2. Click the **Trigger** tab.
3. Click **Create New** to display the Create New Automation Trigger configuration page.

4. Under the **System** section, click **Interface Metrics** to display the configuration editor.

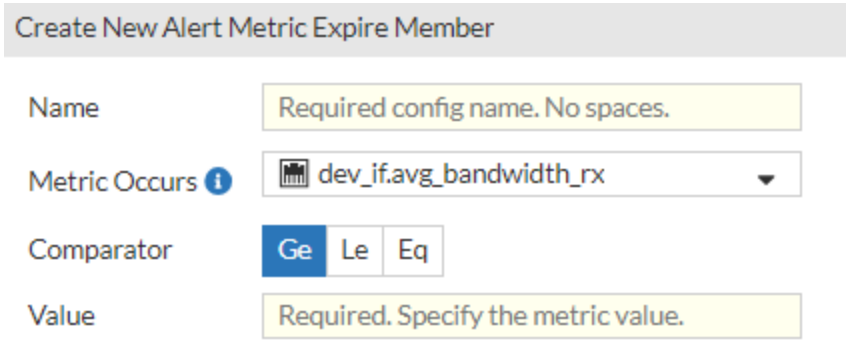


5. Configure the following trigger alert settings:

Setting	Description
Name	Enter a name for the new Interface Metrics trigger alert. The configuration name cannot be edited once it has been saved.
Description	Optionally, you can add a description about this trigger alert configuration.
Instance	Select the network interface on which the Interface Metrics trigger applies.
Duration	Specify the metric duration in seconds. Range: 5-3600 seconds.

6. Click **OK** to commit the Interface Metrics trigger settings.
Once the Interface Metrics trigger alert configuration has been saved, you can then add the alert member configurations under the **Alert Metric Expire Member** section.

7. Under the **Alert Metric Expire Member** section, click **Create New** to display the configuration editor.



8. Configure the following trigger alert member settings:

Setting	Description
Name	Enter a name for the new Interface Metrics trigger alert member. The configuration name cannot be edited once it has been saved.
Metric Occurs	Select the network interface events that will trigger the action.
Comparator	The metric is compared to the Value field according to the selected option: • Ge—greater than • Le—less than • Eq—equal to The action will be triggered if the specified value satisfies the selected option.
Value	Specify the metric value that the **Comparator** uses to determine if the metric triggers an action (for example, 2 milliseconds).

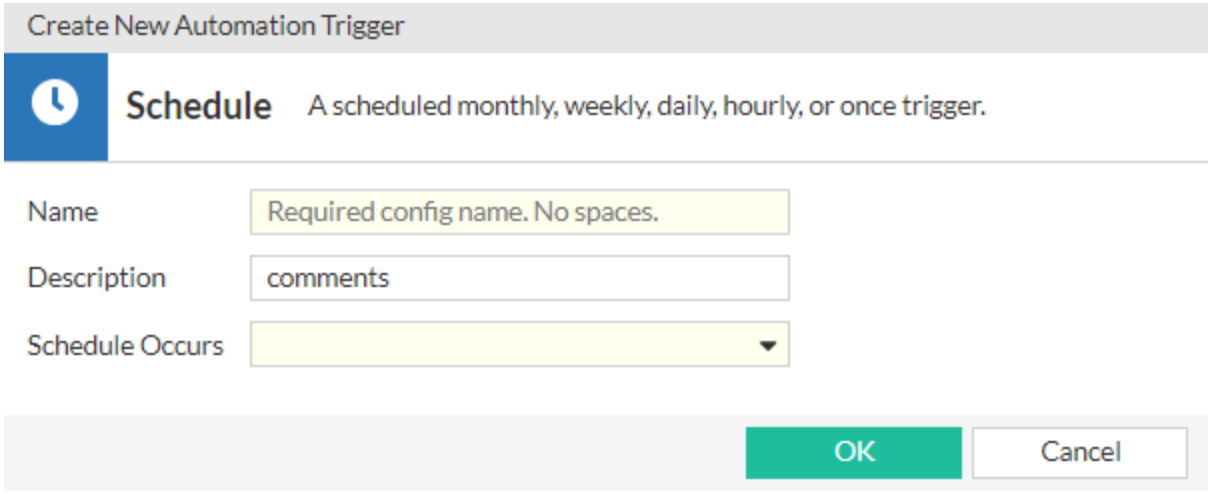
9. Click **OK** to commit the trigger alert member settings.
   The newly created trigger alert member is added under the **Alert Metric Expire Member** section.
10. Click **OK** to save the Interface Metrics trigger alert configuration.

Schedule

To configure a Schedule trigger alert:

1. Go to **Security Fabric > Automation**.
2. Click the **Trigger** tab.
3. Click **Create New** to display the Create New Automation Trigger configuration page.

4. Under the **Miscellaneous** section, click **Schedule** to display the configuration editor.



5. Configure the following trigger alert settings:

Setting	Description
Name	Enter a name for the new Schedule trigger alert. The configuration name cannot be edited once it has been saved.
Description	Optionally, you can add a description about this trigger alert configuration.
Schedule Occurs	Select a user-defined schedule group object or create a new schedule group. For details, see Schedule Group on page 368.

6. Click **OK**.

FortiADC Log

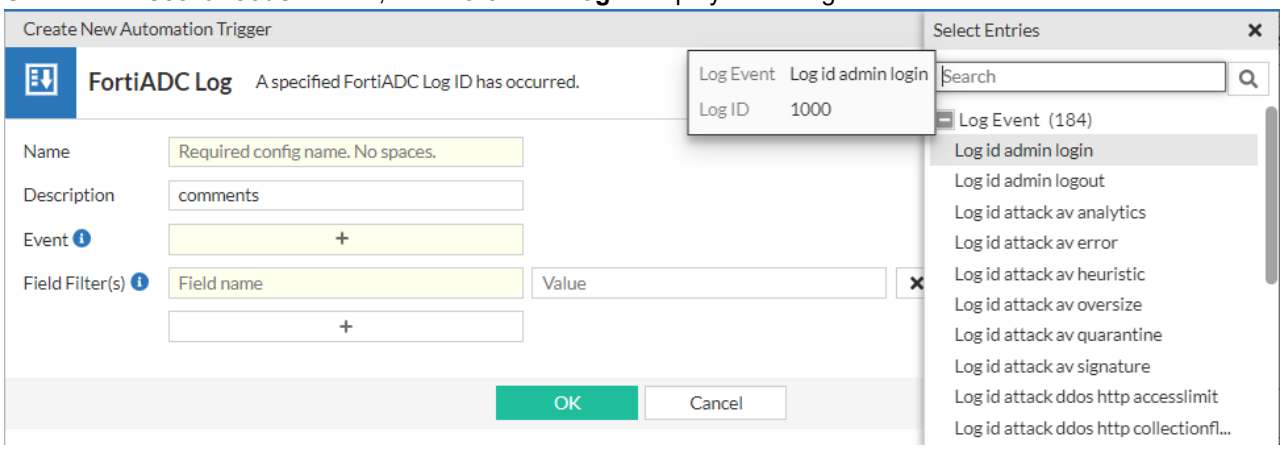
You can configure a FortiADC Log trigger for when a specific event log ID occurs. You can select multiple event log IDs, and apply log field filters.

Exact information from the Event Log is required to configure the FortiADC Log trigger alert, so it is recommended that you download the specific log files to obtain the information prior to beginning the configuration.

To configure a FortiADC Log trigger alert:

1. Go to **Security Fabric > Automation**.
2. Click the **Trigger** tab.
3. Click **Create New** to display the Create New Automation Trigger configuration page.

4. Under the **Miscellaneous** section, click **FortiADC Log** to display the configuration editor.



You can view the last 4 digits of the log ID by hovering over the event entry in the Log Event list.

5. Configure the following trigger alert settings:

Setting	Description
Name	Enter a name for the new FortiADC Log trigger alert. The configuration name cannot be edited once it has been saved.
Description	Optionally, you can add a description about this trigger alert configuration.
Event	Select the FortiADC Log event (identified through the log ID) that will trigger the automation stitch. The Event options correspond to the Message listed in the FortiADC event log. Hover over an entry to view the tooltip that includes the Log Event and Log ID. The Log ID refers to the last 4 digits of the Log ID from the event log. The maximum is 32 log events per trigger configuration.
Field Filters	Enter the Field name and the Value of the logged event that will trigger the automation stitch. The Field name and Value refers to the "key" and its value as it is recorded in the log file. To successfully trigger an automation stitch using the FortiADC Log trigger alert, the field name and value must match what is recorded in the log file exactly. • Field name — Enter the field name as it appears exactly in the log file. • Value — Specify the value of the specified field name as it appears exactly in the log file.

6. Click **OK**.

Example: Create a FortiADC Log trigger alert for failed logins

In this example, you will create a FortiADC Log alert to trigger when a login fails due to incorrect password input.

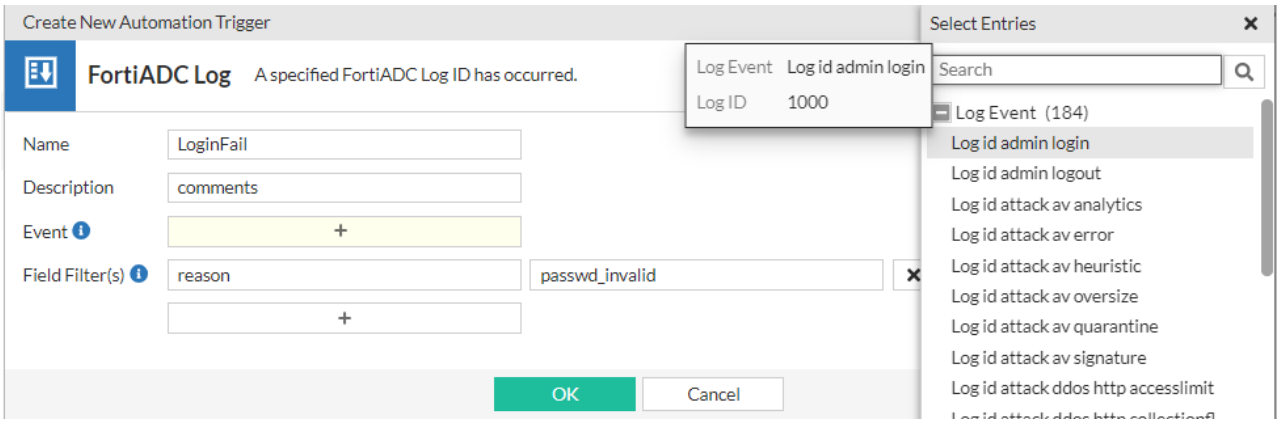
1. Download the required event log file to obtain the specific information to use in the FortiADC Log trigger configuration.
   a. Go to **Log & Report > Event Log**.
   b. Locate the "Admin login" event log.
   c. Download the log file.

2. Open the log file to obtain the required information.

```
date=2023-12-13 time=10:28:41 log_id=0001001000 type=event subtype=admin pri=information
vd=root msg_id=32713167 user=admin ui=172.19.162.70 action=login status=failure
reason=passwd_invalid logdesc=Admin login msg=User admin login successfully from
172.19.162.70
```

   Copy the **log id** and **reason** fields and their values to use in later steps.

3. Go to **Security Fabric > Automation**, click the **Trigger** tab and **Create New**.
4. Click **FortiADC Log** to configure a new FortiADC Log automation trigger using the event log information collected earlier.



   a. Select the **Event** as **Log id admin login**. The last 4 digits of the "log id" from the log file should match the "Log ID" of the selected log event entry.
   b. Enter the **Field name** as **reason** and **Value** as **passwd_invalid**.
   c. Click **OK**.

Predefined automation trigger events

Trigger	Events
Security Events	Bot Detected Brute Force Detected CORS Violate Detected CSRF Violate Detected Data Leak Violate Detected DDoS HTTP Access Limit DDoS HTTP Connection Flood DDoS HTTP Request Flood DDoS IP Fragmentation DDoS SYNFLOOD attack start DDoS SYNFLOOD attack stop DDoS TCP Access Flood DDoS TCP Slow Data Attack Generic Attack Detected Geo Violate Detected

Trigger	Events
	HTML Validation Detected JSON Violate Detected OPENAPI Violate Detected Protocol Constraint Detected Reputation Violate Detected Request Blocked SEC Biometrics Base Detected SEC Threshold Violate Detected SOAP Violate Detected SQL Injection Attack Detected URL Pattern Violate Detected Virtual Server Authentication Fail Web Anti Defacement Detected XML Violate Detected XSS Attack Detected
HA Failover	HA Master Failover HA Peer Lost
System Events	Admin user login failed and blocked IP ARP Conflict Bad Device Fan Bad PSU Fan Certification Expire Config Create Config Delete Config Update CRL Expires Device Rebooted Device Upgrade Completed FW SNAT Port Exhausted Gateway HC Down Gateway HC Up Gateway Inbound Bandwidth Gateway Inbound Spillover Gateway Outbound Bandwidth Gateway Outbound Spillover Gateway Total Spillover GLB GW Available GLB GW Not Available GLB Real Server Available

Trigger	Events
	GLB Real Server Not Available
	GLB Virtual Server Available
	GLB Virtual Server Not Available
	Good Device Fan
	Good PSU Fan
	High CPU Temp
	High CPU Usage
	High Device Temp
	High Disk Usage
	High Memory Usage
	High Power Supply
	High PSU Temp
	High PSU Voltage
	High Voltage
	Link Group HC Down
	Link Group HC Up
	Log Full
	Logical Interface Disabled
	Logical Interface Down
	Logical Interface Up
	Lost Log Disk
	Low Power Supply
	Low PSU Voltage
	Low Voltage
	Normal CPU Temp
	Normal Device Temp
	Normal PSU Temp
	OCSP Response Expires
	PSU Failure
	Real Server Connection Limit Start
	Real Server Connection Limit Stop
	Real Server Connection Rate Start
	Real Server Connection Rate Stop
	Real Server Disabled
	Real Server Enabled
	Real Server HC Down
	Real Server HC Up
	Real Server Maintain Mode
	Slow Device Fan
	Slow PSU Fan

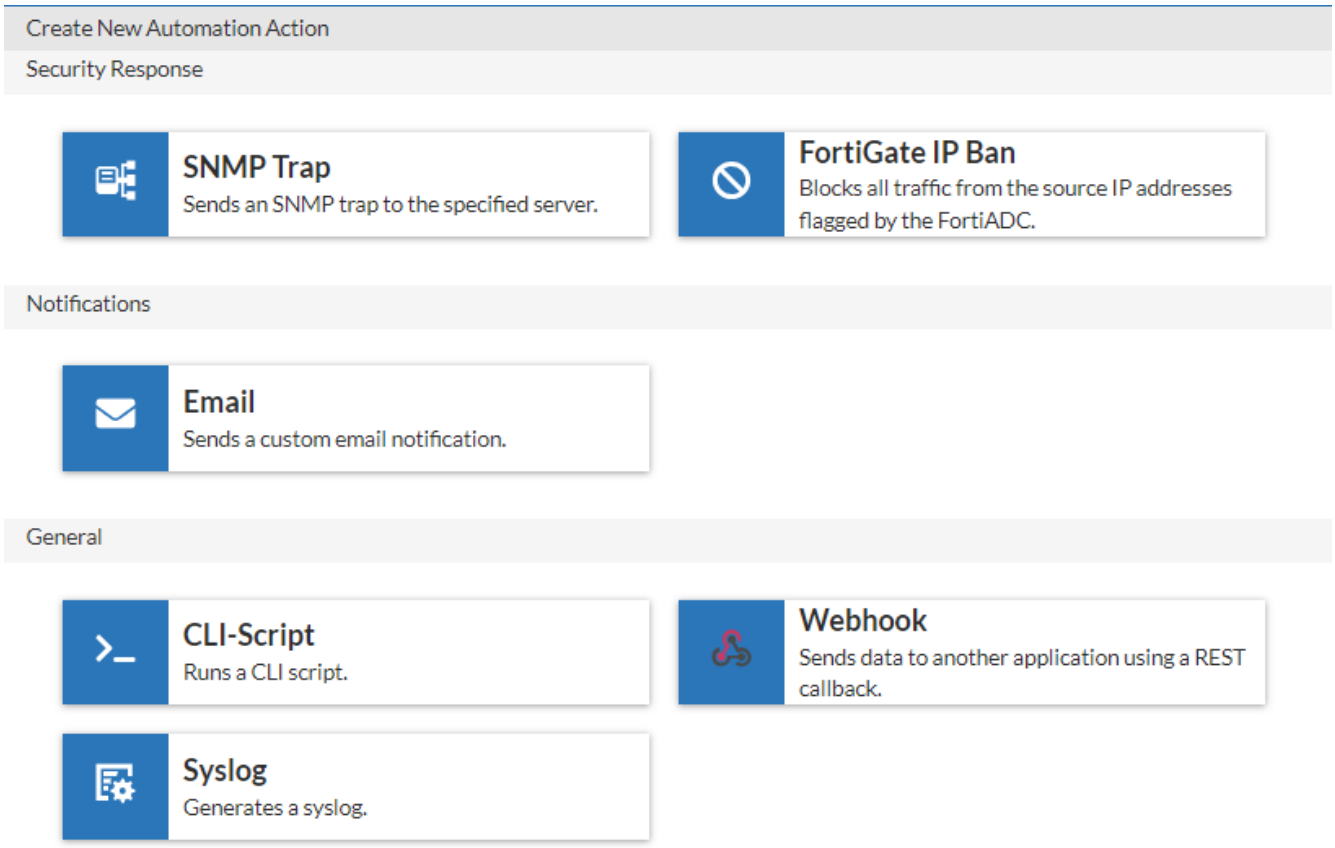
Trigger	Events
	SSD MWI Near Threshold
	SSD MWI Reached Threshold
	SSL Certificate Revoked
	User Login
	User Logout
	Virtual Server Connection Limit Start
	Virtual Server Connection Limit Stop
	Virtual Server Connection Rate Start
	Virtual Server Connection Rate Stop
	Virtual Server Disabled
	Virtual Server Down
	Virtual Server Enabled
	Virtual Server IP Pool Limit
	Virtual Server Maintain Mode
	Virtual Server Transaction Rate Start
	Virtual Server Transaction Rate Stop
	Virtual Server Up

Configuring Automation Actions

On the **Security Fabric > Automation > Action** tab, you can view the list of available automation response actions that have been user-defined. After defining your automation actions, you can combine them with a trigger to create an automation stitch. For details, see Creating automation stitches on page 42

FortiADC supports six response action types:

- SNMP Trap on page 68 — Sends an SNMP trap to the specified server in response to the trigger. This action is not supported for the Schedule trigger.
- FortiGate IP Ban on page 70 — Blocks all traffic from the source IP addresses flagged by the FortiGate in response to the trigger. This action can only be used with the Period Block IP trigger.
- Email on page 71 — Sends a custom email notification in response to the trigger.
- CLI Script on page 73 — Runs a CLI script in response to the trigger. This action is not supported for the Period Block IP trigger.
- Webhook on page 74 — Sends data to another application using a REST callback in response to the trigger.
- Syslog on page 76 — Generates a syslog in response to the trigger.



Before you begin:

- You must have Global Administrator access. Ensure that your admin account settings has **Global Admin** set to **Yes**. For more information, see Creating administrator users on page 249.

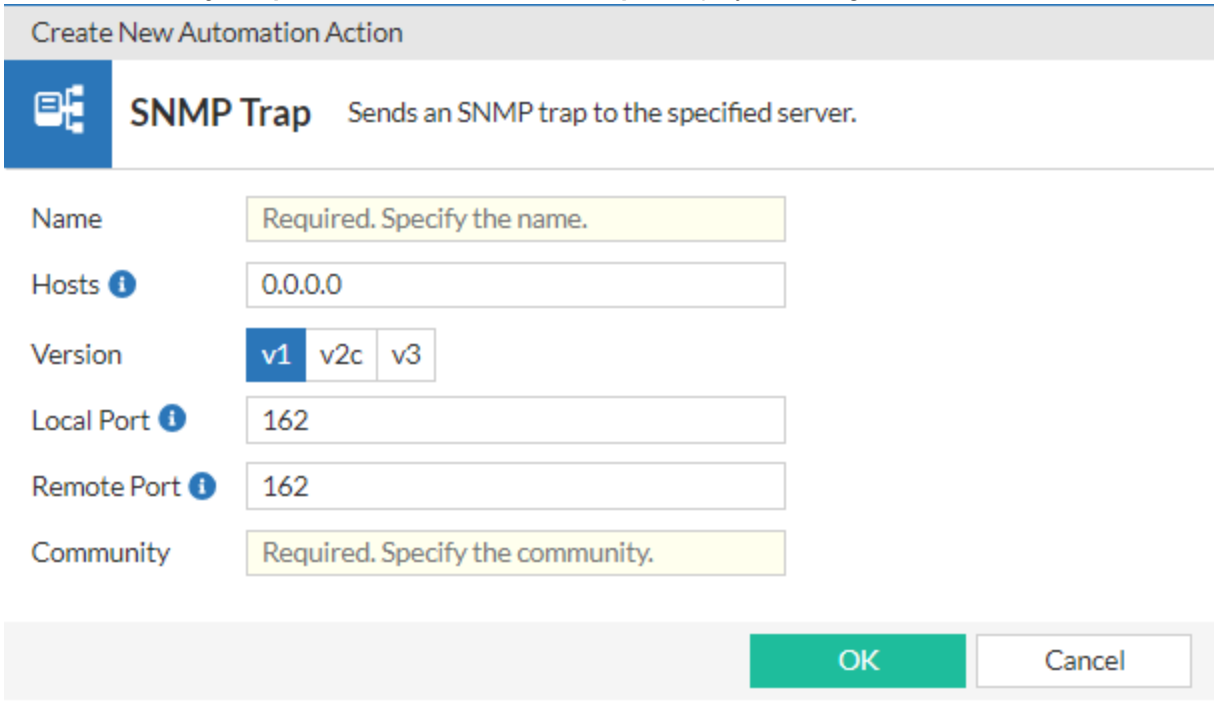
SNMP Trap

Use this action to send SNMP traps to the specified server in response to a trigger event.

To configure an SNMP Trap response action:

1. Go to **Security Fabric > Automation**.
2. Click the **Action** tab.
3. Click **Create New** to display the Create New Automation Action configuration page.

4. Under the **Security Response** section, click **SNMP Trap** to display the configuration editor.



5. Configure the following response action settings:

Setting	Description
Name	Enter a name for the new SNMP Trap action. The configuration name cannot be edited once it has been saved.
Hosts	Specify the IP address that will receive this message.
Version	Select the SNMP version to use • v1 • v2c • v3 **Note**: If using the **System Event** trigger **Admin user login failed and blocked IP**, it is recommended to use SNMPv2 or SNMPv3 only. When using SNMPv1 with the Admin user login failed and blocked IP system event trigger, the event is not recorded on the SNMP client even when the SNMP action is triggered successfully. Whereas the SNMP client records using SNMPv2 or SNMPv3 both properly reflect the triggered action for Admin user login failed and blocked IP event.
Local Port	Specify the source port number. Default: 162 Range: 0-65535
Remote Port	Specify the destination port number. Default: 162 Range: 0-65535
Security Level	The **Security Level** option is available if **v3** is selected for **Version**. The SNMP security level to use: • Auth But no Privacy • Auth And Privacy

Setting	Description
	• No Privacy
Auth Algorithm	The **Auth Algorithm** option is available if **Auth But no Privacy** or **Auth And Privacy** is selected for **Security Level**. The authentication algorithm to use: • SHA1 • MD5
Auth Password	The **Auth Password** option is available if **Auth But no Privacy** or **Auth And Privacy** is selected for **Security Level**. The password to the authentication algorithm.
Private Algorithm	The **Private Algorithm** option is available if **Auth And Privacy** is selected for **Security Level**. The private algorithm to use: • AES • DES
Private Password	The **Private Password** option is available if **Auth And Privacy** is selected for **Security Level**. The password to the private algorithm.
User	Specify the User.

6. Click **OK**.

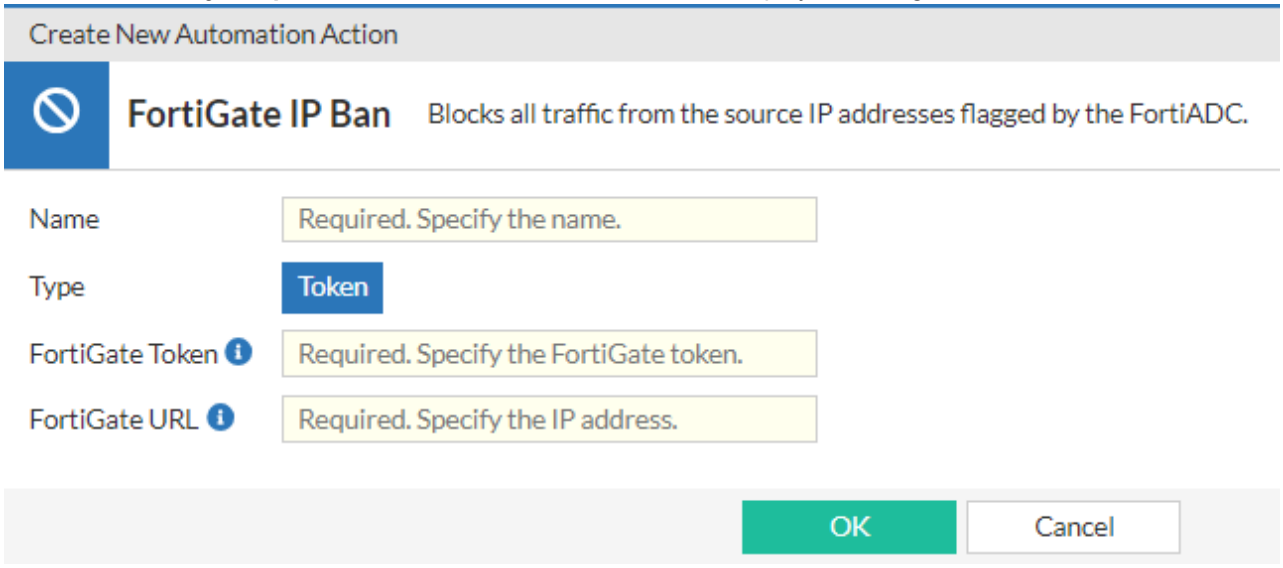
FortiGate IP Ban

Use this action to block all traffic from the source addresses flagged by the FortiGate in response to the Period Block IP trigger. See FortiGate IP Ban action for details.

To configure a FortiGate IP Ban response action:

1. Go to **Security Fabric > Automation**.
2. Click the **Action** tab.
3. Click **Create New** to display the Create New Automation Action configuration page.

4. Under the **Security Response** section, click **FortiGate IP Ban** to display the configuration editor.



5. Configure the following response action settings:

Setting	Description
Name	Enter a name for the new FortiGate IP Ban action. The configuration name cannot be edited once it has been saved.
Type	Token
FortiGate Token	Specify the FortiGate Token. To get the token, log in to FortiGate, go to **System> Administrator**, create a new REST API Administrator, then generate API key.
FortiGate URL	Specify the IP address of the FortiGate URL. For example, https://10.106.155.107

6. Click **OK**.

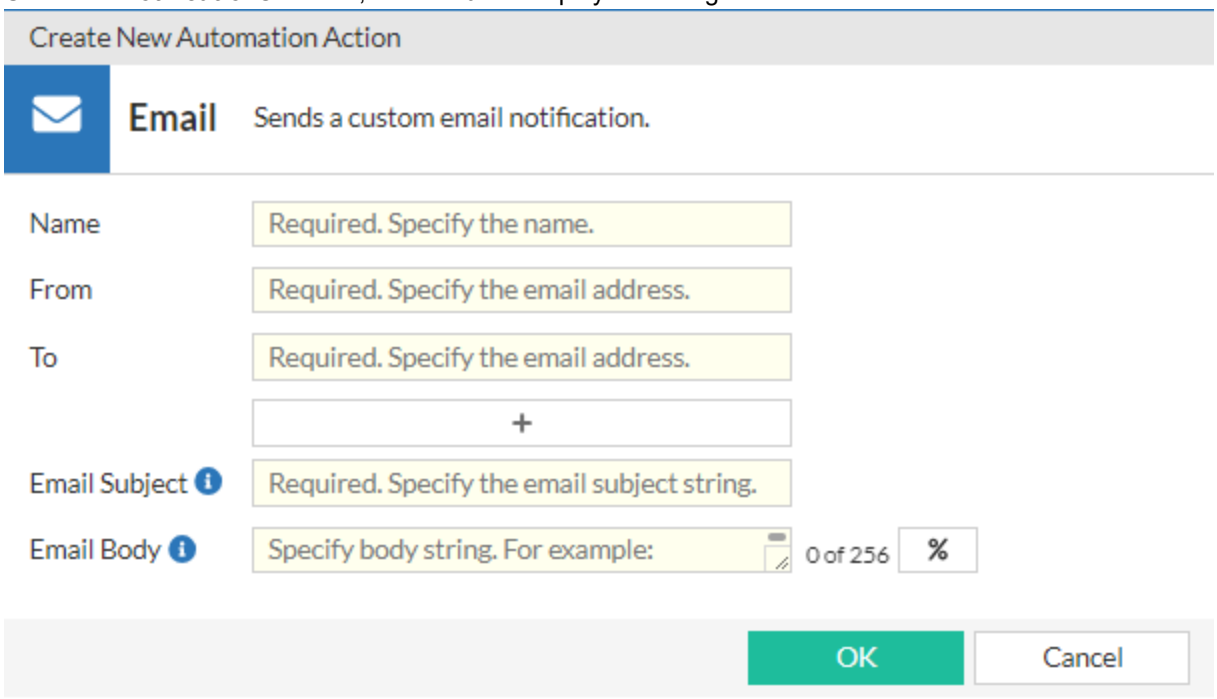
Email

Use this action to send a custom email notification in response to a trigger event.

To configure an Email response action:

1. Go to **Security Fabric > Automation**.
2. Click the **Action** tab.
3. Click **Create New** to display the Create New Automation Action configuration page.

4. Under the **Notifications** section, click **Email** to display the configuration editor.



5. Configure the following settings:

Setting	Description
Name	Enter a name for the new Email action. The configuration name cannot be edited once it has been saved.
From	Specify the sender email address of this notification.
To	Specify the recipient email address of this notification.
Email Subject	Specify the email subject string.
Email Body	Write the email message in the Email Body. Maximum 256 characters. You can insert specific system data, such as parameters from logs or previous action results by wrapping the parameter with %% tags to replace the expression with the JSON value for that parameter. There are 6 available Action Parameters: • %%results%% — In automation stitches with more than one response action, %%results%% inserts the complete result from the previous action, such as a CLI Script action. • %%alert_msg%% — The complete alert string will be inserted from the event source when it happens, such as from the WAF module. • %%metric_obj%% — Applicable to System trigger alerts (SLB Metrics, System Metrics, and Interface Metrics), %%metric_obj%% inserts the metric instance (port) and value that was configured in the trigger alert. • %%block.srcip%% — The quarantined Source IP will be inserted when the WAF module blocks the IP. • %%event.srcip%% — The Source IP of the event will be inserted from the

Setting	Description
	event source when it happens, such as from the WAF module. • %%log_msg%% — Applicable to the FortiADC Log trigger alert, %%log_msg%% inserts the complete log string from the log source when it happens.

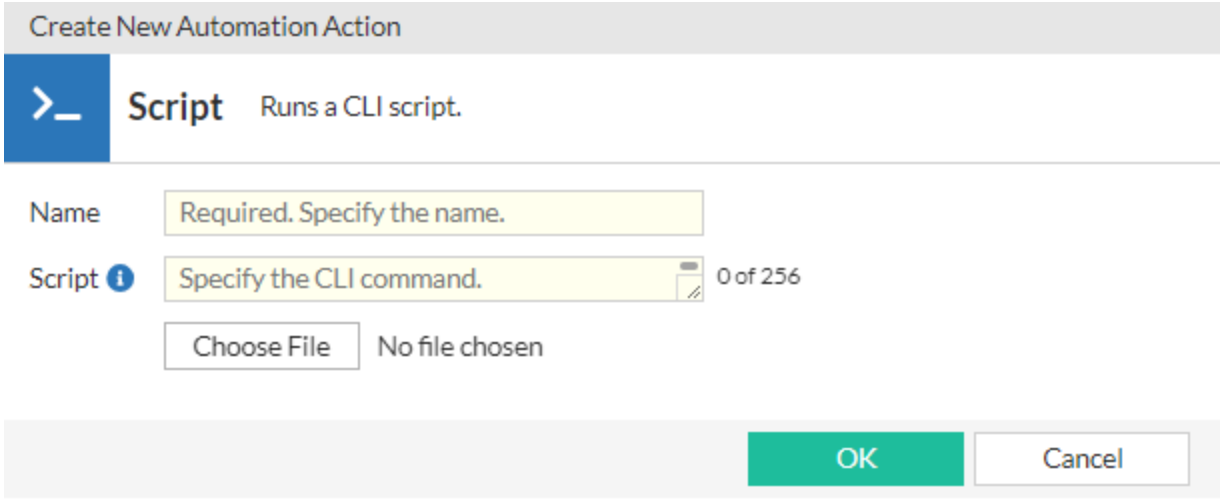
6. Click **OK**.

CLI Script

Use this action to run a CLI script in response to a trigger event, such as to make appropriate configuration changes. The scripts can be manually entered or uploaded as a file.

To configure a CLI Script response action:

1. Go to **Security Fabric > Automation**.
2. Click the **Action** tab.
3. Click **Create New** to display the Create New Automation Action configuration page.
4. Under the **General** section, click **CLI Script** to display the configuration editor.



5. Configure the following response action settings:

Setting	Description
Name	Enter a name for the new CLI Script action. The configuration name cannot be edited once it has been saved.
Script	Manually enter or upload the script. • To manually enter the script, type it into the **Script** field. • To upload a script file, click **Choose File** and locate the file on your management computer. Maximum 256 characters.

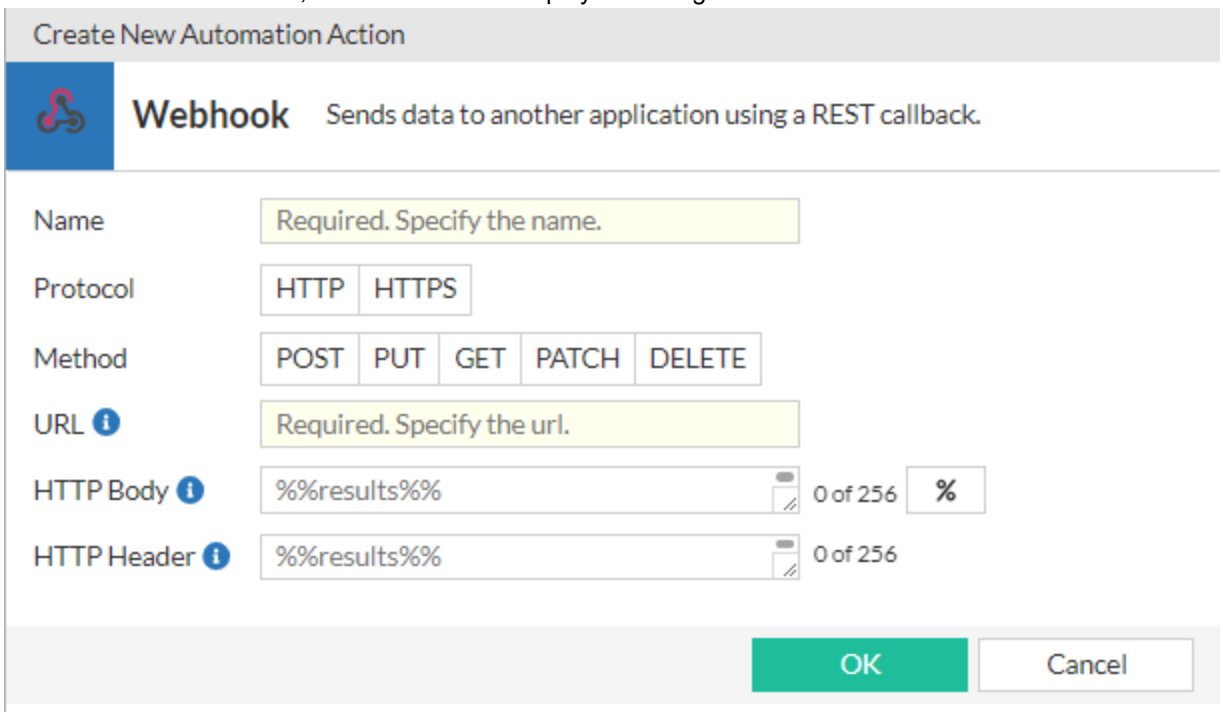
6. Click **OK**.

Webhook

Use this action to send data to another application using a REST callback in response to a trigger event.

To configure a Webhook response action:

1. Go to **Security Fabric > Automation**.
2. Click the **Action** tab.
3. Click **Create New** to display the Create New Automation Action configuration page.
4. Under the **General** section, click **Webhook** to display the configuration editor.



5. Configure the following response action settings:

Setting	Description
Name	Enter a name for the new Webhook action. The configuration name cannot be edited once it has been saved.
Protocol	Select the request protocol to use: • HTTP • HTTPS
Method	Specify the request method: • POST • PUT • GET • PATCH • DELETE
URL	Specify the request URL. For example, 10.106.155.130:90/test
HTTP Body	Specify the request body. For example, 'msg': 'abc', 'user': 'jack' You can insert specific system data, such as parameters from logs or previous action results by wrapping the parameter with %% tags to replace the expression with the JSON value for that parameter. There are 6 available Action Parameters: • %%results%% — In automation stitches with more than one response action, %%results%% inserts the complete result from the previous action, such as a CLI Script action. • %%alert_msg%% — The complete alert string will be inserted from the event source when it happens, such as from the WAF module. • %%metric_obj%% — Applicable to System trigger alerts (SLB Metrics, System Metrics, and Interface Metrics), %%metric_obj%% inserts the metric instance (port) and value that was configured in the trigger alert. • %%block.srcip%% — The quarantined Source IP will be inserted when the WAF module blocks the IP. • %%event.srcip%% — The Source IP of the event will be inserted from the event source when it happens, such as from the WAF module. • %%log_msg%% — Applicable to the FortiADC Log trigger alert, %%log_msg%% inserts the complete log string from the log source when it happens.
HTTP Header	Specify the HTTP request header name and value. For example, customerheader1:value1 customerheader2:value2 Ensure to only use space as the delimiter for multiple headers.
TLS Certificate	The **TLS Certificate** option is available if the **Protocol** is **HTTPS**. Select a TLS certificate to verify by the server to validate the HTTPS connection to the webhook endpoint. A valid TLS certificate is required if the HTTPS server is enabled for two-way authentication. However, a TLS certificate is optional if the HTTPS server is not enabled for two-way authentication.

Setting	Description
Verify Remote Host	The **Verify Remote Host** option is available if the **Protocol** is **HTTPS**. Enable to verify that the remote server matches the host URL using a CA certificate. This option is disabled by default.
CA Certificate	The **CA Certificate** option is available if the **Protocol** is **HTTPS** and **Verify Remote Host** is enabled. Select the CA certificate to use to verify the remote server. FortiADC will verify that the IP or domain name matches in the Remote host field or the Subject alternative name field in the certificate CN.

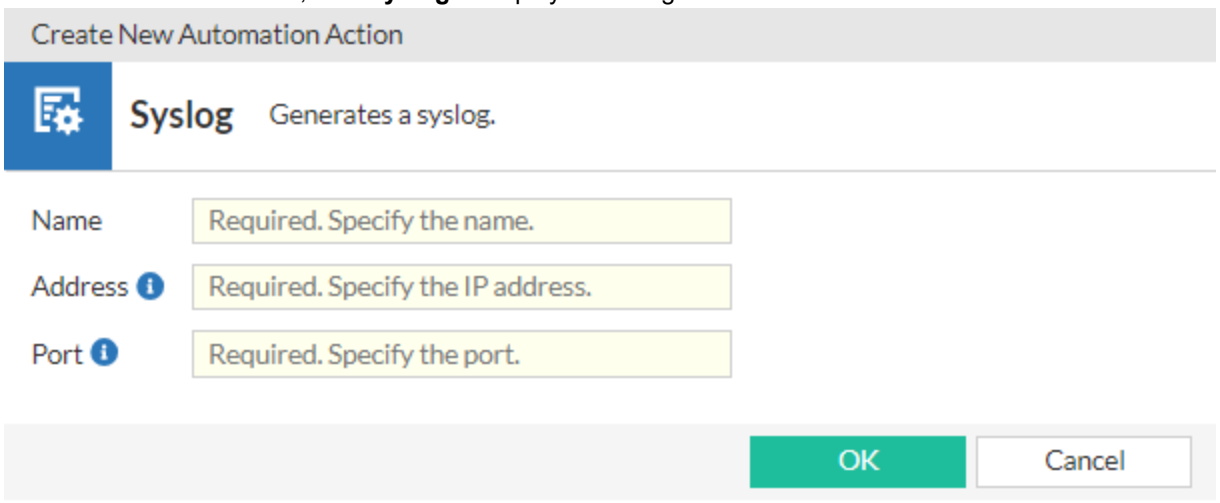
6. Click **OK**.

Syslog

Use this action to generate a syslog message in response to a trigger event.

To configure a Syslog response action:

1. Go to **Security Fabric > Automation**.
2. Click the **Action** tab.
3. Click **Create New** to display the Create New Automation Action configuration page.
4. Under the **General** section, click **Syslog** to display the configuration editor.



5. Configure the following response action settings:

Setting	Description
Name	Enter a name for the new Email action. The configuration name cannot be edited once it has been saved.
Address	Specify the IP address that will receive this message.
Port	Specify the port that will receive this message. Range: 1-65535

6. Click **OK**.

Diagnose commands

To test an automation stitch, use the `diagnose debug module alertd` command.

For more information, see `diagnose debug module` in FortiADC CLI Reference.

Fabric connectors

You can use fabric connectors to integrate FortiADC with other Fortinet Security Fabric solutions:

- FortiSIEM Connector on page 77
- FortiAnalyzer Connector on page 78
- FortiSandbox Connector on page 82
- FortiADC Manager Connector on page 83
- FortiGSLB Connector on page 85
- FortiClient EMS Connector on page 86
- FortiGate Security Fabric Connector on page 89

The fabric connectors define the type of connector and include information for FortiADC to communicate with and authenticate with the products.

FortiSIEM Connector

When you create a connector for FortiSIEM, you are specifying how FortiADC can communicate with FortiSIEM for pushing logs to FortiSIEM.

FortiADC will connect to FortiSIEM by UDP, TCP or TCP SSL depending on FortiSIEM connector setting.

Requirements:

- The FortiSIEM service is required to be exposed on External IP.

To create a FortiSIEM Connector:

1. Go to **Security Fabric > Fabric Connectors**.
2. Click **Create New**.
3. Under **Other Fortinet Products**, select **FortiSIEM**.
4. Configure the following **Syslog Server** options, and then click **Save**.

Status	Toggle on/off to enable/disable the Fabric Connector object.
Address	Type the IP address of the FortiSIEM Log server.
Port	Specify the port that FortiADC uses to communicate with the log server.
Proto	Select the protocol used for log transfer from the following: • UDP • TCP • TCP SSL

TCP Framing	Select one of the following options: • Traditional • Octet Counted This field appears only if **Proto** is **TCP** or **TCP SSL**.
Log Level	Select the severity level of the logs. All the exported logs will be attached with the selected severity level.
CSV	Enable to export the logs in .csv file.
Facility	Select the source facility of the logs. We only support the local use facilities which are not reserved and are available for general use.
Event	Enable to export Event logs.
Traffic	Enable to export Traffic logs.
Security	Enable to export Security logs.

After the connector is created, FortiADC will push the logs to FortiSIEM server. The above configurations are also available in **Log&Report > Log Setting > Syslog Server**.

FortiAnalyzer Connector

When you create a connector for FortiAnalyzer, you are specifying how FortiADC can communicate with FortiAnalyzer for pushing logs to FortiAnalyzer. You can choose between two protocol types for sending logs to FortiAnalyzer: Syslog or OFTP.

Using the Syslog protocol will allow FortiADC to connect to FortiAnalyzer by UDP, TCP or TCP SSL depending on the FortiAnalyzer connector setting.

OFTP (Optimized Fabric Transfer Protocol) is used to synchronize information between FortiAnalyzer and other Fortinet products. So, this is the recommended protocol to use for pushing logs to FortiAnalyzer.

Requirements:

- The FortiAnalyzer service is required to be exposed on External IP.



FortiADC supports integration with FortiAnalyzer versions 7.0.2 or later. As earlier versions of FortiAnalyzer is not optimally compatible with FortiADC, unexpected behavior may occur.

Integration with FortiADC 320F and 420F models requires FortiAnalyzer version 7.6.2 or later.

To create a FortiAnalyzer Connector:

1. Go to **Security Fabric > Fabric Connectors**.
2. Click **Create New**.
3. Under **Other Fortinet Products**, click **FortiAnalyzer** to display the configuration editor.

4. Configure the following **Server Type** settings:

Setting	Description
Type	Select either of the following options: • Syslog — To send logs to FortiAnalyzer using the syslog protocol. • FortiAnalyzer — To send logs to FortiAnalyzer using the OFTP.

Depending on your Type selection, either the Syslog Server or FortiAnalyzer configuration section will appear.

5. If the **Server Type** is **Syslog**, configure the following **Syslog Server** settings:

Setting	Description
Status	Enable/disable the Fabric Connector object.
Address	Specify the IP address of the FortiAnalyzer Log server.
Port	Specify the port that FortiADC uses to communicate with the log server. This is the listening port number of the syslog server. Usually this is UDP port 514.
Proto	Select the protocol used for log transfer from the following: • UDP • TCP • TCP SSL
TCP Framing	If **Proto** is **TCP** or **TCP SSL**, the **TCP Framing** options appear. Select one of the following options: • Traditional • Octet Counted
Log Level	Select the lowest severity to log from the following options: • 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. The exported logs will include the selected severity level and above. 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 severity level Alert and Emergency.
CSV	Enable to export the logs as a CSV file.
Facility	Select the source facility of the logs. We only support the local use facilities which are not reserved and are available for general use.
Event	Enable/disable logging for events.

Setting	Description
Event Category	If **Event** is enabled, the **Event Category** options appear. Select one or more of the following event categories to include in the event logs export: • 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	Enable/disable logging for traffic processed by the load-balancing modules.
Traffic Category	If **Traffic** is enabled, the **Traffic Category** options appear. Select one or more of the following traffic categories to include in the traffic logs export: • 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	Enable/disable logging for traffic processed by the security modules.
Security Category	If **Security** is enabled, the **Security Category** options appear. Select one or more of the following security categories to include in the security logs export: • DDoS — DoS protection logs. • IP Reputation — IP Reputation logs. • WAF — WAF logs. • GEO — Geo IP blocking logs. • AV — AV logs. • IPS — IPS logs. • FW — Firewall logs.

6. If the **Server Type** is **FortiAnalyzer**, configure the following **FortiAnalyzer** settings:

Setting	Description
Status	Enable/disable the Fabric Connector object.
Address	Specify the IP address of the FortiAnalyzer Log server.
Log Level	Select the lowest severity to log from the following options: • Emergency — The system has become unstable.

Setting	Description
	• 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. The exported logs will include the selected severity level and above. For example, if you select **Error**, the system collects logs with severity level Error, Critical, Alert, and Emergency. If you select **Alert**, the system collects logs with severity level Alert and Emergency.
Event	Enable/disable logging for events.
Event Category	If **Event** is enabled, the **Event Category** options appear. Select one or more of the following event categories to include in the event logs export: • 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	Enable/disable logging for traffic processed by the load-balancing modules.
Traffic Category	If **Traffic** is enabled, the **Traffic Category** options appear. Select one or more of the following traffic categories to include in the traffic logs export: • 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	Enable/disable logging for traffic processed by the security modules.
Security Category	If **Security** is enabled, the **Security Category** options appear. Select one or more of the following security categories to include in the security logs export: • DDoS — DoS protection logs. • IP Reputation — IP Reputation logs.

Setting	Description
	• WAF — WAF logs. • GEO — Geo IP blocking logs. • AV — AV logs. • IPS — IPS logs. • FW — Firewall logs.

a. Optionally, click **Test Connectivity** after entering the **Address** to check the FortiAnalyzer OFTP connectivity. The **Connection Status** appears showing the OFTP connection status.
There are three possible OFTP connection statuses:

Icon	OFTP Status	Description
	Connected	The FortiADC has successfully connected to FortiAnalyzer and is authorized by FortiAnalyzer as a Fabric Device. FortiADC can now send log data to FortiAnalyzer.
	Disconnected	The FortiADC cannot connect to FortiAnalyzer. Ensure there are no network connectivity issues.
	Need authorization	The FortiADC has successfully connected to FortiAnalyzer but is not authorized by FortiAnalyzer as a Fabric Device. This status may indicate the authorization is either denied or pending. If pending authorization, the status will change to **Connected** once authorization is successful on the FortiAnalyzer server.

If the status is not Connected, edit the FortiAnalyzer connector accordingly to troubleshoot the connection issue.

7. Click **Save**.

After the connector is created, FortiADC will push the logs to the FortiAnalyzer server. The above configurations are also available in **Log&Report > Log Setting > Syslog Server** tab or **FortiAnalyzer** tab.

FortiSandbox Connector

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 enable a FortiSandbox Connector:

1. Go to **Security Fabric > Fabric Connectors**.
2. Double-click **FortiSandbox**.

3. Configure the following **Fabric Device Settings**:

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.

4. Click Save.

FortiCloud Sandbox file upload limits

FortiCloud Sandbox file upload limit on page 83 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

FortiADC Manager Connector

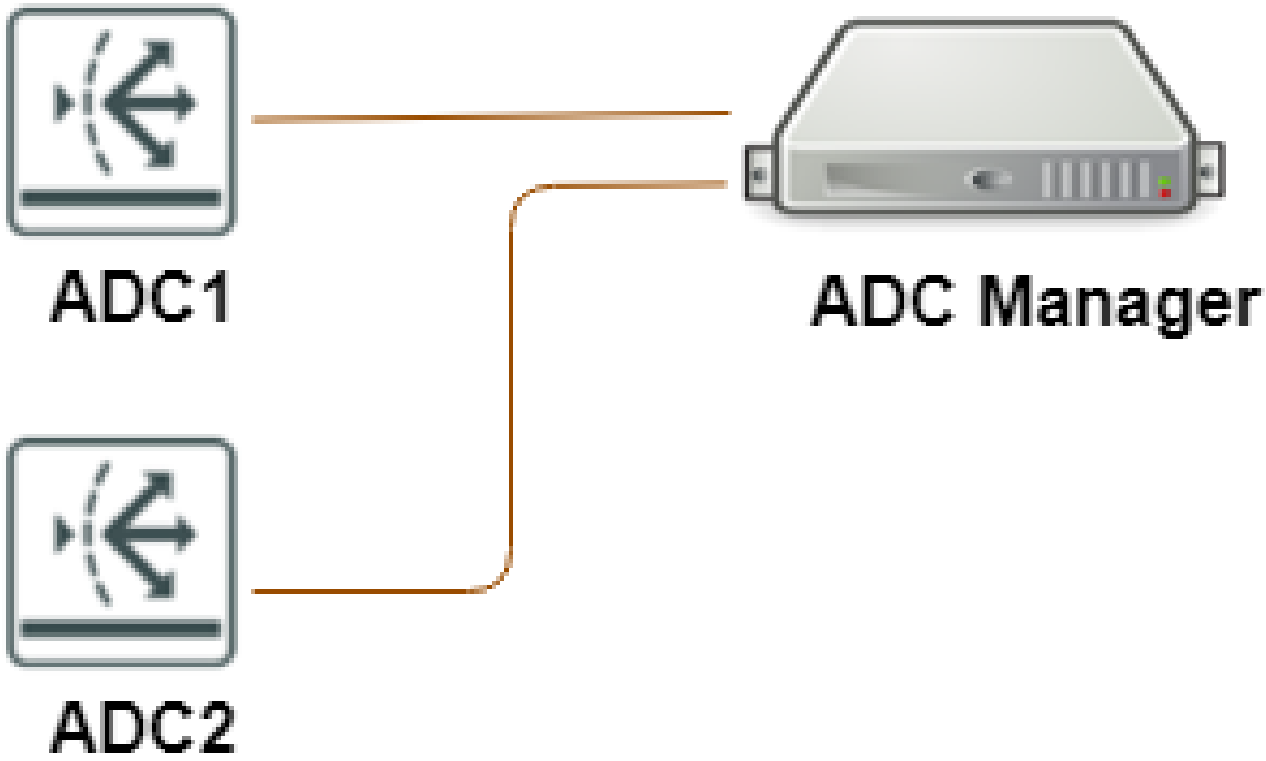
Central Management allows the FortiADC to be connected to a FortiADC Manager. Multiple FortiADCs can be managed by the FortiADC Manager. If you have large networks with multiple FortiADCs, with the FortiADC Manager you can simplify the configuration of these FortiADCs (for example, setting multiple FortiADCs to the same configuration), and view all of their logs and statistics together.

The FortiADC Manager is a powerful tool that gives you more effective control over your FortiADCs.

This guide will show you how to enable central management on your particular FortiADC by connecting to the FortiADC Manager as a Fabric Connector. You will enter the IP address of your manager, then enable Central Management, therefore allowing the FortiADC Manager to manage your FortiADC.

See the FortiADC Manager handbook.

Basic configuration of two FortiADCs linked to a Manager



To enable a FortiADC Manager Connector:

1. Go to **Security Fabric > Fabric Connectors**.
2. Double-click **FortiADC Manager**.

3. Configure the following **Fabric Device Settings**:

Settings	Description
Type	The type of the Central Management None — Initial State of CM Agent. FortiADC Manager — The FortiADC is connected to the Manager. **Note**: The Type is None by default.
Address	The IP address or hostname of the FortiADC Manager. **Note**: The IP address should be empty by default.
Interval	How often the FortiADC tries to connect to the Manager. Default 10 seconds. Range 10-120.
Register	Enable/disable registration to FortiADC Manager. This will enable/disable the connection to the FortiADC Manager. This is disabled by default.
Management Status	The connection status of the FortiADC. • Online — FortiADC Manager successfully connects to CM Server. • Offline — FortiADC Manager failed to connect CM Server. It can happen at the first connection trial or if FortiADC Manager lost the connection. **Note**: FortiADC 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 FortiADC Manager tries to connect with 'State is not None' and CM Server does not have the record of this FortiADC (identified by license). The connection will be rejected by CM Server.

4. Click **Save**.

FortiGSLB Connector

When you enable the FortiGSLB function, FortiADC will connect to the FortiGSLB Cloud server and sync up the configuration to the cloud.

To enable a FortiGSLB Connector:

1. Go to **Security Fabric > Fabric Connectors**.
2. Double-click **FortiGSLB**.
3. Configure the following **Fabric Device Settings**:

Settings	Description
Status	Enable/disable the FortiGSLB function.
Interval	Specify how often FortiADC should try to connect to the FortiGSLB. (Range: 10-1800; default: 15).
Cloud Server URL	Specify the URL of the cloud server.

4. Click **Save**.

After the FortiGSLB connector is enabled, FortiADC will push information to the FortiGSLB Cloud and show the assigned DNS server.

FortiClient EMS Connector

The FortiADC Security Fabric device can link to FortiClient Endpoint Management Server (EMS) for endpoint connectors. Up to three EMS servers can be added to the Security Fabric. EMS settings are synchronized between all Fabric members. Once the FortiADC is authorized as a Fabric device in FortiClient EMS, FortiClient EMS automatically synchronizes ZTNA tags, the EMS CA certificate, and FortiClient endpoint information to the FortiADC.

The FortiClient EMS connector is an integral part of the Zero Trust Network Access (ZTNA) functionality. For more information, see Zero Trust Network Access (ZTNA) on page 854 and How device identity and trust context is established with FortiClient EMS on page 858.

Requirements:

- FortiClient EMS running version 7.0.3 or later
- FortiClient running 7.0.1 or later
- FortiADC hardware, VM, or cloud platform that support FortiClient EMS.



FortiClient EMS is supported in most FortiADC platforms but not all of them. The following lists the hardware models, cloud platforms, and VM environments that support FortiClient EMS.

Hardware models:

- FAD-120F, FAD-220F, FAD-300F, FAD-400F, FAD-1200F, FAD-2200F, FAD-4200F, FAD-5000F

Cloud platforms with BYOL (PAYG FortiADC does not support FortiClient EMS):

- AWS (Amazon Web Services), Microsoft Azure, GCP (Google Cloud Platform), OCI (Oracle Cloud Infrastructure), Alibaba Cloud

VM environments:

- VMware, Microsoft Hyper-V, KVM, Citrix Xen, Xen Project Hypervisor
  Note: The most recent certificate embedded license is required. If your license was issued prior to April 2021, please obtain a new certificate embedded license for your VM through Fortinet Customer Service & Support.

- Read-Write access permission for FortiADC Systems settings

To create and configure a FortiClient EMS connector:

1. Go to **Security Fabric > Fabric Connectors**.
2. Click **Create New**.
3. Under **Core Network Security**, click **FortiClient EMS** to display the configuration editor.
4. Configure the following **FortiClient EMS** settings:

Setting	Description
Name	Specify the FortiClient Enterprise Management Server (EMS) name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces.
IP/Domain name	Specify the server IPv4 address or the domain name of the FortiClient EMS FQDN. For example: 192.0.2.1

Setting	Description
HTTPS Port	Specify the FortiClient EMS HTTPS access port number. Range: 1-65535, default: 443

5. Click **Save**.
   The **Verify EMS server certificate** dialog displays the following message:
   In order for the FortiClient EMS and FortiADC to communicate, the following certificate provided by the FortiClient EMS must be reviewed for correctness, and accepted if deemed valid.
   Do you wish to Accept the certificate as detailed below?
6. After you have verified the EMS server certificate information displayed, click **OK** to accept the EMS server certificate.
   The **Verify completed** dialog displays the following message:
   This FortiADC is not authorized on FortiClient EMS yet. Please let FortiClient EMS to authorize it.
   Note: This message will only appear if the FortiADC device has not yet been authorized as a Fabric Device through FortiClient EMS.
7. Click **OK**.

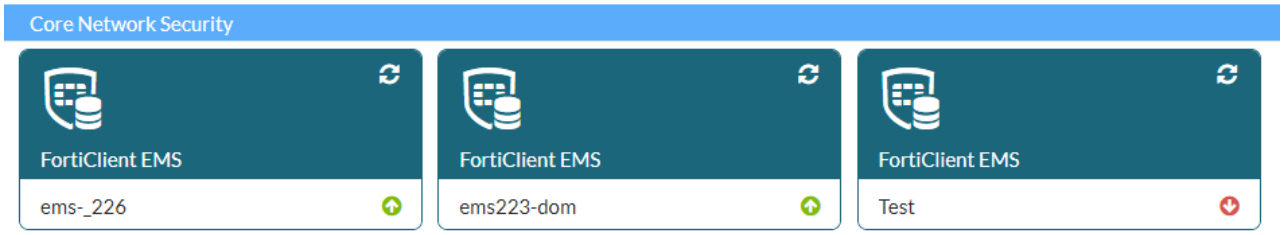
The newly created FortiClient EMS connector is added to the **Security Fabric > Fabric Connectors** page, under the **Core Network Security** section. The FortiClient EMS connector will not be connected until the FortiADC has been authorized as a Fabric Device in FortiClient EMS.

To authorize the FortiADC as a Fabric Device in FortiClient EMS:

1. Login to FortiClient EMS.
2. From the FortiClient EMS landing page, the **Fabric Device Authorization Requests** pop-up displays the Serial Number and IP information of the FortiADC device. Click **Authorize**.
3. Alternatively, you can go to **Administration > Fabric Devices** and select the Fabric device you want to authorize.

To check and troubleshoot the FortiClient EMS connector connection:

1. Go to **Security Fabric > Fabric Connectors**.
2. Under the **Core Network Security** section, locate the FortiClient EMS connector configurations.



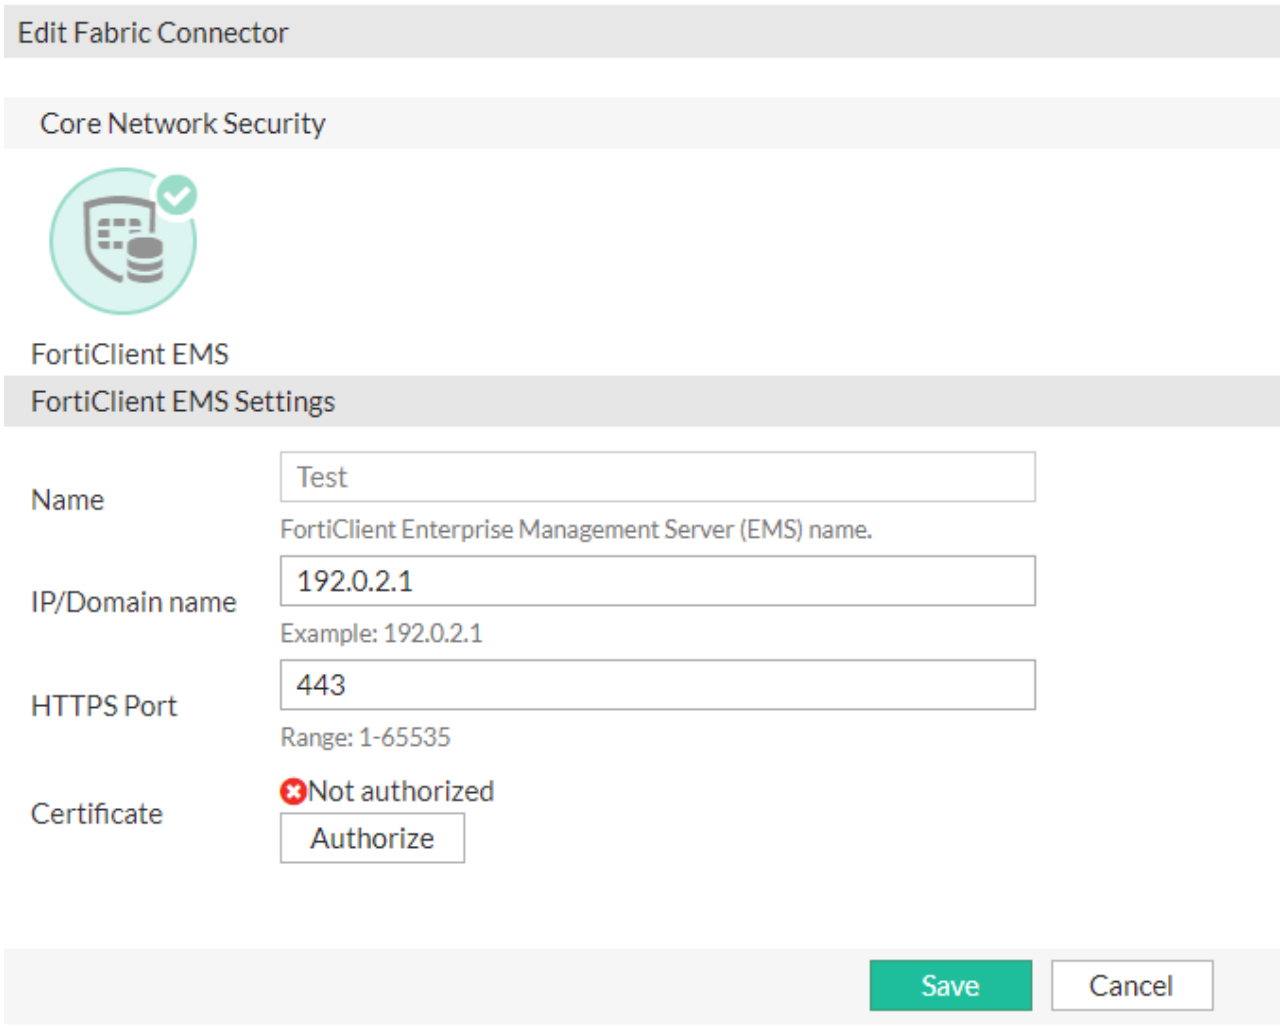
3. The and icons indicate whether FortiClient EMS has successfully authorized the FortiADC Fabric Device for the corresponding FortiClient EMS connector. Hover over the FortiClient EMS connector to see the status details. The table below lists the possible connection statuses for the FortiClient EMS connector.

Icon	EMS Status	Description
	Connected	The FortiADC has been successfully authorized as a Fabric Device through FortiClient EMS.

Icon	EMS Status	Description
	Cert unauthorized	FortiADC does not verify the EMS server's CA certificate. You can edit the FortiClient EMS connector configuration and restart the verification to accept the EMS CA certificate.
	Auth failed	The EMS server does not authorize the FortiADC, indicating the request is either denied or pending authorization. If pending authorization, the status will change to **Connected** once authorization is successful on the EMS server.
	Not reachable	The EMS server was not reachable. Ensure the EMS server IP and system router is properly configured.
	EMS server connection failed	The EMS server connection failed with unknown issue. For example, an incorrect EMS server port may cause this issue.
	No compatible	The EMS server connection failed because the server is not compatible with FortiADC.
	Not sent	The EMS domain name cannot resolve. Ensure proper configuration for the DNS server setting, domain name, and system router.

If the status is not Connected, edit the FortiClient EMS connector accordingly to troubleshoot the connection issue.

4. Locate the newly created FortiClient EMS connector, click the FortiClient EMS connector configuration then click **Edit**, or double click the configuration object to display the configuration editor.

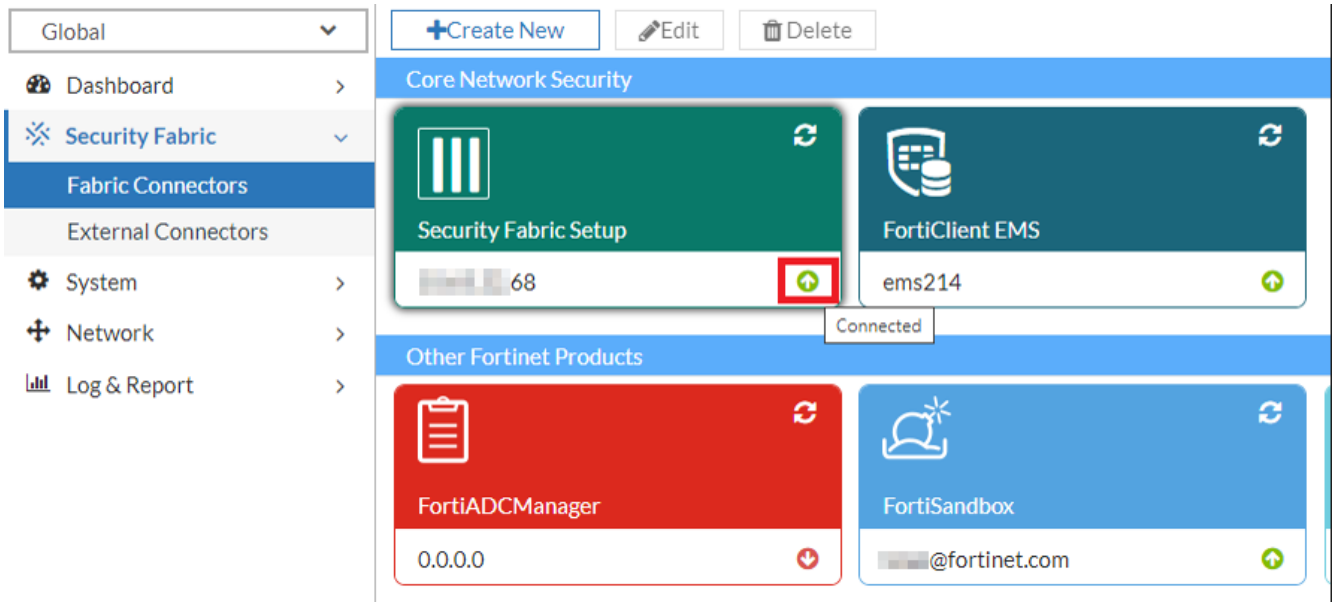


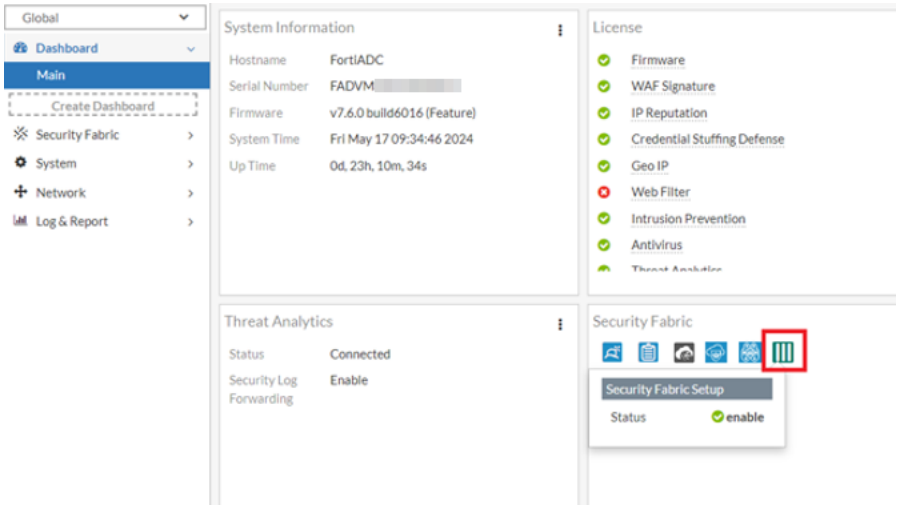
5. Edit the configuration to troubleshoot the connection issue then click **Authorize** to restart the verification to accept the EMS CA certificate.
   A request is resent to the FortiClient EMS to authorize the FortiADC as a Fabric Device in FortiClient EMS. The FortiClient EMS connector will not be connected until the FortiADC has been authorized as a Fabric Device in FortiClient EMS.

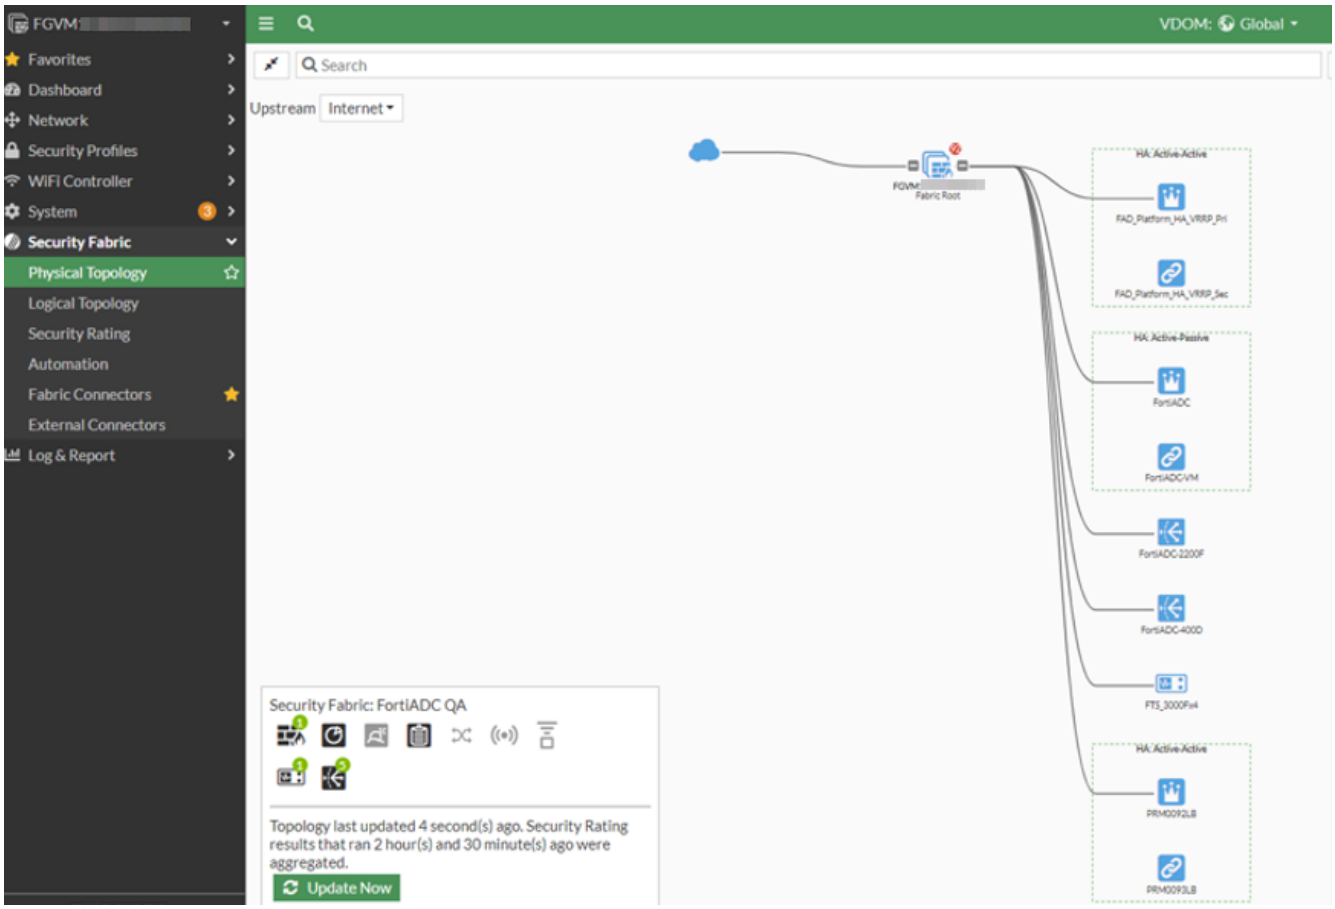
FortiGate Security Fabric Connector

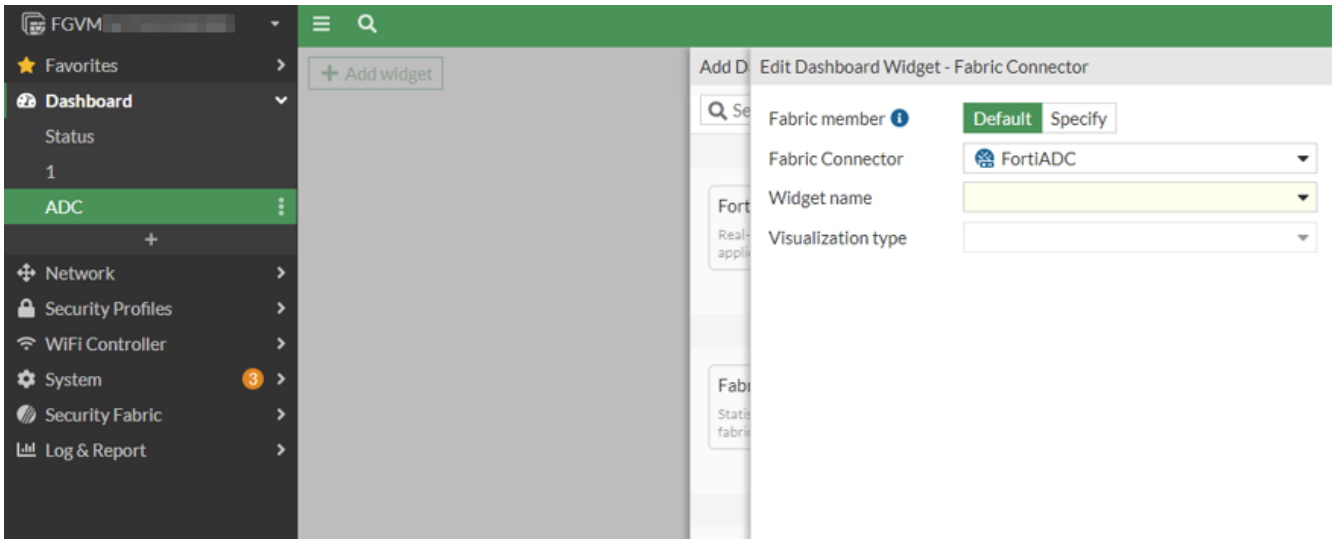
The FortiGate Security Fabric Setup connector facilitates the seamless integration of FortiADC into the FortiGate Security Fabric through the root FortiGate unit. Upon successful integration, the FortiADC device becomes a visible and interactive element within the Security Fabric topology views on the FortiGate interface. Furthermore, administrators can leverage the Fabric Device widget on the FortiGate dashboard to monitor FortiADC statuses, ensuring a comprehensive and unified security posture across the network.

After configuring the Security Fabric Setup connector and successfully connecting to the FortiGate Security Fabric, from the FortiGate, you will be able to view the FortiADC on topology pages, and create a dashboard Fabric Device widget to view FortiADC data.







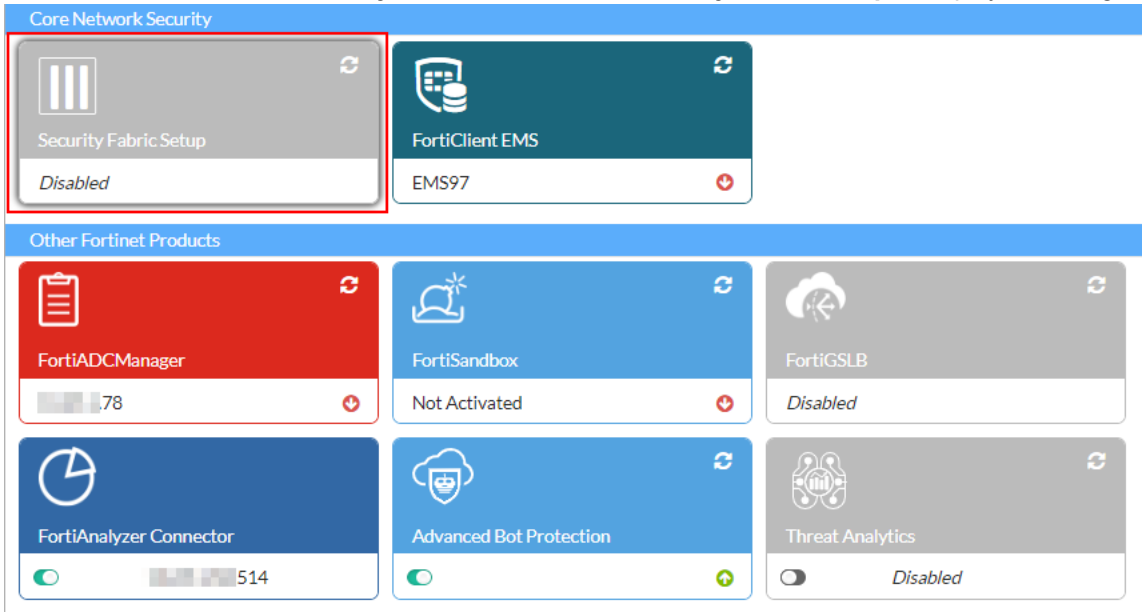


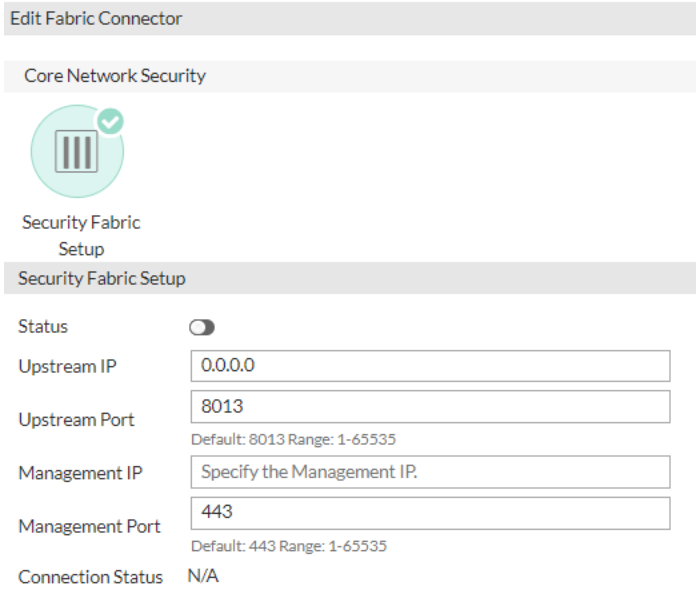
Before you begin:

- You must have Read-Write access permission for FortiADC Systems settings

To configure a FortiGate Security Fabric Setup connector:

1. Go to **Security Fabric > Fabric Connectors**.
2. Under the **Core Network Security** section, double-click **Security Fabric Setup** to display the configuration editor.





3. Configure the following settings:

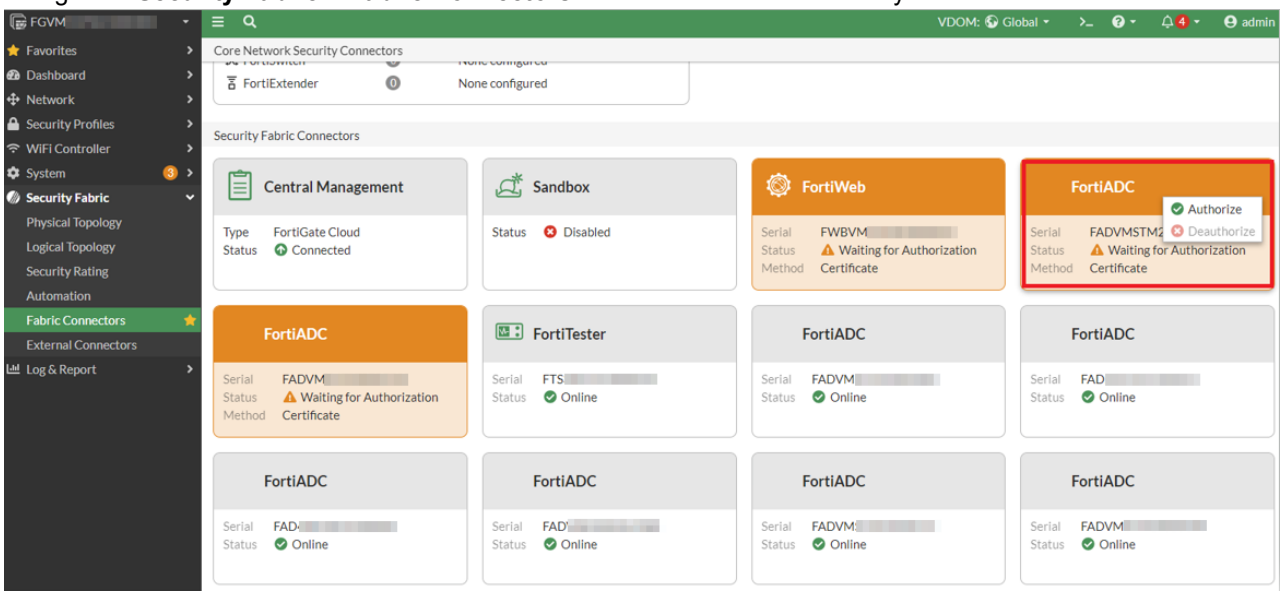
Setting	Description
Status	Enable the FortiGate Security Fabric Setup connector. This is disabled by default.
Upstream IP	Specify the Upstream FortiGate IP address.
Upstream Port	Specify the Upstream port for the FortiGate. The default port is 8013. The valid

Setting	Description
	range is 1-65535.
Management IP	Specify the Management IP address of this FortiADC appliance to join the Security Fabric.
Management Port	Specify the Management port of this FortiADC appliance to join the Security Fabric. The default port is 443. The valid range is 1-65535.

4. Click **Save**.
   The newly enabled Security Fabric Setup connector will display a Connection Status of **Pending Authorization**. The Security Fabric Setup connector will not be connected until the FortiADC has been authorized as a Fabric Device in FortiGate.

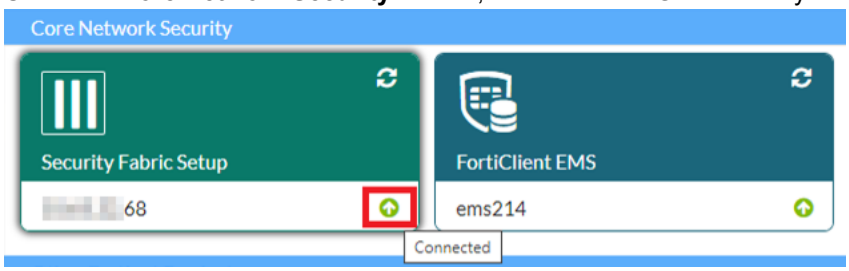
To authorize the FortiADC as a Fabric Device in FortiGate:

1. Login to FortiGate.
2. Navigate to **Security Fabric > Fabric Connectors** and select the Fabric device you want to authorize.



To check and troubleshoot the FortiGate Security Fabric Setup connector connection:

1. Go to **Security Fabric > Fabric Connectors**.
2. Under the **Core Network Security** section, locate the FortiGate Security Fabric Setup connector configuration.



3. The and icons indicate whether FortiGate has successfully authorized the FortiADC Fabric Device. Hover over the FortiGate Security Fabric Setup connector to see the status details. The table below lists the possible connection statuses for the FortiGate Security Fabric Setup connector.

Icon	EMS Status	Description
	Connected	The FortiADC has been successfully authorized as a Fabric Device through FortiGate.
	Pending Authorization	The FortiADC is waiting to be authorized as a Fabric Device in the FortiGate Security Fabric.
	Connecting	The FortiADC has successfully connected to the FortiGate server, but has not yet been authorized as a Fabric Device in the FortiGate Security Fabric.
	Authorization Rejected	The FortiADC has been Deauthorized as a Fabric Device in the FortiGate Security Fabric.
	Not Available	The FortiGate server was not reachable. Ensure the Upstream IP and system router is properly configured.
	Not Connected	The connection failed with unknown issue.

If the status is not Connected, edit the FortiGate Security Fabric Setup connector accordingly to troubleshoot the connection issue.

To further troubleshoot the FortiGate Security Fabric Setup connector issues, you can use the `diagnose debug module csfd all` CLI command to view debug logs.

External connectors

Cloud SDN connectors provide integration and orchestration of Fortinet products with public and private cloud solutions. In a typical cloud environment, resources are dynamic and often provisioned and scaled on-demand. By using an SDN connector, you can ensure that changes to cloud environment attributes are automatically updated in the Security Fabric.

To protect the East-West or North-South traffic in these environments, the FortiADC uses the SDN connector to sync the dynamic addresses that these volatile environments use. You can then configure the dynamic address objects as sources or destinations for firewall policies. When you make changes to cloud environment resources, such as moving them to a new location or assigning different IP addresses to them, you do not need to modify the policy in FortiADC, as the SDN connector syncs changes to the cloud address objects.

The following external connectors are available in the Security Fabric.

Public SDN:

- Amazon Web Services (AWS) Connector on page 95
- Oracle Cloud Infrastructure (OCI) Connector on page 98

Private SDN:

- Kubernetes Connector on page 100
- Splunk Connector on page 104
- SAP Connector on page 105

Authentication:

- FortiAuthenticator (FAC) Connector

Threat Feeds:

- IP Address Connector on page 106



If VDOMs are enabled, SDN and Threat Feeds connectors are in the global settings, and Authentication connectors are per VDOM.

Amazon Web Services (AWS) Connector

When you create an Amazon Web Services (AWS) connector, you are authorizing FortiADC to periodically (every 30 seconds by default) get information from AWS instances and dynamically populate it in the server pool configuration.

Before you begin:

- Ensure the system time is synchronized between AWS EKS and FortiADC. You can enable NTP on FortiADC to correct the time clock. For details, see Configuring system time on page 150.
- If you want to access AWS EKS objects through the Metadata IAM role for the FortiADC EC2 instance, you must have permissions enabled. For details, see Accessing AWS EKS objects through Metadata IAM role on page 96.

To create an AWS Connector:

1. Go to **Security Fabric > External Connectors**.
2. Click **Create New**.
3. Under **Public SDN**, select **Amazon Web Services (AWS)** to display the configuration editor.

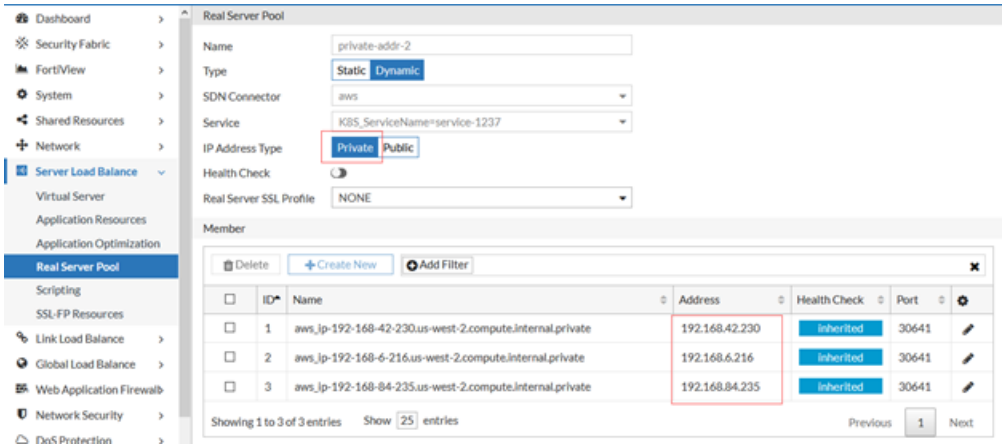
4. Configure the following settings:

Setting	Description
Name	Type a name for the external connector object.
Status	Toggle on to enable the external connector object. Toggle off to disable the external connector object.
Update Interval (s)	Specify the update interval for the connector to get AWS objects and dynamically populate the information in the server pool configuration.
Access Key ID	Specify the access key ID.
Secret Access Key	Specify the secret access key.
Region Name	Select the region where your instances are deployed from the drop-down menu.
Use Metadata IAM	When FortiADC is deployed on AWS, you can assign IAM role for it to access EC2 instances and EKS objects.

5. Click **Save**.

After the connector is created, you can select this connector when creating a server pool. FortiADC will then get the IP addresses of the instances from AWS and dynamically populate the objects in the server pool configuration.

You can use the **IP Address Type** option to get the private address or public address of the instance. This option is supported only when the FortiADC is deployed on AWS.



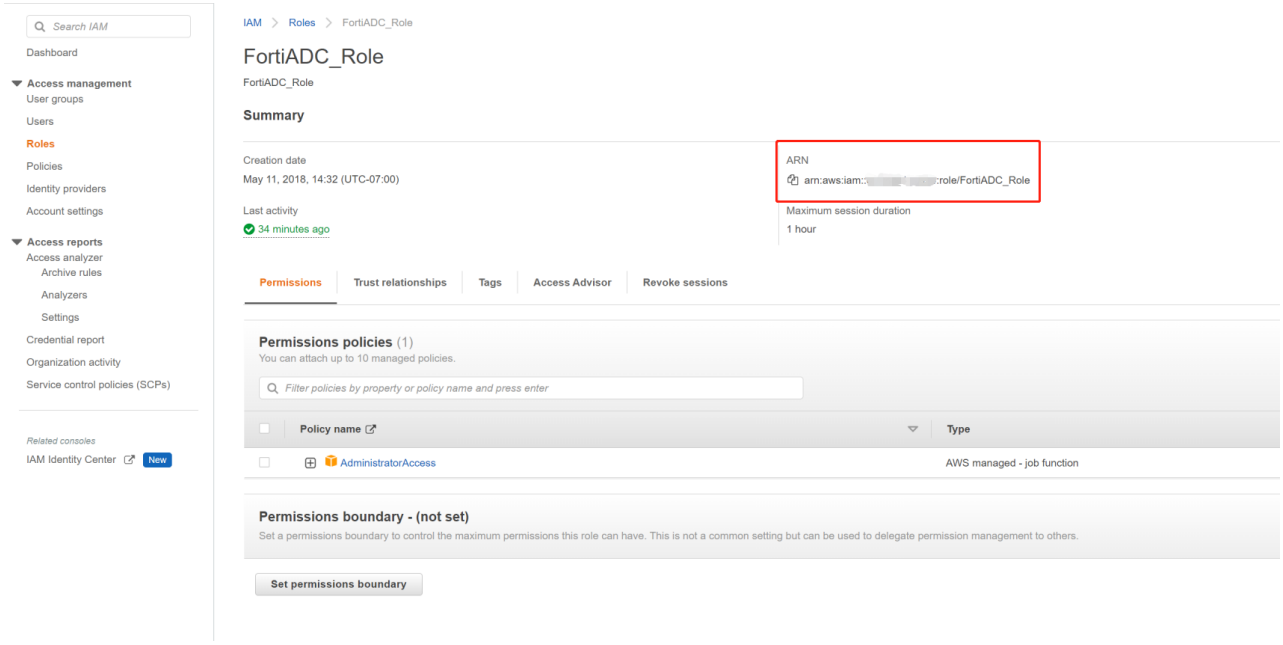
Accessing AWS EKS objects through Metadata IAM role

If you want to use the Metadata IAM role for the FortiADC EC2 instance to access the AWS EKS objects, follow the steps below to enable permission before you configure the AWS SDN connector.

Note: If you have already configured the AWS SDN connector with Metadata IAM enabled, it must be disabled and re-enabled in order for the below steps to take effect in the new configuration.

1. Go to the **AWS Dashboard** and navigate to **IAM > Access management > Role**.
2. Create a role and grant the role with the **AdministratorAccess** permission policy to allow the FortiADC EC2 instance to call AWS EKS services on your behalf.

3. Record the Role **ARN** information to be used later.



4. Assign the newly created IAM role to the FortiADC EC2 instance.
5. Add role-based access control (RBAC) access to the IAM role using the `aws-auth ConfigMap`.
   a. Check the current `aws-auth ConfigMap` and copy the **roleARN** information.
   The roleARN may appear differently depending on the way the EKS cluster node group is created. In this context, the EKS cluster node group is created with the Amazon EKS vended AWS CloudFormation templates, which makes the **NodeInstanceRole** the roleARN.

   ```
   kubectl describe configmap -n kube-system aws-auth
   ```

   ```
   ubuntu@[redacted]:~$ kubectl describe configmap -n kube-system aws-auth
   Name:         aws-auth
   Namespace:    kube-system
   Labels:       <none>
   Annotations:  <none>

   Data
   ====
   mapRoles:
   ----
   - groups:
     - system:bootstrappers
     - system:nodes
     rolearn: arn:aws:iam::[redacted]:role/eksctl-qa-cluster-nodegroup-ng-16-NodeInstanceRole-[redacted]
     username: system:node:{{EC2PrivateDNSName}}


   BinaryData
   ====

   Events:  <none>
   ```

   b. Download the configuration map template.

   ```
   curl -o aws-auth-cm.yaml https://amazon-eks.s3.us-west-
   2.amazonaws.com/cloudformation/2020-10-29/aws-auth-cm.yaml
   ```

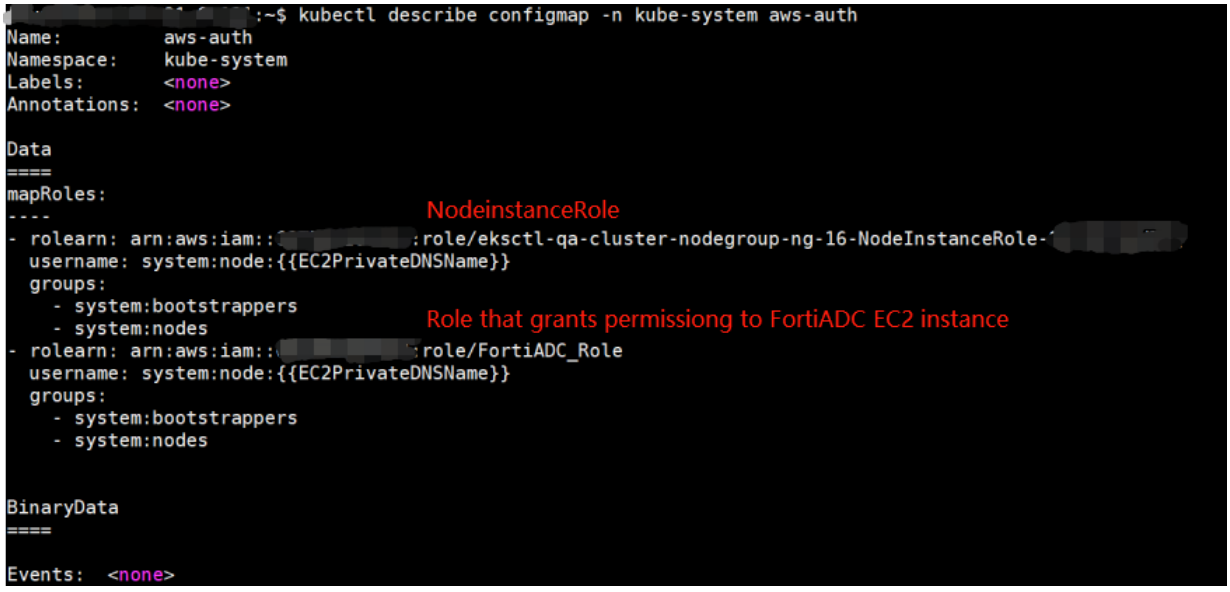
   c. Create a new `aws-auth` configuration by adding the IAM role created previously and the NodeInstanceRole into `aws-auth-cm.yaml`.

```
        apiVersion: v1
       kind: ConfigMap
       metadata:
         name: aws-auth
         namespace: kube-system
       data:
         mapRoles: |
           - rolearn: arn:aws:iam::xxxx:role/eksctl-qa-cluster-nodegroup-ng-16-
NodeInstanceRole-yyyy
             username: system:node:{{EC2PrivateDNSName}}
             groups:
               - system:bootstrappers
               - system:nodes
           - rolearn: arn:aws:iam::xxxx:role/FortiADC-role
             username: system:node:{{EC2PrivateDNSName}}
             groups:
               - system:bootstrappers
               - system:nodes
```

d. Apply the `aws-auth-cm.yaml` by `kubectl apply` command.

```
kubectl apply -f aws-auth-cm.yaml
```

e. Use the `kubectl describe` command to check the `aws-auth ConfigMap`.



For more information on AWS IAM user and role access to the EKS cluster, refer to AWS official documentation.

Oracle Cloud Infrastructure (OCI) Connector

When you create an Oracle Cloud Infrastructure (OCI) connector, you are authorizing FortiADC to periodically (default 30s) get information from OCI instances and dynamically populate it in the server pool configuration.

To create an OCI Connector:

1. Go to **Security Fabric > External Connectors**.
2. Click **Create New**.
3. Under **Public SDN**, select **Oracle Cloud Infrastructure (OCI)** to display the configuration editor.

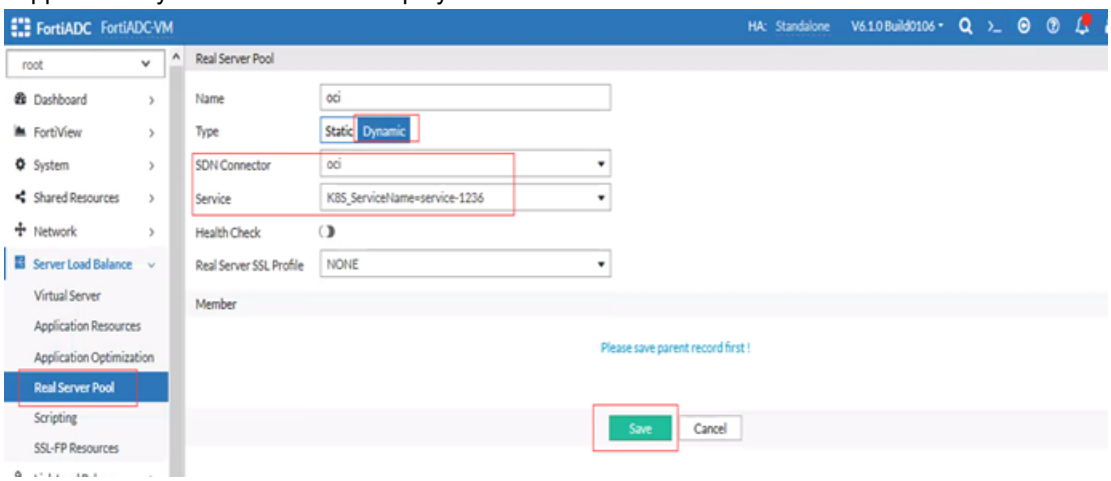
4. Configure the following settings:

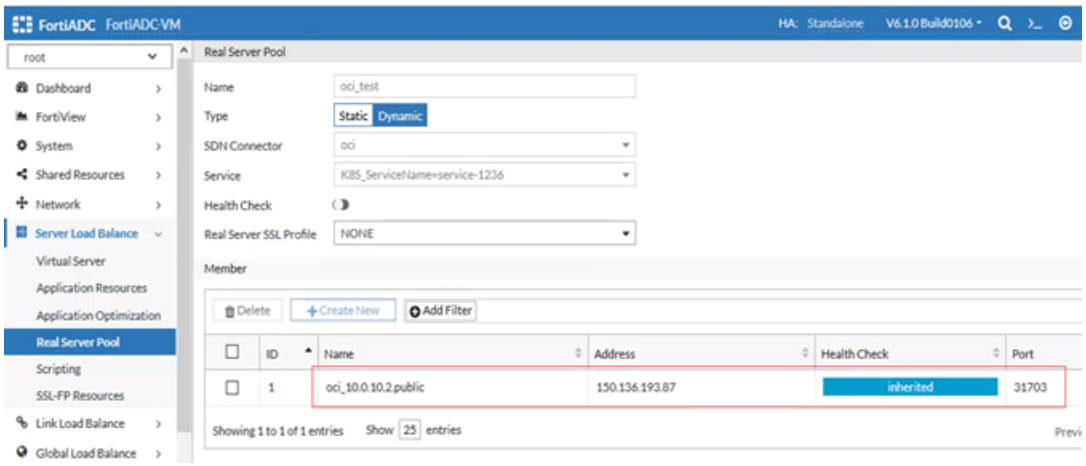
Setting	Description
Name	Type a name for the external connector object.
Status	Toggle on to enable the external connector object. Toggle off to disable the external connector object.
Update Interval (s)	Specify the update interval for the connector to get OCI objects and dynamically populates the information in the server pool configuration.
OCI region type	Specify the OCI region type.
OCI region	Specify the OCI region where your compute instances are located.
User ID	The user's OCID.
Tenant ID	The tenancy's OCID. Refer to this article on how to get the user's OCID and tenancy's OCID: https://docs.oracle.com/en-us/iaas/Content/API/Concepts/apisigningkey.htm#five
Compartment ID	The OCID of the Compartment in which your compute instances are deployed.
Certificate	The certificate that FortiADC uses to build connections with OCI instances. You can select an existing one or create a new one. Refer to Manage Certificates.
OCI HA status	Enable this option so that the system will use this connector to get the HA members' information if this FortiADC is deployed in HA mode and is the primary node.
Use Metadata IAM	Enable this option to assign IAM role for FortiADC to access OCI objects. **Note:** It must be enabled if the connector is used for OCI HA.

5. Click **Save**.

After the connector is created, you can select this connector when creating a server pool. FortiADC will then get the IP addresses of the instances from OCI and dynamically populate the objects in the server pool configuration. This step is not required if you have enabled **OCI HA status** because in this case the connector will be used by the system to get the information of the HA members instead of the server pool members.

You can use the **IP Address Type** option to get the private address or public address of the instances. This option is supported only when FortiADC is deployed on OCI.





Kubernetes Connector

When you create a fabric connector for Kubernetes, you are specifying how FortiADC can communicate with Kubernetes.

FortiADC will be authenticated to periodically (default 30s) get Kubernetes objects (services, nodes) and dynamically populates and updates the related objects, including pool member and real server in its server pool configuration.

Requirements:

- The Kubernetes service is required to be exposed with NodePort type.

To obtain the IP address, port, and secret token in Kubernetes:

When configuring the Kubernetes connector in FortiADC, you must provide the IP address and port that the Kubernetes deployment is running on.

1. On the primary node of your Kubernetes cluster, run `kubectl config view` to get the IP address.
   The following is an example. Take note of the IP address.

```
root@master-node:~#
root@master-node:~# kubectl config view
apiVersion: v1
clusters:
- cluster:
    certificate-authority-data: DATA+OMITTED
    server: https://10.106.170.70:6443
  name: kubernetes
contexts:
- context:
    cluster: kubernetes
    user: kubernetes-admin
  name: kubernetes-admin@kubernetes
current-context: kubernetes-admin@kubernetes
kind: Config
preferences: {}
users:
- name: kubernetes-admin
  user:
    client-certificate-data: REDACTED
    client-key-data: REDACTED
root@master-node:~#
```

2. Run `kubectl get services` to get the port number. FortiADC only supports "NodePort" service type.
   The following is an example:

```
root@master-node:~# kubectl get services
NAME           TYPE        CLUSTER-IP      EXTERNAL-IP   PORT(S)          AGE
kubernetes     ClusterIP   [redacted]      <none>        443/TCP          26h
service-1236   NodePort    [redacted]6     <none>        1236:32663/TCP   50s
```

   Take note of the port number of this service, i.e. service-1236 in the above example.

3. Create a cluster role to grant the FortiADC permission to perform operations and retrieve objects.
   a. Run `cat > <filename>.yaml` to create a yaml file specifying the cluster role.
      For example, running `cat > access_clusterrole.yaml` will create the file "access_clusterrole.yaml".
      Then, type the following to insert it in the file. In this example, the role is named as `psn-reader`. You can give it other names as you desire. Remember to Type `Ctrl-d` at the end to save the file.

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
   # "namespace" omitted since ClusterRoles are not namespaced
   name: psn-reader
rules:
- apiGroups: [""]
   resources: ["pods", "services", "nodes"]
   verbs: ["get", "watch", "list"]
```

   b. Run `cat > <filename>.yaml` to create a yaml file, then insert the following to attach the cluster role to a service account.
      In the following example, the file "cluster_role_bind.yaml" is created, and the role "psn-reader" is attached to the service account "default" for it to read pods, node, or services in default namespace.
      If you want to attach the role to a new service account, use `kubectl create serviceaccount <Service_account_name>` to create one, then attach the role to it.
      Remember to Type `Ctrl-d` at the end to save the file.

```
~# cat > cluster_role_bind.yaml
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
   name: read-psn-global
subjects:
   - kind: ServiceAccount
      name: default #name is case sensitive
      namespace: default
roleRef:
      kind: ClusterRole
      name: psn-reader
      apiGroup: rbac.authorization.k8s.io
```

c. Run `kubectl apply -f access_clusterrole.yaml` to execute the configurations in this file.

d. Run `kubectl apply -f cluster_role_bind.yaml` to execute the configurations in this file.

4. Get secret token.



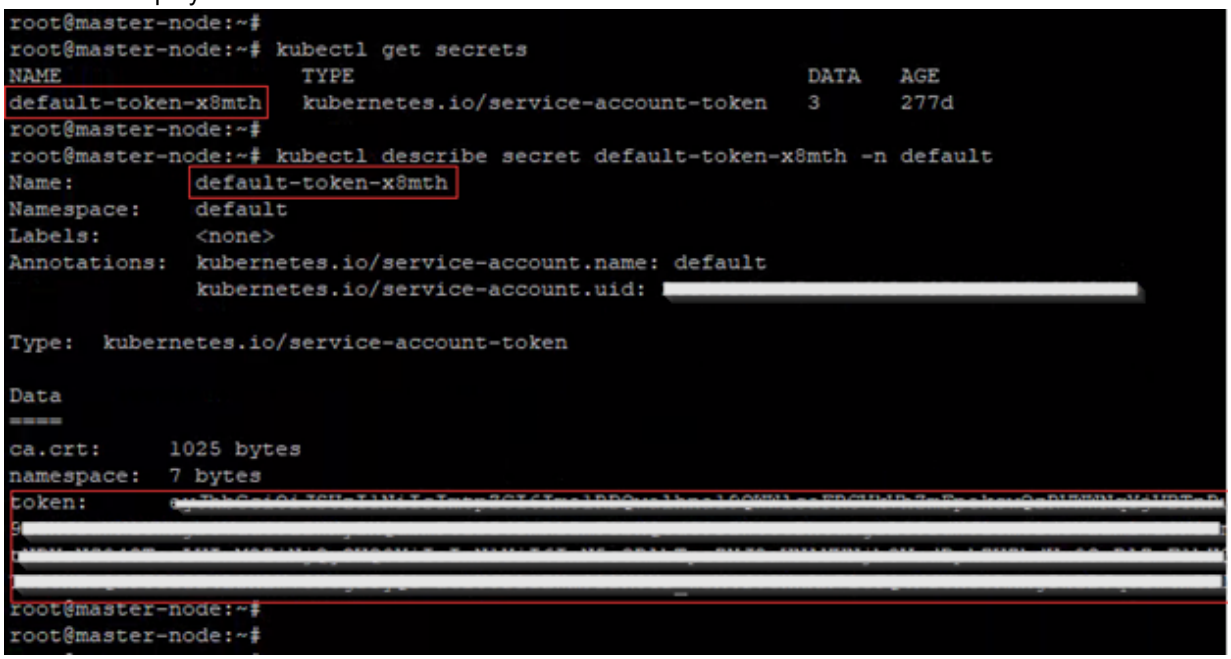
Versions of Kubernetes before v1.22 automatically created long term credentials for accessing the Kubernetes API. This older mechanism was based on creating token Secrets that could then be mounted into running Pods. In more recent versions, including Kubernetes v1.28, API credentials are obtained directly by using the TokenRequest API, and are mounted into Pods using a projected volume. The tokens obtained using this method have bounded lifetimes, and are automatically invalidated when the Pod they are mounted into is deleted.

You can still manually create a service account token Secret; for example, if you need a token that never expires. However, using the TokenRequest subresource to obtain a token to access the API is recommended instead.

a. Run `kubectl get secrets` to view the secrets.

b. Run `kubectl describe secrets <secret_token_name> -n <service_account_name>` to view the secret token. Take note of the token.
In the following example, the information of the secret token "default-token-x8mth" stored in "default" service

account is displayed.

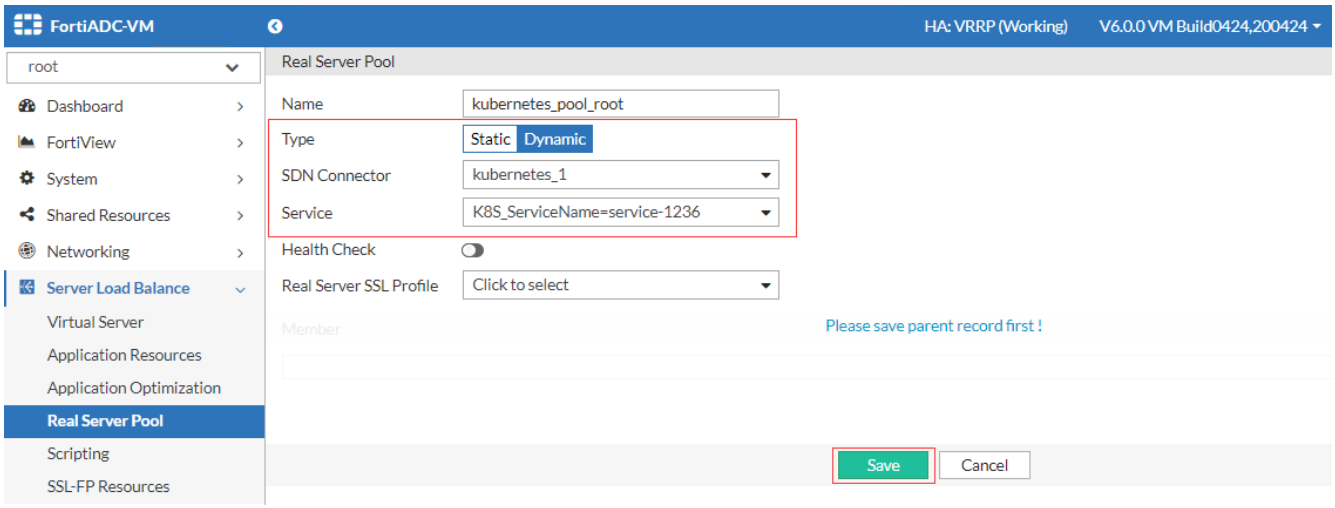


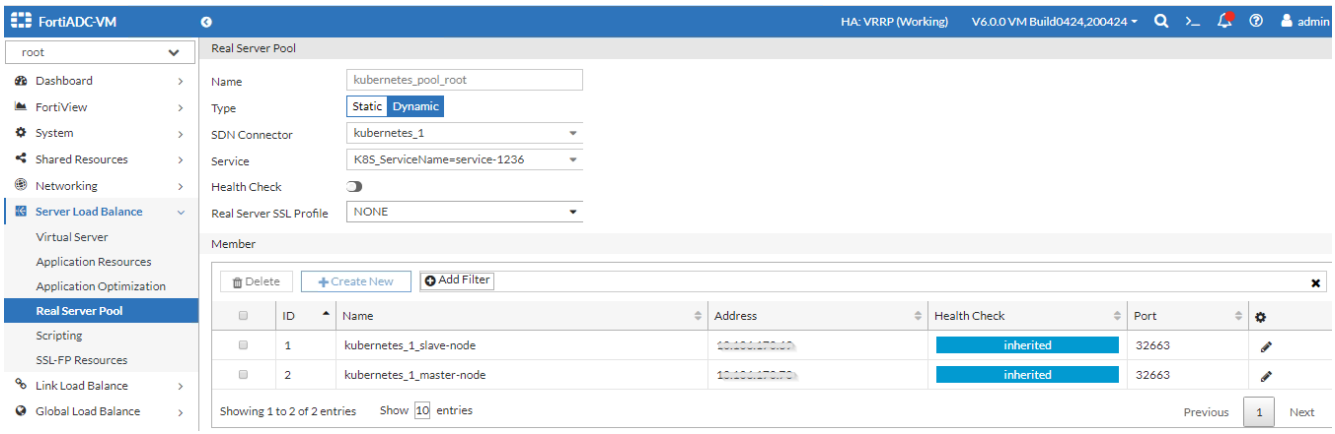
To create a Kubernetes Connector:

1. Go to **Security Fabric > External Connectors**.
2. Click **Create New**.
3. Under **Private SDN**, select **Kubernetes**. The Kubernetes screen is displayed.
4. Configure the following options, and then click Save. You will be required to provide the IP address, port, and the secret token you have obtained in the above section: To obtain the IP address, port, and secret token in Kubernetes:

Name	Type a name for the external connector object.
Status	Toggle on to enable the external connector object. Toggle off to disable the external connector object.
Update Interval (s)	Specify the update interval for the connector to get Kubernetes objects and dynamically updates the IP addresses.
IP	Type the IP address of the Kubernetes API server.
Port	Specify the port that FortiADC uses to communicate with the Kubernetes API server.
Secret Token	Specify the secret token.

After the connector is created, you can select this connector when creating a server pool. FortiADC will then get the IP addresses of the real servers from the Kubernetes deployment and dynamically populate the objects in the server pool configuration.





Splunk Connector

When you create a connector for Splunk, you are specifying how FortiADC can communicate with Splunk for pushing logs to Splunk.

FortiADC will connect to Splunk by UDP, TCP or TCP SSL depending on Splunk connector setting.

Requirements:

- The Splunk service is required to be exposed on External IP.

To create a Splunk Connector:

1. Go to Security Fabric > External Connectors.
2. Click **Create New**.
3. Under Private SDN, select Splunk. The Splunk screen is displayed.
4. Configure the following options, and then click Save.

Status	Toggle on to enable the external connector object. Toggle off to disable the external connector object.

Address	Type the IP address of the Splunk Log server.
Port	Specify the port that FortiADC uses to communicate with the log server.
Proto	Select the protocol used for log transfer.
Log Level	Select the severity level of the logs. All the exported logs will be attached with the selected severity level.
CSV	Enable to export the logs in .csv file.
Facility	Select the source facility of the logs. We only support the local use facilities which are not reserved and are available for general use.
Event	Enable to export Event logs.
Traffic	Enable to export Traffic logs.
Security	Enable to export Security logs.

After the connector is created, FortiADC will push the logs to Splunk server. The above configurations are also available in Log&Report > Log Setting > Syslog Server.

SAP Connector

When you create an SAP connector, you are authorizing FortiADC to periodically (default 30s) get information of SAP instances and dynamically populates it in server pool configuration.

To create an SAP Connector:

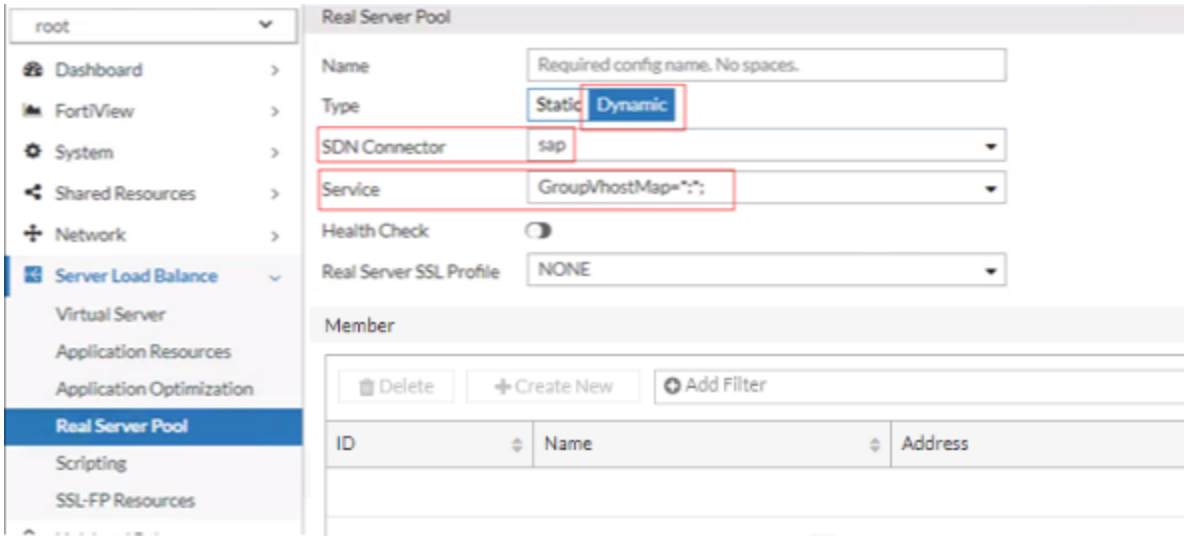
1. Go to Security Fabric > External Connectors.
2. Click **Create New**.
3. Under Public SDN, select SAP. The SAP screen is displayed.
4. Configure the following options, and then click Save.

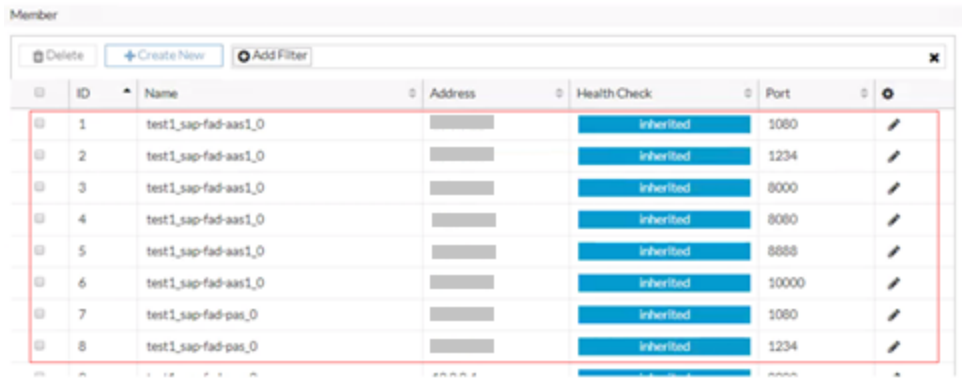
Name	Type a name for the external connector object.
Status	Toggle on to enable the external connector object. Toggle off to disable the external connector object.
Update Interval (s)	Specify the update interval for the connector to get SAP objects and dynamically populates the information in the server pool configuration.
IP	Type the IP address, FQDN, or hostname of the SAP server.
SAP DNS Suffix	Enter the DNS name of the SAP system is required. This option is required if hostname is used for the SAP server.
SAP MS HTTP Port	Specify the SAP MS HTTP port that FortiADC uses to communicate with the SAP server.
SAP ICM HTTP Port	Specify the ICM HTTP Port.

SAP SID Admin	Specify the SID admin account that FortiADC uses to access the resources in this account.
SAP Password	Specify the password.

After the connector is created, you can select this connector when creating a server pool. FortiADC will then get the IP addresses of the compute instances from SAP and dynamically populates the objects in server pool configuration, as shown in the following screenshots.

You can use the **IP Address Type** option to get the private address or public address of the instance.





IP Address Connector

Creating an IP Address connector allows you to dynamically import an external block list from an HTTP/HTTPS server in the form of a plain text file. Block lists can be used to enforce special security requirements, such as long term policies to always block access to certain websites, or short term requirements to block access to known compromised locations. The lists are dynamically imported, so that any changes are immediately imported by FortiADC.

After you have imported your external block list through the IP Address connector, you can apply the IPs as the source or destination address for IPv4 and IPv6 firewall policies.



- You cannot delete an IP Address connector or modify its status if the external resource is being used in an IPv4 or IPv6 firewall policy.
- Up to 512 external resources can be supported, however, large numbers of external resources may affect system performance.

Requirements:

- The external block list must be accessible from an HTTP/HTTPS server.
- The import file must be in plain text and each line must contain an IP, IP Range, or Subnet in the below formats:

IP/ IP Range/ Subnet	Example
IPv4	192.168.2.100
IPv4 Range	172.200.1.4/16
IPv4 Subnet	172.16.8.1-172.16.8.100
IPv6	2001:0db8::eade:27ff:fe04:9a01
IPv6 Range	2001:0db8::eade:27ff:fe04:9a01/120
IPv6 Subnet	2001:0db8::eade:27ff:fe04:aa01-2001:0db8::eade:27ff:fe04:ab01

- The maximum import file size is 1 MB (which is about 5000 line entries).

To create and configure an IP Address connector:

1. Go to **Security Fabric > External Connectors**.
2. Click **Create New**.
3. Under **Threat Feeds**, click **IP Address** to display the configuration editor.
4. Configure the following **IP Address** settings:

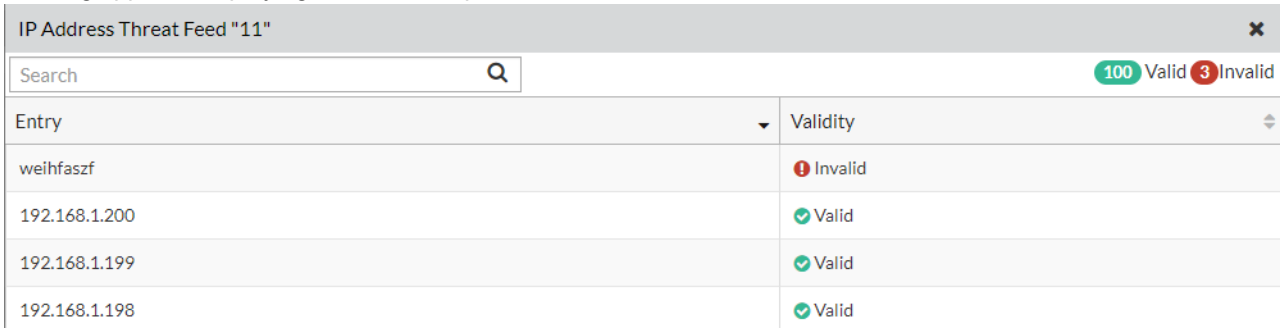
Setting	Description
Name	Specify the name of the IP Address connector. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces.
URI of External Resource	Specify the URI of the HTTP/HTTPS server where the IP address list is stored.
HTTP Basic Authentication	Enable/disable HTTP Basic Authentication to require username and password to access the IP address list.
Username	The **Username** option is available if **HTTP Basic Authentication** is **enable**. Specify the username to be used to access this IP address list.
Password	The **Password** option is available if **HTTP Basic Authentication** is **enable**. Specify the password to be used to access this IP address list.
Refresh Rate	Specify the refresh rate in minutes. (Default: 5. Range: 1-43200 minutes). FortiADC will retrieve the data from the HTTP/HTTPS server periodically according to the refresh rate.

Setting	Description
Comments	Optionally, enter comments about the IP Address connector.
Status	Enable/disable the IP Address connector.

5. Click **Save**.
   The newly created IP Address connector appears on the External Connectors page under Threat Feeds.

To view the external block list IP entries and the resource update status:

1. Go to **Security Fabric > External Connectors**.
2. Under **Thread Feeds**, double-click the **IP Address Connector** to display the configuration editor.
3. From the **Last Update** field, you can see the date of when the resource was last updated.
4. Click **View Entries** to display the IP address list entries.
   A dialog appears displaying the entries imported for the IP Address Connector.



   The imported file has been parsed line by line and marked as valid or invalid based on whether the entry meets format requirements for IP, IP Range, or Subnet.

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.

This section is organized into the following sub-menu topics:

- Logical Topology on page 109
- Server Load Balance on page 112
- Security on page 121
- System on page 139
- Global Load Balance on page 141
- Link Load Balance on page 142
- ZTNA FortiClient endpoint on page 143

Logical Topology

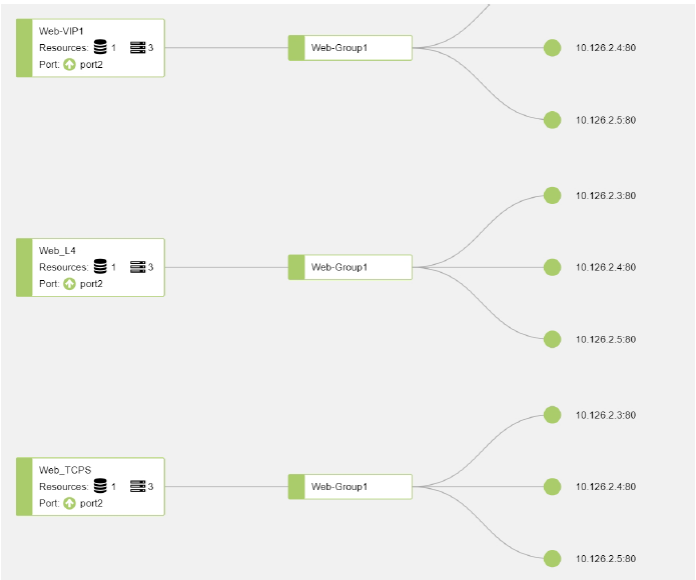
This page displays the logical topology of your FortiADC network structure. It shows your FortiADC appliance or appliances identified by serial number and the real servers connected to it

Note: This page is read-only.

Server Load Balance Logical Topology

The **FortiView > Logical Topology > Server Load Balance** 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 touting, schedule pools, real-server pools, and real-server pool members configured on a virtual server, as illustrated in Server Load Balance Logical topology on page 109.

Server Load Balance 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 113
- Real-server pool—Opens the page with details of the real-server pool. See Real server pool details on page 116
- Real server —Opens the page showing details of the real server. See Real server pool member details on page 118

Link Load Balance Logical Topology

The **FortiView > Logical Topology > Link Load Balance** 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 605
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.

Global Load Balance Logical Topology

The **FortiView > Logical Topology > Global Load Balance** 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 619
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.

Server Load Balance

The **FortiView > Server Load Balance** sub-menu shows server load-balancing configurations on your FortiADC.

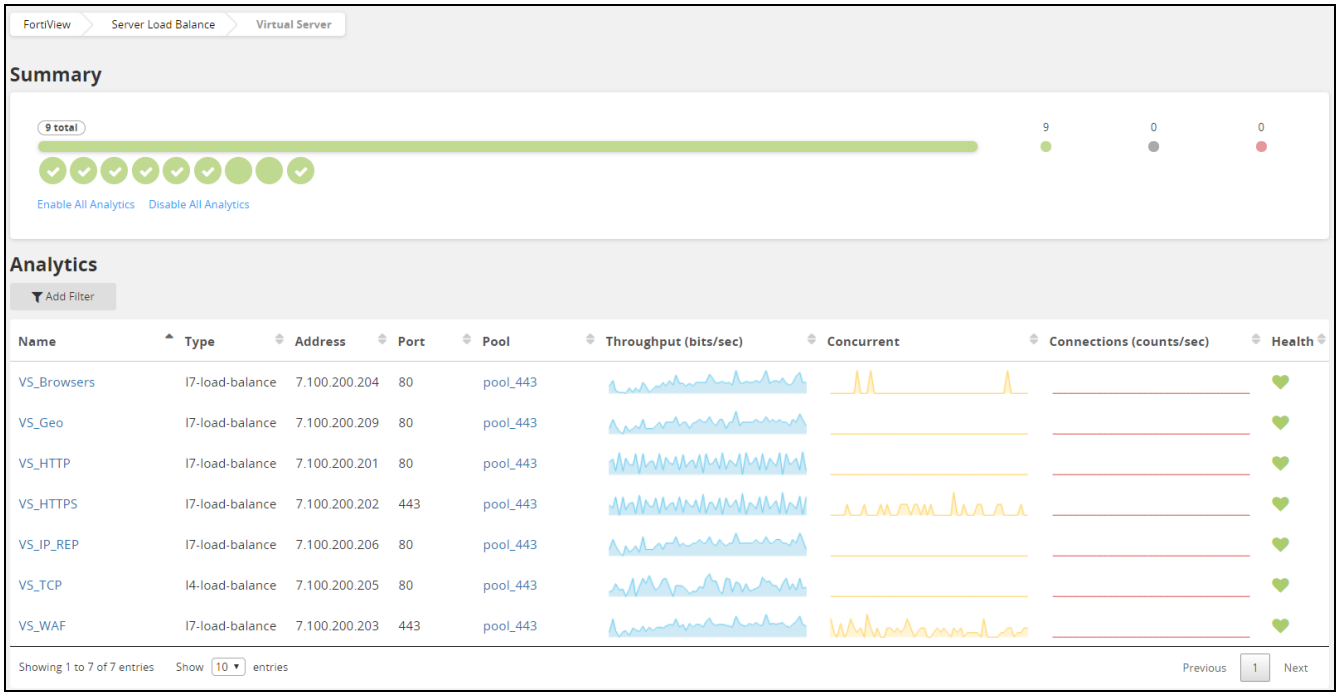
It has the following pages:

- Virtual Server
- Data Analytics
- All Sessions on page 121
- Traffic Logs

Virtual Servers

Under the **FortiView Server Load Balance** sub-menu, the Virtual Server page (Virtual server on page 112) 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 113 describes the information on the FortiView > Server Load Balance > Virtual Server page.

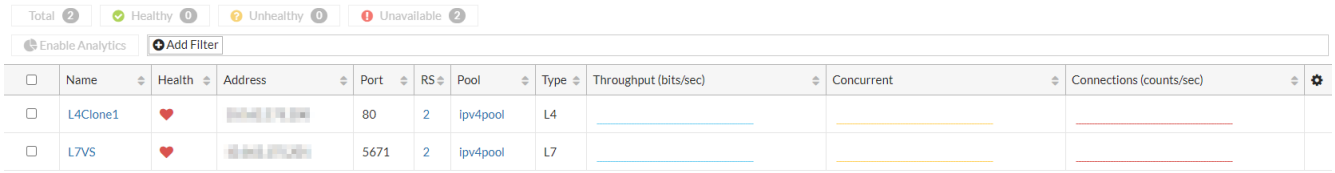
Virtual Server table

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.
Type	The type of virtual servers, which can be one of the following: • l2 = Layer 2 • l4 = Layer 4 • l7 = 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 selected on the **Virtual Server** page under the **FortiView Server Load Balance** sub-menu.

On the **Virtual Server** page, click the virtual server name to view its details.

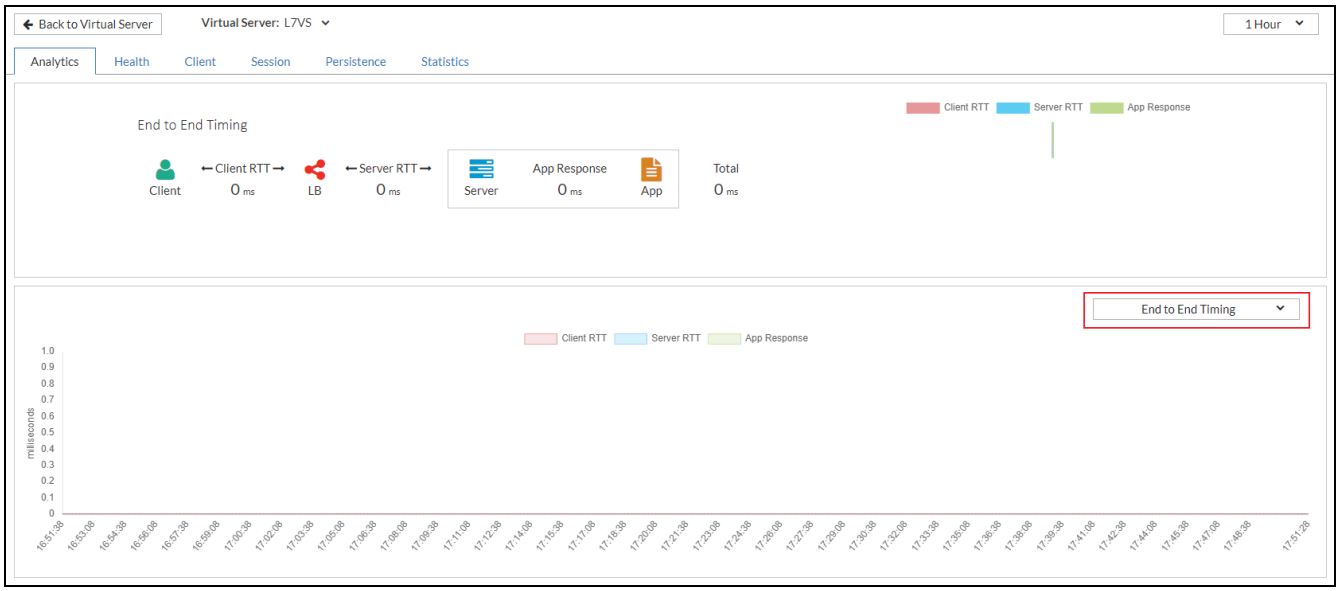


The virtual server details include the following pages:

- 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:

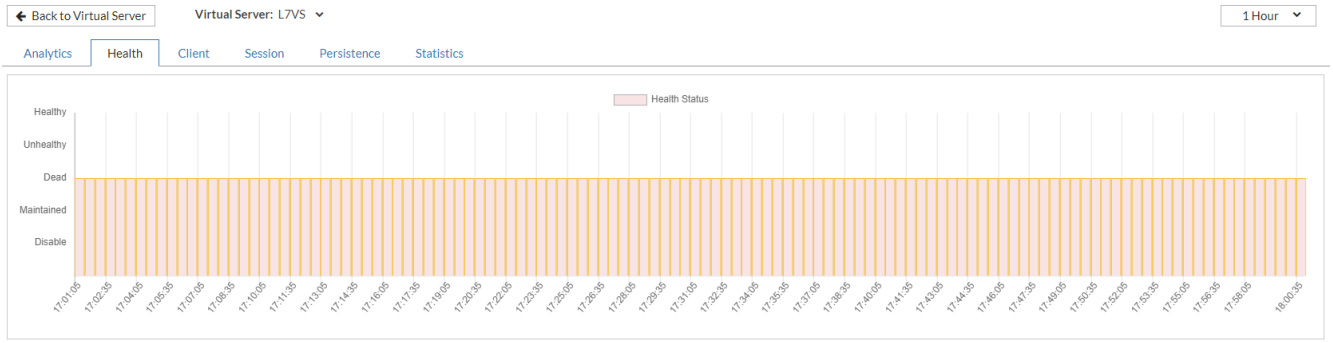
- 1 Hour
- 6 Hour
- 1 Day
- 1 Week
- 1 Month
- 1 Year

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)
  Note: End to End Timing graphs may not deliver precise timings below the 100 ms threshold.
- Throughput
- Concurrent Connections
- Connections per Second
- Request

Health

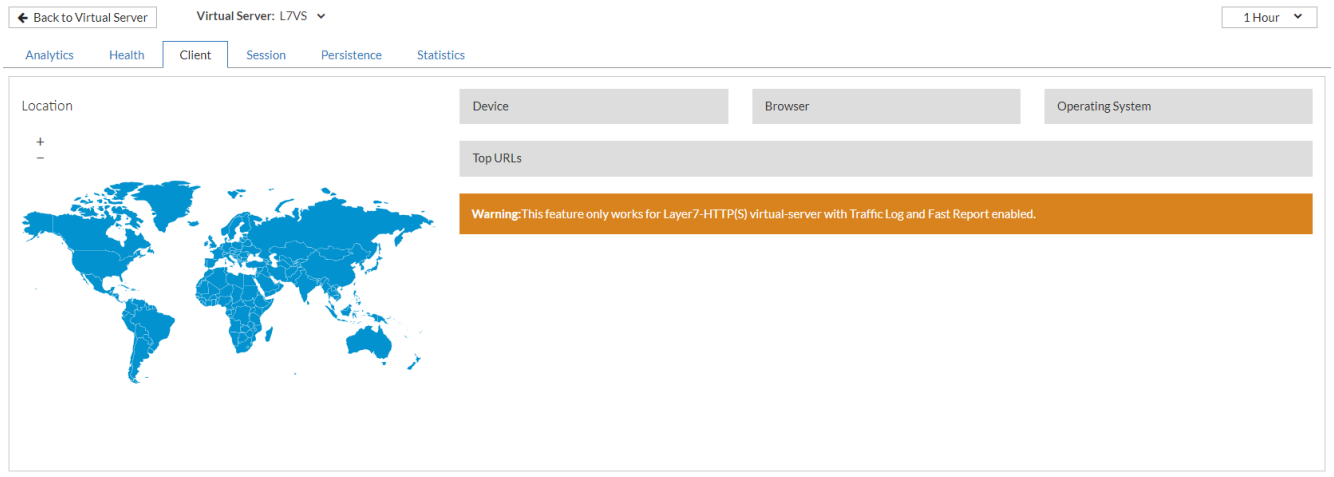
This page uses a bar graph to show the virtual server's health status in a specific time frame, as shown in the following figure:



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.
  The following device types are supported: Mobile and PC.
- 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.
  The following web browsers are supported: IE, Firefox, Safari, Opera, Edge, and Nintendo. If the client is using a browser that is not on the supported list, then the Browser will display as "Unknown".
- 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.
  The following operating systems are supported: Windows, Open BSD, Linux, Mac OS, and iOS. If the client is using an operating system that is not on the supported list, then the Operating System will display as "Unknown".
- 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 121

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 121

Statistics

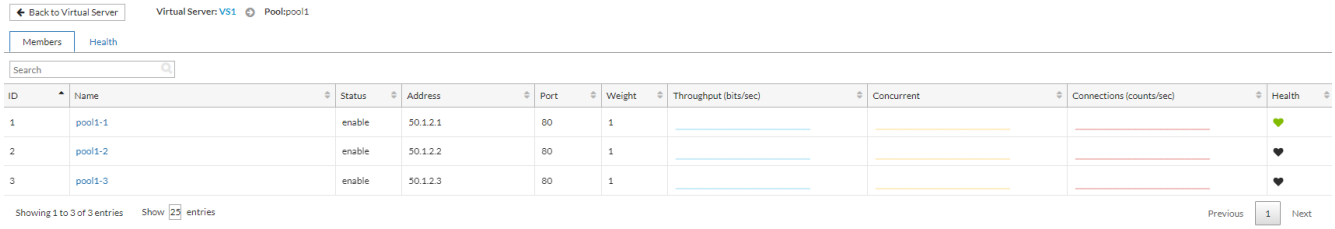
This page shows the statistics for this particular VS, its counters. After reboot, all the counters will restart on their own. It is limited to HTTP/HTTPS virtual servers.

Real server pool details

The real server pool details page shows detailed information about the real server pool you selected on the **Virtual Server** page under the **FortiView Server Load Balance** sub-menu.

The top of the page shows the name of the virtual server and the real server pool to which it is assigned. There 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.

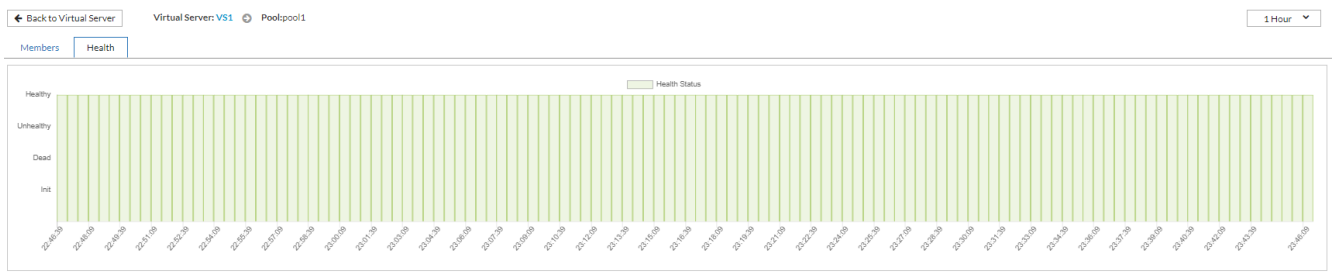
Member



The **Member** page shows key information about the real servers in a real server pool, as described in the table below.

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).
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



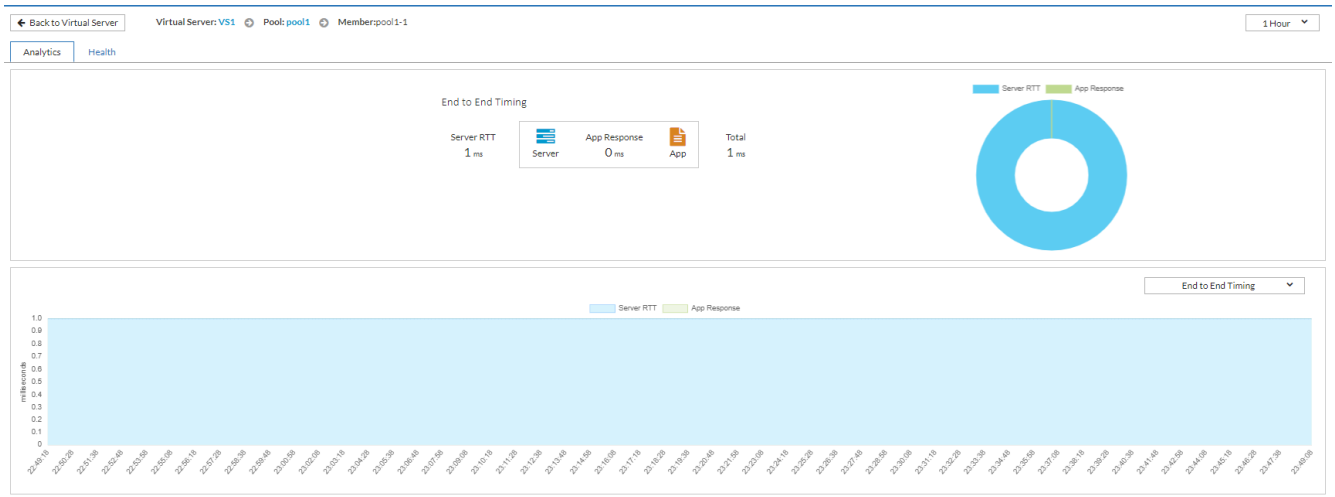
This graph shows the overall health status of the real server pool.

Real server pool member details

This page shows detailed information about the real server pool member selected on the **Virtual Server > Real server pool details** page under the **FortiView Server Load Balance** sub-menu.

The top of the page shows the name of the virtual server, the real server pool, and the real server pool member to which it is assigned. There are two tabs—Analytics and Health. The former shows the analytical information about the real server pool member, whereas the latter shows the health state of the real server in general.

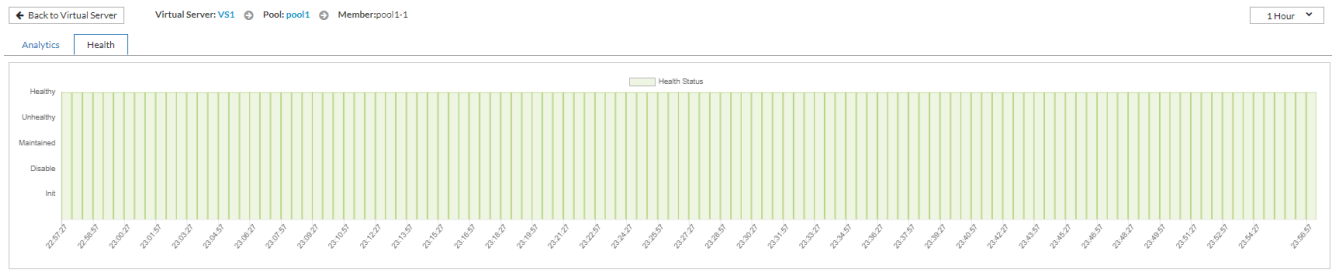
Analytics



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.

Health



This graph shows the overall health status of the real server.

Data Analytics

Under the **FortiView Server Load Balance** sub-menu, the **Data Analytics** page shows, initially, the **Dynamic Charts** page.

This is among three tabs:

- Dynamic Charts on page 119
- Static Charts on page 120
- Statistics on page 120

First we will speak of the **Dynamic Charts** page.

Dynamic Charts

In the Dynamic Charts 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.

By default, Dynamic Charts is enabled to automatically refresh with new data to add to the charts. The data is tied to the **Fast Stats** log setting, in which the **Traffic** option is automatically enabled/disabled when Dynamic Charts is enabled/disabled. For details, see Configuring fast stats log settings.

To add a widget (chart):

1. Go to **FortiView > Data Analytics**.
   The configuration page displays the **Dynamic Charts** tab.
2. Click the **Add Widget** to open the Fast Report dialog.
3. Make the entries and selections as described in Data Analytics Widget on page 119.
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—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 (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.

All Sessions

Under the **FortiView Server Load Balance** sub-menu, the **All Sessions** page contains two tabs, which open the **Session Table** and **Persistence Table**, respectively.

Viewing the Session or Persistence 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. Go to **FortiView > 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.

Security

The **FortiView > Security** sub-menu shows network security information captured by FortiADC.

It has the following pages:

- OWASP Top 10 on page 121
- OWASP Top 10 Compliance on page 124
- Threat Map on page 129
- Data Analytics on page 130
- Security Logs on page 133
- Blocked IP on page 138

OWASP Top 10

Through the **FortiView > OWASP Top 10** page, you can monitor threats by OWASP Top 10 to analyze the 10 most critical attacks targeted to your application.

You can see the total number of threats, the types of actions FortiADC carries out in response to specific types of attacks, and how severe attacks are. You can then use this insight to tune your FortiADC configurations to best address specific threats your environment faces.



To view the OWASP Top 10 data on this page, you need to first enable the security WAF log. Go to **Log & Report > Log Setting** to enable the corresponding log.

The OWASP Top 10 is a standard awareness document for web application security that represents a broad consensus about the most critical security risks to web applications. The project is managed by the Open Web Application Security Project (OWASP), a nonprofit foundation dedicated to improving the security of software. The OWASP Top 10 is updated periodically to reflect the evolving landscape of web application security.

The following summarizes the most recent OWASP Top 10, which was last updated in 2021:

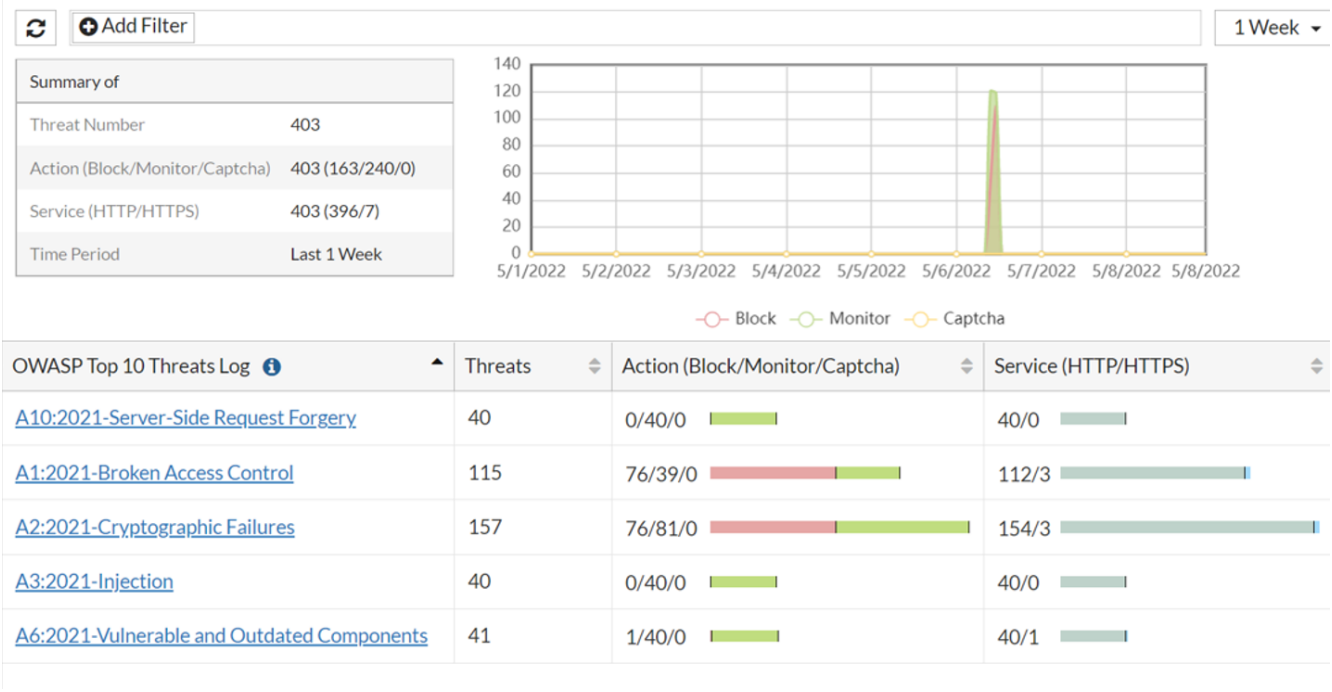
OWASP Top 10	Description	Example
Broken Access Control (A01:2021)	Flaws in access control allow unauthorized users to gain access to resources or perform actions beyond their intended permissions.	Bypassing access controls by modifying URL parameters, exploiting unprotected APIs, and accessing data intended for other users.
Cryptographic Failures (A02:2021)	Weaknesses in cryptographic processes can lead to exposure of sensitive data.	Using outdated or weak encryption algorithms, poor key management, and lack of encryption for sensitive data.
Injection (A03:2021)	Injection flaws, such as SQL, NoSQL, OS, and LDAP injection, occur when untrusted data is sent to an interpreter.	SQL injection, command injection, and other injection attacks that can execute unintended commands or access data without proper authorization.
Insecure Design (A04:2021)	Insecure design refers to the lack of security considerations in the design and architecture of an application.	Designing systems without proper threat modeling, security requirements, and design reviews.
Security Misconfiguration (A05:2021)	Security misconfigurations can occur at any level of an application stack, such as improper configurations of servers, databases, and frameworks.	Default configurations, error messages revealing stack traces, and unnecessary features enabled.
Vulnerable and Outdated Components (A06:2021)	Using components with known vulnerabilities can compromise application security.	Outdated libraries, frameworks, and other software modules that contain known security flaws.
Identification and Authentication Failures (A07:2021)	Flaws in authentication and session management can lead to compromised credentials or sessions.	Weak password policies, improper session management, and poor implementation of multi-factor authentication (MFA).
Software and Data Integrity Failures (A08:2021)	Failures related to software updates, critical data, and CI/CD pipeline integrity can compromise system integrity.	Untrusted software packages, unvalidated software updates, and insecure deployment processes.
Security Logging and Monitoring Failures (A09:2021)	Insufficient logging and monitoring can delay or prevent the detection of security breaches.	Lack of proper logging mechanisms, ineffective monitoring, and failure to log security-relevant events.
Server-Side Request	SSRF vulnerabilities occur when a web	Using URL input fields to fetch internal

OWASP Top 10	Description	Example
Forgery (SSRF) (A10:2021)	application is tricked into making requests to unintended locations.	resources, such as metadata services, or accessing internal services not meant to be exposed externally.

From FortiADC 7.1.0, the OWASP Top 10 list has been updated to the latest 2021 version. The OWASP Top 10 Wizard is automatically updated to the 2021 list, and the OWASP Top 10 2021 log data will be displayed through FortiView.

Log data from OWASP Top 10 2017 can still be accessed through the Security log.

From the **FortiView > OWASP Top 10** page, you can see the total threat data that FortiADC has detected for each OWASP Top 10 threat:



The summary OWASP Top 10 threats shows the total number of threats, actions, and service used according to the threat type.

The OWASP Top 10 Threats log analysis is based on the WAF log, so the data may not match the OWASP Top 10 Threats on the Dashboard with 100% accuracy.

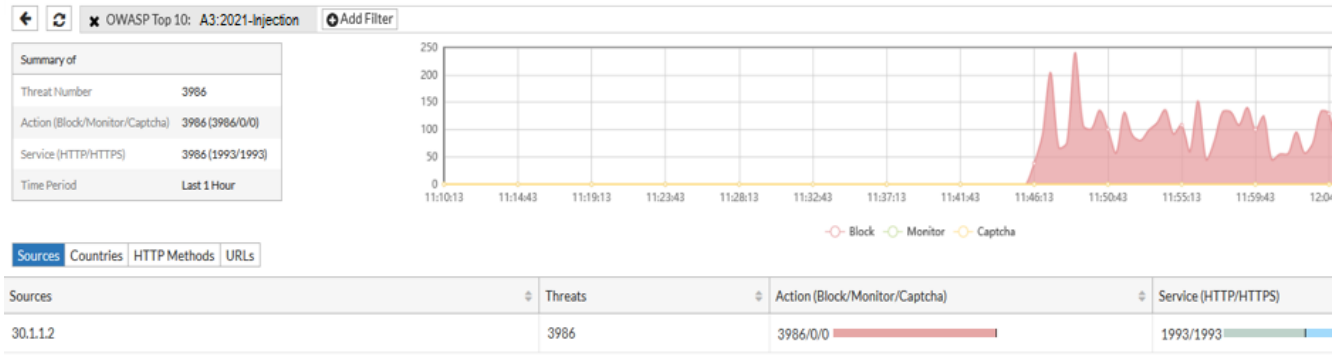
The reason for this inconsistency is due to the way the data is obtained for the Dashboard and FortiView. While the Dashboard obtains the data directly from the FortiADC, the FortiView statistics are calculated and re-aggregated by the Log module. Another cause for data inconsistencies is when a WAF action is predefined as "silent-deny", which will not be sent or recorded in the WAF log when triggered.

Viewing individual OWASP Top 10 threats

There are two ways to drill down into the key elements about a specific threat:

- Double-click the threat from the OWASP Top 10 Threats Log.
- Click the **Add Filter** icon and select the OWASP Top 10 threat from a drop-down menu.

From here, you can view information about the source IP of the attacks, countries from which the attacks are launched, the HTTP methods used, and the targeted URLs under the **Sources**, **Countries**, **HTTP Methods**, and **URLs** for the specified OWASP Top 10 threat.



OWASP Top 10 Compliance

Through the **FortiView > OWASP Top 10 Compliance** page, you can view your compliance rate against the OWASP Application Security Top 10 by assessing the security configuration of each application and categorizing them into the top 10 categories. Based on the assessment results, you can refine configurations to enhance the security posture of your virtual server.

The OWASP Top 10 Compliance dashboard currently only supports HTTP, HTTPS, HTTP2, and Explicit-HTTP virtual servers; no other virtual server types will be included in the OWASP Top 10 Compliance rating.

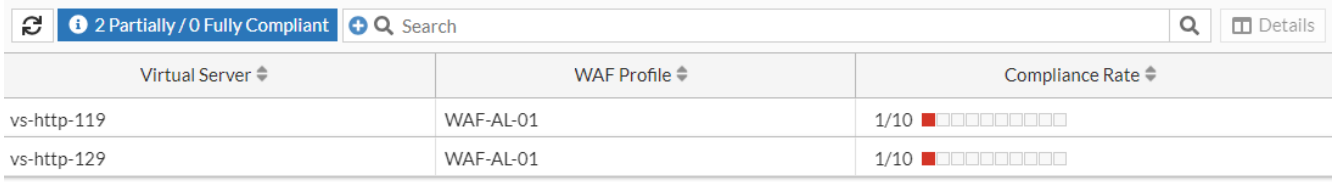


To view the OWASP Top 10 Compliance dashboard, you need to first enable the **OWASP Top10 Compliance** option from system settings. Go to **System > Settings**, and enable **OWASP Top10 Compliance** in the **Basic** settings.

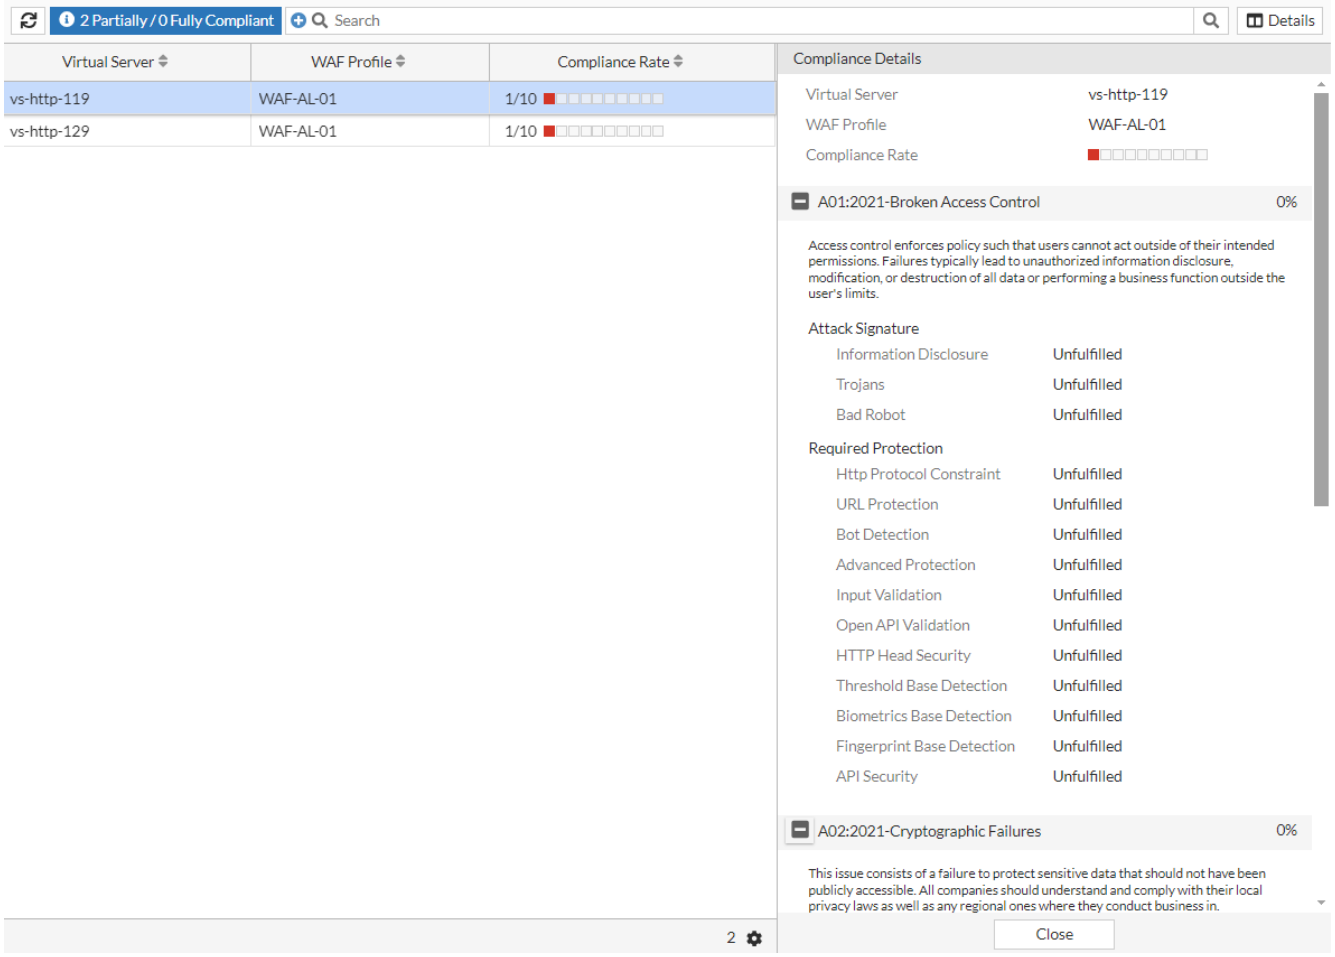
The OWASP Top 10 represents the most critical security risks to web applications as identified by OWASP (Open Web Application Security Project), an open community dedicated to enabling organizations to develop, purchase, and maintain applications and APIs that can be trusted. Organizations strive to comply with the OWASP Top 10 to ensure that the most critical security vulnerabilities in their web applications are addressed, allowing them to develop stronger defenses against cyber threats and reduce the risk of data breaches.

FortiADC's OWASP Top 10 Compliance dashboard empowers users to achieve full compliance by providing critical insights and evaluating the protection level of each server policy. The **Compliance Rate** of each virtual server is determined by analyzing their current security-related configurations—including WAF modules, Antivirus, FortiView, and Traffic log statuses—and assessing their level of protection against each OWASP Top 10 threat. This comprehensive analysis allows you to assess the effectiveness of your server policies and make the necessary adjustments to address OWASP Top 10 security risks effectively. To learn more about how the Compliance Rate is calculated, see Interpreting the Compliance Rate on page 125.

From the **FortiView > OWASP Top 10 Compliance** page, you can view the Compliance Rate of your virtual servers:



To view the **Compliance Details**, you can either select the Virtual Server and click **Details** or you can double-click the Virtual Server.



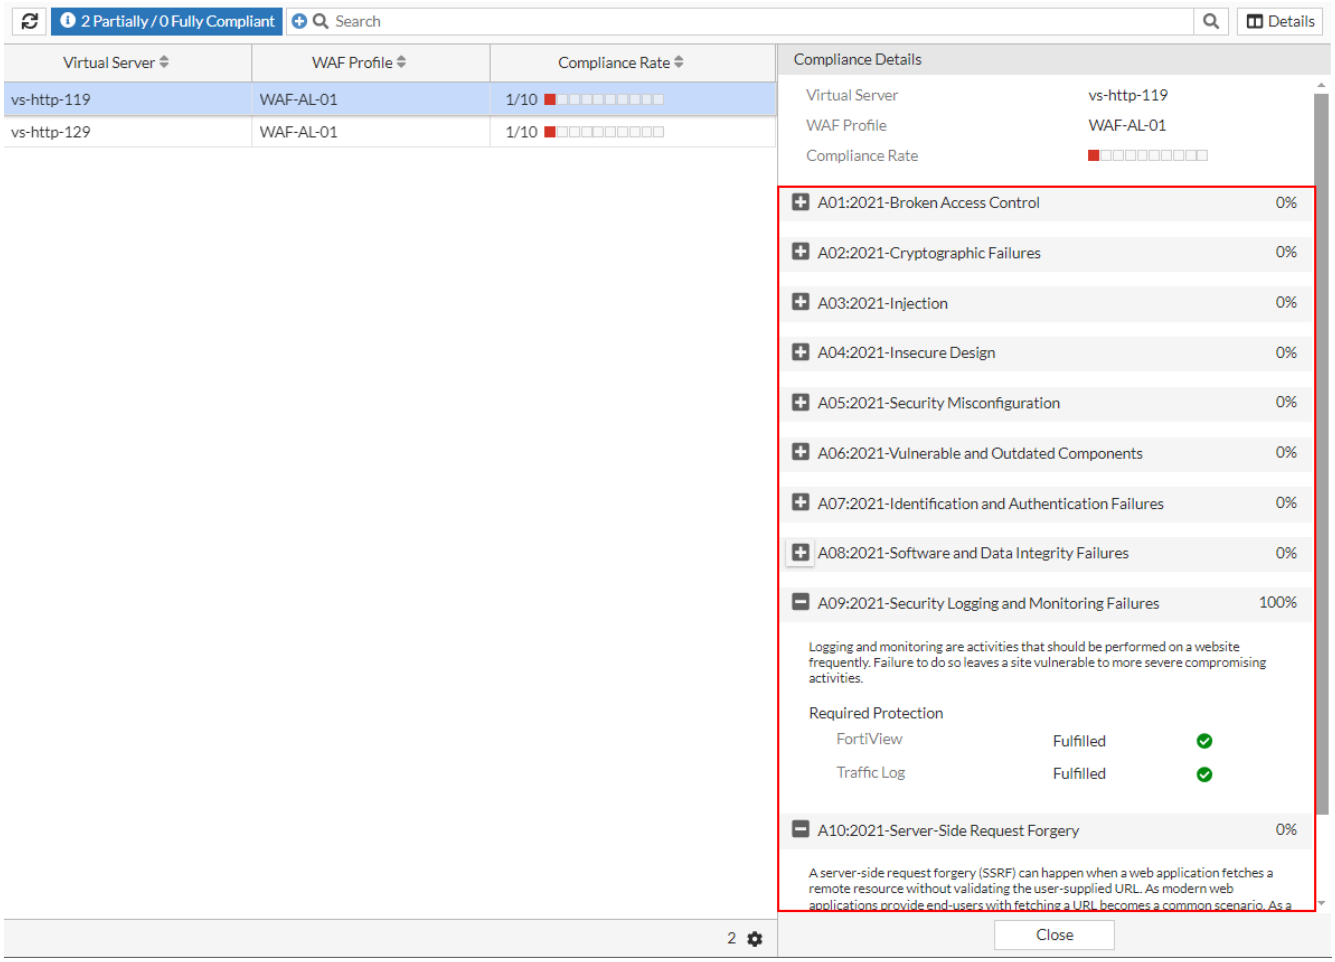
Interpreting the Compliance Rate

The **Compliance Rate** evaluates how well your virtual server policy aligns with the best practices recommended by OWASP for mitigating the Top 10 vulnerabilities. It assesses the configuration and rules in place to protect against these risks.

The Compliance Rate is calculated based on the percentage of compliance that has been reached for each OWASP Top 10 threat. For example, if the Compliance Rate is 1/10, then your virtual server configurations satisfies the protection requirements against 1 out of 10 OWASP Top 10 threats.

For each OWASP Top 10 threat, to reach 100% protection compliance, required protections must be fulfilled by completing specific security settings that target against the particular Top 10 vulnerability. The Compliance Rate is then calculated by aggregating the compliance percentage for all 10 threats.

In the example below, virtual server vs-http-119 has a Compliance Rate of 1/10, with only the **A09:2021-Security Logging and Monitoring Failures** vulnerability having fulfilled 100% of the protection requirements by enabling FortiView and Traffic Log.



The following lists the protection requirements for each OWASP Top 10 threat.

OWASP Top 10	Category	Configuration Setting
A01:2021-Broken Access Control Access control enforces policy to ensure that users cannot act outside of their intended permissions. Failures in access control typically lead to unauthorized information disclosure, modification or destruction of data, or performing business functions beyond the user's limits.	Attack Signature	Information Disclosure
		Trojans
		Bad Robot
	Required Protection	Http Protocol Constraint
		URL Protection
		Bot Detection
		Advanced Protection
		Input Validation
		Open API Validation
		HTTP Head Security
		Threshold Base Detection
		Biometrics Base Detection
		Fingerprint Base Detection
		API Security
A02:2021-Cryptographic Failures This issue involves a failure to protect sensitive data that should not be publicly accessible. Organizations should understand and comply with their local privacy laws as well as any regional regulations where they conduct business.	Attack Signature	Known Exploits
		Credit Card Detection
		Information Disclosure
	Required Protection	Cookie Security
		Data Loss Prevention
A03:2021-Injection An injection occurs when an attacker sends invalid data to a web application with the intention of making it perform actions it is not designed or programmed to do. Common types of injections include SQL, NoSQL, OS command, ORM, LDAP, and EL or OGNL injections.	Attack Signature	Generic Attacks
		Cross Site Scripting
		Cross Site Scripting(Extended)
		SQL Injection
		SQL Injection(Extended)
	Required Protection	SQL Injection Detection
		XSS Injection Detection
		XML Validation
		JSON Validation

OWASP Top 10	Category	Configuration Setting
A04:2021-Insecure Design Insecure design is a broad category representing various weaknesses, characterized by "missing or ineffective control design". An insecure design cannot be fixed by a perfect implementation because, by definition, the necessary security controls were never created to defend against specific attacks.	Attack Signature	Bad Robot
		Credit Card Detection
		Cross Site Scripting
		Generic Attacks
		Information Disclosure
		Known Exploits
		SQL Injection
		Trojans
	Required Protection	JSON Validation
		XML Validation
		Open API Validation
		API Gateway
		API Security
		Bot Detection
		Threshold Base Detection
		Biometrics Base Detection
		Fingerprint Base Detection
		Antivirus
A05:2021-Security Misconfiguration Security misconfigurations are design or configuration weaknesses resulting from configuration errors or shortcomings. Without a consistent and repeatable application security configuration process, systems are at a higher risk.	Attack Signature	Generic Attacks
	Required Protection	XML Validation
		JSON Validation
		CORS Protection
A06:2021-Vulnerable and Outdated Components Components with known vulnerabilities, such as CVEs, should be identified and patched, while outdated or malicious components should be evaluated for viability and the risks they may introduce. Attackers actively seek out websites using vulnerable components and aggressively exploit them.	Attack Signature	Generic Attacks
		Generic Attacks(Extended)
		Known Exploits

OWASP Top 10	Category	Configuration Setting
A07:2021-Identification and Authentication Failures Confirmation of the user's identity, authentication, and session management is critical to protect against authentication-related attacks. A broken authentication vulnerability can allow an attacker to use manual and/or automatic methods to try to gain control over any account they want in a system.	Attack Signature	Generic Attacks
	Required Protection	CSRF Protection
		Brute Force Attack Detection
		Credential Stuffing Defense
		API Gateway
A08:2021-Software and Data Integrity Failures Software and data integrity failures occur when code and infrastructure do not protect against integrity violations. These failures can take many forms, especially as the web evolves and the use of third-party code and services in web applications becomes increasingly common.	Attack Signature	Generic Attacks
A09:2021-Security Logging and Monitoring Failures Logging and monitoring should be performed frequently on a website. Failure to do so leaves the site vulnerable to more severe compromise.	Required Protection	FortiView
		Traffic Log
A10:2021-Server-Side Request Forgery Server-side request forgery (SSRF) occurs when a web application fetches a remote resource without validating the user-supplied URL. With modern web applications allowing end-users to fetch URLs, SSRF incidents are increasingly common.	Attack Signature	Generic Attacks

Threat Map

Under the **FortiView Security** sub-menu, the **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 come.
- 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.

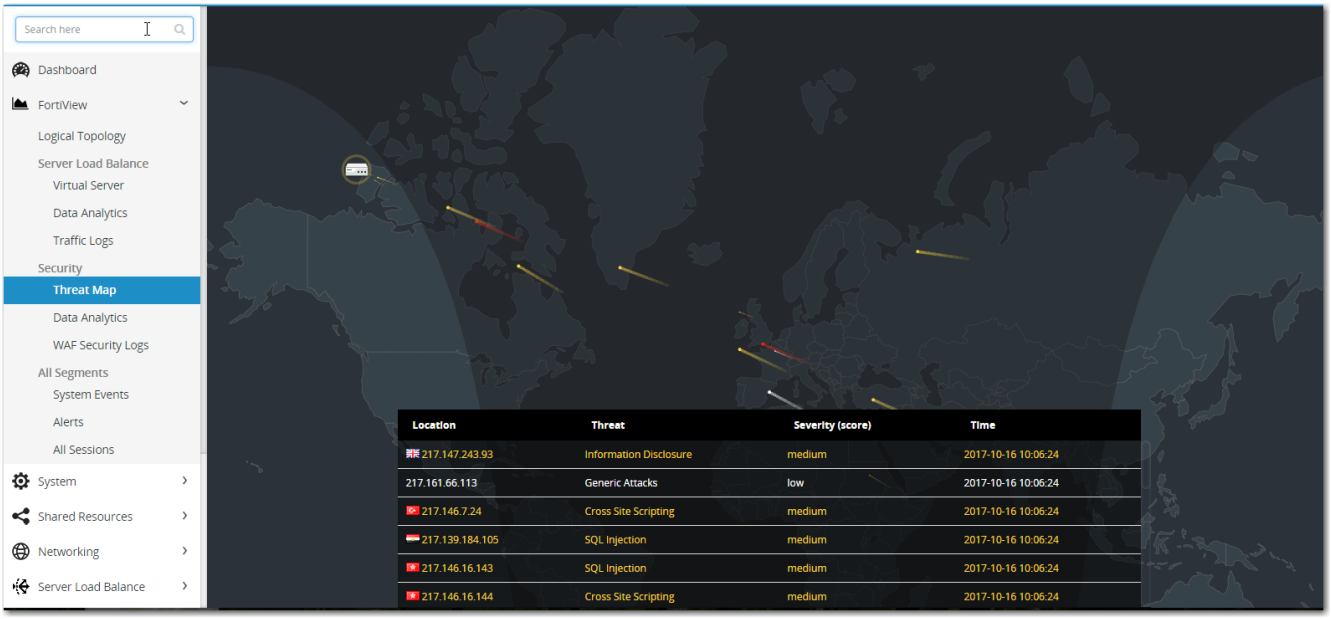
The severity of threats are color-coded:

- 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 130 shows the Threat Map with only one FortiADCappliance.

Threat map



Data Analytics

Under the **FortiView Security** sub-menu, the 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 130.
4. Click Save when done,

Data Analytics widget

Chart/Graph	Description
Name	Enter a unique name for a chart.

Chart/Graph	Description
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 • 1 Hour • 1 Day • 1 Week • 1 Month

Chart/Graph	Description
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.

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 **FortiView > Security > Data Analytics > 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.

Security Logs

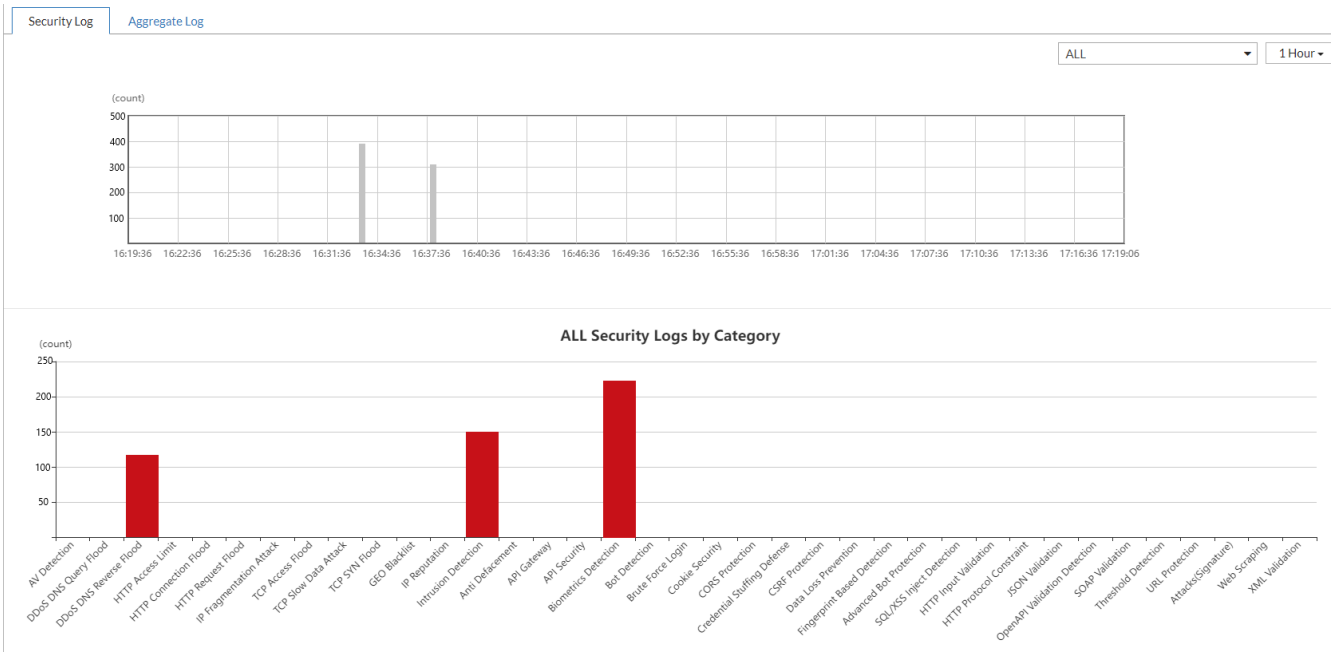
The **FortiView > Security Logs** page provides you with graphical analysis tools to view and analyze the statistical data collected from **Log & Report > Security Log**. All security logs from Log & Report > Security Log can be accessed from FortiView > Security Logs except for logs related to the Firewall module.

There are two types of FortiView logs:

- Security Log on page 133 — displays a bar graph of the security log event count against a specific time-period, from where you can drill down to a detailed view of particular logs.
- Aggregate Log on page 136 — displays a doughnut chart and bar graph that provide an aggregate view of security logs within a selected time-frame.

Security Log

From the **Security Log** tab, you can generate a bar graph of the log count and time-period of your choosing. The default selection is **ALL**, which generates a second bar graph of the log count of all security logs by category.



To view and filter the security log data:

1. Navigate to the settings along the top of the window.
2. Select the Security Log Category. The table below lists the available log options and their associated security module.

Security Log Category	Security Module
AV Detection	Anti Virus

Security Log Category	Security Module
DDoS DNS Query Flood	DoS Protection
DDoS DNS Reverse Flood	
HTTP Access Limit	
HTTP Connection Flood	
HTTP Request Flood	
IP Fragmentation Attack	
TCP Access Flood	
TCP Slow Data Attack	
TCP SYN Flood	
GEO Blocklist	Geo IP Blocklist
IP Reputation	IP Reputation
Intrusion Detection	Intrusion Prevention System (IPS)
Anti Defacement	Web Application Firewall (WAF)
API Gateway	
Bot Detection	
Brute Force Login	
Cookie Security	
CORS Protection	
Credential Stuffing Defense	
CSRF Protection	
Data Leak Prevention	
SQL/XSS Inject Detection	
HTTP Input Validation	
HTTP Protocol Constraint	
JSON Validation	
OpenAPI Validation Detection	
SOAP Validation	
URL Protection	
Attacks(Signature)	
Web Scraping	
XML Validation	

3. Select the time-period from which the selected security logs should be included to generate the graph. You have the following options:
   - 1 Hour
   - 6 Hours
   - 1 Day
   - 1 Week
   - 1 Month
   - 1 Year

From each graph, you can click on any data point to view the associated logs for further analysis. The log columns displayed depends on the security log category. For additional detail, click the  (Detail icon) to show the log details. For further description of each log message, see the FortiADC Log Reference.

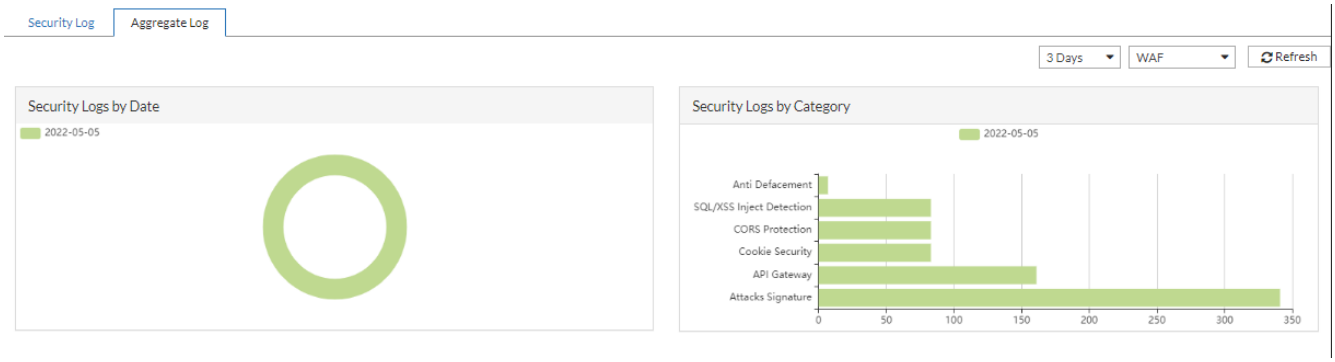
The following table describes the columns for each security log.

Column	Description
Date	Log date.
Time	Log time.
Count	The **Count** column is only available for security logs related to **DoS Protection**, **Geo IP Blocklist**, and **IP Reputation**. Rule match count.
Source	Source IP address.
Destination	Destination IP address.
Action	Action type that was taken as a result.
Destination	Destination IP address.
Service	The **Service** column is only available for security logs related to **Anti Virus** and **IPS**. Specifies the service type.
Severity	The **Service** column is only available for security logs related to **Anti Virus**, **Geo IP Blocklist**, **IPS** and **WAF**. Specifies the security level.
Virus Category	The **Virus Category** column is only available for security logs related to **Anti Virus**. Specifies the virus category.
Rule Name	The **Rule Name** column is only available for security logs related to **IPS**. Specifies the security rule name.
WAF Subcategory	The **WAF Subcategory** column is only available for security logs related to **WAF**. Specifies the Web Application Firewall subcategory.
Action	Action type that was taken as a result.

Column	Description
(Detail icon)	Click the (Detail icon) for the log details. For further description of each log message, see the FortiADC Log Reference. For **WAF** related security logs, the following actions may be performed directly from the log details: • Add Exception — You can add WAF Exceptions directly from the WAF log. This option appears only for WAF subcategories that support WAF Exceptions. For details, see Configuring WAF Exception objects on page 673. • Disable Signature — You can disable WAF signature profiles directly from the WAF log. This option appears only for Attacks Signature WAF subcategories. Disable Signature can only be successful if the WAF signature profile exists, otherwise the disable will fail with the error message "Entry not found". • View Signature — You can view the WAF signature status and information directly from the WAF log. This option appears only for Attacks Signature WAF subcategories.

Aggregate Log

From the **Aggregate Log** tab, you can generate two graphs, a doughnut chart of the security logs by date and a horizontal bar graph of the security logs by category. these graphs provide an aggregate view of security logs within the time-period of your choosing.



To view and filter the aggregate log data:

1. Navigate to the settings along the top of the window.
2. Select the security logs from the following options:
   - IP Reputation — Traffic logged by the IP Reputation feature.
   - DDoS — Traffic logged by the DoS Protection feature.
   - WAF — Traffic logged by the Web Application Firewall feature.
   - GEO — Traffic logged by the Geo IP block list feature.
   - AV — Traffic logged by the Anti Virus module.
   - IPS — Traffic logged by the IPS feature.
3. Select the time-frame from the following options:

- 3 Days
- 5 Days
- 7 Days

From each graph, you can click on any data point to view the associated logs for further analysis. The log columns displayed depends on the security log category. For additional detail, click the  (Detail icon) to show the log details. For further description of each log message, see the FortiADC Log Reference.

The following table describes the columns for each security log.

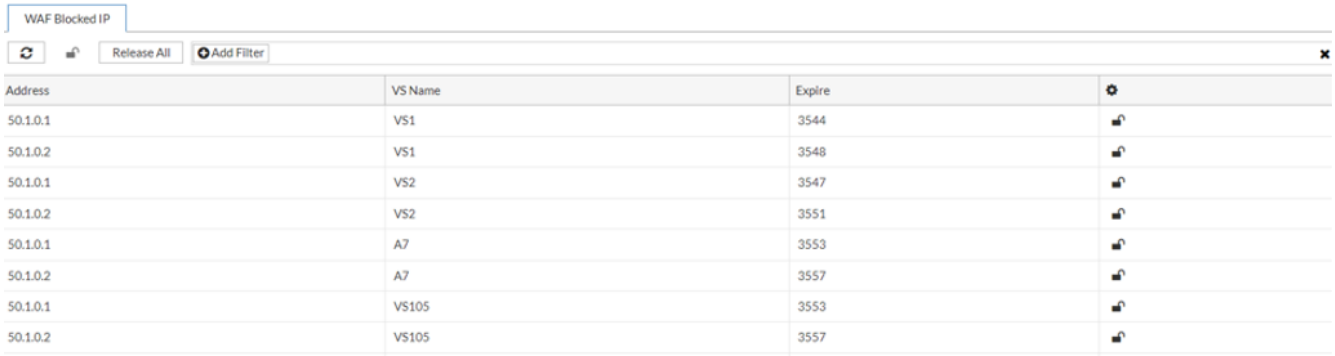
Column	Description
Date	Log date.
Time	Log time.
Count	The **Count** column is only available for **DDoS**, **GEO**, and **IP Reputation**. Rule match count.
Source	Source IP address.
Destination	Destination IP address.
Action	Action type that was taken as a result.
Destination	Destination IP address.
Service	The **Service** column is only available for **AV** and **IPS**. Specifies the service type.
Severity	The **Service** column is only available for security logs related to **AV**, **GEO**, **IPS** and **WAF**. Specifies the security level.
Virus Category	The **Virus Category** column is only available for security logs related to **AV**. Specifies the virus category.
Rule Name	The **Rule Name** column is only available for security logs related to **IPS**. Specifies the security rule name.
WAF Subcategory	The **WAF Subcategory** column is only available for security logs related to **WAF**. Specifies the Web Application Firewall subcategory.
Action	Action type that was taken as a result.
 (Detail icon)	Click the  (Detail icon) for the log details. For further description of each log message, see the FortiADC Log Reference. For **WAF** security logs, the following actions may be performed directly from the log details: • Add Exception — You can add WAF Exceptions directly from the WAF log. This option appears only for WAF subcategories that support WAF Exceptions. For details, see Configuring WAF Exception objects on page 673.

Column	Description
	• Disable Signature — You can disable WAF signature profiles directly from the WAF log. This option appears only for Attacks Signature WAF subcategories. Disable Signature can only be successful if the WAF signature profile exists, otherwise the disable will fail with the error message "Entry not found". • View Signature — You can view the WAF signature status and information directly from the WAF log. This option appears only for Attacks Signature WAF subcategories.

Blocked IP

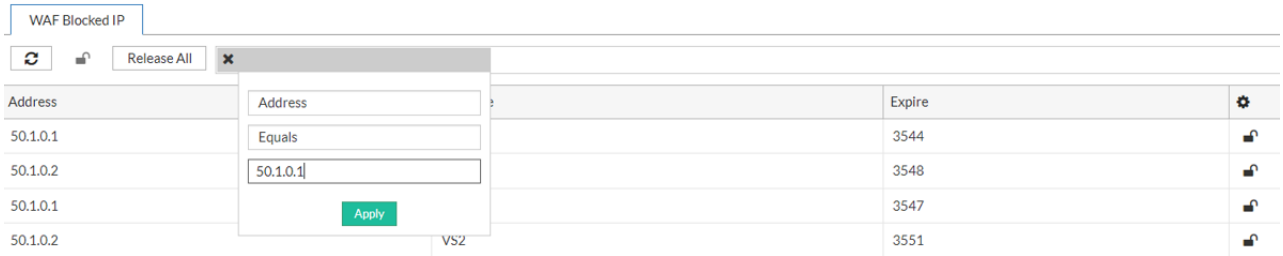
The **FortiView > Blocked IP** page displays all client IP addresses that are currently blocked by WAF modules through the **Block** or **Period Block** actions. Through the **FortiView > Blocked IP** page, you can view and release IP addresses prior to the block expiry period.

From the **WAF Blocked IP** tab, you can filter through the list of WAF blocked IP addresses and release any or all of the IP addresses that match the filter criteria.



To filter the blocked IP list:

1. Click **Add Filter** to display the filter editor.
2. Select either the **Address** or **VS Name** from which to filter the WAF Blocked IP list.
   a. To filter by **Address**, enter the IP address to complete the filter equation.
   b. To filter by **VS Name**, select the virtual server name from the drop-down menu to complete the filter equation.
3. Click **Apply**.

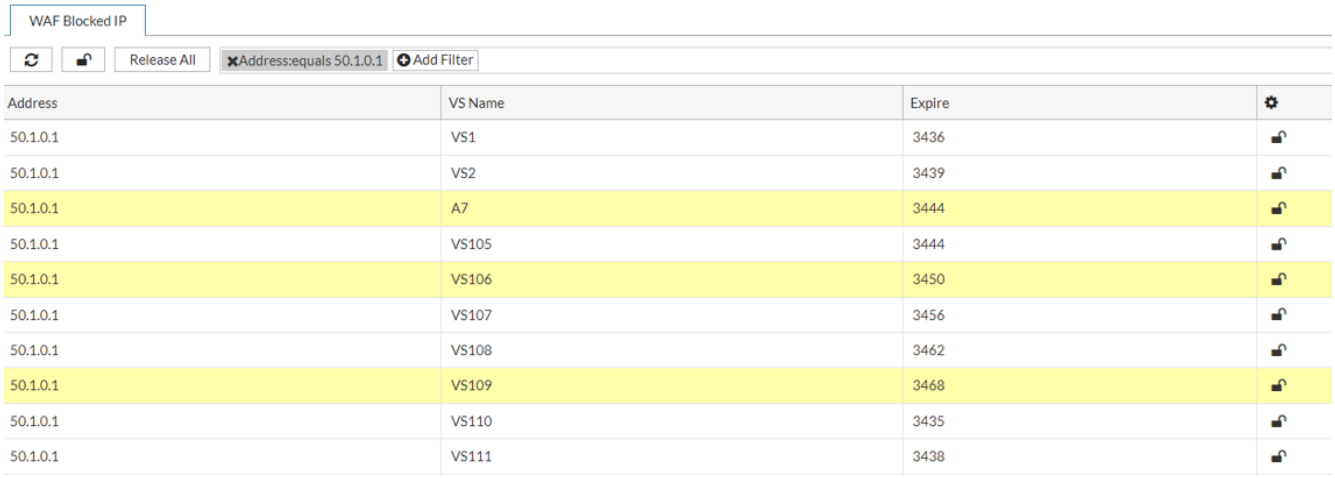


The blocked IP addresses that match the filter criteria display in the list.

To release the blocked IP address:

You can release blocked IP addresses through any of the following options:

- Release all listed IP addresses — Click **Release All**.
- Release multiple selected IP addresses — Select the IP addresses from the list and click the  (**Release button**).
- Release a single IP address — Locate the IP address from the list and click the  (**Release icon**) in that row.



The following table describes the columns of the blocked IP address list:

Column	Description
Address	The blocked IP address.
VS Name	The virtual server that has blocked the IP address.
Expire	The remaining time the IP address is blocked in seconds.
	Click the  (Release icon) to release the IP address.

System

The **FortiView > System** sub-menu shows system information captured by FortiADC.

It has the following pages:

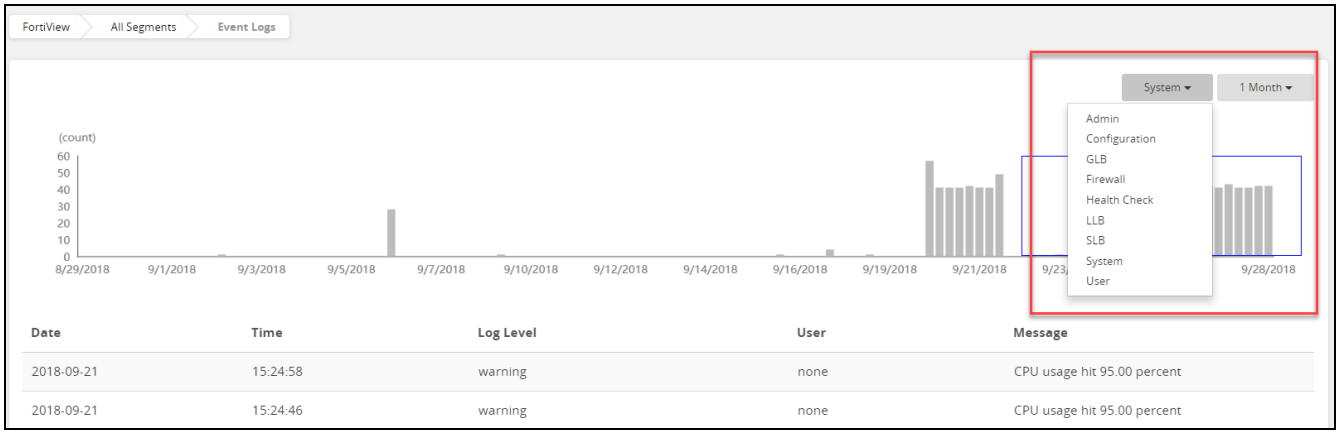
- Event Logs on page 139
- Automation on page 140

Event Logs

Under the **FortiView Security** sub-menu, the **Event Logs** page 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.

Automation

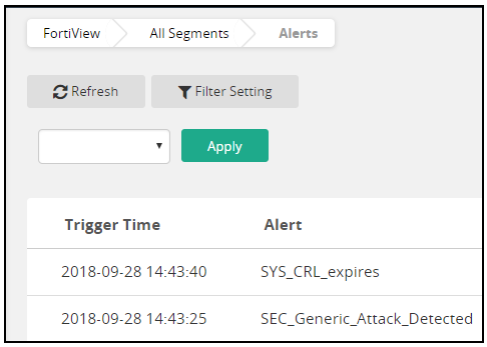
Under the FortiView System sub-menu, the Automation page shows the automation alert messages that the system has generated.

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.



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.

Global Load Balance

The **FortiView > Global Load Balance** sub-menu shows global load-balancing configurations on your FortiADC.

It has two pages:

- Host on page 141
- Data Analytics on page 142

Host

Under the **FortiView Global Load Balance** sub-menu, the **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 619

Click the number next to arrow icon in Virtual Server Pool to show the virtual servers inside the pool.

Data Analytics

Under the **FortiView Global Load Balance** sub-menu, the **Data Analytics** page shows global load-balancing information in charts called "widgets". You can create Dynamic Charts using the Add Widget button at the top right. The **Refresh** button will refresh the charts.

1. 1. After clicking **Add Widget**, the Analytic dialogue box will open.
2. 2. Select the **Type**. Based on what type you select, the type will appear as the box below it. So if you select type Virtual Server Pool, in the second box, you will be able to select a preexisting virtual server pool, or create a new one.
3. 3. Select the range.

Data Analytics Widget

Chart/Graph	Description
Type	Data Center, Virtual Server Pool, Server, Link
Range	1HOUR, 6HOURS, 1DAY, 1WEEK, 1MONTH, 1YEAR

Link Load Balance

The **FortiView > Link Load Balance** sub-menu shows link load-balancing configurations on your FortiADC.

It has one page:

- Gateway on page 142

Gateway

Under the **FortiView Link Load Balance** sub-menu, the Gateway page shows gateway 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

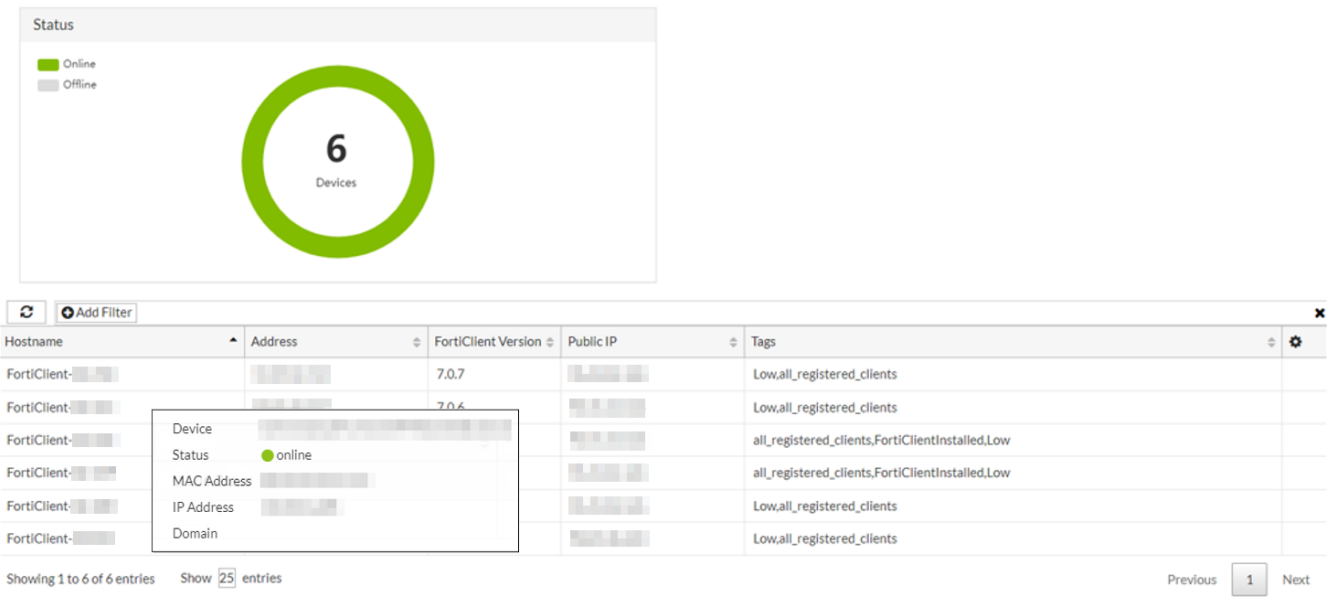
Editing gateway configuration

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 607

ZTNA FortiClient endpoint

Under the **FortiView > FortiClient** sub-menu, you can monitor the **ZTNA** page for the real-time status of the endpoints registered to FortiClient EMS. From here, you can view the FortiClient endpoint information and status that are synchronized to FortiADC from FortiClient EMS.

The FortiClient endpoint information is integral to the Zero Trust Network Access (ZTNA) functionality. For more information, see Zero Trust Network Access (ZTNA) on page 854 and How device identity and trust context is established with FortiClient EMS on page 858.



The following describes the information displayed in each column:

Column	Description
Hostname	The hostname of the endpoint device. This column is displayed by default. **Note**: You can hover over the **Hostname** column to view the device details synchronized from the FortiClient EMS.
Address	The IP address of the endpoint device. This column is displayed by default.
FortiClient Version	The FortiClient version installed on the endpoint device. This column is displayed by default.
Public IP	The public IP address of the endpoint device. This column is displayed by default.
Tags	The ZTNA tags assigned to the endpoint device. This column is displayed by default.

Column	Description
MAC	The MAC address of the endpoint device. This column is not displayed by default and must be added manually.
OS Type	The operating system installed on the endpoint device. This column is not displayed by default and must be added manually.
OS Version	The operating system version installed on the endpoint device. This column is not displayed by default and must be added manually.

You can add or remove columns from display by clicking the  **(gear icon)** to open the **Column Configuration** editor.

System

The **System** menu contain features and configurations that allow you to manage the system.

This section is organized into the following sub-menu topics:

- Settings on page 145
- Cloud Auto Scaling on page 163
- Azure LB Backend on page 166
- Virtual Domain on page 171
- High Availability on page 180
- Traffic Group on page 246
- Administrator on page 247
- WCCP on page 258
- SNMP on page 259
- Replacement Messages on page 266
- FortiGuard on page 266
- Debug on page 274
- Certificate on page 275

Settings

The **System > Settings** sub-menu includes the following topics:

- Configuring basic system settings on page 145
- Configuring system time on page 150
- Updating firmware on page 152
- Configuring an SMTP mail server on page 155
- Pushing/pulling configurations on page 156
- Backing up and restoring configuration on page 158

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. Go to **System > Settings**.
   The configuration page displays the Basic tab.
2. Complete the configuration as described in Basic settings configuration on page 146.
3. Save the configuration.

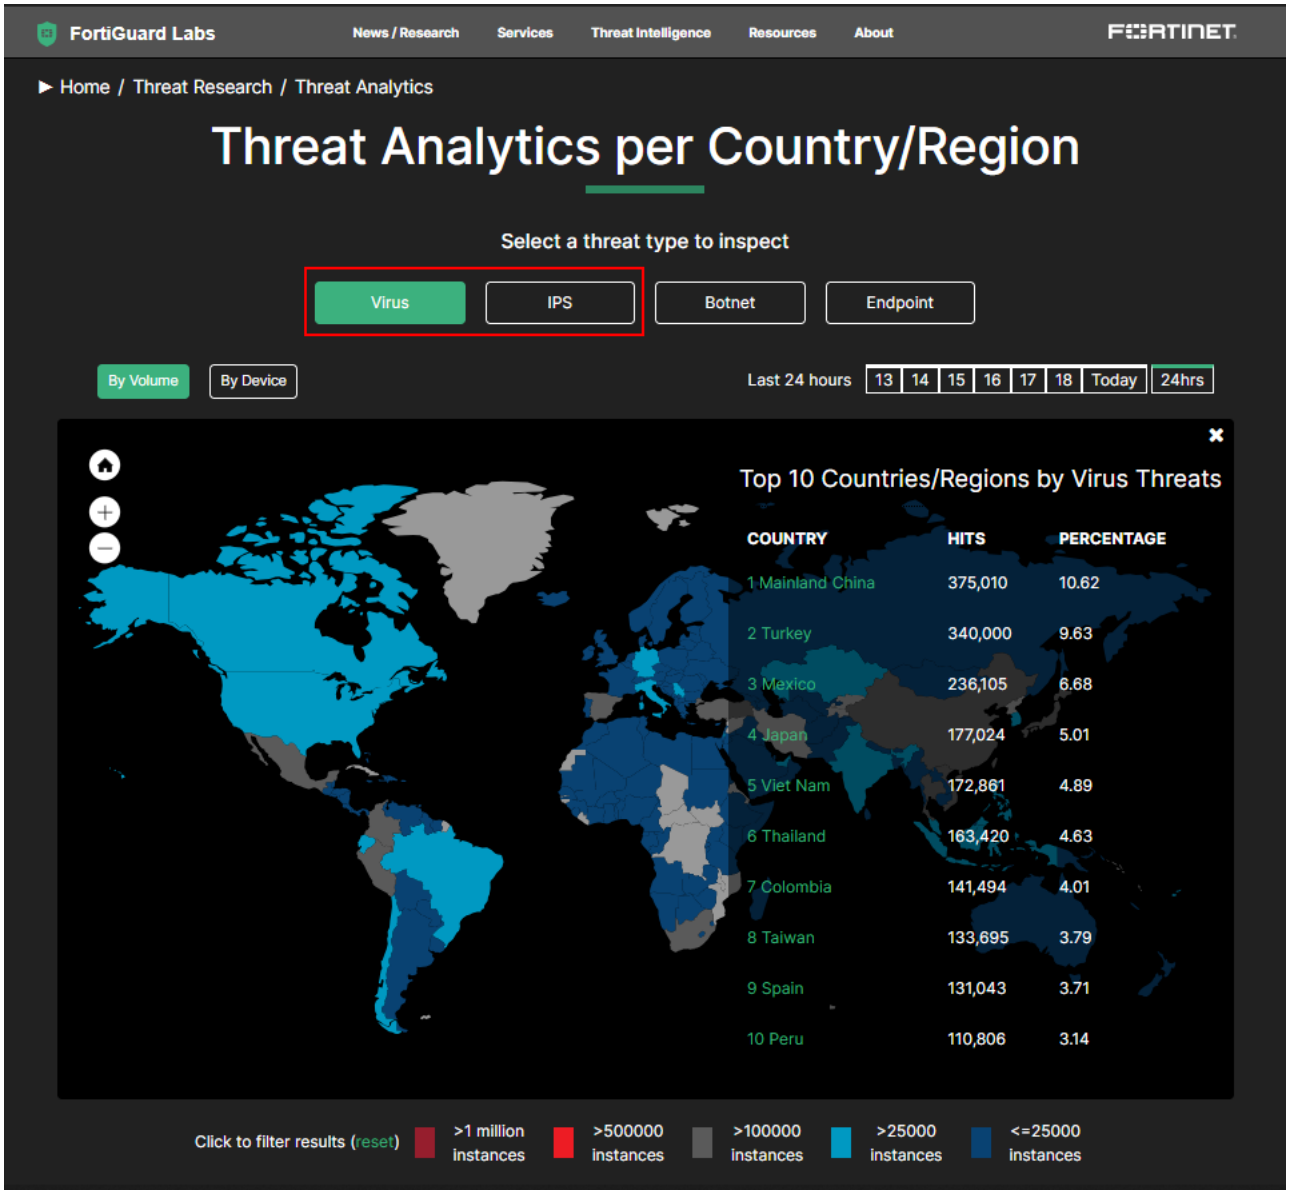
Basic settings configuration

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 `get system status` 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	Select from the following supported languages: • English • Simplified Chinese • Japanese • Spanish • Traditional Chinese • Portuguese
Idle Timeout	Log out an idle administrator session. The default is 30 minutes.
HTTPS Server Cert	Select a certificate object.
Default Intermediate CA Group	Select an Intermediate CA group.
SSH Port	Specify the port for the SSH service. Usually, SSH uses port 22.
Telnet Port	Specify the port for the Telnet service. Usually, Telnet uses port 25.
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, 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.

Settings	Guidelines
Virtual Domain	Enables the virtual domain feature. Before you enable it, make sure you understand how the system implements virtual domains. See Virtual Domain.
Virtual Domain Mode	The **Virtual Domain Mode** option is available if **Virtual Domain** is enabled. Select the virtual domain mode: • Independent Network — each VDOM functions independently within its own network, unaffected by activity from other VDOMs on the system. • Share Network — VDOMs function as administrative domains (ADOMs), sharing the same network interface and routing between all ADOMs. For details, see Enabling the Virtual Domain feature and selecting the Virtual Domain Mode on page 174.
HTTP Port	Specify the port for the HTTP service. Usually, HTTP uses port 80.
Redirect to HTTPS	When enabled, all HTTP connections to FortiADC will be redirected to HTTPS. HTTPS-Redirect switch is enabled by default.
HTTPS Port	Specify the port for the HTTPS service. Usually, HTTPS uses port 443.
Config Sync Enable	Enable/disable the configuration synchronization feature. This feature is related to Pushing/pulling configurations, not HA synchronization. Disabled by default.
Pre Login Banner	Enable/disable the pre-login banner feature to show login disclaimer messages. Disabled by default.
Admin Bypass VDOM Check	The **Admin Bypass VDOM Check** option is available if **Virtual Domain** is **enabled** and the **Virtual Domain Mode** is **Independent Network**. Once enabled, all non-root VDOM administrators can login through the root VDOM interface without needing root VDOM privileges. From the root VDOM interface, the non-root VDOM administrator can access and modify the settings relating to their designated non-root VDOM. This is disabled by default. **Note**: When the **Admin Bypass VDOM Check** is enabled, admin event logs are recorded solely under the root VDOM, regardless of the VDOM assigned to the user. This is due to the way the log API handles event logging, allowing event logging for only one VDOM at a time.
OWASP Top10 Compliance	Enable OWASP Top10 Compliance to view the security compliance rate for each SLB virtual server in FortiView. This is disabled by default. For details, see OWASP Top 10 Compliance on page 124.
Feedback Options	
Upload detection statistics to FortiGuard	Enable or disable FortiADC detection statistics upload to FortiGuard. This is enabled by default. For details, see Sending FortiADC Threat Telemetry to FortiGuard on page 148.

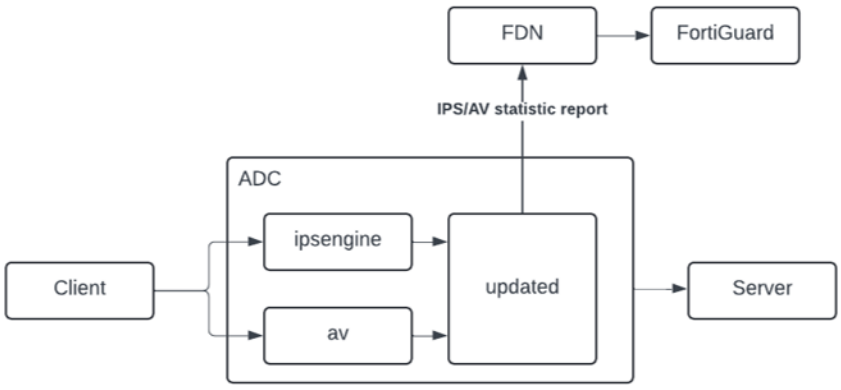
Sending FortiADC Threat Telemetry to FortiGuard

FortiADC is designed to send threat telemetry data to FortiGuard to enhance threat intelligence and enable product-specific threat analysis. FortiGuard processes this data for analysis, including metrics on Antivirus, IPS, and other security statistics. The AV/IPS analysis has been standardized, allowing FortiADC to send IPS and Antivirus data to FortiGuard for comprehensive threat assessment.



FortiADC Threat Telemetry Workflow

The **FortiADC** Threat Telemetry Workflow outlines the systematic process through which the IPS and AV engines detect rule matches, collect relevant statistics, and transmit this data to FortiGuard for comprehensive analysis.

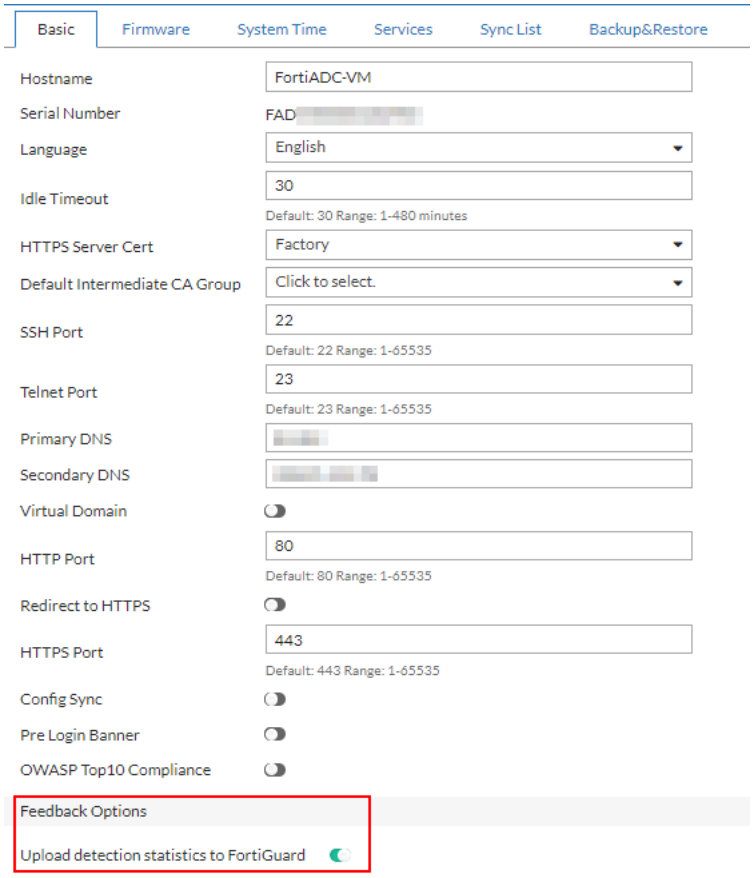


1. **Detection** — The IPS or AV engine detects rule matches based on the configured threat protection policies.
2. **Data Transmission to Daemon** — After detection, the IPS/AV engine sends the matched rule information to the updated daemon for further processing.
3. **Statistics Collection** — If the option is enabled to send statistics, the daemon collects IPS/AV statistics, including details of the detected rule matches.
4. **Data Transmission to FortiGuard** — The daemon sends the collected statistics to FortiGuard every 60 minutes by default. This reporting interval can be configured by the user through CLI.
5. **Data Processing** — Once FortiGuard receives the IPS/AV statistics, it processes the data for analysis and generates a report.
6. **Portal Display** — The processed data is then made available and displayed on the FortiGuard portal for user review.

To enable FortiADC to transmit threat telemetry data to FortiGuard:

1. Navigate to **System > Settings**.
2. From the **Basic** tab, scroll down to locate the **Feedback Options** section.
3. Enable the **Upload detection statistics to FortiGuard** option.

Note: This option is enabled by default after upgrading to FortiADC version 7.6.1.



Users can configure the statistics collection interval from the CLI.

```
config system global
   set fds-statistics {enable|disable}
   set fds-statistics-period <integer>
end
```

`fds-statistics`	Enable or disable FortiADC detection statistics upload to FortiGuard. This is enabled by default.
`fds-statistics-period`	Specify the FortiGuard statistics collection period in minutes. The default value is 60 minutes, and the valid range is 1-1440 minutes.

Configuring system time

The system time must be accurate for many features to work, including scheduling, logging, and SSL/TLS-related features.

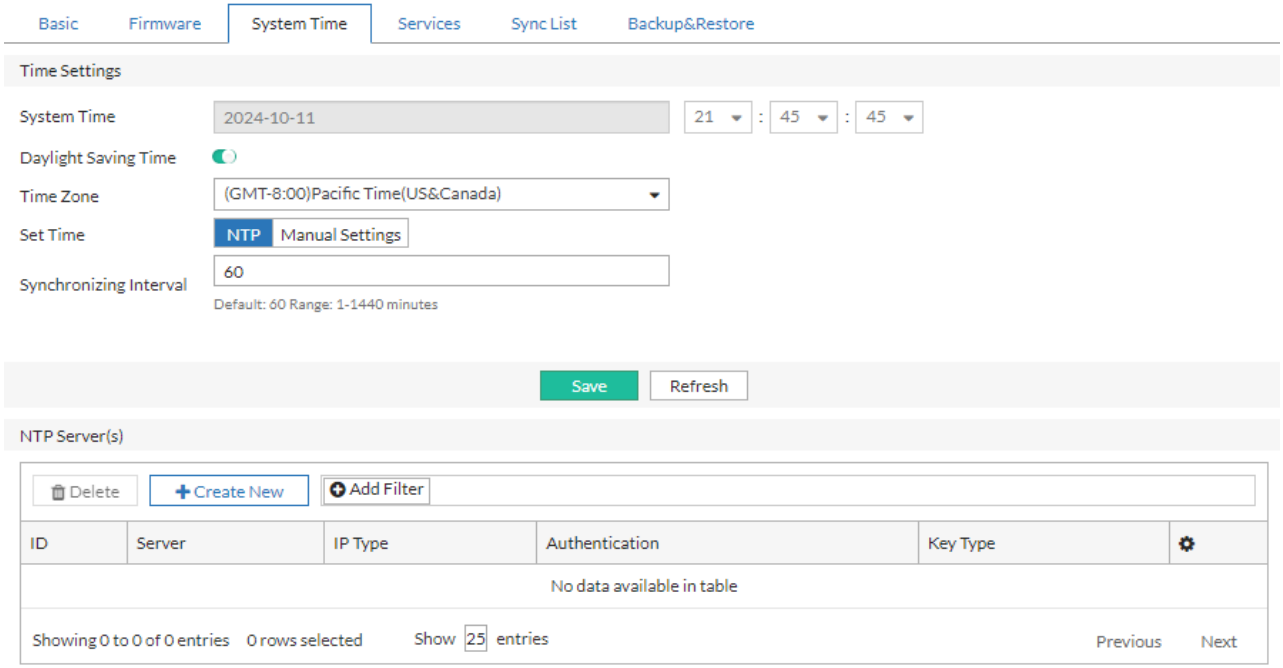
We recommend that you use Network Time Protocol (NTP) to maintain the system time. As an alternative when NTP is not available or is impractical, you can set the system time manually.

Before you begin:

- You must have Read-Write permission for System settings.

To configure the system time:

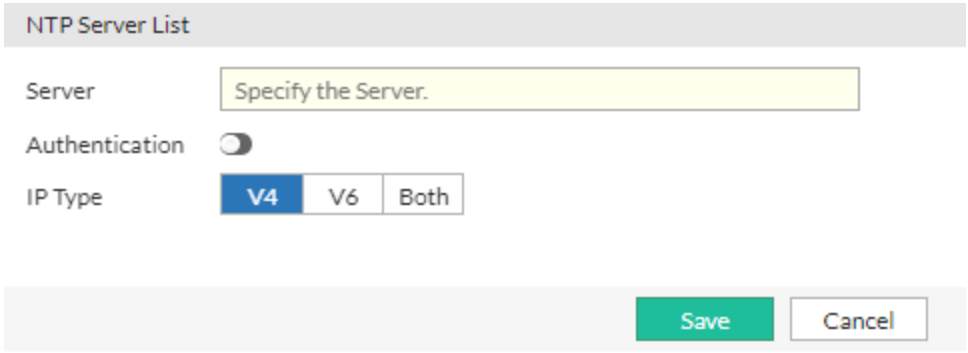
1. Go to **System > Settings**.
2. Click the **System Time** tab.
3. Configure the Time Settings:



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.
Set Time	Select either of the following options: • NTP — Set the system time based on the Network Time Protocol. • Manual Settings — Set the system time manually. The default option is NTP.
Synchronizing Interval	The **Synchronizing Interval** option is available if **Set Time** is **NTP**. Specify how often the system synchronizes its time with the NTP server. The default is 60 minutes. The valid range is 1-1440.

4. Click **Save**.
If the system time is based on NTP, the **NTP Server(s)** section becomes configurable once the Time Settings configuration is successfully saved.

5. Under the NTP Server(s) section, click **Create New** to display the configuration editor.
6. Configure the NTP Server List settings:



Setting	Guidelines
Server	Specify a space-separated list of IP addresses or FQDNs for an NTP server or pool, such as `pool.ntp.org`. To find an NTP server, go to http://www.ntp.org.
Authentication	Enable to apply authentication keys to secure the NTP server. This is disabled by default.
IP Type	Select the IP type from the following: • V4 • V6 • Both The default option is V4.
Key Type	The **Key Type** option is available if **Authentication** is **enabled**. Select the key type from the following: • SHA1 • SHA256 • AES128 • AES256 The default option is SHA1.
Key	The **Key** option is available if **Authentication** is **enabled**. Specify the Key in hexadecimal format. The maximum length is 127 digits or characters.
Key ID	The **Key ID** option is available if **Authentication** is **enabled**. Specify the Key ID. The valid range is 0-65536

7. Click **Save**.

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 may need to reconfigure some settings.



Firmware versions are tagged to indicate the maturity level of its feature development:

- Feature (F) — New features are included in this version.
- Mature (M) — This version primarily includes bug fixes and/or vulnerability patches, and no new or major features.

Important: Read the release notes for release-specific upgrade considerations.

Updating firmware using the web UI

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.

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 **Firmware** tab.
3. Click **Upgrade Firmware** to display the configuration editor.
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.

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 156.
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. SMTPS is not supported by FortiADC, so if you want to build up secure connections with FortiADC, STARTTLS can work as an substitution to SMTPS.
Username	Username for authentication to the SMTP server.
Password	Password for authentication to the SMTP server.

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. Go to **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 `config config`, `config system` (except `config system mailserver`), `config user`, and `config vdom` commands.
	• Networking—Includes `config router` commands.
	• SLB—Includes `config load-balance` commands.
	• Log—Includes `config log` commands and `config system mailserver`.
	• LLB—Includes `config link-load-balance` commands.
	• GLB—Includes `config global-load-balance` and `config global-dns-server` commands.
	Note: The Shared Resource type is automatically selected when GLB is selected. If any GLB configurations reference a shared resource, that shared resource configuration must be synced together with the GLB configuration, otherwise the GLB configuration will not work with the missing referenced object.
	• Security—Includes `config security waf` commands.
	• Shared Resource—Includes the shared resources that are used across different modules.
	These configurations include the following commands:
	`config system health-check`
	`config system health-check-script`
	`config system schedule-group`
	`config system address`
	`config system address6`
	`config system addrgrp`
	`config system addrgrp6`
	`config system isp-addr`
	`config system service`
	`config system servicegrp`
	• User—Includes `config user` 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	• system global • system tcpdump • system accprofile • system admin

Module	Commands
	• system ha • system snmp sysinfo • system snmp community • system snmp user • system alert-snmp-trap • system fortiguard • system hsm info • system hsm partition • config sync-list
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 setting • global-dns-server general
Log & Report	• system mailserver

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 Global Administrator access. Ensure that your admin account settings has **Global Admin** set to **Yes**. For more information, see Creating administrator users on page 249.

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 159.
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.)

Actions	Guidelines
	• 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 162.
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 yourFortiADC 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 sever:

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

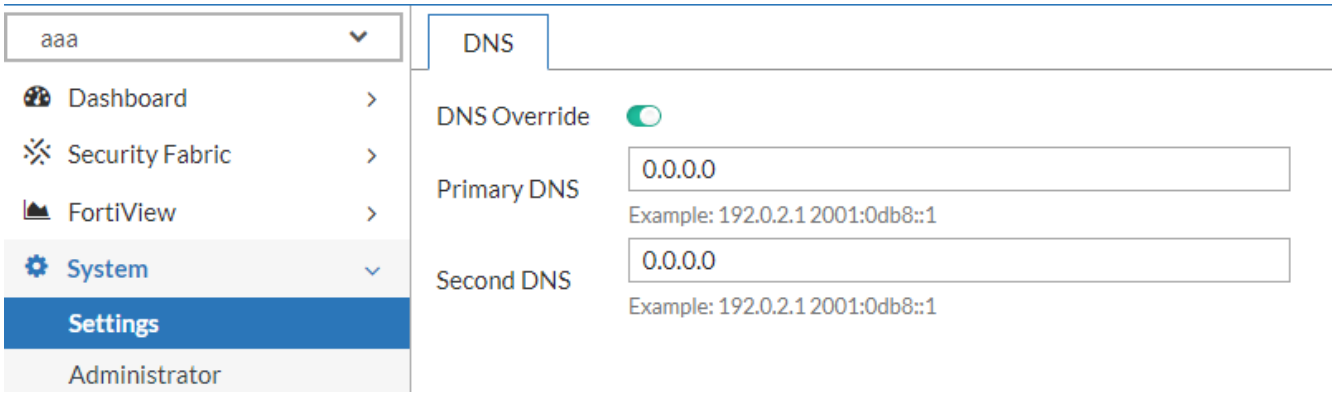
Hardware model	Maximum total backup file size
FortiADC 60F	50 MB
FortiADC 100F	50 MB
FortiADC200D	50 MB
FortiADC 200F	50 MB
FortiADC 300D	100 MB
FortiADC 400D	100 MB
FortiADC 700D	100 MB
FortiADC 1000F	100 MB
FortiADC 1500D	100 MB
FortiADC 2000D	100 MB
FortiADC 2000F	200 MB
FortiADC 4000D	200 MB
FortiADC 4000F	200 MB
All FortiADC VMs	100 MB

Configuring DNS settings for non-root VDOM

For FortiADC systems with VDOM enabled, you can configure the System DNS resolver for non-root VDOMs and override Global DNS settings to set a DNS server IP per VDOM for more flexibility.

This functionality is only available in the non-root VDOM. The root VDOM defaults to inherit the Global DNS settings.

Note: The DNS settings per VDOM is also not available when VDOMs are in Share Network mode (ADOM mode) as all ADOMs will share the same network namespace.



DNS settings per VDOM does not affect the following global modules which will continue to use Global DNS settings:

- Security Fabric external and fabric connectors
- URL filters
- FortiGuard
- SMTP for logs and alerts
- NTP
- Admin FortiToken Cloud
- AntiVirus cloud

Before you begin:

- You must have VDOM enabled and have access to a non-root VDOM.
- 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 system settings.

To configure DNS settings for a non-root VDOM:

1. Go to **System > Settings**.
   The configuration page displays the **DNS** tab.
2. Configure the following settings:

Setting	Description
DNS Override	Enable to apply the DNS settings configured for this non-root VDOM instead of inheriting from the Global DNS settings.
Primary DNS	Specify the IP address for the primary DNS server.
Second DNS	Specify the IP address for the secondary DNS server.

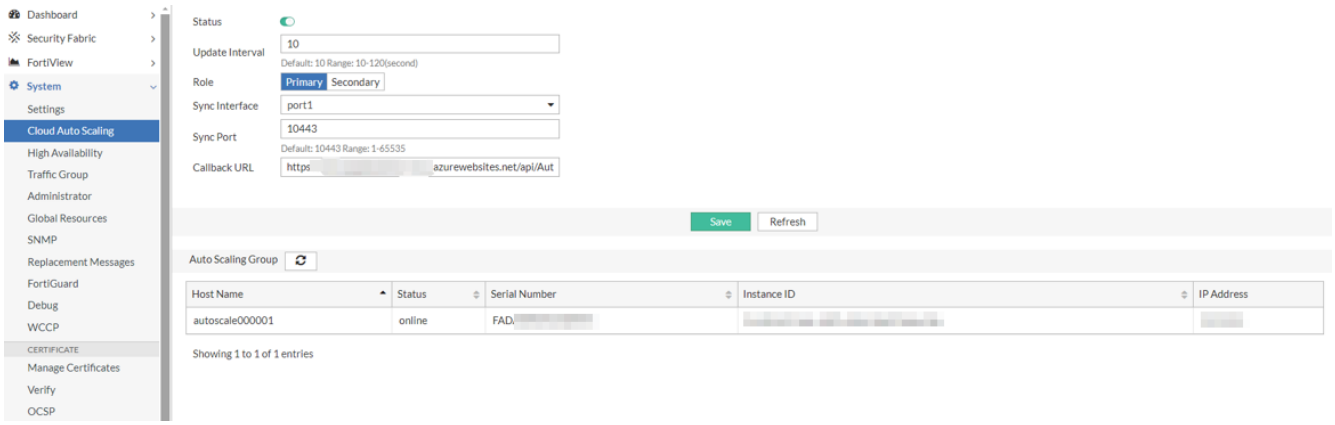
3. Click **Save**.

Cloud Auto Scaling

FortiADC supports cloud autoscaling for Azure and AWS platforms. Once autoscaling is deployed on the cloud, the status of FortiADC devices can be monitored through the GUI of the primary node. FortiADCs can be in one of two states:

- **init** — In this state, the FortiADC establishes a connection with the primary node and performs a full configuration synchronization.
- **online** — This state indicates the synchronization process has successfully completed.

Before connecting to any secondary node via GUI, console, or SSH, ensure its status has transitioned to "online."



Any updates to the Cloud Auto Scaling configuration in for are not automatically synchronized with the autoscaling group on the cloud platform. To prevent configuration discrepancies between nodes, the autoscaling group must be manually updated to maintain synchronization consistency.

Azure Autoscaling

The Azure Autoscaling feature facilitates dynamic scaling of FortiADC Virtual Machine Scale Sets (VMSS) based on real-time CPU utilization metrics. This architecture comprises a primary FortiADC and several secondary FortiADCs that distribute web traffic from the Azure Load Balancer in front of the VMSS. This setup optimizes resource allocation and enhances system resilience by automatically adjusting active instances to meet fluctuating traffic demands, ensuring high availability and efficient load management.

The autoscaling feature, available in FortiADC 7.6.1 for On-demand (PAYG) instances, can be deployed on the Azure platform using ARM templates. A server-less web application running on Azure App Service enables authorized access to resources within the auto-scaling cluster. This application selects the primary node in the FortiADC Virtual Machine Scale Set (VMSS) and communicates its IP address and VMID to the secondary nodes. Configuration synchronization occurs unidirectionally from the primary node to secondary nodes, ensuring consistent settings across the cluster.

For detailed steps on deploying Azure autoscaling, refer to the FortiADC Azure Deployment Guide.

Other considerations when Azure autoscaling is enabled include:

- The Azure Load Balancer (LB) Backend tab will be hidden from the user interface.
- The virtual server address will also be concealed, with the virtual server operating using the interface IP address and port instead.

Troubleshooting

Debug information can be accessed through the Azure Function App's AutoscaleHandler log or the FortiADC CLI.

In the AutoscaleHandler log, you can verify whether the FortiADC VMs in the VMSS are sending heartbeat callbacks in a timely manner and whether they are maintaining a healthy state. Check both the logs and DynamoDB records for the deployed resources.

On the FortiADC-VM, you can utilize CLI commands to monitor the status of the cloud autoscaling daemon. Use the following commands:

- To review the heartbeat callback results and failover information in case a primary election is triggered:

```
# diagnose debug cloud-autoscale autoscaled
```

- To check the synchronization between the primary and secondary FortiADCs, as well as view any crash logs if they exist:

  ```
  # diagnose debug cloud-autoscale autoscale-tunnel
  ```

Limitations

Synchronization

The configuration synchronization for FortiADC is consistent with the settings applied when HA VRRP is enabled. Note that synchronization is not managed by the HA module.

When the FortiADC is operating as the primary node, it listens on the IP address `ip:10443`, where ip is the interface IP defined in the sync-interface under system auto-scale settings. When a secondary FortiADC connects to the primary, it initiates a full configuration synchronization to ensure consistency with the primary.

Configuration changes can only be initiated by the primary node. Any modifications made on a node assigned the secondary role will not be synchronized to other nodes within the VMSS.

High Availability

When auto-scaling is enabled, the HA mode must operate in standalone mode. Switching the HA mode to VRRP is prohibited while auto-scaling is active.

AWS Autoscaling

You can deploy FortiADC virtual machines (VMs) to support Autoscaling on AWS. This requires a manual deployment incorporating AWS CloudFormation Templates (CFTs).

Multiple FortiADC-VM instances form an Autoscaling group (ASG) to provide highly efficient clustering at times of high workloads. FortiADC-VM instances can be scaled out automatically according to predefined workload levels. When a spike in traffic occurs, the Lambda script is invoked to automatically add FortiADC-VM instances to the ASG. Autoscaling is achieved by using FortiADC Cloud Autoscaling features such as system autoscale that synchronize operating system (OS) configurations across multiple FortiADC-VM instances at the time of scale-out events.

FortiADC Autoscale for AWS is available with FortiADC 7.2.0 and supports On-demand (PAYG) instances.

For detailed steps on deploying AWS autoscaling, refer to the FortiADC AWS Deployment Guide.

Troubleshooting

Debug information can be accessed via AWS CloudWatch or the FortiADC CLI. AWS CloudWatch provides logs to verify if FortiADC-VMs within the ASG are sending timely heartbeat callbacks and to check the health status of each FortiADC instance.

On the FortiADC-VM, the CLI can be used to monitor the cloud autoscaling daemon's status:

- Use `diagnose debug cloud-autoscale autoscaled` to view heartbeat callback results and failover details if a primary node election is initiated.
- Use `diagnose debug cloud-autoscale autoscale-tunnel` to examine the synchronization process between primary and secondary FortiADCs, and to identify any crash logs if they are present.

Azure LB Backend

The **Azure LB Backend** configuration object supports FortiADC integration with Azure Load Balancers (ALBs) for high-availability deployments. This feature ensures seamless failover and reduces configuration errors by enabling independent backend IP address assignments on each FortiADC node. It is exclusive to FortiADC instances deployed in Azure using the ARM template for VRRP HA. For details, see the FortiADC Document on Deploying High Availability on Azure.

Azure Load Balancers provide a consistent front-end IP address as the single access point for clients, even during failover. They distribute inbound traffic to FortiADC nodes in the backend pool based on load-balancing rules and health probes. These rules determine how traffic on specific front-end IP addresses and ports is routed to backend addresses and ports.

The **Azure LB Backend** configuration synchronizes across FortiADC nodes in an HA cluster, but the backend IP address is not included in the synchronization. On the secondary node, the backend IP is set to `0.0.0.0` by default. Administrators must manually update this value with the appropriate **Azure LB Backend** IP address for the secondary node. This design ensures that IP configurations remain independent and eliminates the need for IP migration during failover.

This feature leverages both external and internal Azure Load Balancers to manage traffic and maintain consistent client access during normal operation and failover scenarios. By decoupling backend IP configurations from the failover process, FortiADC ensures uninterrupted service and reliable traffic distribution.



As the IP address in the **Azure LB Backend** configuration object would not be synchronized in HA mode, you must configure it on both the FortiADCs in the cluster.

This feature is only available in the FortiADC Azure based images.

Configuring the Azure LB Backend and Virtual Server for HA VRRP Mode Using Azure Load Balancer

To enable proper health check functionality after deploying HA VRRP with Azure Load Balancer, you must configure the probing virtual server to respond to health check packets on the FortiADC side.

This process involves two key steps:

- **Configure Azure LB Backend** — Use the Azure LB Backend configuration to define the backend pool IP information on both FortiADC-VMs. Each FortiADC must have its own backend IP configured independently, as the IP address field is not synchronized across HA nodes.
- **Configure the Probing Virtual Server** — Set up the probing virtual server to reference the backend IP pool configured in the Azure LB Backend. This ensures that the virtual server responds to health check packets, allowing the Azure Load Balancer to monitor the FortiADC nodes' availability and distribute traffic accordingly.

By completing these configurations, the Azure Load Balancer can maintain reliable traffic distribution during normal operation and failover scenarios in an HA VRRP deployment.

Before you begin:

- You must have already deployed FortiADC instances in Azure via the ARM template for VRRP HA.
- You must have obtained the backend pool information from the Azure Load Balancer. This will be used in the FortiADC Azure LB Backend configuration.

- You must have obtained the health probe information from the Azure Load Balancer. This will be used in the probing virtual server configuration.

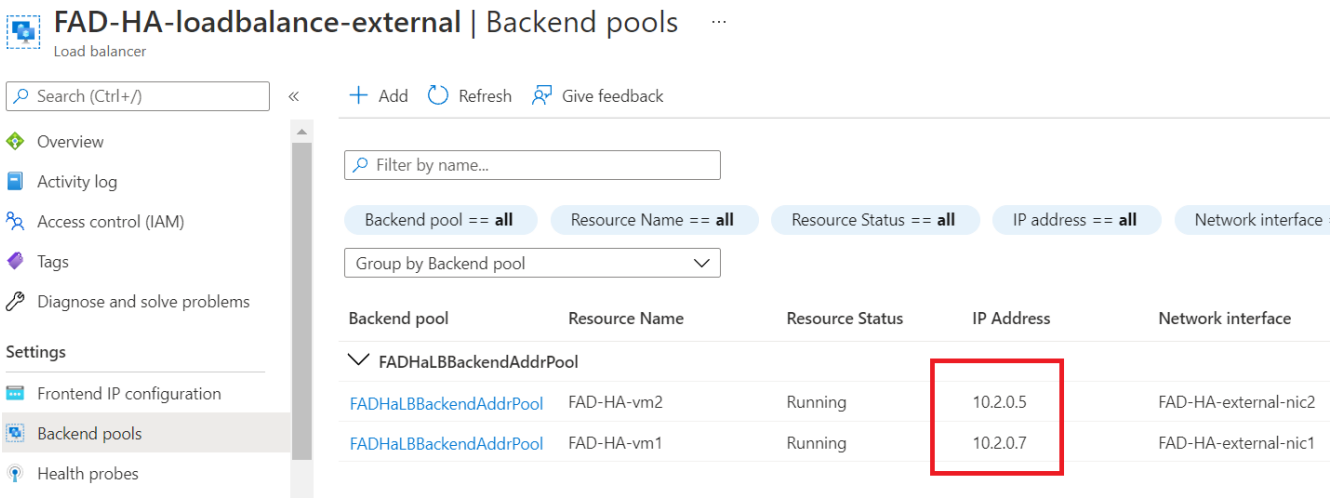
To configure the Azure LB Backend:

1. Navigate to **System > Azure LB Backend**.
2. Click **Create New** to display the configuration editor.
3. Configure the following settings:

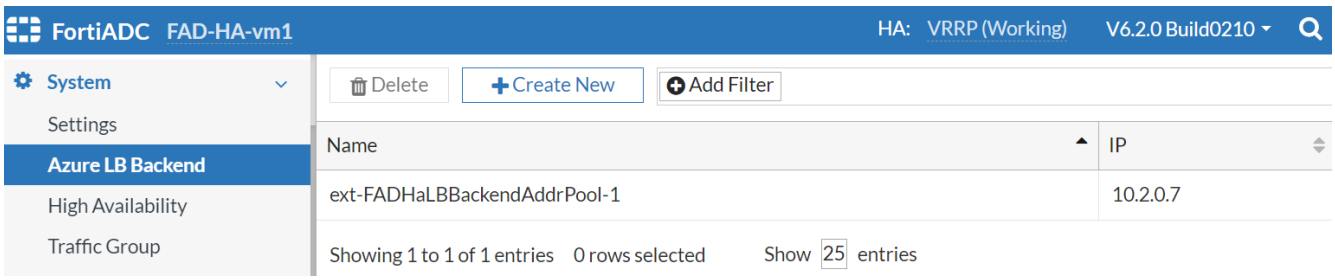
Settings	Guidelines
Name	Azure LB Backend Pool Name. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No spaces. After you initially save the configuration, you cannot edit the name.
IP	IP address used in the Azure LB Backend Pool. **Note**: The IP address would not be synchronized in HA mode.

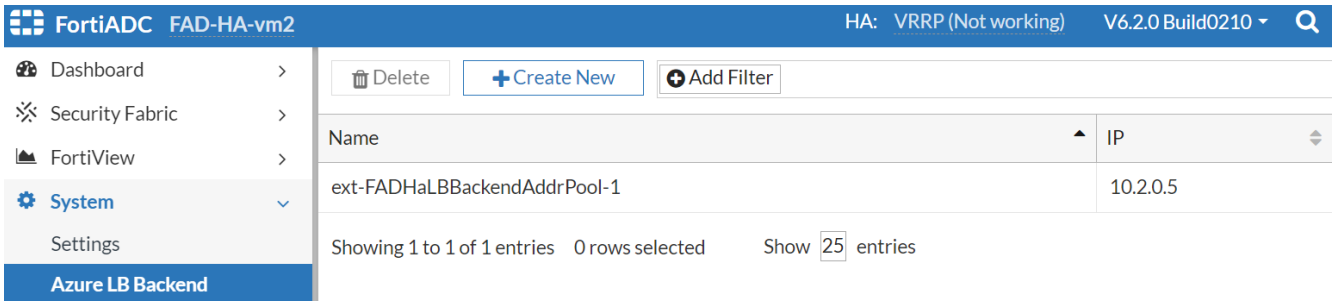
4. Save the configuration. Repeat the steps to configure the Azure LB Backend in the secondary FortiADC.

In this example, the Azure Load Balancer (ALB) backend pools contain both FortiADC-VMs in the `FADHaLBBackendAddrPool`, with IPs allocated from the `FadcOutSideSubnet`. The IP addresses assigned are `10.2.0.5` and `10.2.0.7`, respectively.



The resulting **Azure LB Backend** configurations on both FortiADC nodes will use these IP addresses from the ALB backend pools.





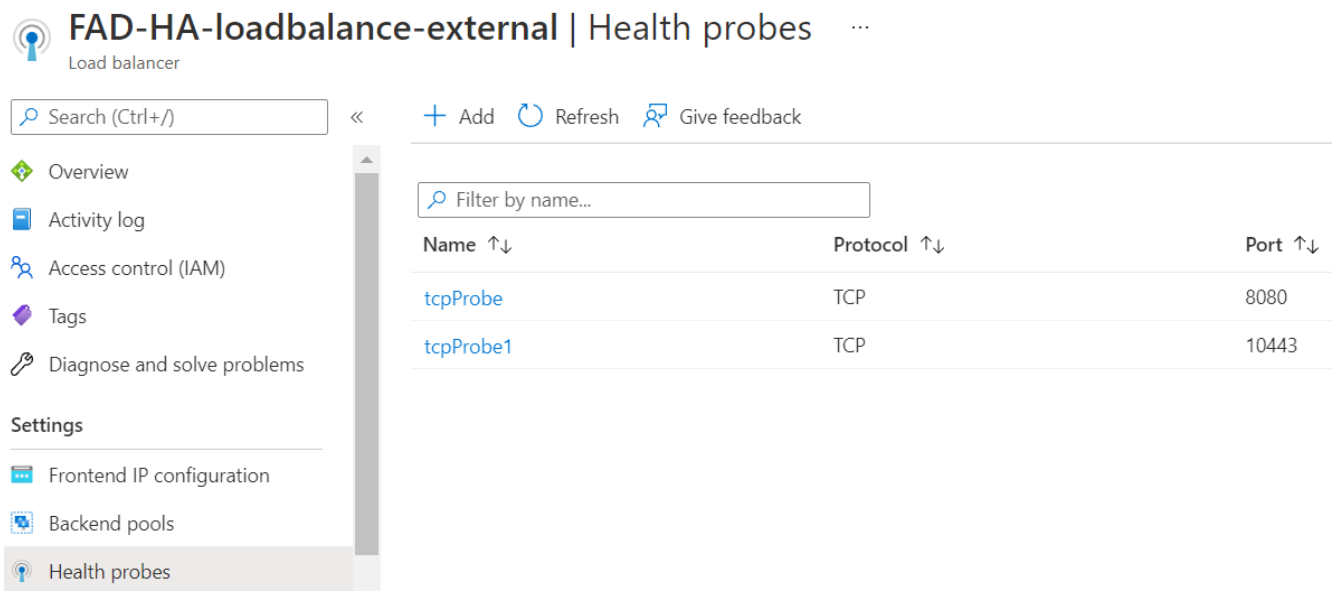
To configure the probing virtual server:

1. Navigate to **Server Load Balancer > Virtual Server**.
   The configuration page displays the **Virtual Server** tab.
2. Click **Create New > Advanced Mode** to open the Advanced Mode configuration editor.
3. Configure the following key settings from the **General** tab:

Settings	Guidelines
Use Azure LB Backend IP	Enable or disable the virtual server to reference Azure LB Backend IP as the virtual server IP.
Azure LB Backend IP	Select the Azure LB Backend IP Configuration used as the virtual server IP.
Port	Specify the port according to the ALB health probe information.
Profile	Select a profile that matches the protocol of the ALB health probe or the service the virtual server is intended to provide. For example, if the ALB health probe uses the TCP protocol, select a TCP profile. If the virtual server is intended to serve HTTPS traffic, select an HTTPS profile.

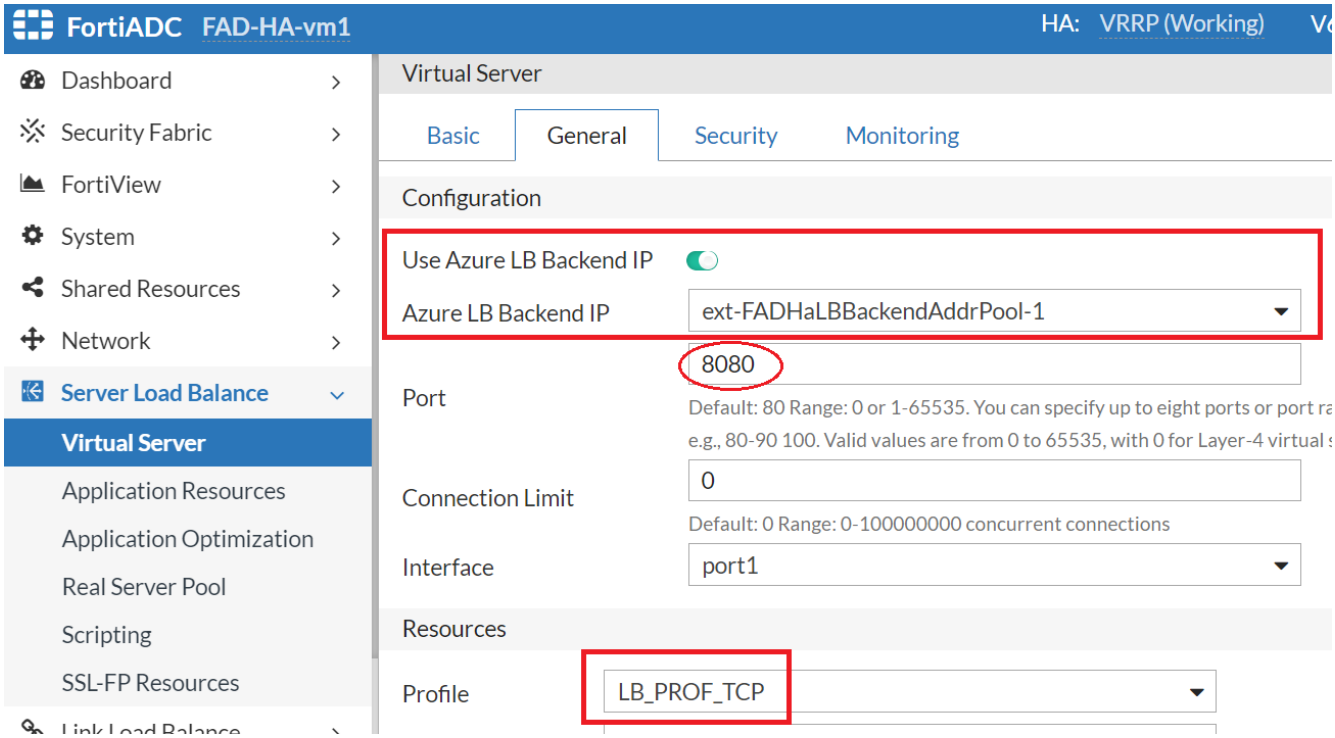
4. Save the configuration.
   The updated virtual server settings will be replicated to the secondary node through HA cluster synchronization.

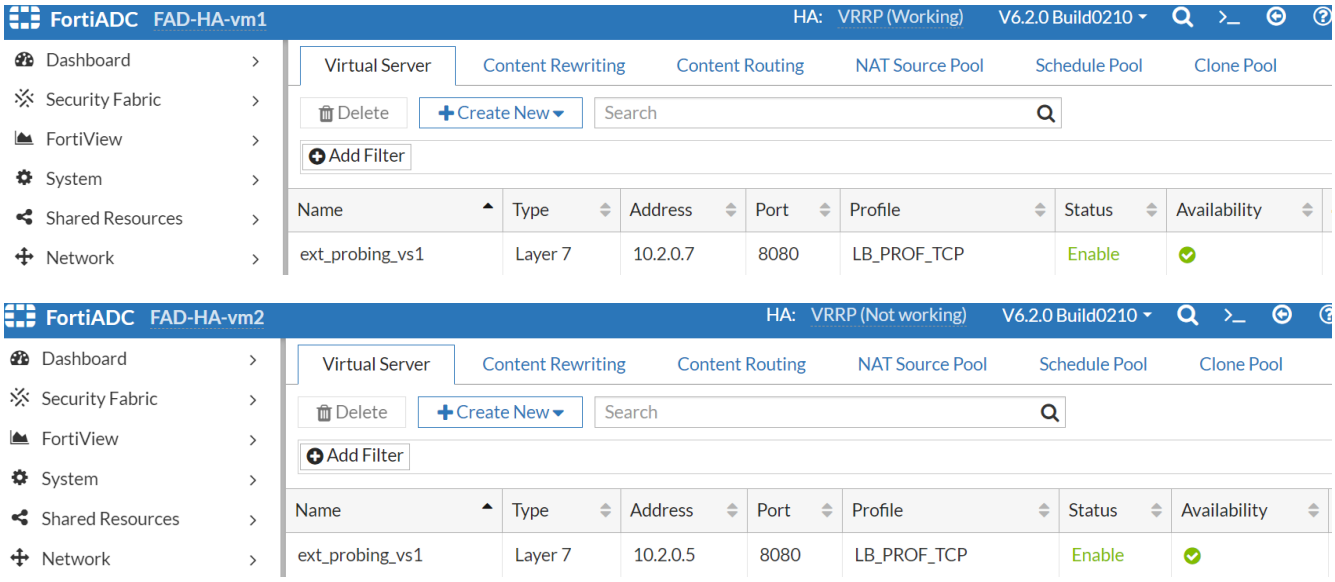
In this example, the Azure Load Balancer (ALB) health probes include two TCP health probes configured on ports `8080` and `10443`. When business traffic is received from users, `FAD-HA-loadbalance-external` directs the traffic to the active FortiADC-VM in the HA VRRP setup, which in this case is `FAD-HA-VM1`. The ALB determines the active node by sending health check packets to both FortiADC-VMs and identifying the node that responds to the health checks.



The resulting probing virtual server configurations on both FortiADC nodes utilize the ports defined by the ALB health probes. The probing virtual server is configured to use `ext-FADHaLBBackendAddrPool-1` as the virtual server IP, with port `8080` and the `LB_PROF_TCP` profile, matching the TCP health probe defined in the ALB configuration. After saving the configuration, the probing virtual server will automatically reference the IP address specified in the `ext-FADHAKBBackendAddrPOOL` entry configured within the Azure LB Backend settings.

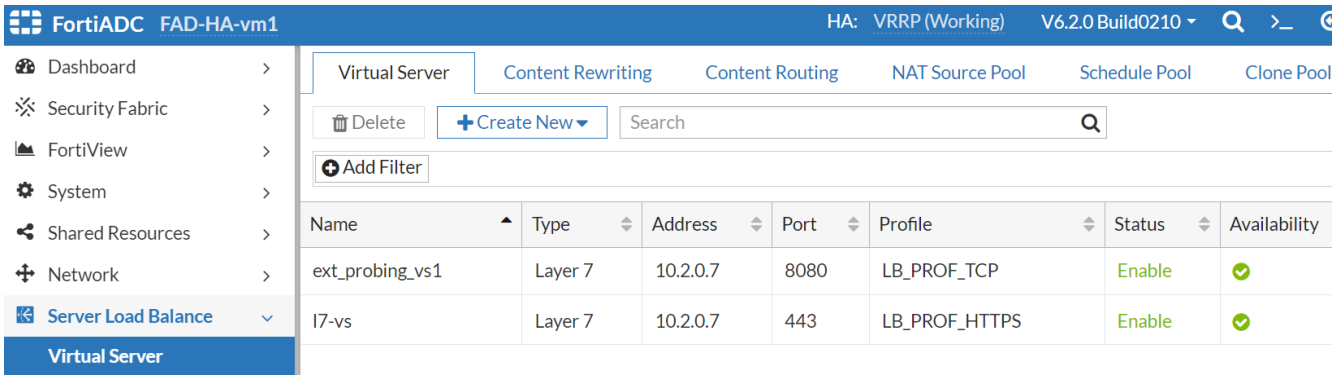
Note: For FortiADC version 6.2.0 or later, LB_PROF_TCP is deprecated. Please use LB_PROF_L7_TCP instead.





After configuring the probing virtual server, you can create a Layer 7 virtual server on `FAD-HA-VM1` to manage HTTPS business traffic. This virtual server is set to use port `443` with the `LB_PROF_HTTPS` profile and references the Azure Load Balancer backend IP. This setup ensures efficient and reliable handling of HTTPS traffic for users.

Note: Only the default route type is supported by the ALB in this setup.



Virtual Domain

The **System > Virtual Domain** sub-menu includes the following topics:

- Virtual Domain (VDOM) and Administrative Domain (ADOM) overview on page 171
- Enabling the Virtual Domain feature and selecting the Virtual Domain Mode on page 174
- Creating a virtual domain on page 175
- Assigning administrator users and network interfaces to VDOMs on page 175
- Virtual domain policies on page 176
- Disabling a virtual domain on page 179

Virtual Domain (VDOM) and Administrative Domain (ADOM) overview

A Virtual Domain (VDOM) is a complete FortiADC instance that runs on the FortiADC platform. VDOM configuration objects contain all of the system and feature configuration options of a full FortiADC instance and can be used to divide a FortiADC into two or more virtual units that function independently, allowing it to support multi-tenant deployments.

The VDOM feature supports two Virtual Domain Modes that allow the VDOMs to function independently with its own networking or as administrative domains (ADOMs) with shared networking between all ADOMs. When the VDOM is in the Independent Network mode, you can provision an administrator account with privileges to access and manage only their assigned VDOM. The VDOM user can then configure their VDOM as desired untethered to other VDOMs. Alternatively, when the VDOM is in Share Network mode, it functions as an ADOM that shares the same networking interfaces and routing between all the ADOMs. The ADOM functionality enables the administrator to constrain access privileges to a subset of server load-balancing servers by defaulting all interface settings to the root ADOM.

The Virtual Domains feature is not enabled by default and requires an administrator with "super admin" or "global admin" access to enable. The **admin** account holder (also known as the "super admin") can enable and configure all VDOMs and provision accounts with "global admin" access that grants administrators permissions to enable and configure VDOMs as well. The super admin and global admin have unrestricted access to all virtual domains that have been created on the system and can provision administrator accounts to access their assigned domains.

After the Virtual Domain feature is enabled, virtual domain administrators can enter their assigned VDOM/ADOM and see a subset of the typical menus or CLI commands appear, allowing access to only the feature configurations, logs and reports specific to their VDOM/ADOM. Unlike super admin and global admin users, VDOM/ADOM administrators do not have access to global settings.

Differences between super admin/global admin, and VDOM/ADOM administrators when virtual domains are enabled:

	Super admin or global admin user	VDOM/ADOM administrators
Access to global settings (`config global`)	Yes	No

	Super admin or global admin user	VDOM/ADOM administrators
Can create administrator accounts	Yes — administrator accounts can be assigned to access other virtual domains on the system.	Yes — administrator accounts can only be assigned access to the VDOM/ADOM administrator's own virtual domain.
Can create and access all VDOMs/ADOMs	Yes	No

Basic steps:

1. Enable the Virtual Domain feature and select the Virtual Domain Mode.
2. Create a VDOM or ADOM configuration object and assign administrators to the domain.
3. If the Virtual Domain Mode is Independent Network, then assign network interfaces and administrators to the VDOM.
   Note: If the Virtual Domain Mode is Share Network (ADOM mode), all network interface settings are restricted to the root settings.

GUI and CLI functional availability for administrators of VDOM, root ADOM, and non-root ADOM

For administrators provisioned to access only their assigned virtual domains, the GUI and CLI functions available to them depend on their Virtual Domain Mode and whether their virtual domain is root or non-root. VDOMs configured in the Independent Network mode function independently within its own network, allowing the VDOM administrator to have full unrestricted access to all configurations within their own VDOM. Administrators of VDOMs in the Independent Network mode have full unrestricted access to all configurations within their own VDOM; as these VDOMs function independently within their own network, modifications can be made without affecting other VDOMs on the system. In contrast, administrators of ADOMs (VDOMs in Share Network mode) do not have full access to all configurations due to all ADOMs sharing the same network interfaces and routing as the root ADOM. As a result, administrators of non-root ADOMs have restricted access, partial access, or completely no access to GUI and CLI functions relating to networking.

The following table lists the difference in GUI/CLI function availability between root and non-root ADOM administrators.

Configuration		Root ADOM	Non-root ADOM
Network	Interface	**Virtual Domain** option is hidden from the **Interface** settings. The interface settings are automatically defaulted to the root ADOM.	Read-only access for **Interface** settings. Data pulled from root ADOM.
	Routing	Read-write access for all configurations.	Read-only access for all configurations. Data pulled from root ADOM.
	NAT	Read-write access for all configurations.	No access to configurations. **NAT** is hidden.
	QoS	Read-write access for all configurations.	No access to configurations. **QoS** is hidden.

Configuration		Root ADOM	Non-root ADOM
Link Load Balance	*All configurations under Link Load Balance*	Read-write access for all configurations.	Read-only access for all configurations. Data pulled from root ADOM.
Global Load Balance	*All configurations under Global Load Balance*	Read-write access for all configurations.	No access to all configurations. **Global Load Balance** is hidden.
Network Security	Firewall	Read-write access for all configurations.	No access to all configurations. **Firewall** is hidden.
DoS Protection	Networking	Read-write access for all configurations.	Partial access: **IP Fragmentation Protection** and **TCP SYN Flood Protection** are hidden.
FortiView	Logical Topology	Read-write access for all configurations.	Partial access: **Global Load Balance** is hidden, and **Link Load Balance** is read-only with data pulled from root ADOM.
	Host	Read-write access for all configurations.	No access to all configurations. **Host** is hidden.
	Data Analytics (under Global Load Balance)	Read-write access for all configurations.	No access to all configurations. **Data Analytics** under **Global Load Balance** is hidden.
	Gateway	Read-write access for all configurations.	Read-only access to **Link Load Balance** data pulled from root ADOM. The **Monitor** option is hidden.
	Interfaces	Read-write access for all configurations.	All FortiADC interfaces are shown. Data pulled from root ADOM.
Log & Report	Log Setting	Read-write access for all configurations.	Partial access: **Link Load Balance (LLB)**, **Global Load Balance (GLB)**, and **Firewall (FW)** options are hidden from the **Local Log** and **Fast Stats** settings.
	Traffic Log	Read-write access for all configurations.	Partial access: **Link Load Balance (LLB)** and **Global Load Balance (GLB)** filter options are hidden.
	Security Log	Read-write access for all configurations.	Partial access: **Firewall** filter option is hidden.

Configuration		Root ADOM	Non-root ADOM
	Event Log	Read-write access for all configurations.	Partial access: **Link Load Balance (LLB)**, **Global Load Balance (GLB)**, and **Firewall** filter options are hidden.
	Report Setting	Read-write access for all configurations.	**DNS-Top-Policy-by-Count** and **DNS-Top-Source-by-Count** are not supported in **Query Set**.

Enabling the Virtual Domain feature and selecting the Virtual Domain Mode

By default, the Virtual Domain feature must be enabled for the Virtual Domain configuration to be visible in the GUI. Once you enable Virtual Domain, you can select the Virtual Domain Mode to determine the VDOM networking options.

There are two virtual domain modes:

- Independent Network — each VDOM functions independently within its own network, unaffected by activity from other VDOMs on the system.
- Share Network — VDOMs function as administrative domains (ADOMs), sharing the same network interface and routing between all ADOMs.



Once configured, switching between the virtual domain modes is not recommended. If you need to switch virtual domain modes, the Virtual Domain feature must first be disabled. For details, see Disabling a virtual domain on page 179.

Before you begin:

- You must have super admin (**admin** administrator) or global admin permission to enable the Virtual Domain feature.

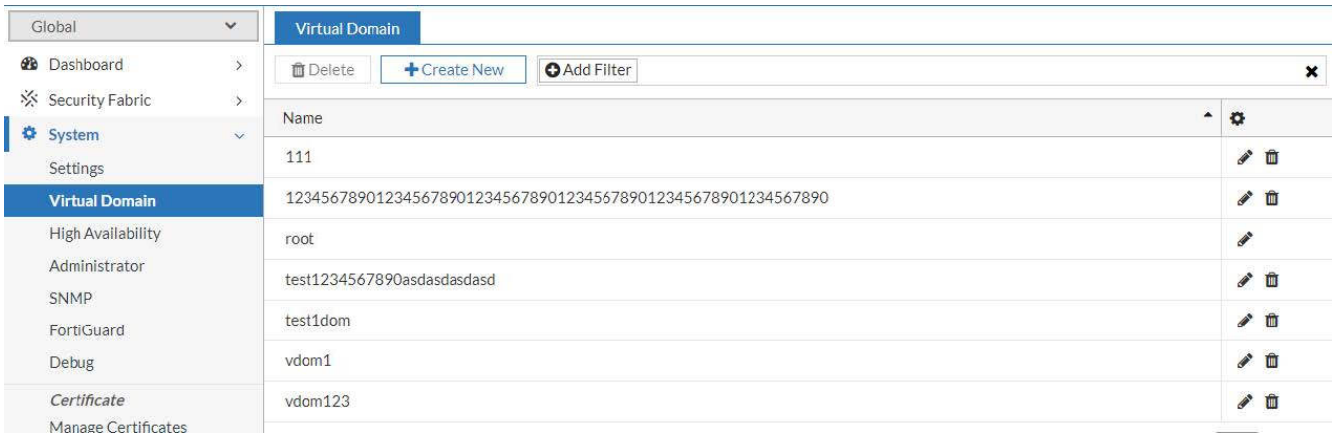
To enable the Virtual Domain and select the Virtual Domain Mode:

1. Go to **System > Settings**.
   The configuration page displays the **Basic** tab.
2. Enable **Virtual Domain**.
   The **Virtual Domain Mode** field appears.
3. In the **Virtual Domain Mode** field, select either **Independent Network** or **Share Network**.
4. Click **Save**.

Super admin login with virtual domain on page 174 shows the landing page after the super admin logs into the system when the Virtual Domain feature is enabled. From here, the super admin 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 super admin 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 super admin and global admin users are able to add, delete, enable, and disable any non-root virtual domains on the system.

Before you begin:

- You must have super admin (**admin** administrator) or global admin permission to create virtual domains.

To create a virtual domain:

1. Go to **System > Virtual Domain**.
2. Click **Create New**, enter a unique name for the virtual domain.
3. Click **Save**.
   The newly created virtual domain is listed under the Virtual Domain page.
4. Optionally, double-click or click the (edit icon) of the virtual domain to create and impose custom policies.
   This function is only applicable for virtual domains in Independent Network mode. For more information, see Virtual domain policies on page 176.
   For more information about the Independent Network and Share Network virtual domain modes, see Virtual Domain (VDOM) and Administrative Domain (ADOM) overview on page 171.

Assigning administrator users and network interfaces to VDOMs

After creating the virtual domain, you can assign administrator users to manage it. These virtual domain administrators can access only the domain they are assigned.

For virtual domains in Independent Network mode, you need to assign network interfaces to the virtual domain. If the Virtual Domain Mode is Share Network (ADOM mode), all network interface settings are defaulted to the root settings, so assigning network interfaces is unnecessary.

Before you begin:

- You must have super admin (**admin** administrator) or global admin permission to assign administrator users and network interfaces to virtual domains.

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.

When virtual domain administrators log into the FortiADC system, they do not see the Virtual Domain menu in the GUI. Instead, they only see configuration settings and data for the virtual domain they have been assigned to. Furthermore, the difference in what GUI and CLI functions are available to virtual domain administrators also depend on whether the Virtual Domain Mode is Independent Network or Share Network. For more information, see Virtual Domain (VDOM) and Administrative Domain (ADOM) overview on page 171.

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.

Virtual domain policies

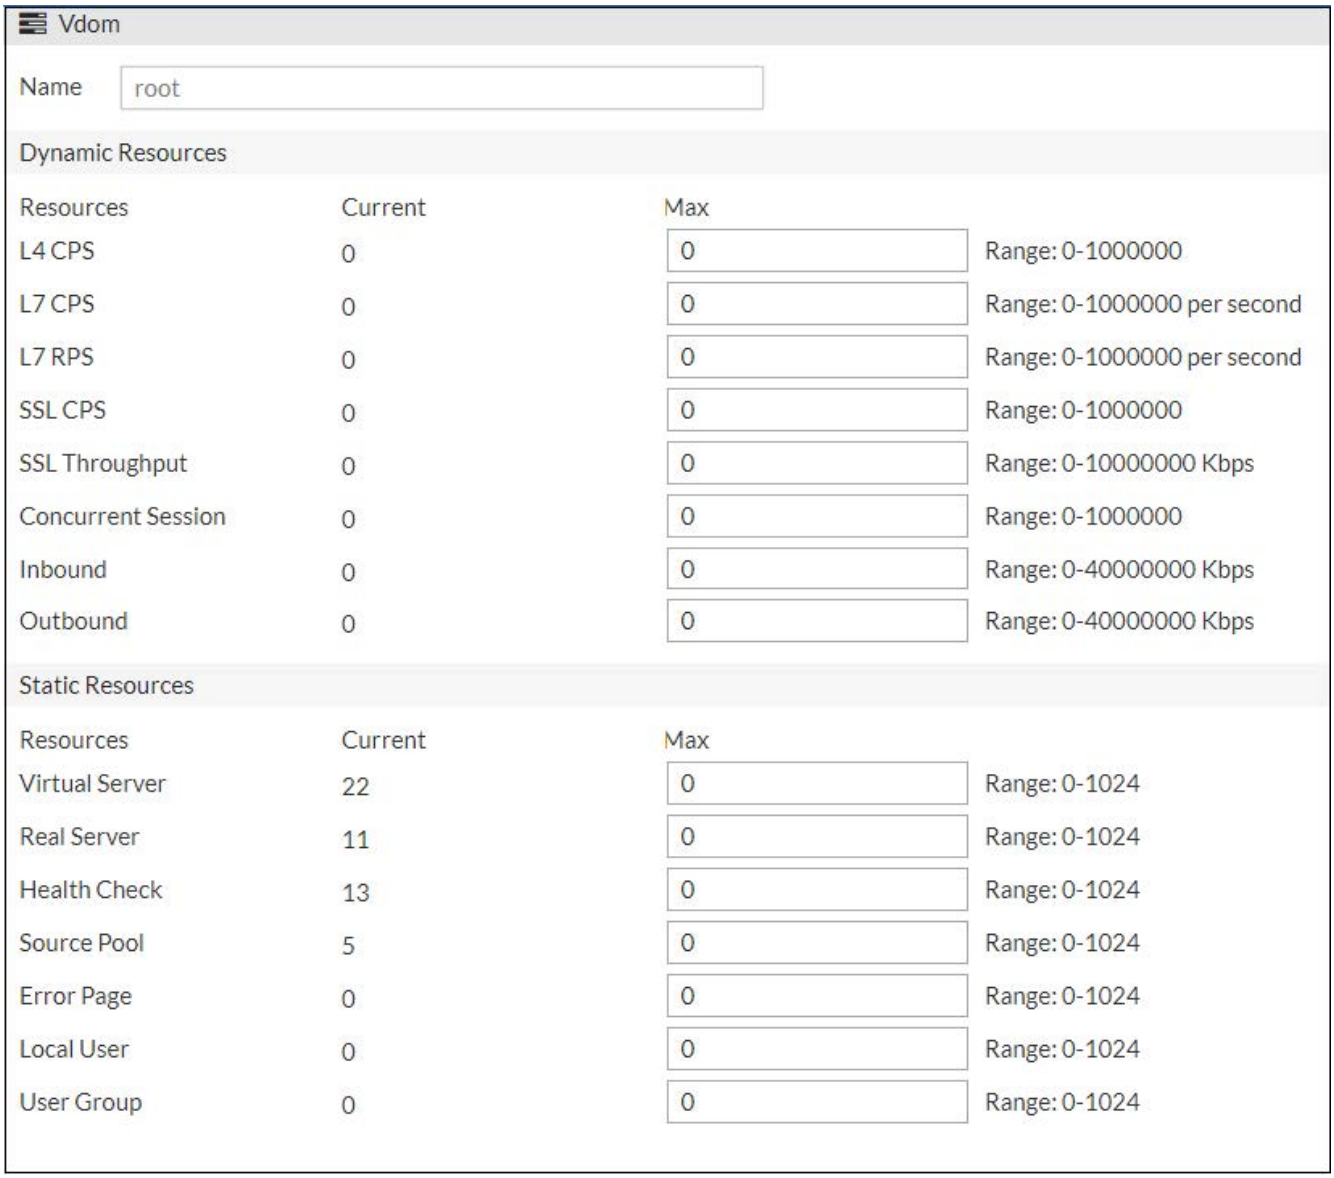
For virtual domains in Independent Network mode, FortiADC allows you to create and apply custom policies or restrictions on each virtual domain you have added. For each virtual domain, you can configure the maximum range for its Dynamic Resources and Static Resources. Dynamic Resources are related to a virtual domain's performance, while Static Resources are related to its configuration. The Vdom configuration dialog on page 176 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. (For more information about the Independent Network and Share Network virtual domain modes, see Virtual Domain (VDOM) and Administrative Domain (ADOM) overview on page 171.)

By default, all per-VDOM resource settings are set to have no limits. This means that any single VDOM can use all of the FortiADC device's resources. This could deprive other VDOMs of the resources that they require, to the point that could be unable to function. We recommend setting maximum values on the resources that are vital to you.



Custom virtual domain policies is not supported for FortiADC 5000F model.

Vdom configuration dialog



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.

Parameter	Description
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 *TCP* 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 *TCP* 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. 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.

Parameter	Description
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

The Virtual Domain feature can be disabled through the Web UI. You may need to disable virtual domains in certain scenarios, such as switching to a different Virtual Domain Mode.

Before you begin:

- You must have super admin (**admin** administrator) or global admin permission to disable the Virtual Domain feature.
- You must have deleted all non-root virtual domains on the system, leaving only the root (default) virtual domain.

To disable the Virtual Domain feature:

1. Go to **System > Settings**.
   The configuration page displays the **Basic** tab.
2. Disable **Virtual Domain**.
3. Click **Save**.
   The FortiADC should refresh automatically.

High Availability

FortiADC appliances can be deployed either as standalone units or within High Availability (HA) clusters.

Overview

In a FortiADC HA cluster, the configuration consists of a minimum of two and up to eight nodes, with each node representing an instance of the FortiADC appliance. The role of each node in traffic handling is dictated by the HA deployment mode. In both Active-Passive and Active-Active HA modes, the cluster elects a single Primary node, while the remaining nodes operate as Secondary. In an Active-Passive deployment, the Primary node exclusively processes all incoming traffic, with Secondary nodes in standby mode. In contrast, in an Active-Active cluster, all nodes concurrently process traffic, distributing the load across the cluster.

Regardless of the deployment mode, nodes continuously monitor each other's operational status via HA heartbeat messages, which are periodically transmitted across all nodes. If the Primary node fails, the HA cluster initiates its election mechanism, selecting a new Primary node from the pool of Secondary nodes to ensure high availability and prevent service disruption.

The election of the Primary node follows a hierarchical set of criteria:

1. Link health — Nodes with failed monitor port links are marked as down.
2. Remote IP monitor health — Assesses the status of remote IP health checks.
3. Override setting — Prioritizes the configured priority setting over uptime.
4. Port availability — The node with the highest number of available ports is preferred.
5. Uptime — The node with the longest uptime is favored.
6. Device priority — Nodes with a lower priority number take precedence (for example, priority 1 over priority 2).
7. Serial number sorting — Serial numbers are evaluated character by character from left to right, with higher alphanumeric values (such as '9' and 'z') given precedence.

Heartbeat monitoring forms the backbone of the failover mechanism in any HA mode. Each node sends periodic heartbeat packets to all other nodes within the cluster. If a node ceases to transmit heartbeat traffic, it is considered non-operational, triggering the failover sequence, which redistributes traffic to the remaining operational nodes to maintain continuity.

Full configuration synchronization in the HA cluster is controlled by the Configuration Source, a designated node responsible for replicating configuration settings across the cluster. During cluster initialization, full configuration synchronization is typically initiated from this node, propagating its configuration to other members, with exceptions defined by the configuration policies. Post-initialization, if session and persistence settings are enabled for HA, nodes continuously synchronize their session and persistence tables in real time. This ensures uniform session and persistence state information across the cluster, which is critical for enabling stateful failover and maintaining load balancing with minimal disruption during node transitions.

HA deployment modes

FortiADC offers three High Availability (HA) deployment modes, each optimized for different operational requirements and deployment scenarios. Select the HA mode that best aligns with your performance requirements and deployment capabilities.

Active-Passive (HA-AP)

In an Active-Passive HA configuration, a single Primary node processes all traffic and handles all data operations, while up to seven Secondary nodes remain in standby mode. If HA specific management IP is not configured, the Primary node's management IP address is the only one active, and administrative access is permitted only when this node is in the Primary role. During failover, if the Primary node becomes unavailable or requires maintenance, a Secondary node automatically assumes the Primary role, ensuring uninterrupted service and seamless transition. To manage a standby node, users must either connect through a console port or configure the Management Interface to access all nodes via the Management IP. HA-AP delivers a highly stable and straightforward deployment with minimal prerequisites, making it suitable for environments where high reliability and ease of management are essential.

For more information and detailed steps to deploy Active-Passive HA, see Deploying an Active-Passive cluster on page 207.

Active-Active (HA-AA)

In HA-AA mode, all nodes actively process traffic, distributing the load and ensuring high availability. This configuration requires a complex setup and strict adherence to specific deployment conditions to achieve optimal performance. It is designed for high-demand environments needing maximum performance and advanced load balancing capabilities. In an Active-Active cluster, each node has a unique IP address for all interfaces, including the management interface. When operating in standalone mode, the physical port IP address is active; in HA mode, the address assigned from the HA node IP list is used. Administrative access is available via the active management IP address for any node. This configuration supports robust performance and scalability, provided all ideal deployment conditions are met.

For more information and detailed steps to deploy Active-Active HA, see Deploying an Active-Active cluster on page 218.

Active-Active-VRRP (HA-VRRP)

HA-VRRP employs a VRRP-like protocol to enable concurrent traffic handling across all FortiADC nodes configured with multiple traffic groups, enhancing throughput and performance through effective load balancing and redundancy. This mode requires fewer deployment conditions than HA-AA, making it suitable for environments needing increased performance and active traffic management. In an Active-Active-VRRP cluster, FortiADC uses the Heartbeat Interface for internal node status communication and supports synchronization of the sessions, persistence, and image through both the Heartbeat Interface and Data Interface. Although named similarly, FortiADC does not implement the VRRP protocol and cannot interact with third-party VRRP devices, while still ensuring robust internal synchronization and management capabilities.

For more information and detailed steps to deploy Active-Active-VRRP HA, see Deploying an Active-Active-VRRP cluster on page 236.

HA system requirements

Ensure the following requirements are met to configure HA in FortiADC.

- Hardware and firmware consistency — All appliances must be of the same hardware model and run the identical firmware version.
- Redundant network topology — The network must be designed to handle node failures, with physical cabling and routes capable of redirecting traffic to the remaining member nodes.
- Physical ports for communication — Each HA appliance must have at least one physical port dedicated to heartbeat and data traffic between cluster members. For Active-Passive configurations, connect the ports directly with a

crossover cable. For Active-Active clusters, connect the nodes using a shared Layer 2 switch.

- Heartbeat and synchronization traffic — Ensure that heartbeat and synchronization traffic between nodes is carried over the designated physical network ports. If switches are used, the interfaces must support Layer 2 multicast for proper communication. When applying Synchronization settings, configure the Heartbeat Interface (responsible for HA status communication) and the Data Interface (used for traffic forwarding) on separate physical ports to prevent potential conflicts and ensure optimal performance. Using distinct ports minimizes the risk of latency and packet loss, which could otherwise disrupt HA synchronization and impact data traffic.
- Licensing requirements — Each appliance must have a valid license. For FortiADC-VM, a paid license is required; trial licenses are not supported.



FortiADC-VM supports High Availability (HA). However, if you prefer not to utilize the native HA functionality, you can leverage your hypervisor or virtual machine environment manager to enhance availability. For instance, in VMware environments, you can utilize vMotion or VMware HA to manage virtual appliances across a hardware cluster, improving overall system resilience and uptime.

HA configuration and monitors

The following sections cover the steps on how to view and configure settings required for HA.

- Configuring HA settings on page 183
- HA configuration synchronization on page 192
- Monitoring an HA cluster on page 194
- Updating firmware for an HA cluster on page 200
- HA Troubleshooting on page 202

Configuring HA settings

Configure the HA settings of the local node. The configuration options are based the supported cluster modes: Standalone, Active-Passive, Active-Active, and Active-Active-VRRP.

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. Click the **Configuration** tab.
3. Click the (edit icon) to launch the configuration editor.
4. Select the **Cluster Mode**: Standalone, Active-Passive, Active-Active, or Active-Active-VRRP.
   If **Standalone** mode is selected, no additional settings are required. Otherwise, configure the **Basic** required settings for the selected mode.

a. Configure **Active-Passive** Basic settings:

Setting	Description
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	Allows you to determine which configuration the system uses when synchronizing the configuration between the HA nodes. 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 primary-secondary 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. Range of acceptable values is 0 to 255. Default is 100.
Monitor Interface	One or more network interfaces that correlate with a physical link. These ports will be monitored for link failure. Port monitoring (also called interface monitoring) monitors physical network ports to verify that they are functioning properly and linked to their networks. You can monitor physical interfaces and 802.3ad aggregated interfaces. **Note:** To prevent an unintentional failover, do not configure port monitoring *until* you configure HA on all appliances and have plugged in the cables to link the physical network ports that will be monitored.
Heartbeat Interface	Set the network interface to be used to receive and send the HA heartbeat packet between the HA group members. Use the same port numbers for all cluster members. **Note**: To avoid heartbeat loss and ensure reliable HA synchronization, configure the Heartbeat and Data Interfaces on separate physical ports. This prevents conflicts and minimizes failover risks.

Setting	Description
Data Interface	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. **Note**: To avoid heartbeat loss and ensure reliable HA synchronization, configure the Heartbeat and Data Interfaces on separate physical ports. This prevents conflicts and minimizes failover risks.

b. Configure **Active-Active** Basic settings:

Setting	Description
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	Allows you to determine which configuration the system uses when synchronizing the configuration between the HA nodes. 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 primary-secondary 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. Range of acceptable values is 0 to 255. Default is 100.
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.
Monitor Interface	One or more network interfaces that correlate with a physical link. These ports will be monitored for link failure. Port monitoring (also called interface monitoring) monitors physical network ports to verify that they are functioning properly and linked to their networks. You can monitor physical interfaces and 802.3ad aggregated interfaces.

Setting	Description
	Note: To prevent an unintentional failover, do not configure port monitoring *until* you configure HA on all appliances and have plugged in the cables to link the physical network ports that will be monitored.
Heartbeat Interface	Set the network interface to be used to receive and send the HA heartbeat packet between the HA group members. Use the same port numbers for all cluster members. **Note**: To avoid heartbeat loss and ensure reliable HA synchronization, configure the Heartbeat and Data Interfaces on separate physical ports. This prevents conflicts and minimizes failover risks.
Data Interface	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. **Note**: To avoid heartbeat loss and ensure reliable HA synchronization, configure the Heartbeat and Data Interfaces on separate physical ports. This prevents conflicts and minimizes failover risks.
Node List	Select the node IDs for the nodes in the cluster. An Active-Active cluster can have up to eight members (0-7).

c. Configure **Active-Active-VRRP** Basic settings:

Setting	Description
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.

Setting	Description
Config Priority	Allows you to determine which configuration the system uses when synchronizing the configuration between the HA nodes. 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 primary-secondary 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. Range of acceptable values is 0 to 255. Default is 100.
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.
Heartbeat Interface	Set the network interface to be used to receive and send the HA heartbeat packet between the HA group members. Use the same port numbers for all cluster members. **Note**: To avoid heartbeat loss and ensure reliable HA synchronization, configure the Heartbeat and Data Interfaces on separate physical ports. This prevents conflicts and minimizes failover risks.
Data Interface	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. **Note**: To avoid heartbeat loss and ensure reliable HA synchronization, configure the Heartbeat and Data Interfaces on separate physical ports. This prevents conflicts and minimizes failover risks.
Heartbeat Type	Select one of the following: • Multicast • Broadcast • Unicast
Local Address	The **Local Address** option is available if the **Heartbeat Type** is **Unicast**. Enter the local IP address. This IP address should be set for the physical interface that is used for the heartbeat interface.
Peer Address	The **Peer Address** option is available if the **Heartbeat Type** is **Unicast**. Enter the peer IP address(es). A minimum of one peer IP address is required, and a maximum of seven peer IP addresses is supported.

5. Configure the **Synchronization** settings:



When applying Synchronization settings, configure the Heartbeat Interface (responsible for HA status communication) and the Data Interface (used for traffic forwarding) on separate physical ports to prevent potential conflicts and ensure optimal performance. Using distinct ports minimizes the risk of latency and packet loss, which could otherwise disrupt HA synchronization and impact data traffic.

Setting	Description
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 automatically synchronized across HA members. • When disabled, 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 automatically synchronized across HA members. • When disabled, 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 disabled, 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.
Note: The Synchronization settings are not applicable to Firmware Upgrade.	

6. Configure the **Advanced** settings:

Setting	Description
Priority	Number indicating priority of the member node when electing the cluster primary node. The smaller the number, the higher the priority. The valid range is 0 to 9. The default is 5. This setting is optional and does not affect the election of the configuration source. **Note**: By default, unless you enable Override, uptime is more important than this setting.
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). **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:

Setting	Description
	• 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. • 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.
Remote IP Monitor	Enable or disable active monitoring of remote beacon IP addresses to determine if the network path is available. This option is disabled by default.
Failover Threshold	The **Failover Threshold** option is available if **Remote IP Monitor** is **enabled**. The number of unreachable remote beacon IP addresses under the Remote IP Monitor List that must be occur to indicate network failure. The default is 5. The valid range is 1-64.
Failover Hold Time	The **Failover Hold Time** option is available if **Remote IP Monitor** is **enabled** If failover occurs due to a remote IP monitor test, and this node's role changes (to primary or secondary), 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.

a. If the **Remote IP Monitor** option is enabled, add the remote beacon IP addresses under the **Remote IP Monitor List** section.

i. Click **Create New** to display the configuration editor and configure the following Remote IP Monitor List settings:

Setting	Description
Name	Enter a unique name for the configuration. Valid characters are A-Z, a-z, 0-9, _, and –. No space is allowed.
Remote Address	Specify the remote IP address to ping.
Source Port	Specify the interface to send the health check ping.
Health Check Interval	Specify the number of seconds between each health check. Should be more than the timeout to prevent overlapping health checks. The valid range is 1-3600 seconds. The default is 10 seconds.
Health Check Timeout	Specify the number of seconds to wait for a reply before assuming that the health check has failed. The valid range is 1-3600 seconds. The default is 5 seconds.
Health Check Retry	Specify the number of retries to confirm up or down. The valid range is 1-10. The default is 3 retries.

ii. Click **Save** to confirm the settings and exit the Remote IP Monitor List configuration editor.

7. Click **Save** to confirm the HA settings.

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.

HA configuration synchronization

In an HA configuration, the node designated as the Configuration Source is responsible for pushing most of its configuration to the other member nodes in the cluster, a process referred to as HA configuration synchronization. When automatic synchronization is enabled, this process occurs seamlessly whenever a new appliance joins the cluster. If automatic synchronization is disabled, synchronization must be initiated manually.

The following components are included in HA configuration synchronization:

- Core configuration file (`fadc_system.conf` in CLI format)
- X.509 certificates, certificate signing requests (CSRs), and private keys
- Layer-7 virtual server error message files
- Layer-4 TCP connection states, Layer-4 persistence tables, and Layer-7 persistence tables (Source Address Persistence only)
- Health check status (specific to Active-Passive deployments)

For most settings, configuration changes are applied directly to the Configuration Source node, and those settings are then propagated to the other member nodes. However, certain configuration settings are not synchronized. These exceptions are outlined in the table below, while all other settings are synchronized across the cluster.

Setting	Explanation
Hostname	Hostnames are intentionally not synchronized, allowing each node to maintain a unique name within the cluster.
SNMP system information	Each member node maintains its own SNMP system information, allowing for accurate and distinct data collection across the cluster. However, the network interfaces of a standby node remain inactive, preventing active SNMP monitoring for those interfaces.
RAID level	RAID settings are hardware-dependent and determined at boot time by inspecting the drives for software RAID or the controller for hardware RAID. Since these settings are not part of the system configuration, they are not synchronized across the cluster.
HA settings	Most HA configuration settings are intentionally not synchronized to ensure proper HA system operations. Specifically: • Priority and Override settings — These parameters are used in the election process for the Primary node and are kept unsynchronized to allow differentiation between nodes. • Group ID — This identifier allows nodes with the same Group ID to form a cluster. Since it determines cluster membership, it must be set manually and is not synchronized. • HA mode — Administrators may prefer flexibility in switching a node between HA and standalone modes without affecting the entire cluster. For this reason, the HA mode setting is not synchronized, allowing a Primary node to be switched to standalone mode without impacting Secondary nodes, or for a Secondary node to operate in standalone mode without being overwritten by periodic synchronization. • Node list and Local Node ID — These settings apply specifically to Active-

Setting	Explanation
	Active mode and are used to uniquely identify each node within an Active-Active load-balancing group. To maintain this distinction, these settings are not synchronized across the cluster.
HA management settings	HA management settings are intentionally not synchronized, allowing you to configure unique, non-synchronized management interface details on each FortiADC appliance. These HA management IP addresses can be used across all HA modes, independent of the appliance's HA role.

In addition to HA configuration settings, the following data is also **not** synchronized across the cluster:

- Log messages — These logs capture events specific to individual nodes. In the event of a failover, gaps may appear in the log files of the original active node during its downtime. Any log messages generated while the standby node assumes the active role (assuming local log storage is configured) are stored locally on the former standby node and are not transferred back to the original node.
- Generated reports — Similar to log messages, reports are based on events logged on a specific appliance. While report configurations are synchronized across the cluster, the report data itself is node-specific and is not propagated to other members.

You can view the status of cluster members from the dashboard of the primary node. You may need to log in to the system for a non-primary member node in the following situations:

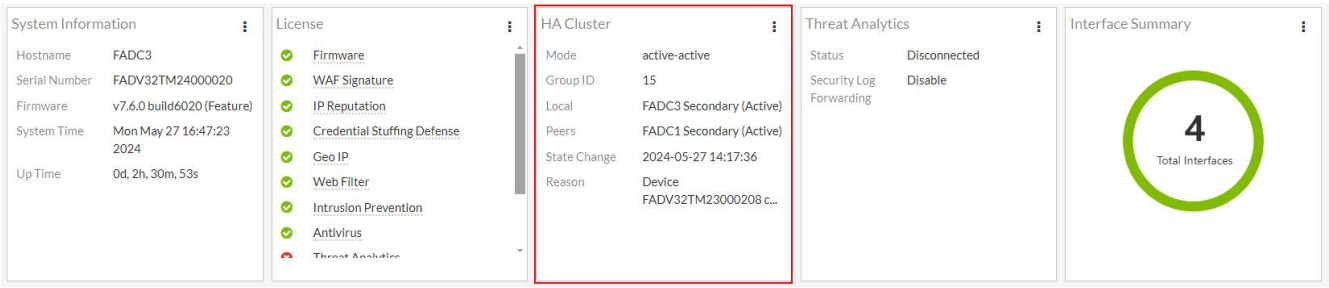
- To configure settings that are not synchronized.
- To view log messages recorded about the member node itself on its own hard disk.
- To view traffic reports for traffic processed by the member node.

Monitoring an HA cluster

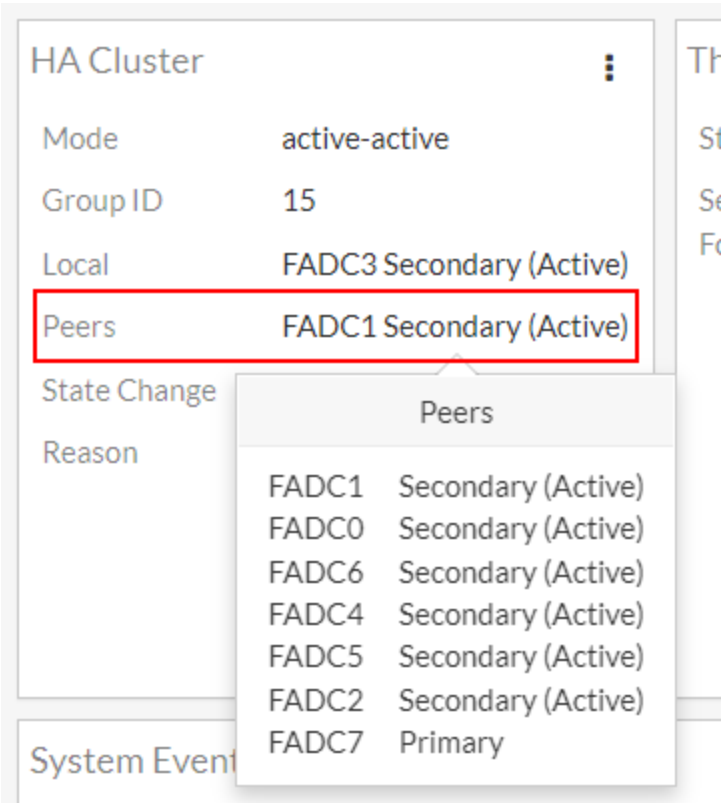
FortiADC provides several tools to monitor the status of your HA cluster from the Dashboard, Top navigation, and the High Availability pages.

HA Cluster Dashboard widget

From the **Dashboard > Main** page, you can view the HA status from the system dashboard HA Cluster widget. You can click the ⋮ (vertical ellipsis icon) and click **See Detail** to view the topology of the HA cluster.



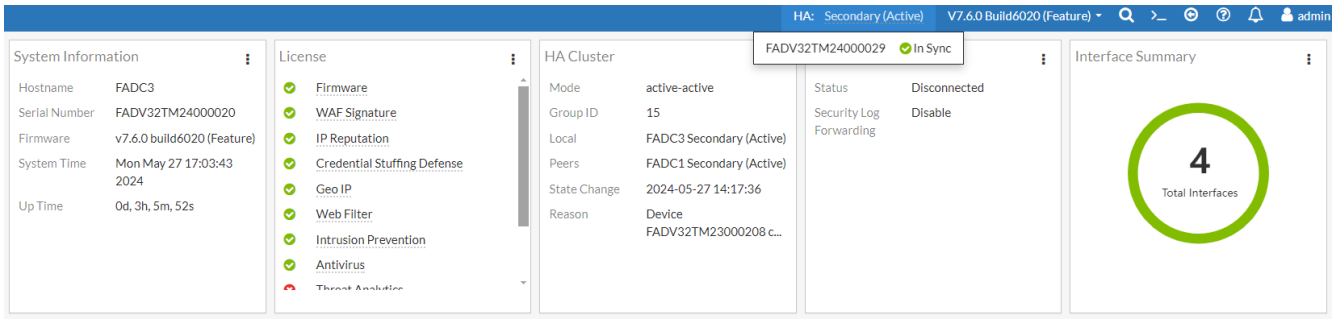
When you hover over the **Peers** field, it displays the list of Peer nodes that are part of the HA cluster and their statuses.



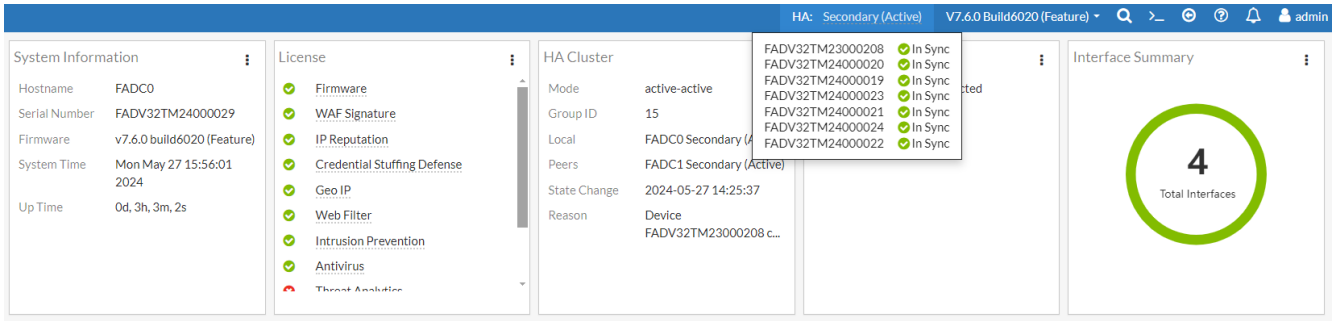
Top Navigation bar

The HA status of the local node is displayed in the Top Navigation bar. When you hover over this status, you can view its HA synchronization status.

In this example, you will see that the local node FADC3 is currently In Sync with another device with the serial number FADV32TM24000029. This means that the local node is not the configuration source, the device that it is currently synchronized to is the configuration source.

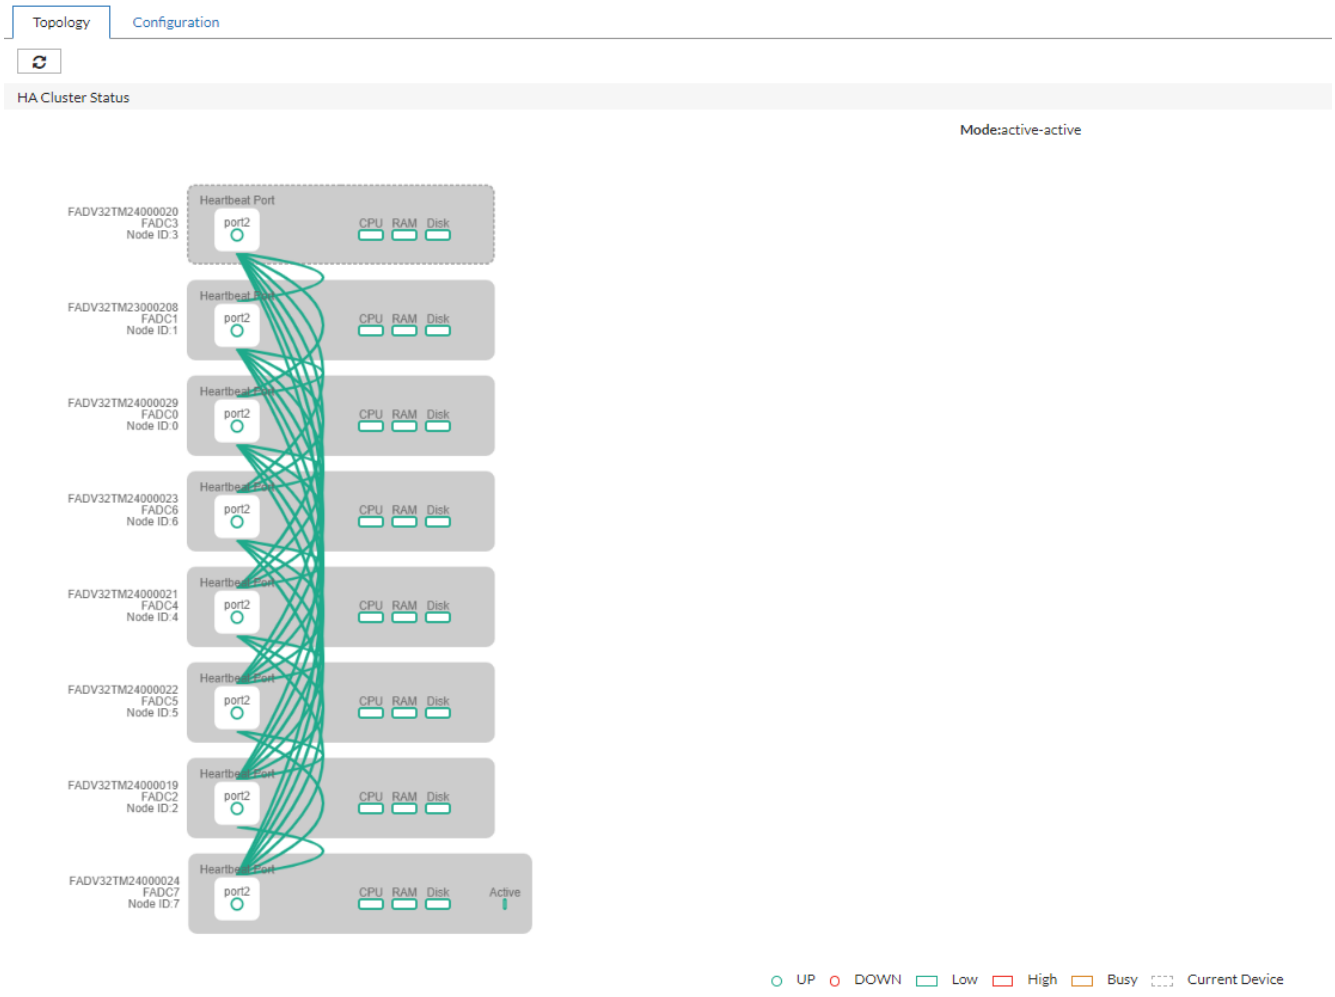


Below is what is displayed if your local node is the configuration source. You will see the list of devices that are currently synchronized to its HA configuration settings, including node FADC3 (you can verify this with the Serial Number from the System Information widget).

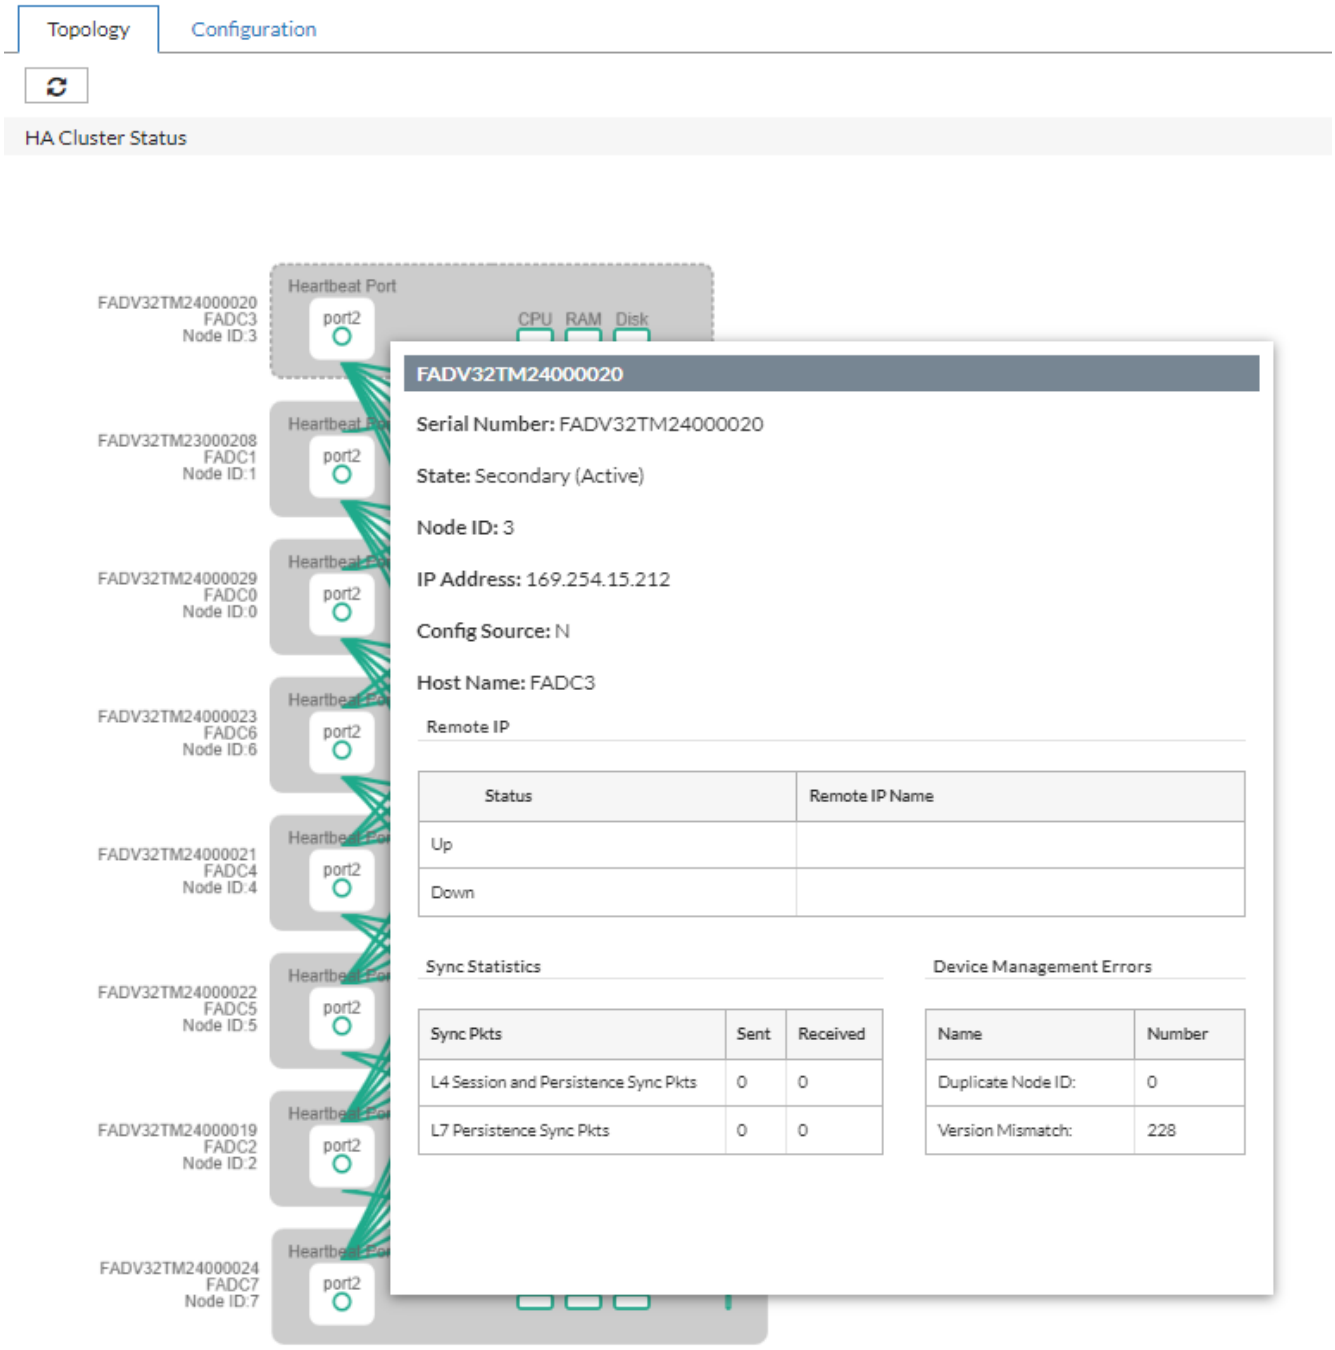


High Availability > Topology page

To view the topology of your HA cluster, navigate to **System > High Availability > Topology**.



Hover over any device to drill down to the detailed information about the HA cluster member.



High Availability > Configuration page

To view the list of member nodes in your HA cluster, navigate to **System > High Availability > Configuration**.

Topology | Configuration

Host Name	State	Serial Number	Node ID	IP Address	Config Source	
FADC0	Secondary (Active)	FADV32TM24000029	0	169.254.255.106	Y	
FADC1	Secondary (Active)	FADV32TM23000208	1	169.254.47.63	N	
FADC3	Secondary (Active)	FADV32TM24000020	3	169.254.15.212	N	
FADC2	Secondary (Active)	FADV32TM24000019	2	169.254.82.86	N	
FADC6	Secondary (Active)	FADV32TM24000023	6	169.254.53.141	N	
FADC4	Secondary (Active)	FADV32TM24000021	4	169.254.10.116	N	
FADC7	Primary	FADV32TM24000024	7	169.254.8.125	N	
FADC5	Secondary (Active)	FADV32TM24000022	5	169.254.132.187	N	

Column	Description
Host Name	The Host Name of the member node.
State	Indicates the HA status of the node and whether it is elected as the HA Primary or Secondary nodes.
Serial Number	The Serial Number of the FortiADC devices in the HA cluster.
Node ID	The Node ID of the member node.
IP Address	The HA in-cluster communication IP address.
Config Source	Indicates whether the member node is the source of configurations for HA synchronization. This is determined by the **Config Priority** set in the HA setting. When the configuration priority values of the nodes are different, the configuration of the device with the lower configuration priority will prevail.

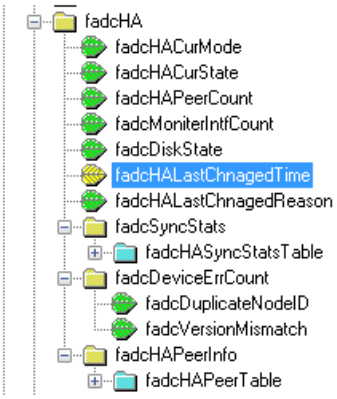
Other HA monitoring methods

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 secondary mode`

The following figure shows FortiADC HA event objects in an SNMP manager.

FortiADC HA event objects in an SNMP manager


fadcHA
fadcHACurMode
fadcHACurState
fadcHAPeerCount
fadcMoniterIntfCount
fadcDiskState
fadcHALastChnagedTime
fadcHALastChnagedReason
fadcSyncStats
fadcHASyncStatsTable
fadcDeviceErrCount
fadcDuplicateNodeID
fadcVersionMismatch
fadcHAPeerInfo
fadcHAPeerTable

Updating firmware for an HA cluster

You can initiate a firmware upgrade from any node within an HA cluster by enabling the **HA Cluster Upgrade** option from the system settings.

The upgrade process begins on the initiating FortiADC device and is sequentially propagated to each of the other nodes in the cluster. For example, when upgrading the firmware on a FortiADC cluster with eight nodes, the current device will sequentially synchronize the new firmware image to each of the remaining nodes, awaiting the completion of the upgrade on one of the other nodes before upgrading itself and rebooting.

The following process occurs when you perform the HA cluster upgrade from the current FortiADC device:

1. The current node pushes the firmware image to the member nodes.
2. The current node notifies the member nodes of the upgrade and handles their user traffic during the process.
3. The upgrade command is executed on the member nodes. They reboot and send an acknowledgment to the current node confirming the completion of the upgrade.
4. The upgrade command is then executed on the current node, which also undergoes a reboot. During this phase, the cluster may perform a new election process to designate a new Primary node based on the established election rules.

Upon completing the firmware upgrade, the system determines the role of the original node based on the HA Override setting:

- **If Override is enabled**, the cluster uses the Device Priority setting to decide which node should be the Primary. Typically, this results in a secondary failover to restore the original roles of the nodes.
- **If Override is disabled**, the cluster prioritizes uptime. Since the original Primary node will have a reduced uptime due to the reboot sequence during the firmware upgrade, it will not reclaim its active role. Instead, the node with the highest uptime remains as the Primary, and no additional failover occurs.

Reboot times can vary depending on the appliance model and the differences between the original and the new firmware versions being installed.

The procedure for upgrading firmware on an HA cluster closely mirrors that of a standalone appliance. To minimize service disruption, follow these steps. This procedure is applicable to both Active-Active and Active-Passive clusters.



If you need to **downgrade** to a previous firmware version, do **not** use the HA cluster upgrade function.

The HA daemon on a member node may recognize the primary node's older firmware and attempt to automatically upgrade it to synchronize with the other nodes, effectively undoing your downgrade. Instead, exit HA mode, perform the downgrade on each node individually, and then re-enable HA mode.

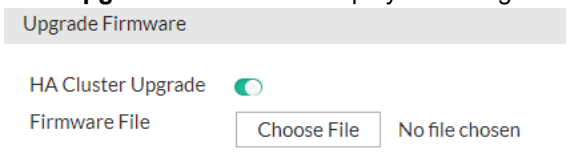
Before you begin:

- 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 in to the web UI of any node within the HA cluster as the `admin` administrator.
2. Go to **System > Settings**.
3. Click the **Firmware** tab.
4. Click **Upgrade Firmware** to display the configuration editor.



5. Enable **HA Cluster Upgrade**.
6. Click **Choose File** to locate and select the file.
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.

HA Troubleshooting

Effective troubleshooting of HA configurations is crucial for ensuring high availability and minimizing downtime. This section provides key strategies and techniques for diagnosing and resolving common issues in HA deployments.

This section covers the following troubleshooting topics:

Common Issues in HA Deployments on page 202

- HA Management Interface on page 202
- HA Configuration Synchronization Issues on page 203
- HA Abnormal State on page 203
- Remote IP Monitor in HA-AP on page 204

HA Event Log on page 204

HA Debug on page 205

Common Issues in HA Deployments

HA Management Interface

In HA-AP mode, direct access to the Secondary device is not possible unless management access has been explicitly enabled. If management access is not configured, you can only access the Secondary device's CLI from the Primary device using the command: `execute ha manage`. To avoid this limitation, it is recommended to enable the Management Interface for all nodes.

To configure the management interface, use the `mgmt-interface` option under the `config system ha` command.

When configuring the management interface, ensure that the IP address is not in the same subnet as the original port1 or mgmt interface to avoid conflicts. It is advisable to have console access during this process. First, clear the old management IP address from the original interface (typically port1 or mgmt), and then configure it under `config system ha`.

For example:

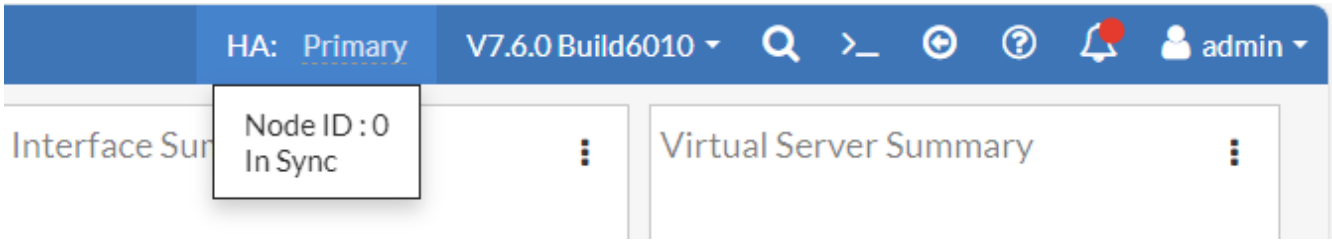
```
FAD2 # config system ha
FAD2 (ha) # set mgmt-status enable
FAD2 (ha) # set mgmt-interface port1
FAD2 (ha) # set mgmt-ip 10.106.188.42/23
FAD2 (ha) # set mgmt-ip-allowaccess http https ping snmp ssh telnet
FAD2 (ha) # end
```

Configure the default route accordingly.

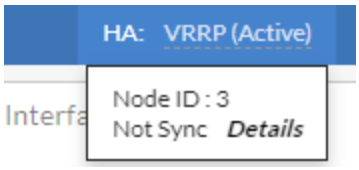
In HA-AA and HA-VRRP modes, configuring the HA Management Interface is not necessary. In HA-AA mode, each device uses the IP address associated with its own local ID. In HA-VRRP mode, interfaces retain their individual IP addresses, as interface addresses are not synchronized across devices.

HA Configuration Synchronization Issues

Once HA peers are established, all required configurations should be synchronized from the Configuration Source. Two types of synchronization occur: incremental synchronization and full synchronization. Incremental synchronization happens whenever a configuration change is made on one of the HA nodes, propagating the changes to the other peers. Full synchronization occurs when a new HA peer with a different configuration joins the cluster or when the command `execute ha force sync-config` is executed.



Under normal circumstances, from the GUI, you should see **In Sync** when you hover your cursor over the HA device from the top navigation, indicating that the configurations are aligned.



However, if an unexpected issue causes a configuration mismatch, the status may indicate an out-of-sync condition. In such cases, you can click **Details** to view a line-by-line comparison of the your current out-of-sync HA configuration against configuration of the device designated as the configuration source. You can manually correct the discrepancies on each device or run the `execute ha force sync-config` command on the device with the correct configuration to perform a full synchronization across the HA cluster.

If manually correcting discrepancies or forcing synchronization with the `execute ha force sync-config` command does not resolve the issue, you should back up the full configuration file from the correctly configured HA peer and restore it to the out-of-sync peer.

HA Abnormal State

If you encounter an abnormal HA state, such as dual Primary nodes or prolonged synchronization delays, begin by checking for errors using the `get system ha-status` command and reviewing the system event log.

The `get system ha-status` command provides an overview of the HA node's status and can reveal potential issues such as duplicate licenses or HA version mismatches. If an error persists, the relevant error count will increase each time you run `get system ha-status`.

For optimal operation, heartbeat and data ports should be directly connected or, at a minimum, placed within the same VLAN via switches. When configuring multiple heartbeat or data ports, ensure that each pair resides in a separate VLAN. If the connections are correct and the issue persists, enable debugging to diagnose the root cause of the abnormal HA state.

For example:

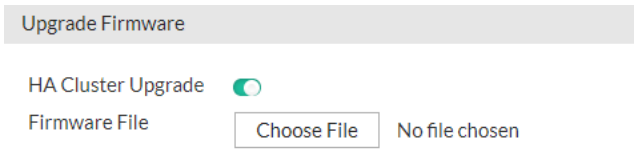
You can enable the "heartbeat" debug option to verify whether heartbeat messages are being sent and received successfully. If all heartbeat messages are properly transmitted and received, proceed by enabling the "errors" debug option to check for any errors. Record any errors that appear and work to resolve them.

In some cases, HA packet delays due to network conditions can cause HA timeouts, even when ports are correctly configured and connected. To mitigate this, consider increasing the `hb-lost-threshold` or `hb-interval` settings in `config system ha`.

For further details about using HA debug options, see HA Debug on page 205.

Firmware Upgrade in HA Cluster

Users can upgrade all FortiADC units within the HA cluster with a single click by enabling the **HA Cluster Upgrade** option in the **Upgrade Firmware** GUI.



For details, see Updating firmware for an HA cluster on page 200.

Remote IP Monitor in HA-AP

The Remote IP Monitor condition is a key factor in the election process when appliances join an HA cluster and determine the Primary node. If devices have identical disk states and monitor interface statuses, the appliance with the fewest failed Remote IP Monitor health checks will be elected as the Primary. At this stage, the Remote IP Monitor failover threshold does not influence the election.

Once an appliance assumes the role of HA-AP Secondary, it no longer processes traffic. Therefore, as long as the current Primary's count of failed Remote IP Monitor checks remains below the failover threshold, changes in the failure count will not trigger a failover.

A failover occurs when the current Primary's failed Remote IP Monitor count exceeds the failover threshold. In this case, one of the Secondary appliances will become the new Primary since all Secondary appliances are considered to have a zero failure count while inactive. Once a Secondary appliance becomes Primary, its Remote IP Monitor checks become active, potentially triggering another failover if its failover threshold is exceeded. If all appliances exceed the failover threshold, continuous failovers could occur. The Failover Hold Time setting determines how long an appliance retains its previous HA state in such scenarios to prevent excessive failover frequency.

HA Event Log

Event logs for all modules are accessible via the GUI under **Log & Report > Event Log**. In the event of HA issues, it is critical to review these logs for any relevant failure events.

The configuration event log is essential for tracking configuration changes on the FortiADC. It logs timestamps for both direct configurations made via the GUI or CLI and incremental synchronizations received from HA peers. When an incremental synchronization completes successfully, the event is logged across all HA members' configuration event logs. On the originating FortiADC, it appears as a direct CLI or GUI change, while on other FortiADC units, it is recorded as an HA synchronization event.

The system event log provides key insights into system-level operations, such as HA state transitions and the initiation of full synchronizations. It also captures indirect factors that could impact HA performance, such as the failure of a heartbeat interface, which may trigger a failover.

For more detailed, real-time HA logs, refer to HA Debug on page 205.

HA Debug

Logs generated by `diagnose debug` CLI commands are output in real-time to the console after debugging is enabled. Since each node in an HA cluster operates as an individual entity, their debug logs will not be identical. To effectively diagnose HA communication issues within the cluster, it is recommended to enable HA debug commands on all nodes in the HA cluster.

To enable HA debug:

1. Enable the system debug via `diagnose debug` command.

   ```
   (M) FAD1 # diagnose debug enable
   ```

2. To output specific traffic flows, specify the number of debug lines required. Certain HA diagnose debug commands necessitate this configuration to activate flow-related debug log output.
   Start debugging. The `[<count>]` option specifies the number of debug lines to output.

   ```
   (M) FAD1 # diagnose debug flow start [<count>]
   ```

   To enable detailed debugging for specific traffic flows within the HA cluster, use the `diagnose debug flow start` command with the relevant `diagnose debug` commands listed below.

   The following `diagnose debug` commands require the `diagnose debug flow start` command to be specified to generate output:

   - `diagnose debug ha heartbeat` — outputs debug information for HA heartbeat messages.
   - `diagnose debug ha arp` — outputs ARP-related debug logs within the HA cluster.
   - `diagnose debug ha layer4` — captures Layer 4 debug logs for HA-related traffic.

   The following `diagnose debug` commands can include flow-related information when `diagnose debug flow start` is specified:

   - `diagnose debug ha basic` — provides basic HA debug information, including flow details.
   - `diagnose debug ha error` — outputs error-specific debug logs, potentially containing flow information.

   For more information, see the FortiADC CLI Reference on diagnose debug flow and diagnose debug ha.

3. Toggle the appropriate HA debug options on or off as needed to suit your diagnostic requirements.
   In the example below, we will enable debugging to view both basic HA information and error-specific logs:

   ```
   (M) FAD1 # diagnose debug ha basic
   ha debug basic enabled
   (M) FAD1 # diagnose debug ha errors
   ha debug errors enabled
   ```

   Note: Entering the same option again will toggle it off.

   ```
   (M) FAD1 # diagnose debug ha basic
   ha debug basic enabled
   ```

```
(M) FAD1 # diagnose debug ha basic
ha debug errors disabled
```

4. Enter the list option to view all HA debug settings.
   The example below shows that the **basic** and **errors** debug options have been successfully enabled from the previous step.

```
(M) FAD1 # diagnose debug ha list
basic: enabled
configuration: disabled
errors: enabled
file: disabled
health-check: disabled
heartbeat: disabled
layer4: disabled
layer7: disabled
message: disabled
state: disabled
sync-status: disabled
upgrade: disabled
arp: disabled
```

5. Enable or disable the **all** HA debug option to view the debug information for each enabled setting.
   The example below shows the HA debug information for the enabled **basic** and **errors** options.

```
(M) FAD1 # diagnose debug ha all
enabled all ha debugs
(M) FAD1 # [10-09 10:00:58] [kernel]Hello pkt: mode 2 group id 14 local_node_id 0 SN
FADV040000146260 sended
[10-09 10:00:58] [kernel]Hello pkt: mode 2 group id 14 local_node_id 0 SN
FADV040000146260 sended
[10-09 10:00:58] [kernel]Hello pkt: mode 2 group id 14 local_node_id 1 SN
FADV040000146261 received
[10-09 10:00:58] [kernel]Hello pkt: mode 2 group id 14 local_node_id 1 SN
FADV040000146261 received
(M) FAD1 # diagnose debug ha all
disabled all ha debugs
```

Use the `diagnose debug disable` command to turn off debugging when you are finished.

Deploying an Active-Passive cluster

In an Active-Passive cluster, one node actively processes traffic, while the other remains in standby mode, ready to take over if the active node becomes unavailable. This makes the HA-AP mode the most stable deployment option, suitable for implementation across various platforms, provided they are correctly configured.

This section covers the following information about the Active-Passive HA deployment:

- About HA Active-Passive Mode in FortiADC on page 207
- Basic HA-AP configuration steps on page 213
- Best practices on page 217

About HA Active-Passive Mode in FortiADC

In an HA Active-Passive (HA-AP) configuration for FortiADC, the active node (designated as the Primary) is responsible for processing all traffic under normal operating conditions. Failover occurs if the active node encounters a critical issue, such as a failure of a monitored interface (e.g., an interface directly connected to an ISP) or a hardware malfunction. Upon detecting such a failure, the standby node (designated as the Secondary) automatically transitions to the active role, assuming full responsibility for traffic processing.

Figure 1: HA Active-Passive Mode Deployment

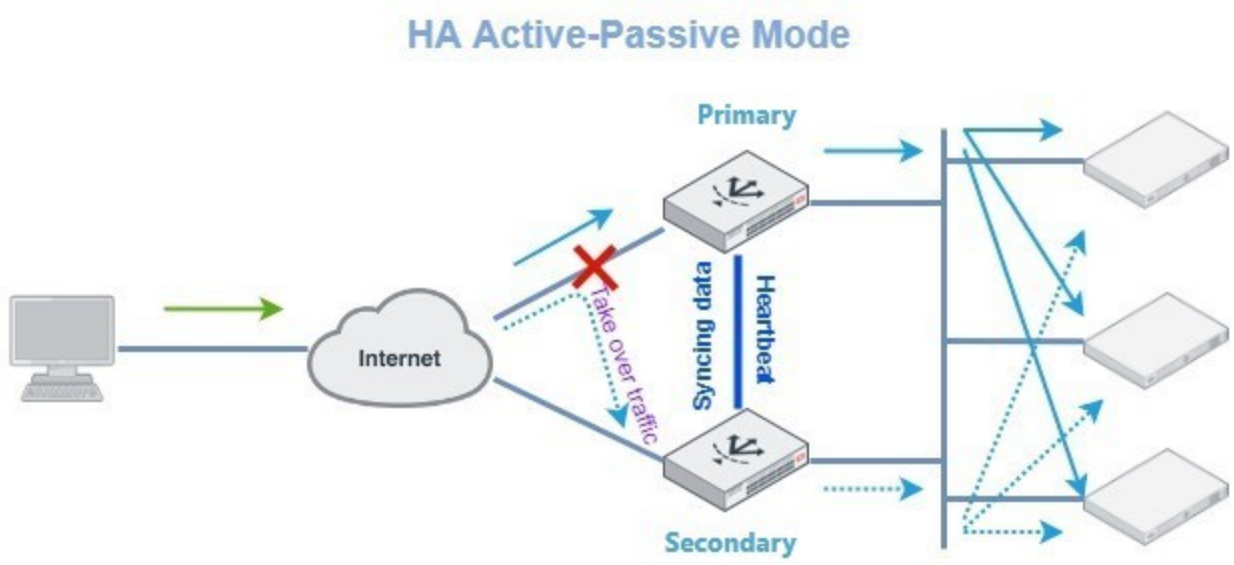
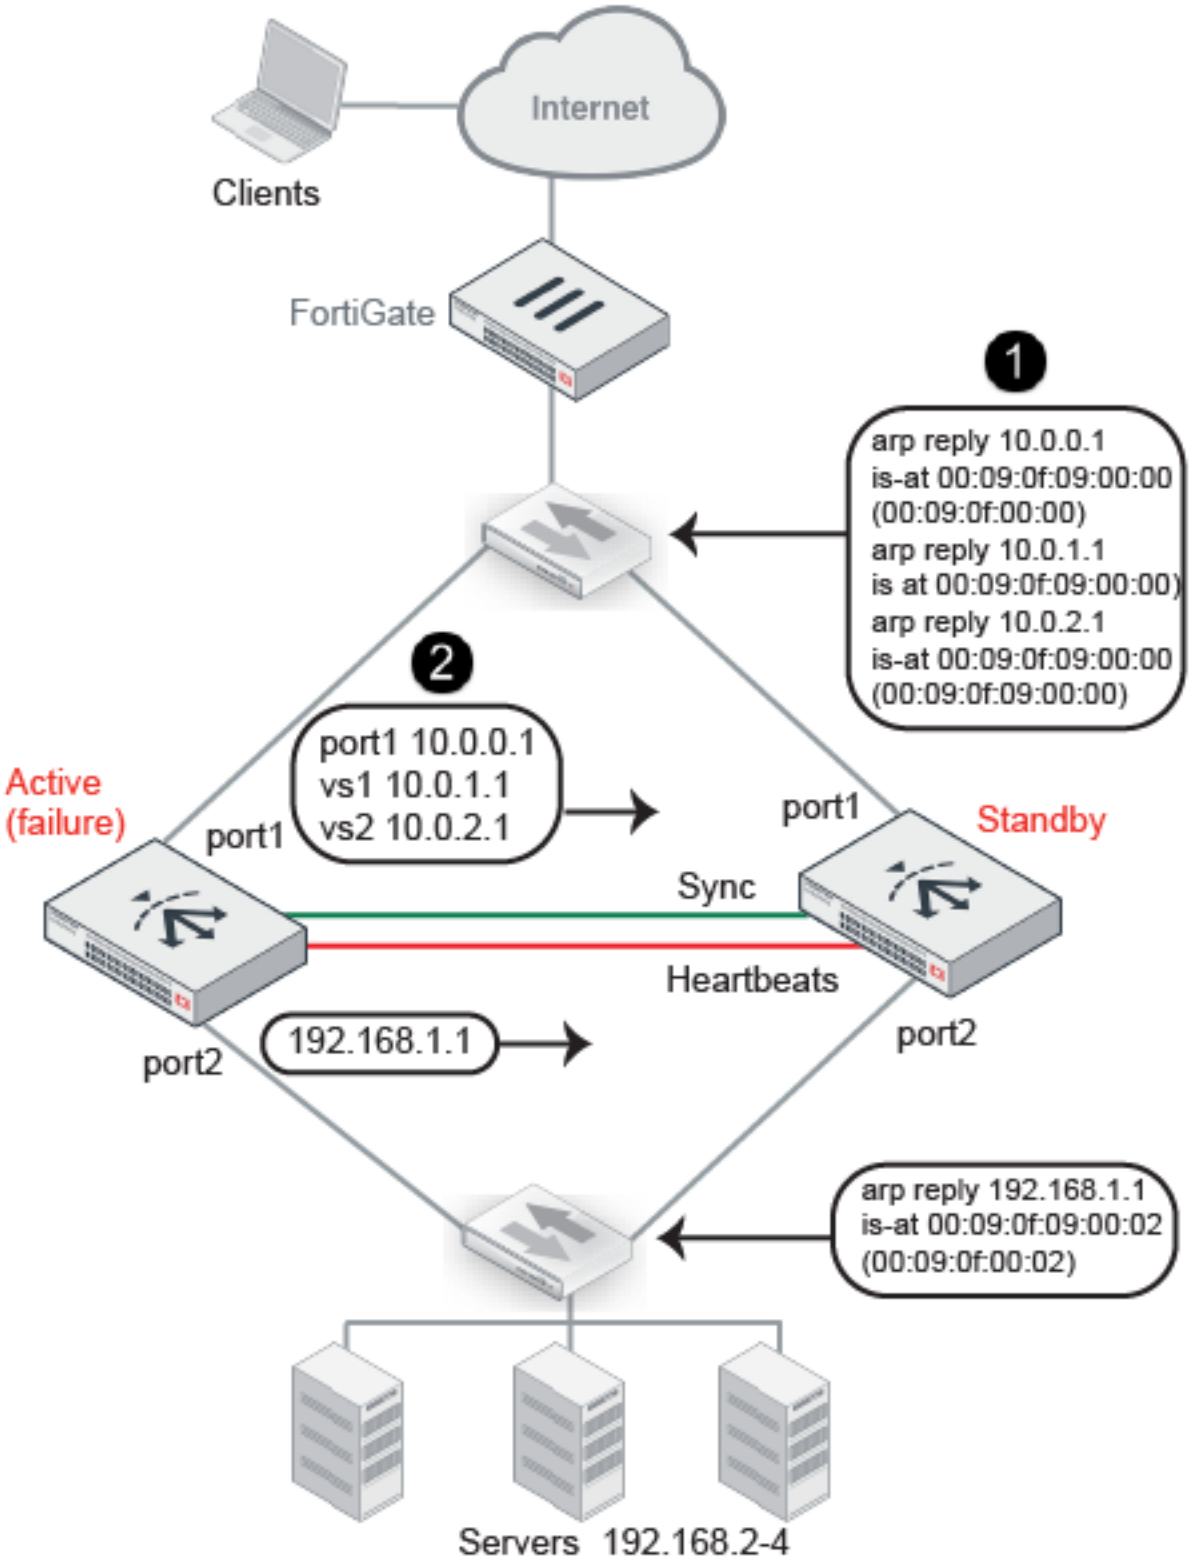


Figure 1 above details the HA-AP mode deployment. In this configuration, the Secondary node remains in standby, with all client-side and server-side traffic exclusively managed by the Primary node. Despite not handling traffic, the Secondary node continuously synchronizes with the Primary, replicating incremental configuration changes, Layer 4 session and persistence tables, Layer 7 persistence data, and health-check statuses.

Figure 2: An Active-Passive cluster at failover—IP address transfer to the new active member



The system is configured to exchange heartbeat packets between nodes to continuously monitor their operational status. These packets are polled at regular intervals to ensure the active appliance remains responsive. If the active

appliance fails to respond within the defined timeout, a failover is initiated, and the standby node transitions to the active role, assuming the role of the primary node. The failover process (illustrated in **Figure 2**) involves two critical actions: (1) The standby node broadcasts a gratuitous ARP to adjacent routers, prompting them to update their ARP tables and reroute traffic for the virtual MAC addresses (vMAC) to its interfaces (provided the virtual-mac option is enabled); (2) The standby node then assumes the IP addresses associated with the unresponsive appliance, enabling uninterrupted network service continuity.



As of FortiADC version 7.6.0, the introduction of the CLI `virtual-mac` option in `config system ha` allows you to enable or disable the use of virual MAC addresses in HA-AP.

By default, virtual MAC addresses are assigned to the Primary FortiADC's interfaces. In the event of a failover, the Secondary node, upon assuming the Primary role, inherits these virtual MAC addresses, thereby reducing traffic disruption. This is particularly advantageous in network environments where security devices, such as firewalls, require MAC address binding for traffic processing continuity.

If the `virtual-mac` option is disabled, the new primary node will instead use its physical MAC addresses for traffic redirection.

It is crucial to note that in Microsoft Hyper-V environments, HA-AP mode defaults to using physical MAC addresses due to platform-specific limitations.

Figure 3: Basic Active-Passive Cluster

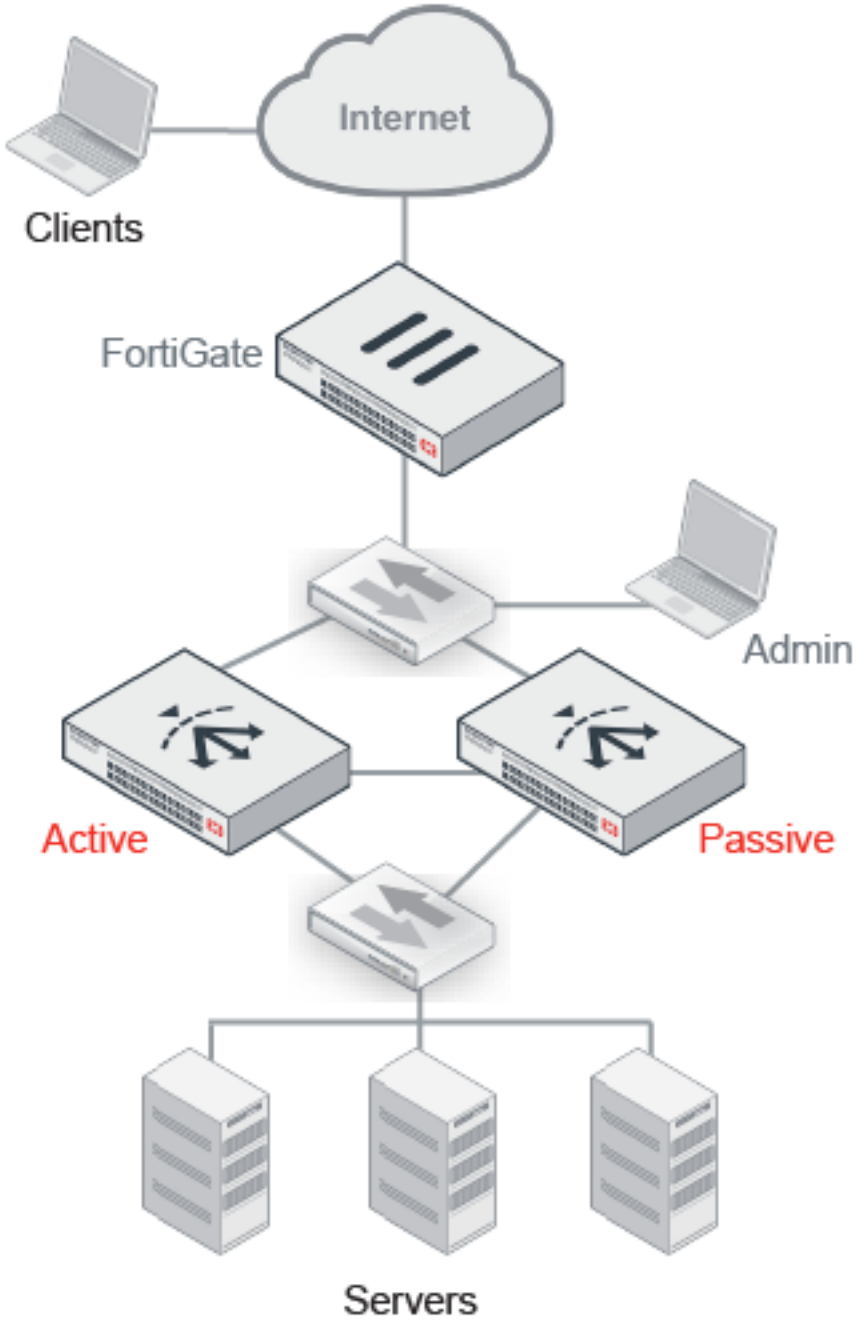
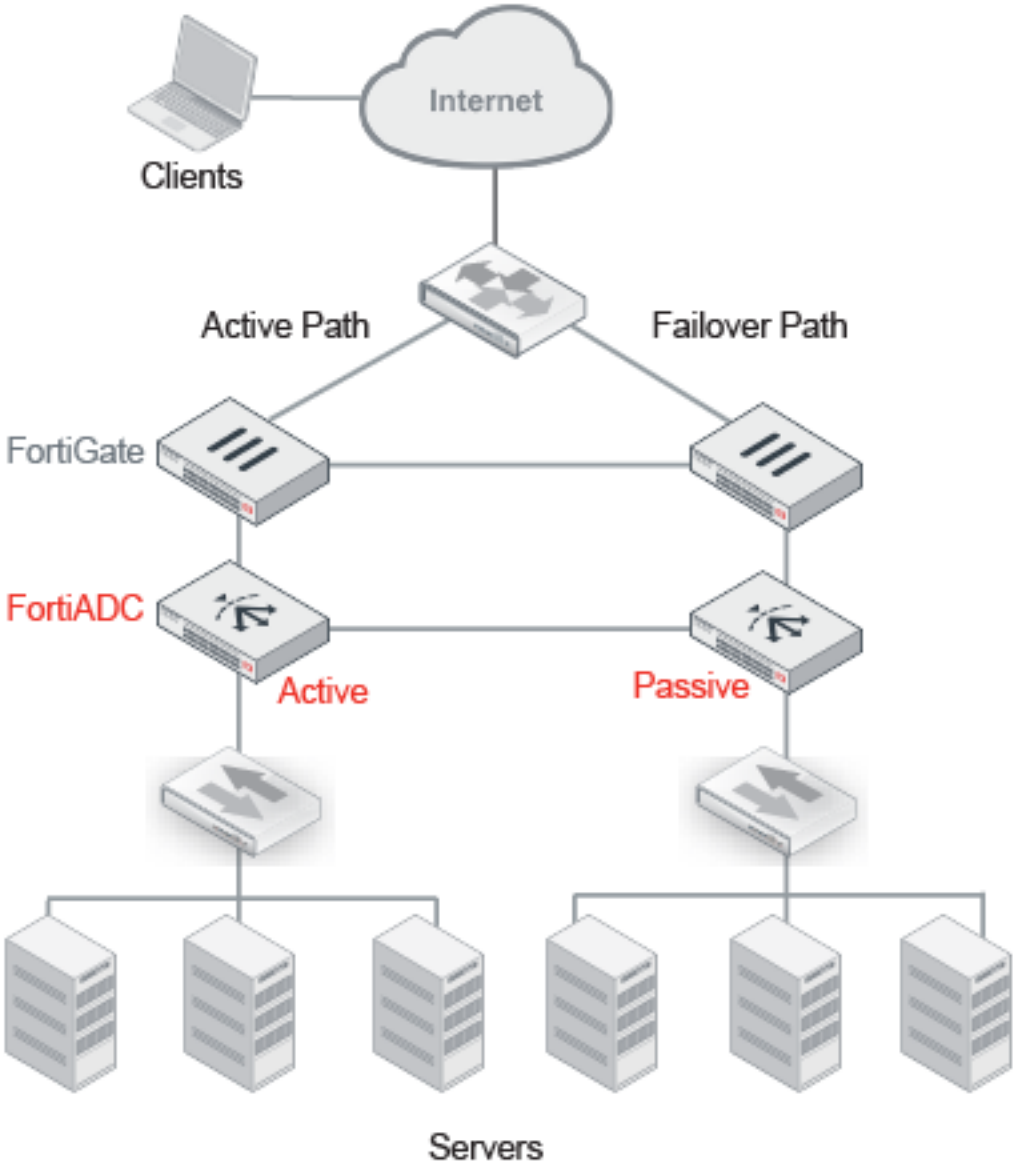


Figure 4: Redundant Path Active-Passive Cluster



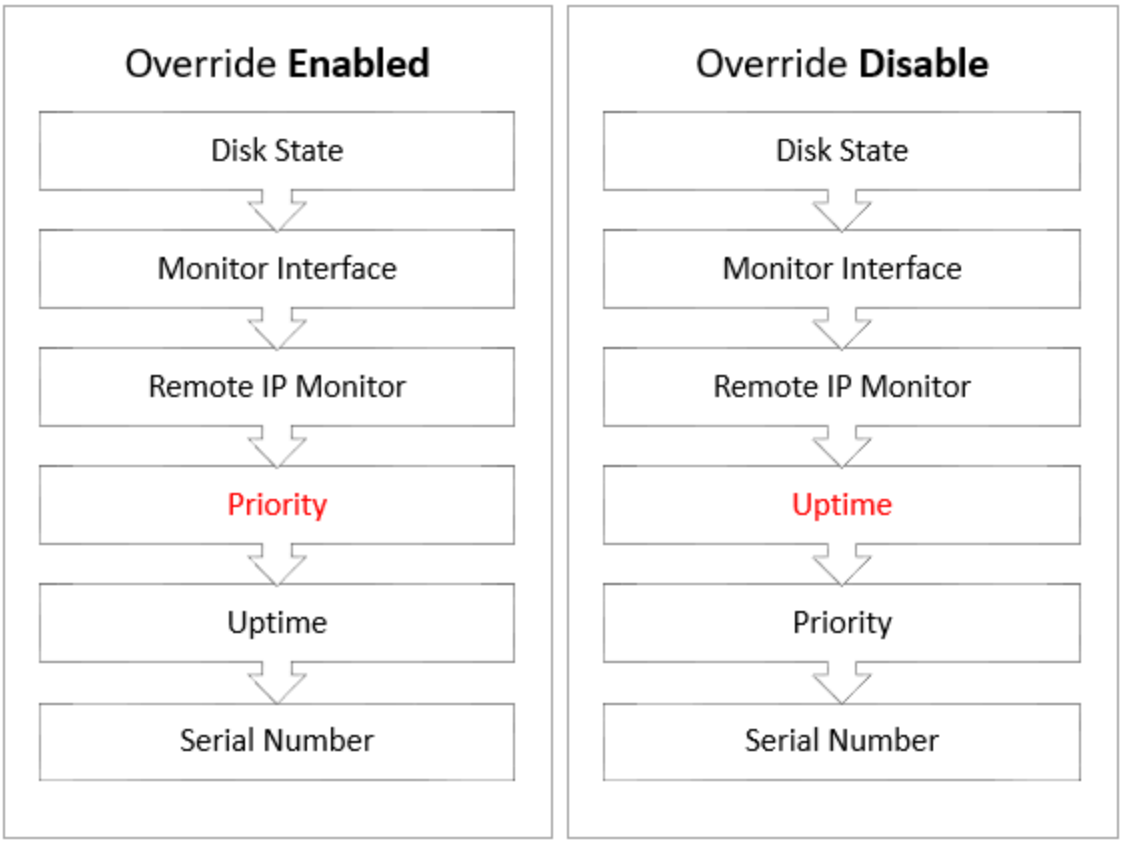
Compare the topologies of a basic Active-Passive cluster configuration (as illustrated in **Figure 3**) and a redundant path Active-Passive cluster configuration (as illustrated in **Figure 4**). In the basic configuration, a single network path is used; if the primary node encounters hardware failure or requires maintenance, failover is triggered, and the standby node becomes the primary, taking over traffic processing along the same network path. In contrast, the redundant path configuration eliminates single points of failure. Should the gateway, load balancer, or switch fail, traffic is seamlessly rerouted through an alternate path. Due to its increased resilience, the redundant path configuration is recommended as a best practice.

Primary election process in an Active-Passive cluster

When the former active appliance rejoins the HA cluster, it will reassume the primary role if it meets the criteria and wins the election process. The system determines the primary device by evaluating and prioritizing specific operational

conditions and configuration parameters. Depending on whether the Override setting has been enabled or disabled, the order each condition is considered differs slightly (illustrated in **Figure 5**).

Figure 5: Priority Criteria for Primary Device Election



Criteria	Description
Disk State	The “disk state” refers to the operational status of the device's hard disk. A device without disk errors takes precedence. If all devices have disk errors, the system proceeds to evaluate the next criterion.
Monitor Interface	The “monitor interface” refers to the count of monitored interfaces that are currently up. The device with the higher number of active interfaces is prioritized. If all devices have the same number of active interfaces, the system proceeds to the next evaluation criterion.
Remote IP Monitor	The "remote IP monitor" actively checks beacon IP addresses to assess network path availability. The device with fewer failed Remote IP Monitor health checks is prioritized. If all devices have the same number of failures, the system evaluates the next criterion.
Priority	The "priority" refers to the value set in the HA configuration; the device with the lowest priority value is favored. If multiple devices share the same priority value, the system moves on to the next evaluation criterion.
Uptime	The "uptime" refers to the duration the device has been operational. The device

Criteria	Description
	with the longest uptime is given precedence. If multiple devices have identical uptime, the system proceeds to the next criterion.
Serial Number	The device with the higher serial number will be designated as the Primary.

Basic HA-AP configuration steps

The following steps outline the deployment of an Active-Passive cluster where the Primary node serves as both the active device and the configuration source, with override enabled. To ensure successful negotiation, all HA devices must use identical heartbeat and data ports, and share the same Group ID.

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. Enable the management interface. In this example, CLI commands are executed from the FortiADC console.
   It is advisable to enable the management interface when deploying HA-AP mode. In HA-AP mode, the primary node is solely responsible for handling traffic, which restricts direct access to the secondary node. This limitation can be operationally inconvenient. The management interface, a virtual interface bound to a physical interface, provides consistent management access across all modes, including standalone. Ensure to execute the following configuration steps on all HA nodes to maintain consistent management access.
   a. The management interface operates as a virtual interface within the system, following similar routing mechanisms as other interfaces. As such, there must be no overlapping subnets within the system. To avoid conflicts, it is standard practice to clear the original IP address from the physical interface. However, this action may result in a temporary loss of connectivity, so it's crucial to have console access before proceeding.

      ```
      FAD2 # config system interface
      FAD2 (interface) # edit port1
      FAD2 (port1) # unset ip
      FAD2 (port1) # end
      ```

   b. Configure the management interface.

      ```
      FAD2 # config system ha
      FAD2 (ha) # set mgmt-status enable
      FAD2 (ha) # set mgmt-interface port1
      FAD2 (ha) # set mgmt-ip 10.106.188.42/23
      FAD2 (ha) # set mgmt-ip-allowaccess http https ping snmp ssh telnet
      FAD2 (ha) # end
      ```

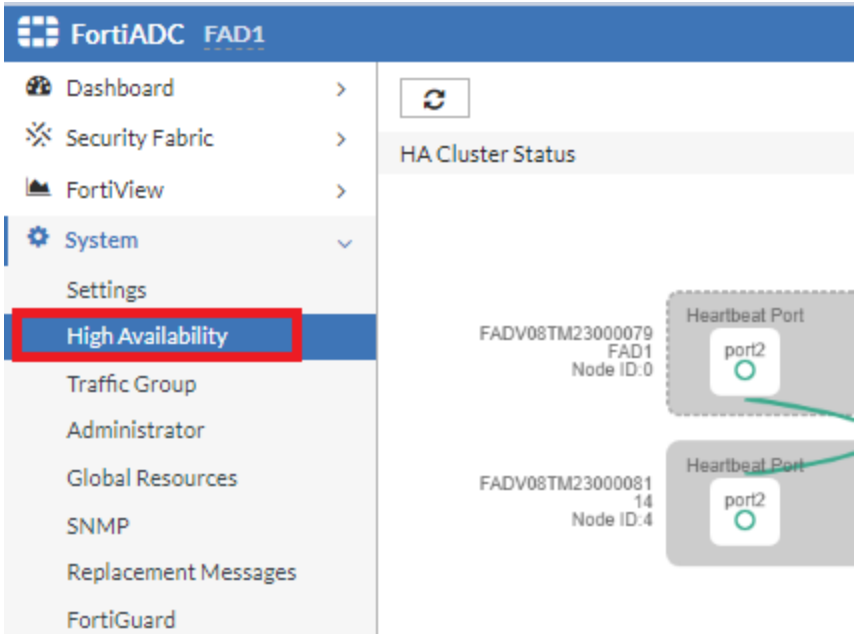
   c. Configure the default route accordingly.

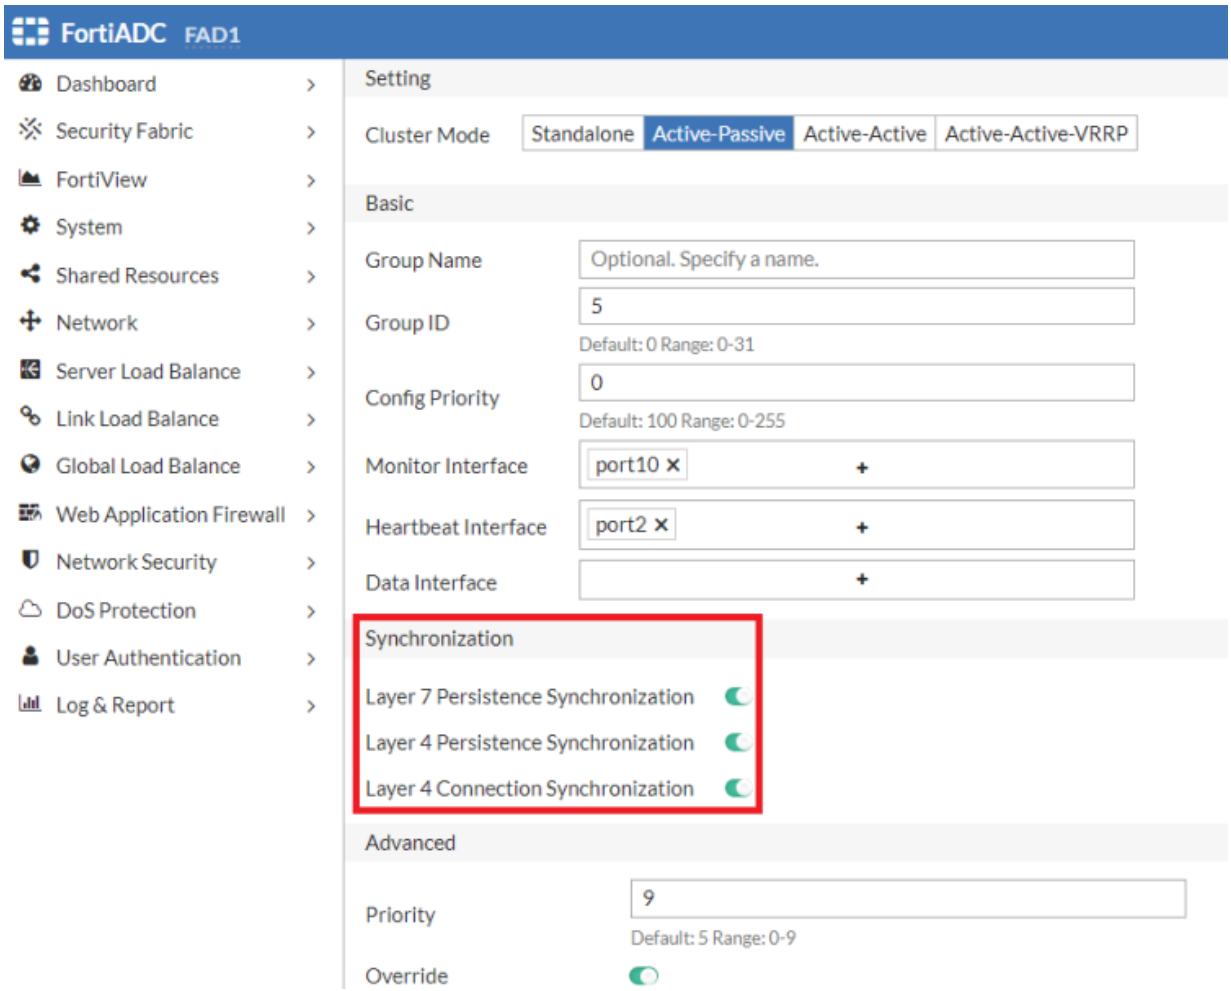
      ```
      FAD2 # config router static
      FAD2 (static) # edit 1
      FAD2 (1) # set gateway 10.106.189.254
      FAD2 (1) # end
      ```

   Note: On VMware ESXi and Hyper-V, specific settings may be required for the VM interface to which the management interface is bound. Please refer to their respective documentation on HA settings.
3. 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:

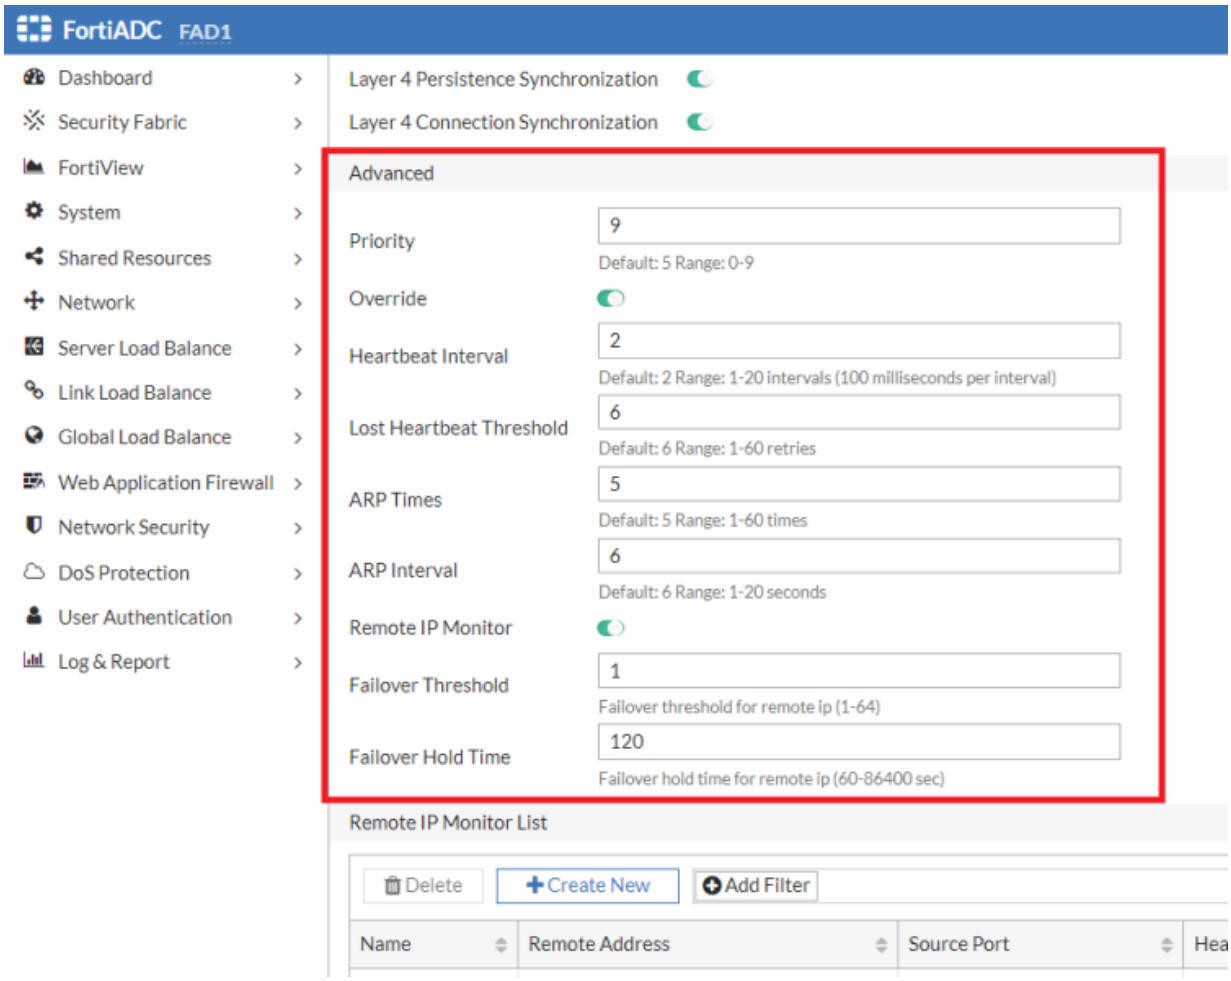
- Connect the two appliances directly with a crossover cable.
- Link the appliances through a switch. If connected through a switch, the heartbeat interfaces must be reachable by Layer 2 multicast.

4. Configure the HA settings for each node.
   Ensure that the Group ID of each HA node is the same and the heartbeat interfaces be connected directly or in the same VLAN.
   a. Log in to the appliance as the **admin** user.
   b. Complete the HA settings as described in Configuring HA settings.
   The example below shows the configuration for the Primary device FAD1.





Note: The configuration source is determined by the **Config Priority** set in the HA setting. When the configuration priority values of the nodes are different, the configuration of the device with the lower configuration priority will prevail. For the secondary node, ensure the Config Priority is set to a number higher than the device you want as the configuration source.



To configure in CLI:

```
config system ha
  set mode active-passive
  set hbdev port6 port7
  set group-id 14
  set group-name group1
  set priority 1
  set config-priority 10
  set override enable
  set l7-persistence-pickup enable
  set l4-persistence-pickup enable
  set l4-session-pickup enable
  set monitor port2 port3 port4 port5
end
```

After saving the HA configuration changes, cluster members will initiate the join or rejoin process within the cluster. Upon saving the configuration on the primary node, it will automatically disseminate the updated configuration to the secondary node through the HA synchronization mechanism.

Best practices

The following best practices should be observed:

- Maintain heartbeat links — Ensure the heartbeat link(s) are consistently operational. If a heartbeat is interrupted—such as by a temporary network cable disconnection—other nodes may incorrectly perceive the primary node as failed, triggering failover in an active-passive deployment. This could result in both nodes operating as active nodes simultaneously if no actual failure has occurred.
- Improve fault tolerance — When connecting HA appliances through switches, link the ports through two separate switches to enhance fault tolerance and reliability. Avoid connecting these switches to the broader network to mitigate potential security vulnerabilities and prevent network load from introducing latency in the heartbeat, which could inadvertently trigger failover.

Deploying an Active-Active cluster

In an Active-Active cluster, traffic from the upstream router can be load-balanced among up to eight member nodes.

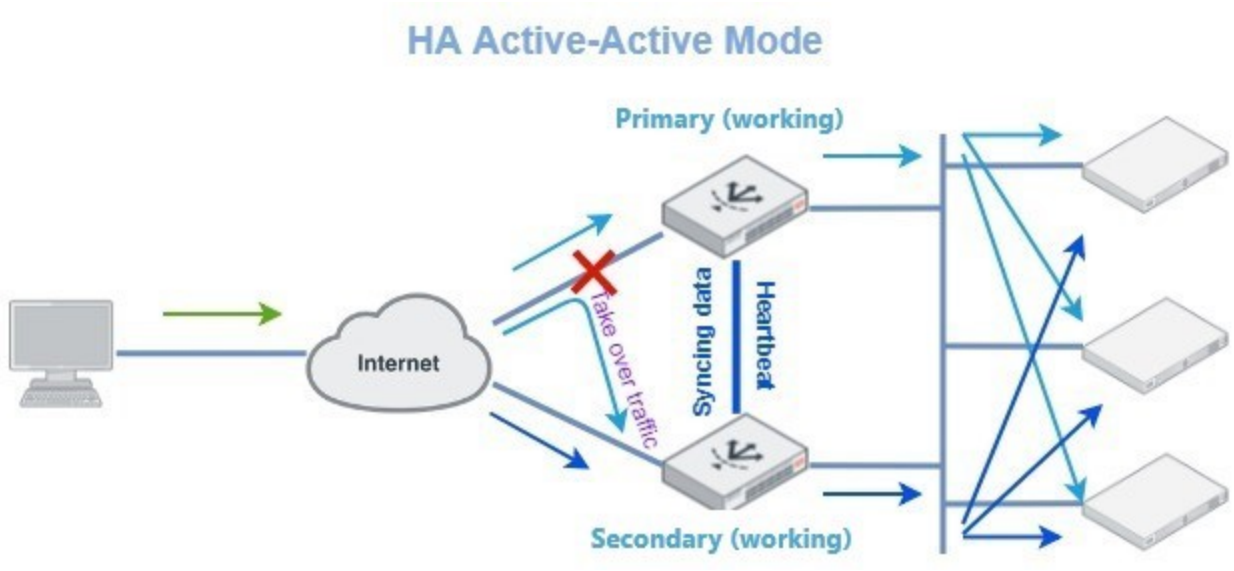
This section covers the following information about the Active-Passive HA deployment:

- About HA Active-Active Mode in FortiADC on page 218
- HA Active-Active configuration overview on page 220
- HA-AA configuration steps on page 222
- Expected behavior on page 226
- How the cluster uses synchronized data on page 226
- Best practices on page 234

About HA Active-Active Mode in FortiADC

In HA Active-Active mode, both primary and secondary nodes can handle traffic concurrently. However, it is important to note that FortiADC only synchronizes Layer 4 virtual server sessions across its peers. Consequently, if Layer 4 traffic is returned from a real server and arrives at a different FortiADC device than the one that initially processed the inbound traffic, the new FortiADC device can still forward the traffic to the client. For traffic routed by FortiADC, a similar issue may arise, potentially impacting performance. To mitigate this, it is ideal to place a routing device between FortiADC and the real servers, which can ensure that return traffic is routed back to the original FortiADC device. For Layer 7 virtual servers, FortiADC typically establishes sessions with real servers using its own interface IP address, allowing traffic to be returned directly to the originating FortiADC device, unless the **Source Address** option is enabled.

Figure 1: HA Active-Active Mode Deployment



If a monitored link on a device fails or if the entire device experiences a failure, its HA peer can assume full traffic handling responsibilities.

In HA-AA mode, each interface on the FortiADC devices must be assigned a unique operational IP address, which should be specified under `config ha-node-ip-list`. This configuration is propagated across all HA peers, with

each node operating under its designated IP based on its Local Node ID. It is essential that each HA node has a distinct Local Node ID to ensure proper identification and avoid conflicts.

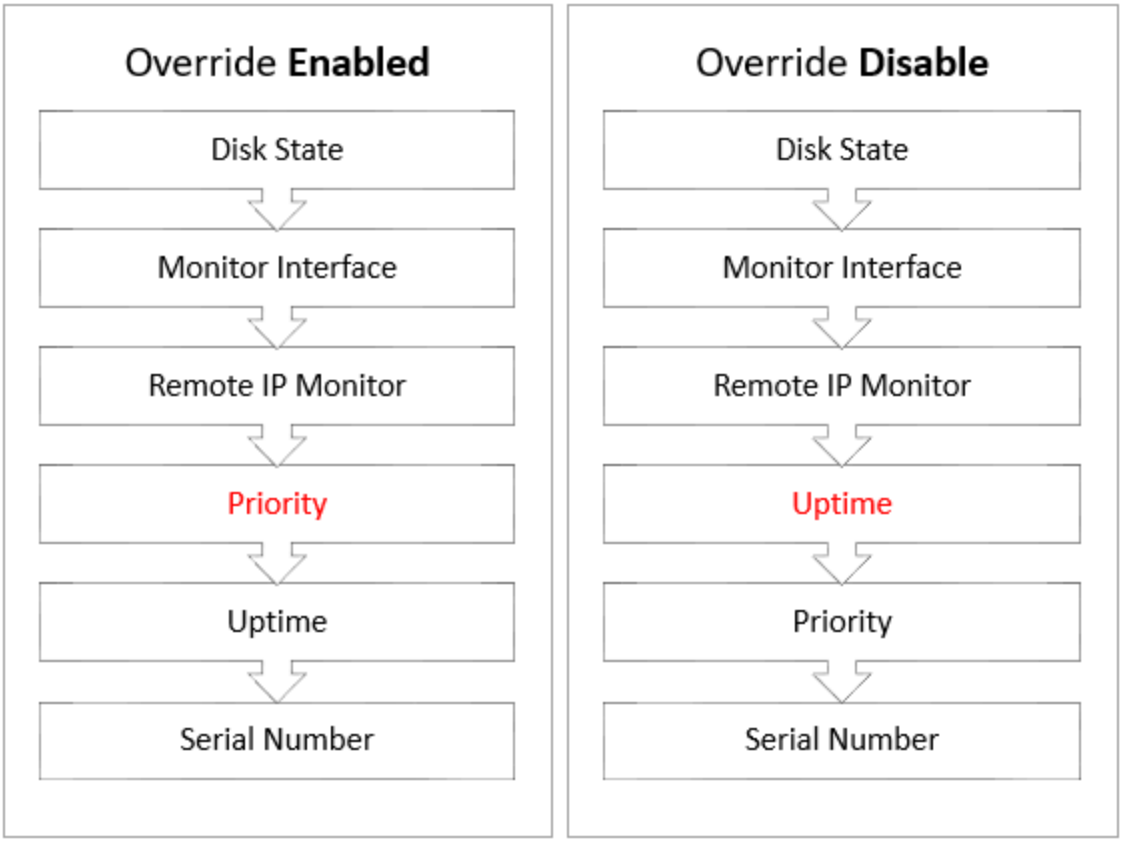
In HA-AA mode, the IP address configured on the interface is no longer utilized for traffic routing. Instead, only the IP addresses listed in `config ha-node-ip-list` are actively used for node communication.

To ensure optimal performance and stability, it is advisable to implement a routing device (typically a router) between the FortiADC and the real server pool. The routing device should support "reverse-route" functionality, which facilitates the routing of return traffic from the real servers back to the FortiADC node that initially distributed the client requests. For example, if FortiADC1 processes requests from client1 and forwards them to real-server1, the router must direct the return packets from real-server1 back to FortiADC1.

Primary election process in an Active-Active cluster

When a previously active appliance rejoins the HA cluster, it will reassume the primary role if it meets the criteria and wins the election process. The system determines the primary device by evaluating and prioritizing specific operational conditions and configuration parameters. Depending on whether the Override setting has been enabled or disabled, the order each condition is considered differs slightly (illustrated in Figure 3).

Figure 3: Priority Criteria for Primary Device Election



Criteria	Description
Disk State	The “disk state” refers to the operational status of the device's hard disk. A device without disk errors takes precedence. If all devices have disk errors, the system

Criteria	Description
	proceeds to evaluate the next criterion.
Monitor Interface	The “monitor interface” refers to the count of monitored interfaces that are currently up. The device with the higher number of active interfaces is prioritized. If all devices have the same number of active interfaces, the system proceeds to the next evaluation criterion.
Remote IP Monitor	The "remote IP monitor" actively checks beacon IP addresses to assess network path availability. The device with fewer failed Remote IP Monitor health checks is prioritized. If all devices have the same number of failures, the system evaluates the next criterion.
Priority	The "priority" refers to the value set in the HA configuration; the device with the lowest priority value is favored. If multiple devices share the same priority value, the system moves on to the next evaluation criterion.
Uptime	The "uptime" refers to the duration the device has been operational. The device with the longest uptime is given precedence. If multiple devices have identical uptime, the system proceeds to the next criterion.
Serial Number	The device with the higher serial number will be designated as the Primary.

HA Active-Active configuration overview

The FortiADC HA Active-Active configuration encompasses the following components:

- Primary node system and feature configuration — Defines the system settings and features for the primary node.
- Interface configuration — Includes the HA node IP list, specifying the IP addresses for each node in the cluster.
- HA configuration — Configures High Availability settings, including failover and synchronization parameters.

Figure 2: Basic Active-Active cluster

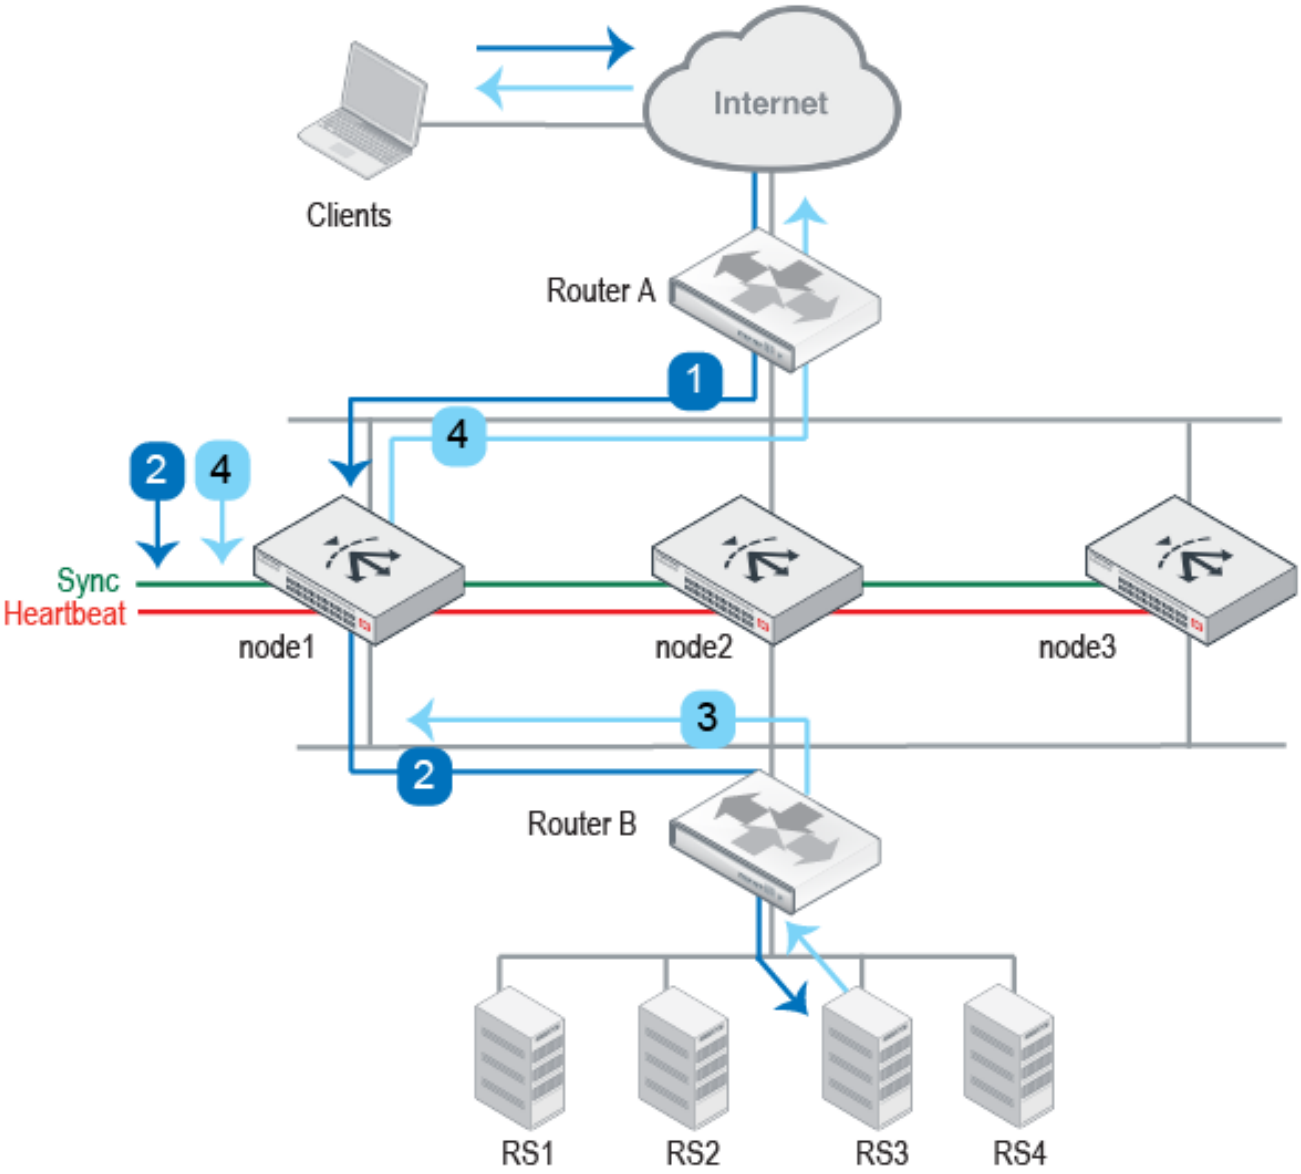


Figure 2 illustrates an example of an Active-Active cluster configuration. In this setup, traffic from the upstream router is load-balanced across up to two member nodes, enabling both nodes to handle traffic concurrently.

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 provide failover support. Each node in the cluster, including the primary node, serves as a backup for the other nodes. In the event of a node failure, the primary node assumes the failed node's IP address and broadcasts gratuitous ARP to adjacent routers. This ensures that traffic for the virtual MAC address (vMAC) is redirected to the primary node's network interfaces.

In an Active-Active cluster, one node is designated as the primary node, while the remaining nodes are member nodes. In this setup, node1 is the primary node, and node2 and node3 are member nodes. Upon cluster formation, the configuration from node1 is propagated to node2 and node3.

When configuring network interfaces for nodes in an Active-Active cluster, you must specify not only the primary IP address for each interface but also define an HA node IP list. This list details the unique HA IP addresses assigned to each node in the cluster. For example, the HA node IP list for port2 includes 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, configure the HA settings for each node while it is in standalone mode. Once all nodes are configured, set the HA mode to Active-Active to form the cluster. The system state will automatically update as each node joins the cluster.

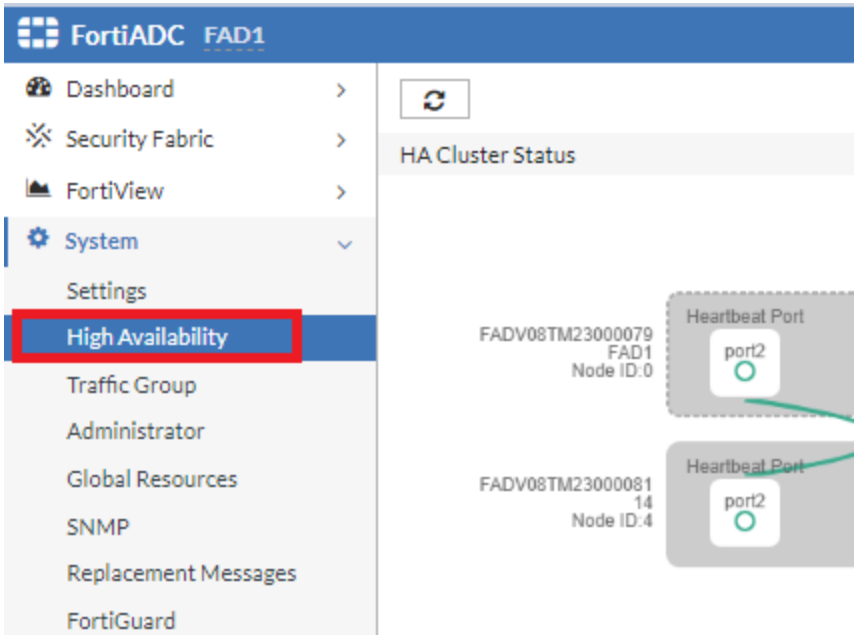
Note: The example illustrates a FortiADC cluster with routers positioned on both sides. If your deployment lacks a router between the FortiADC cluster and the real server pool, and if your real servers support load balancing methods such as Equal-Cost Multi-Path (ECMP), the behavior will align with the described scenario. If ECMP is not supported, the real servers should route reply traffic back to the FortiADC node that originally sent the client traffic.

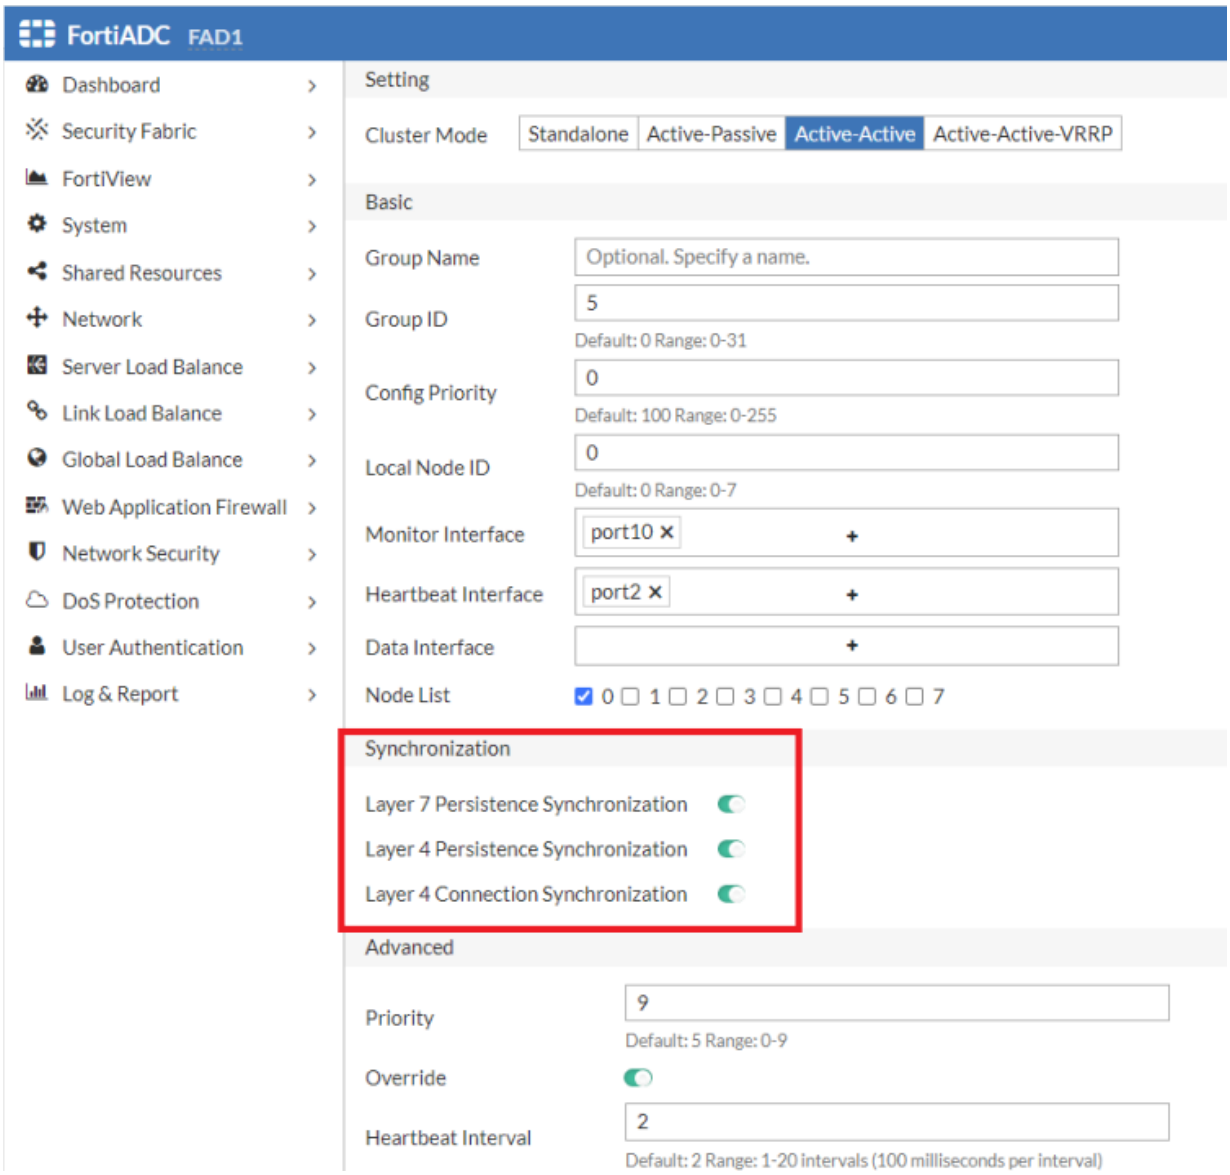
HA-AA configuration steps

The following steps outline the deployment of an Active-Active cluster where the Primary node serves as both the active device and the configuration source, with override enabled. To ensure successful negotiation, all HA devices must use identical heartbeat and data ports, share the same Group ID and Node List, and have different Local Node IDs.

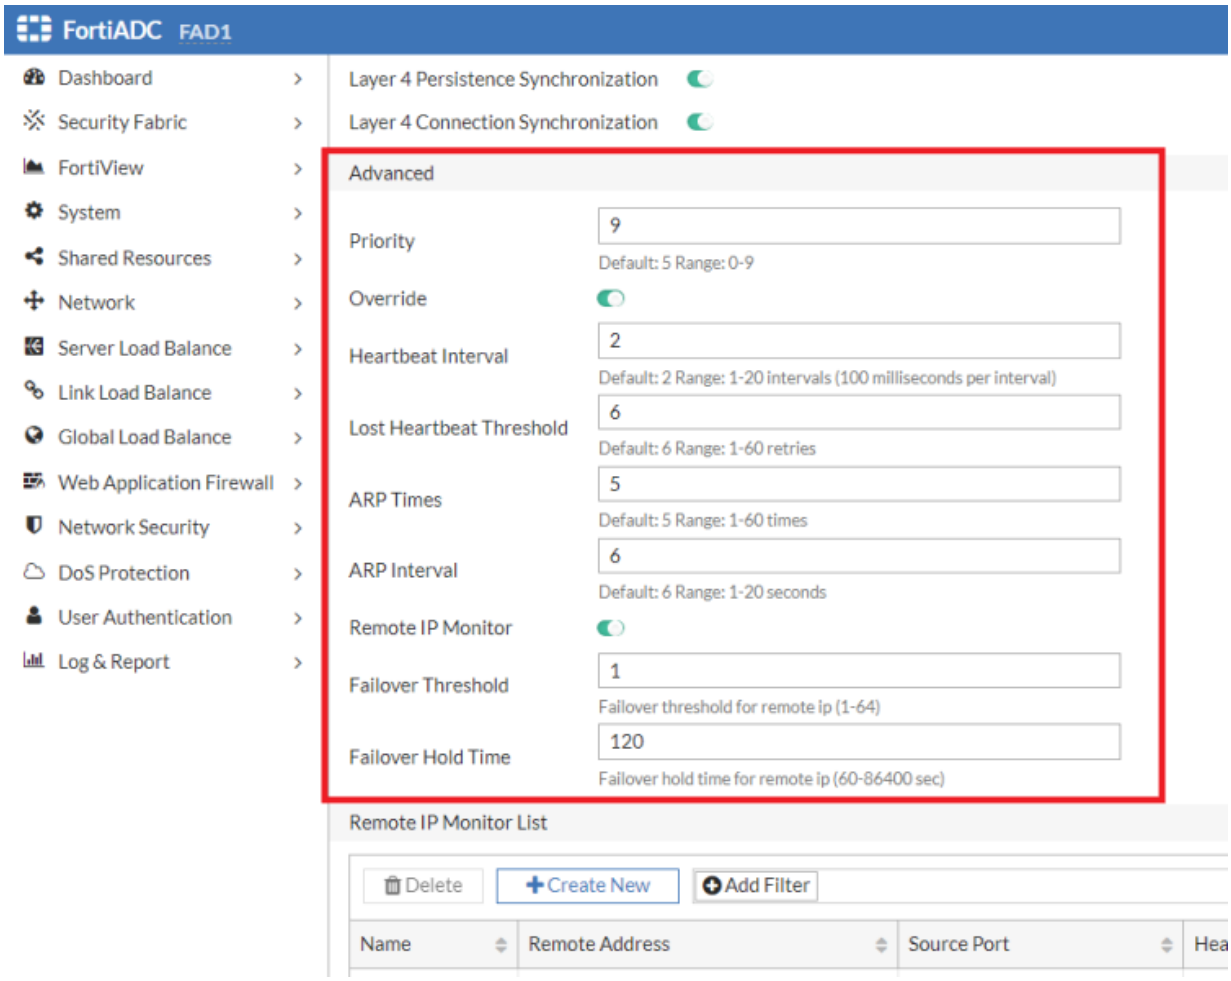
To deploy an Active-Active cluster:

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:
   - 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 HA settings for each node.
   a. Log in to the appliance as the **admin** user.
   b. Complete the HA settings as described in Configuring HA settings.
   The example below shows the configuration for the Primary device FAD1.





Note: The configuration source is determined by the **Config Priority** set in the HA setting. When the configuration priority values of the nodes are different, the configuration of the device with the lower configuration priority will prevail. For the secondary node, ensure the Config Priority is set to a number higher than the device you want as the configuration source.



For the Primary node, ensure to 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 configured in standalone mode, the system utilizes the physical port IP address for network communication. In contrast, when HA is configured in Active-Active mode, the system relies on the HA node IP addresses specified in the `config ha-node-ip-list` for routing and inter-node communication.

To configure in CLI:



```
config system ha
  set mode active-active
  set hbdev port6 port7
  set group-id 14
  set node-list 0 1
  set group-name group1
  set priority 3
  set config-priority 40
  set override enable
  set l7-persistence-pickup enable
  set l4-persistence-pickup enable
  set l4-session-pickup enable
  set monitor port2 port3 port4 port5
end
```

After saving the HA configuration changes, cluster members will initiate the join or rejoin process within the cluster. Upon saving the configuration on the primary node, it will automatically disseminate the updated configuration to the secondary node through the HA synchronization mechanism.

Expected behavior

In Active-Active deployments, it is crucial to enable data synchronization to ensure consistent operation across all nodes. Specifically, configure the following settings:

- Layer 4 Connection Synchronization — Ensures synchronization of TCP connection state data across nodes.
- Layer 4 Session Synchronization — Synchronizes the source IP address table utilized for session persistence to backend servers.
- Layer 7 Session Synchronization — Synchronizes the source IP address table used for application-level persistence to backend servers.



When applying Synchronization settings, configure the Heartbeat Interface (responsible for HA status communication) and the Data Interface (used for traffic forwarding) on separate physical ports to prevent potential conflicts and ensure optimal performance. Using distinct ports minimizes the risk of latency and packet loss, which could otherwise disrupt HA synchronization and impact data traffic.

How the cluster uses synchronized data

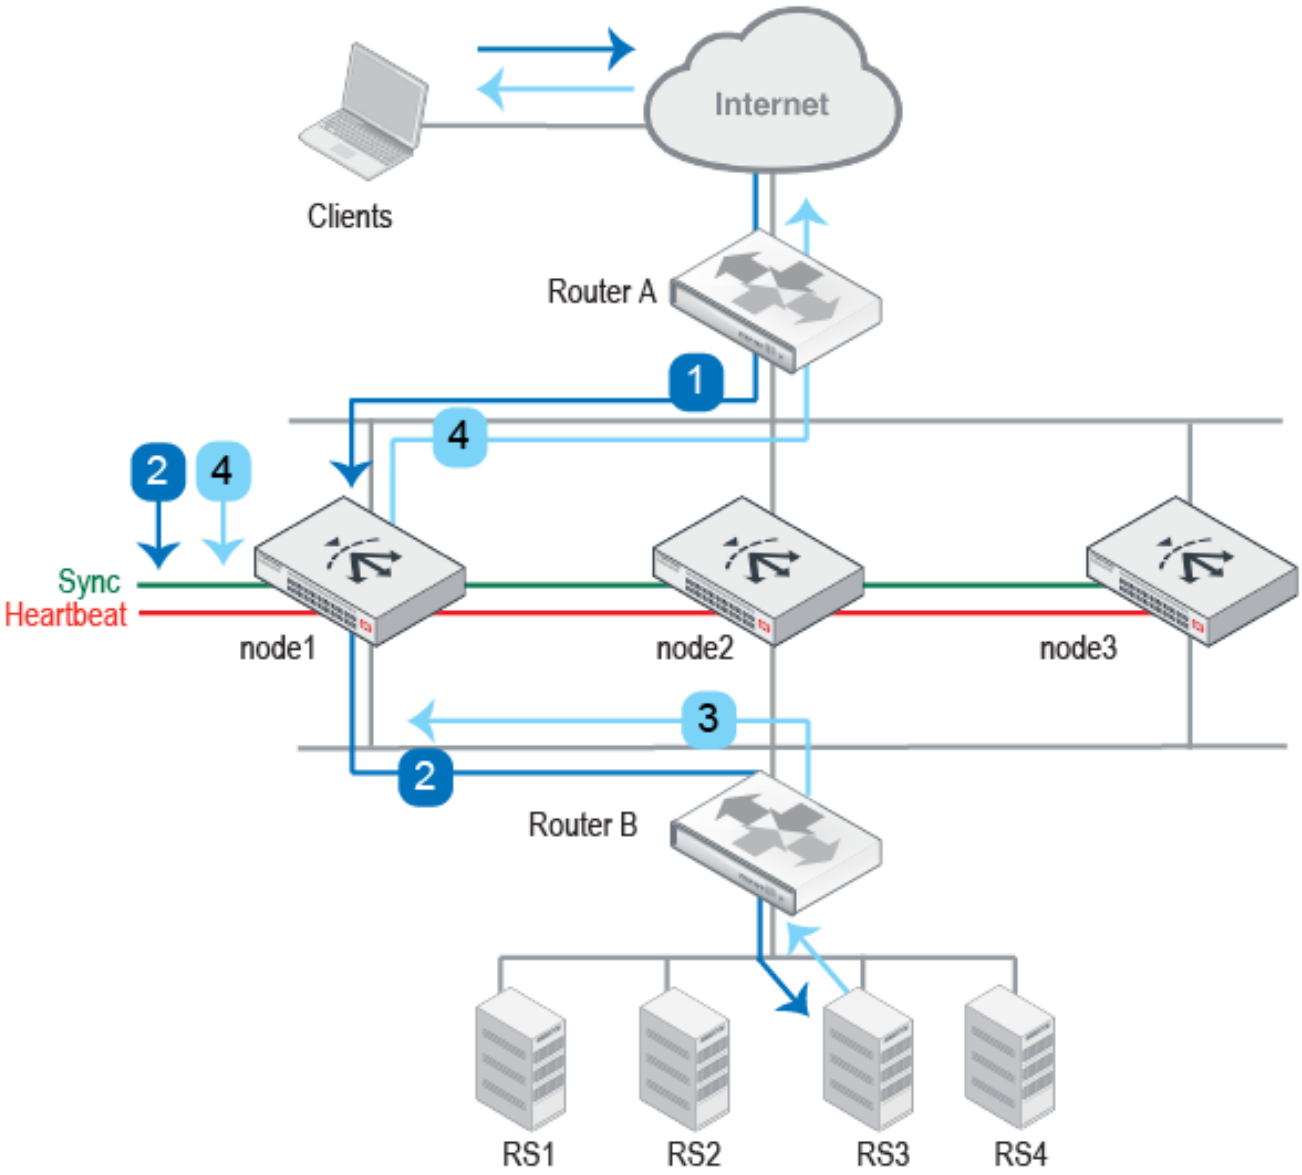
Traffic to TCP virtual servers

When Layer 4 synchronization is enabled, all cluster nodes share TCP connection states and Layer 4 source IP address data for traffic directed to Layer 4 virtual servers. This includes Layer 2 TCP and Turbo HTTP virtual servers, which are

packet-based. The node that initially receives the SYN packet forwards the traffic to the designated real server and concurrently multicasts the session data to the other nodes in the cluster to maintain synchronized session states.

Figure 3 illustrates the sequence of traffic flow where both client-to-server and server-to-client session traffic are consistently routed through the same node.

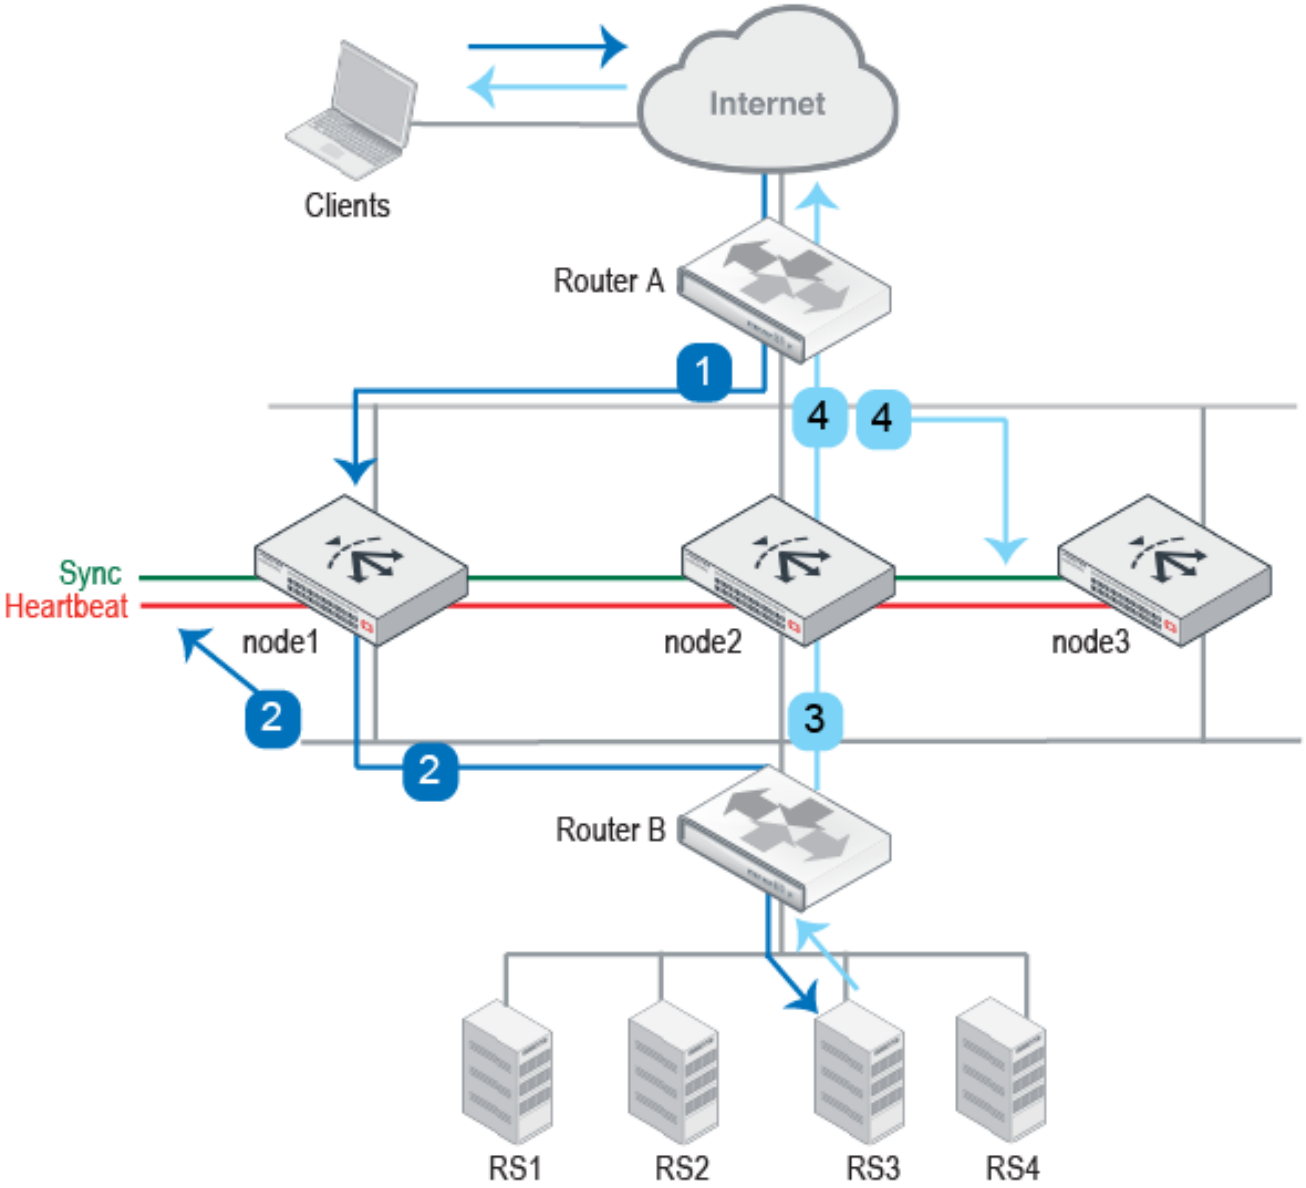
Figure 3: TCP Traffic Flow with ECMP Ensuring Consistent Node Forwarding



1. Router A employs ECMP to direct a client connection request to a specific cluster node—in this case, node1.
2. The cluster node (node 1) processes the traffic by forwarding it to the designated real server and simultaneously multicasts the session data to the entire cluster via the data port.
3. Router B uses ECMP to route the server's response back to the same cluster node—also node1.
4. The cluster node (node1) then forwards the response to the client and again multicasts the session data to the cluster.

Figure 4 illustrates the traffic flow sequence where client-to-server and server-to-client session traffic are routed through different nodes. The diagram shows that synchronization has already taken place before the second node handles the response traffic, ensuring seamless session continuity across the cluster.

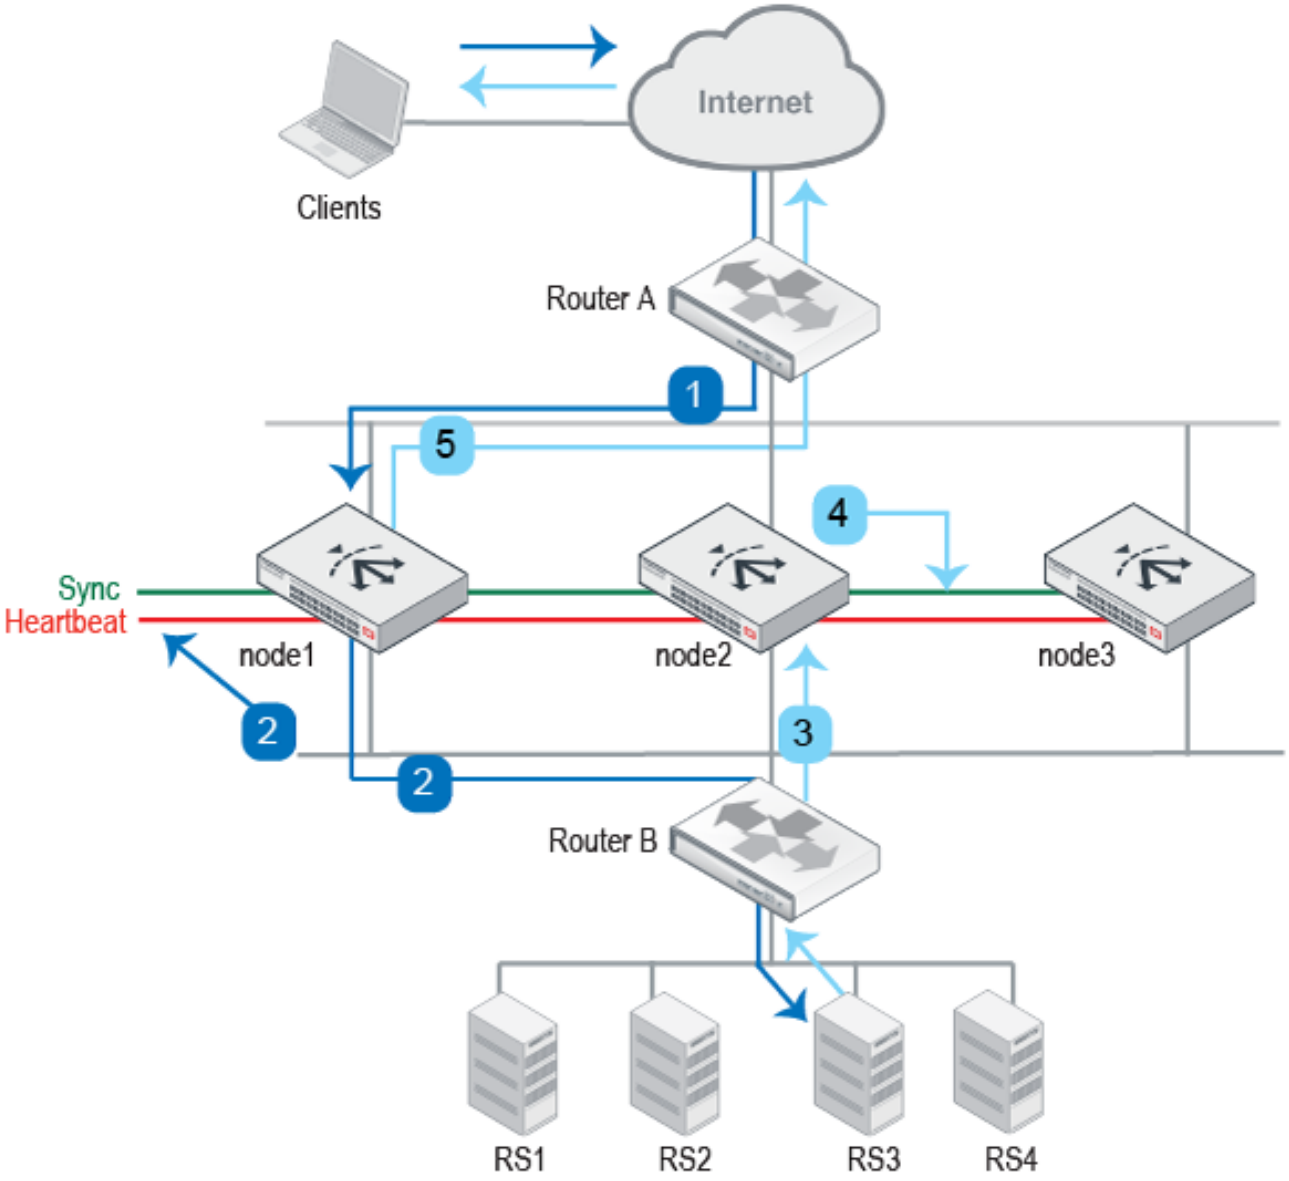
Figure 4: TCP traffic flow post-synchronization



1. Router A employs ECMP to direct a client connection request to a selected cluster node—in this case, node1.
2. The cluster node (node1) processes the request by forwarding traffic to the real server and multicasting session data to the cluster via the data port.
3. Router B utilizes ECMP to route the server's response traffic to a different cluster node, selecting node2.
4. If session synchronization between node1 and node2 has occurred, node2 forwards the response traffic to the client and multicasts the updated session data to the cluster.

Figure 5 illustrates the traffic flow sequence when client-to-server and server-to-client session traffic are routed through different nodes, with synchronization not yet completed by the time the second node receives the response traffic.

Figure 5: TCP traffic flow when synchronization has not yet occurred



1. Router A uses ECMP to forward a client connection request to a selected cluster node—in this case, node1.
2. The cluster node (node1) forwards the traffic to the real server and multicasts the session data to the cluster.
3. Router B employs ECMP to direct the server's response traffic to a different cluster node, which is node2.
4. Since session synchronization between node1 and node2 has not occurred, node2 multicasts the traffic to the cluster.
5. Upon receiving the traffic from node2, node1 forwards it to the client and multicasts the updated session data.

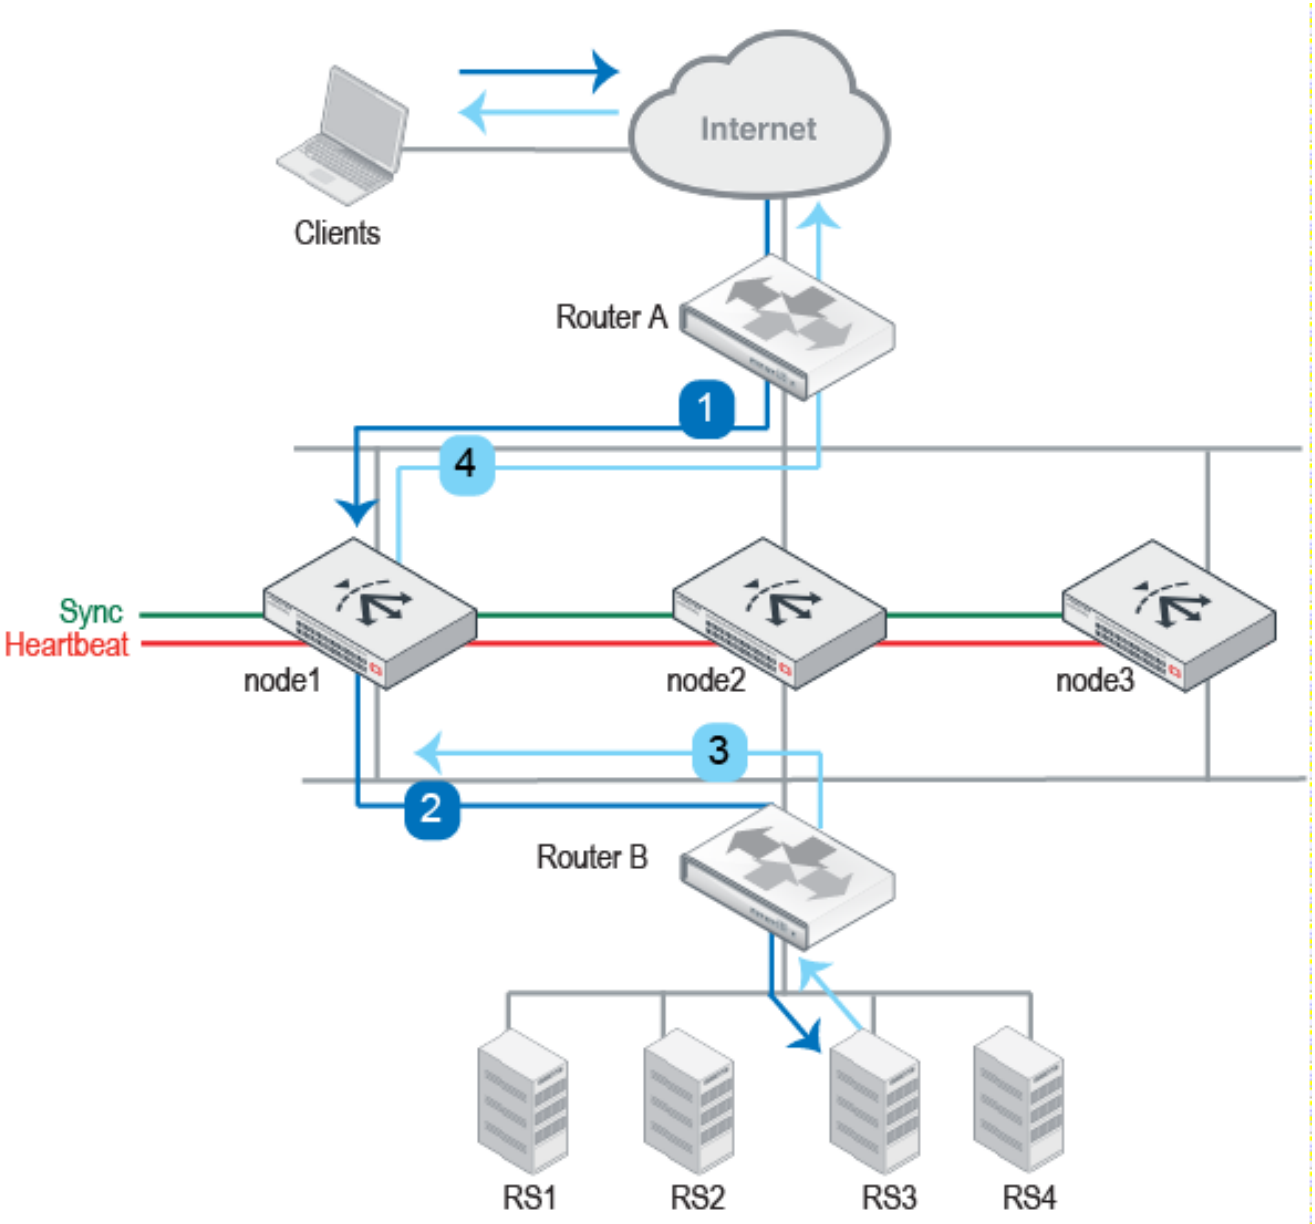
Traffic to HTTP virtual servers

When Layer 7 synchronization is enabled, cluster nodes handle source IP data for traffic to HTTP virtual servers based on the virtual server profile's Source option. If the Source option is enabled, FortiADC forwards traffic to the real server

with the client's source IP address. If the Source option is disabled, the traffic is forwarded with the FortiADC HA cluster node's IP address.

Figure 6 illustrates the sequence of the traffic flow when the Source option is not enabled.

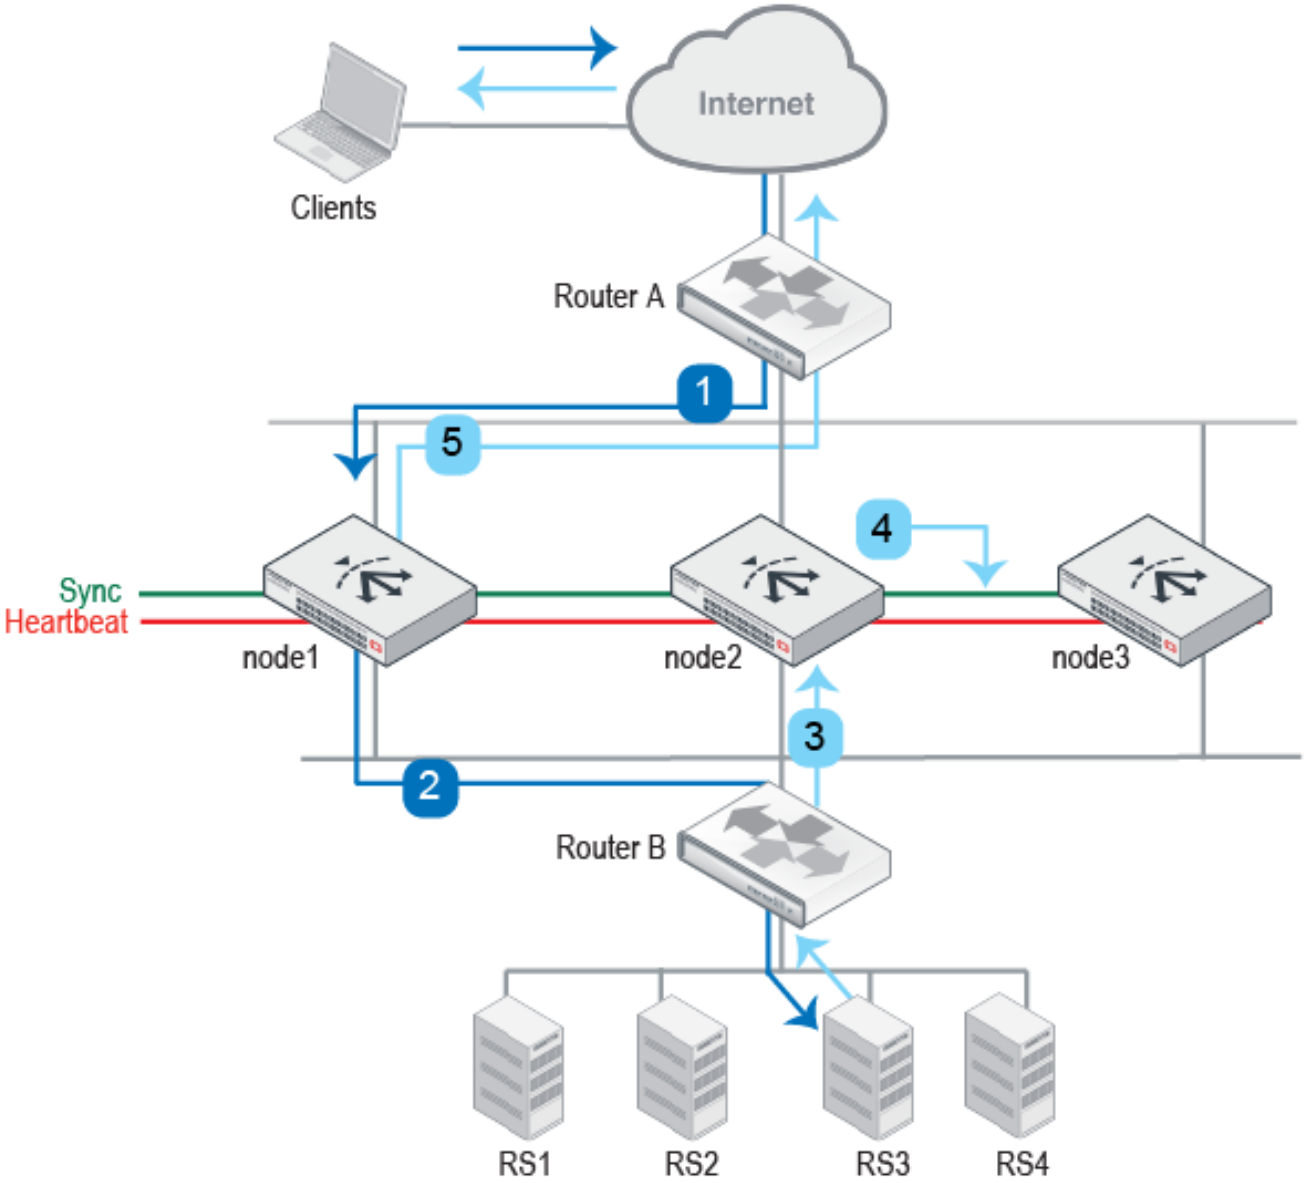
Figure 6: HTTP traffic flow when the Source option is not enabled



1. Router A uses ECMP to select a cluster node for forwarding a client connection request—in this case, node1.
2. The cluster node (node 1) forwards the traffic to the real server. Since the Source option is disabled, the source IP address in the FortiADC-to-real-server traffic is the node1 HA cluster node IP address, which becomes the destination IP address for the response traffic.
3. Router B does not use ECMP; it forwards the traffic directly to the node1 HA cluster IP address.
4. The cluster node (node1) retrieves the real client IP address from its session table and forwards the traffic to the client.

Figure 7 illustrates the sequence of the traffic flow when the Source option is enabled.

Figure 7: HTTP traffic flow when the Source option is enabled



1. Router A uses ECMP to select a cluster node for forwarding a client connection request—in this case, node1.
2. The cluster node (node1) processes the request and forwards the traffic to the real server. With the Source option enabled, the source IP address in the FortiADC-to-real-server traffic is the true client IP address, which then becomes the destination IP address for the server-to-client traffic.
3. Router B uses ECMP to forward the server response traffic. In this example, it routes the traffic to node2.
4. Since node2 was not expecting the server-to-client response, it multicasts the response traffic to the cluster.
5. Node1 receives the response from node2 and forwards it to the client.

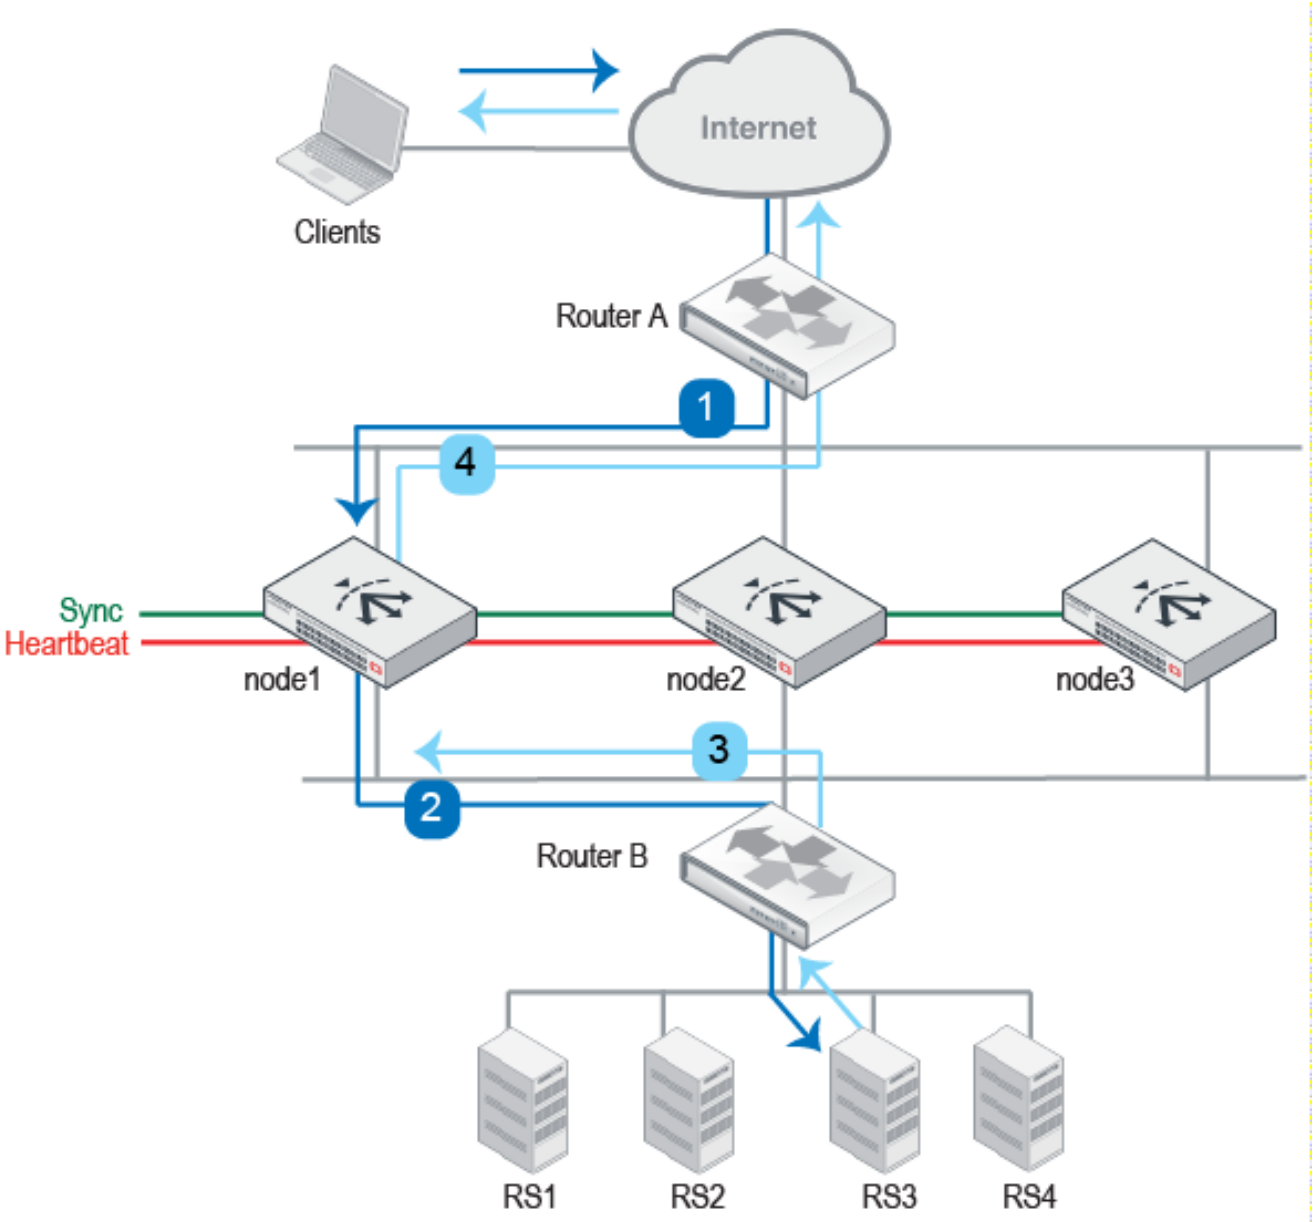
Note: In an Active-Active deployment, the virtual server profile Source option adds latency to the transaction. To reduce latency, use an alternative to the Source option, such as the X-Forwarded-For option, if you require the original client IP to be logged by the real server.

FTP traffic and traffic processed by firewall rules

In an Active-Active deployment, FTP and firewall traffic are exclusively routed through the primary node. FTP communication involves both a control connection and a data connection between the client and server, which complicates the handling of asymmetric routes in an Active-Active cluster. Moreover, traffic subject to stateful firewall rules is not distributed across nodes and is processed solely by the primary node.

Figure 8 illustrates the traffic flow sequence when ECMP routing directs traffic through the primary node.

Figure 8: FTP or firewall traffic flow when ECMP selects the primary node

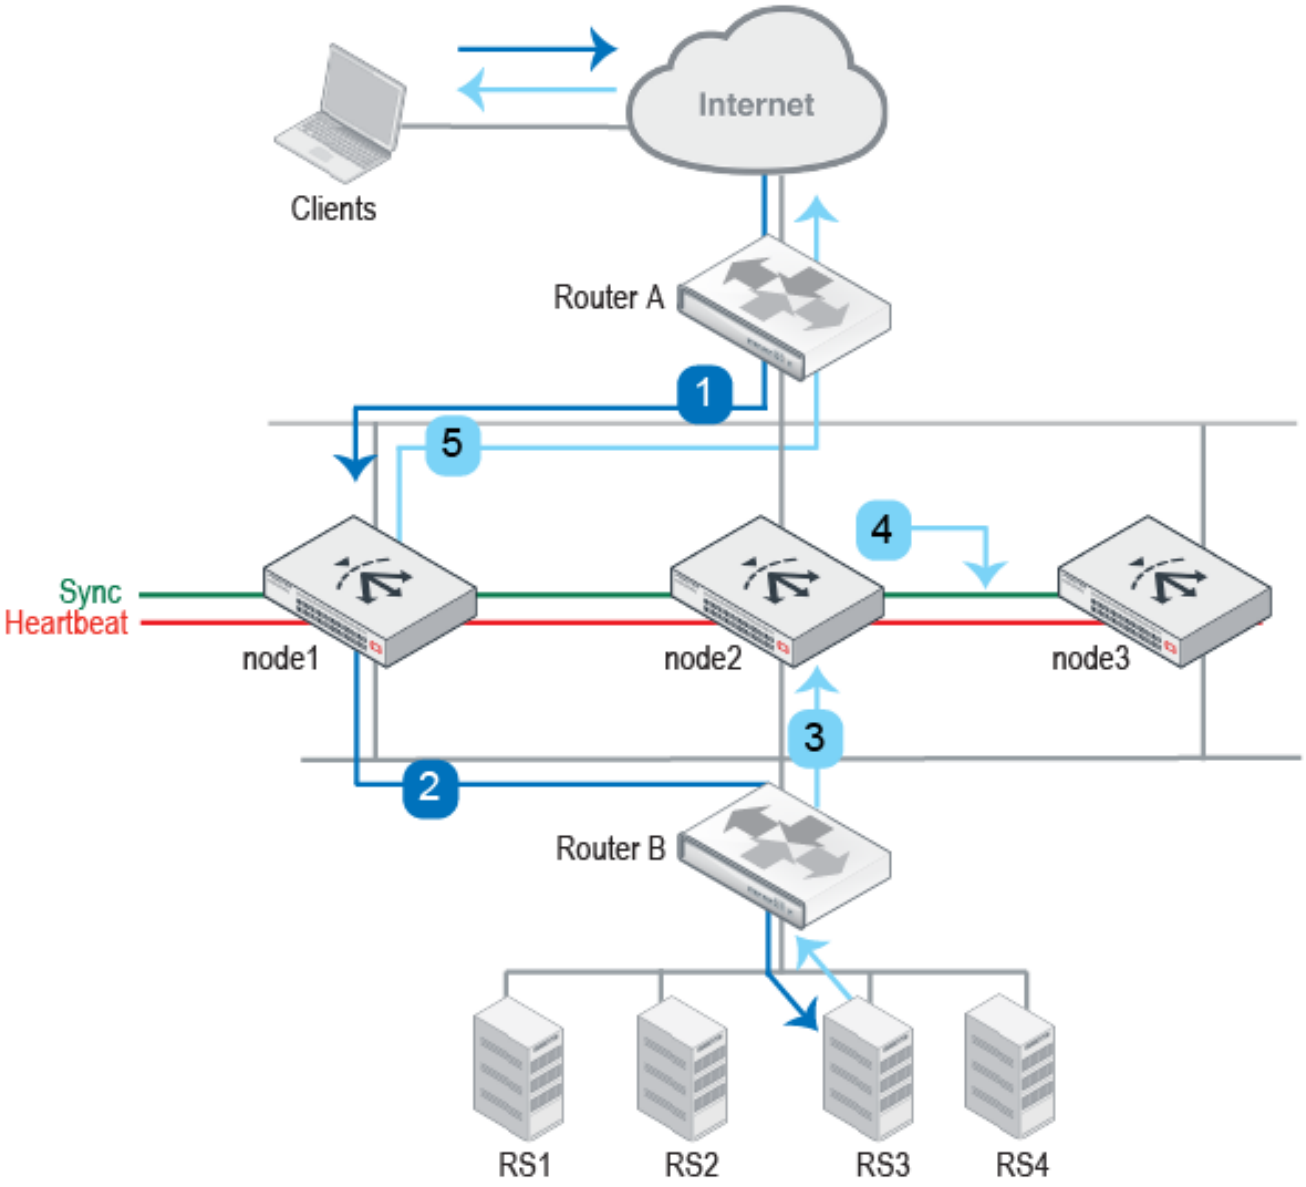


1. Router A uses ECMP routing to select a cluster node for forwarding the client connection request, choosing node1 as the primary node.
2. The primary node (node1) forwards the client request to a real server.

3. Router B employs ECMP to select a cluster node for routing the server response traffic, again choosing node1.
4. The primary node (node1) forwards the server response traffic back to the client.

Figure 9 illustrates the sequence of the traffic flow when ECMP results in an asymmetric route.

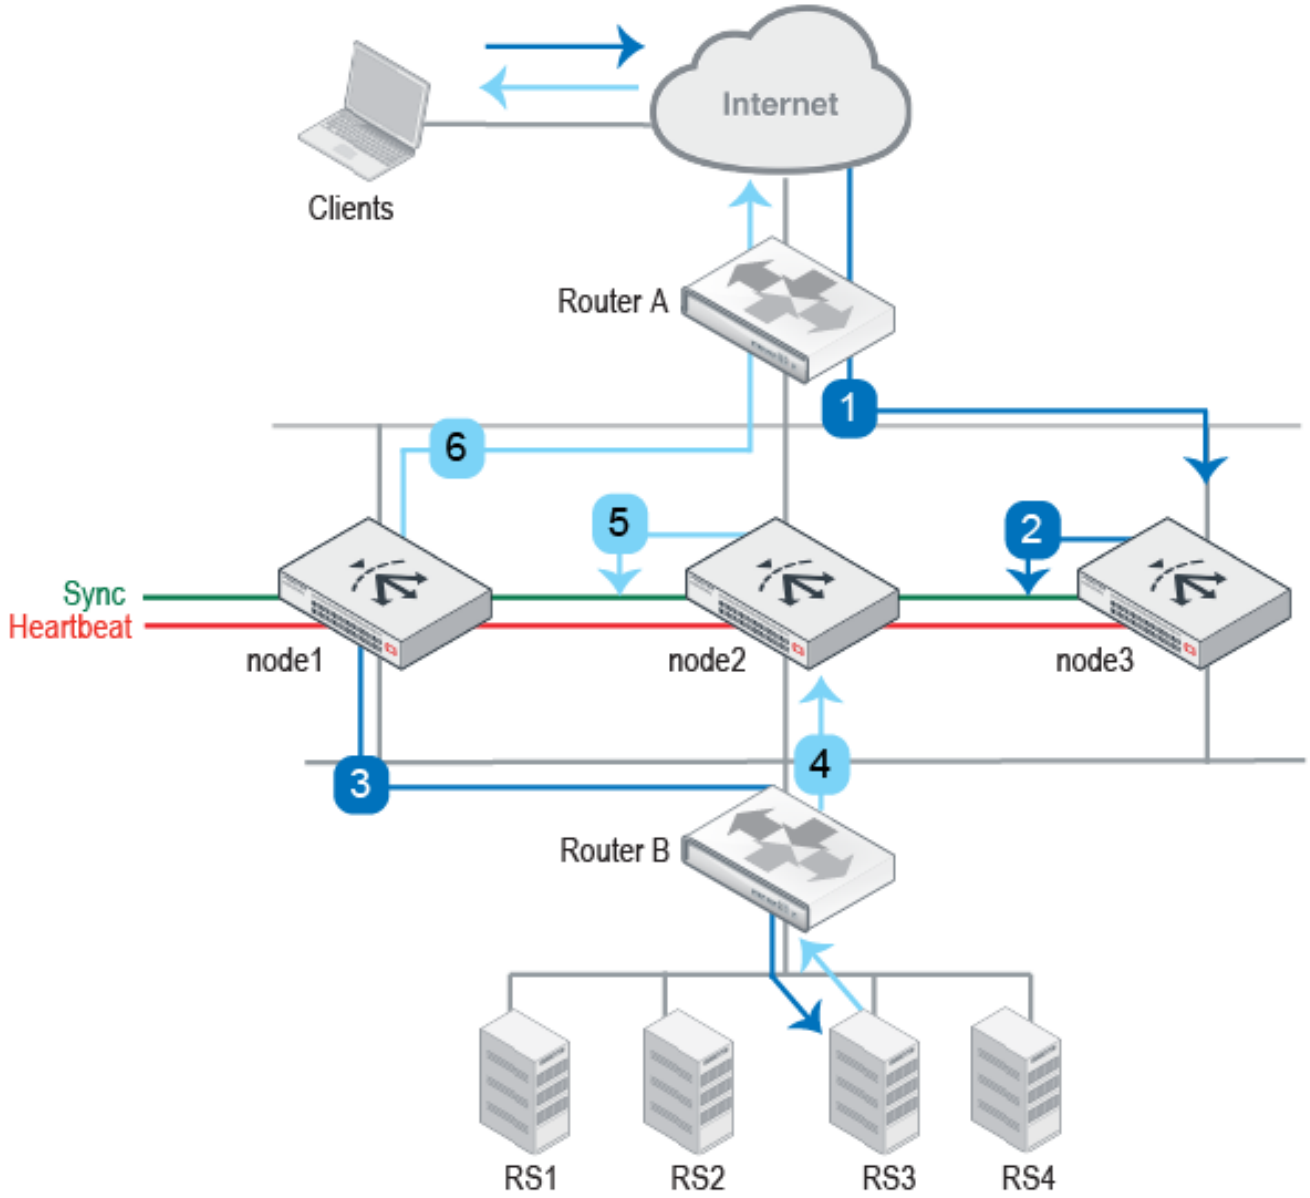
Figure 9: FTP or firewall traffic flow when ECMP results in an asymmetric route



1. Router A uses ECMP routing to choose a cluster node for forwarding the client connection request, selecting node1 as the primary node.
2. The cluster node (node1) forwards the client request to the real server.
3. Router B applies ECMP to select a cluster node for the server response traffic, choosing node2.
4. Since node2 did not expect the server response, it forwards the traffic to the cluster.
5. The primary node (node1) receives the server response traffic from node2 and forwards it to the client.

Figure 10 illustrates the traffic flow sequence when ECMP routing directs client-to-server traffic to a non-primary node.

Figure 10: FTP or firewall traffic flow when ECMP results in traffic sent to a non-primary node



1. Router A uses ECMP routing to select a cluster node for forwarding the client connection request to a real server destination IP address, choosing node3.
2. Since firewall traffic is handled exclusively by the primary node, node3 multicasts the session data to the cluster.
3. The primary node processes the forwarded traffic and sends it to the real server.
4. Router B uses ECMP to choose a cluster node for forwarding the server response traffic, selecting node2.
5. Because node2 did not expect the server response, it forwards the traffic to the cluster.
6. The primary node receives the traffic from node2 and forwards it to the client.

Best practices

The following best practices should be observed:

- Maintain heartbeat links — Ensure the heartbeat link(s) are consistently operational. If a heartbeat is interrupted—such as by a temporary network cable disconnection—other nodes may incorrectly perceive the primary node as failed, triggering failover in an active-passive deployment. This could result in both nodes operating as active nodes simultaneously if no actual failure has occurred.
- Improve fault tolerance — When connecting HA appliances through switches, link the ports through two separate switches to enhance fault tolerance and reliability. Avoid connecting these switches to the broader network to mitigate potential security vulnerabilities and prevent network load from introducing latency in the heartbeat, which could inadvertently trigger failover.

Deploying an Active-Active-VRRP cluster

HA Active-Active-VRRP employs a VRRP-like protocol to enable concurrent traffic handling across all FortiADC nodes configured with multiple traffic groups, enhancing throughput and performance through effective load balancing and redundancy. This mode requires fewer deployment conditions than HA Active-Active, making it suitable for environments needing increased performance and active traffic management.

Please note that FortiADC supports VRRP configurations only between FortiADC units. Integration with VRRP groups that include third-party VRRP devices is not supported.

This section covers the following information about the Active-Active-VRRP HA deployment:

- About HA Active-Active-VRRP Mode in FortiADC on page 236
- HA Active-Active-VRRP configuration overview on page 237
- HA-AA-VRRP configuration steps on page 241

About HA Active-Active-VRRP Mode in FortiADC

The Virtual Router Redundancy Protocol (VRRP) is designed to eliminate the single point of failure in static default routed environments. It employs an election protocol to dynamically assign the responsibility of a virtual router to one of the VRRP routers on a LAN. The router that controls the IP addresses associated with the virtual router is designated as the primary, and it forwards packets sent to these IP addresses. The election process ensures dynamic failover of forwarding responsibilities if the primary becomes unavailable. This setup allows any of the virtual router's IP addresses on a LAN to serve as the default first-hop router for end-hosts. The primary advantage of VRRP is its ability to provide a highly available default path without the need to configure dynamic routing or router discovery protocols on each end-host.

A virtual router is characterized by its Virtual Router Identifier (VRID) and a set of associated IP addresses. Each VRRP router can associate a virtual router with its real interface address and may be configured with additional virtual router mappings and priorities for backup purposes. Coordination of VRID-to-address mappings is required among all VRRP routers on a LAN.

FortiADC implements a VRRP-like high availability (HA) mode rather than the exact VRRP protocol. Consequently, FortiADC's HA Active-Active VRRP mode is referred to as a VRRP-like HA mode.

VRRP configurations serve as an HA solution to maintain network connectivity with the Internet or other networks in the event of a default router failure. VRRP allows designation of primary and backup routers, where the primary router handles traffic, and backup routers monitor the primary, taking over traffic forwarding if the primary router fails.

For detailed specifications, refer to RFC 3768.

Advantages of HA Active-Active-VRRP

When comparing FortiADC's HA Active-Active-VRRP mode with traditional HA Active-Passive and Active-Active clusters, several distinct advantages emerge. These advantages highlight the efficiency, flexibility, and performance benefits that HA Active-Active-VRRP mode offers over the other HA configurations:

Device-Level Failover in HA Active-Active Mode

In HA Active-Active mode, failover occurs at the device level, meaning that the entire device is switched over even if only one monitored port fails.

Traffic Load Design with HA Active-Active-VRRP

In FortiADC's HA Active-Active-VRRP mode, virtual servers can be manually assigned to specific traffic groups, allowing for customized traffic load distribution based on virtual servers.

Efficient Session Synchronization

HA Active-Active-VRRP mode only synchronizes the session and persistence tables to the next available device within the same traffic group using the "failover-order" command. In clusters with more than two devices, this targeted synchronization is more efficient than in HA Active-Passive or Active-Active modes, where the session/persistence table is synchronized across the entire HA group. This approach effectively supports N+M hot-backup functionality.

Simplified Traffic Distribution

HA Active-Active mode requires an external router configured with Equal-Cost Multi-Path (ECMP) routing to distribute traffic among Active-Active nodes. In contrast, HA Active-Active-VRRP mode eliminates the need for an external router for traffic distribution. Both client and server gateways can simply point to the VRRP floating IP.

Consistent Traffic Handling

In HA Active-Active-VRRP mode, all devices within the same traffic group share the same HA status. When both client-side and server-side gateways are pointed to the same floating IP, incoming and outgoing traffic is directed to the same device. As a result, there is no need to multicast traffic across the HA group, enhancing network performance and efficiency.

Load Balancing Under Stress

In HA Active-Active mode, the AA-Primary node handles the traffic of all AA-Passive nodes. If multiple AA devices fail, the AA-Primary may become overloaded, leading to potential issues under high traffic stress.

Direct Server Access

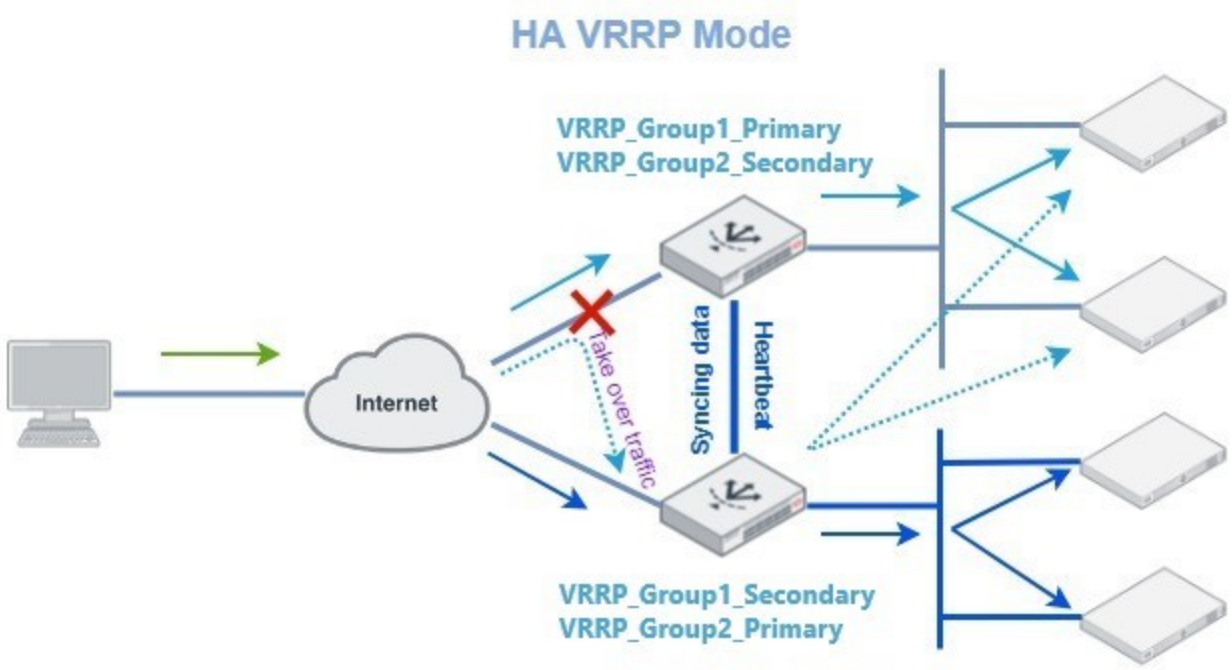
In HA Active-Active mode, direct access to the real server's IP address from the client is not possible. However, this limitation does not exist in HA Active-Active-VRRP mode, allowing for more flexible network configurations.

HA Active-Active-VRRP configuration overview

The FortiADC HA Active-Active-VRRP configuration encompasses the following components:

- Traffic groups and feature configuration — For details, see Traffic Group on page 246.
- Interface and virtual server configuration — This involves configuring the floating IPs and associating them with the appropriate traffic groups.
- HA configuration — Configures High Availability settings, including failover and synchronization parameters.

Figure 1: HA Active-Active-VRRP Mode Deployment



HA-VRRP mode allocates resources into distinct groups, allowing the creation of multiple VRRP groups and the assignment of public IP resources to these groups. This configuration enables an alternative Active-Active mode known as Active-Active-VRRP, commonly referred to as "HA-VRRP." In this mode, each HA node has its own interface IP, while the floating IP is defined on an interface within one of the VRRP groups.

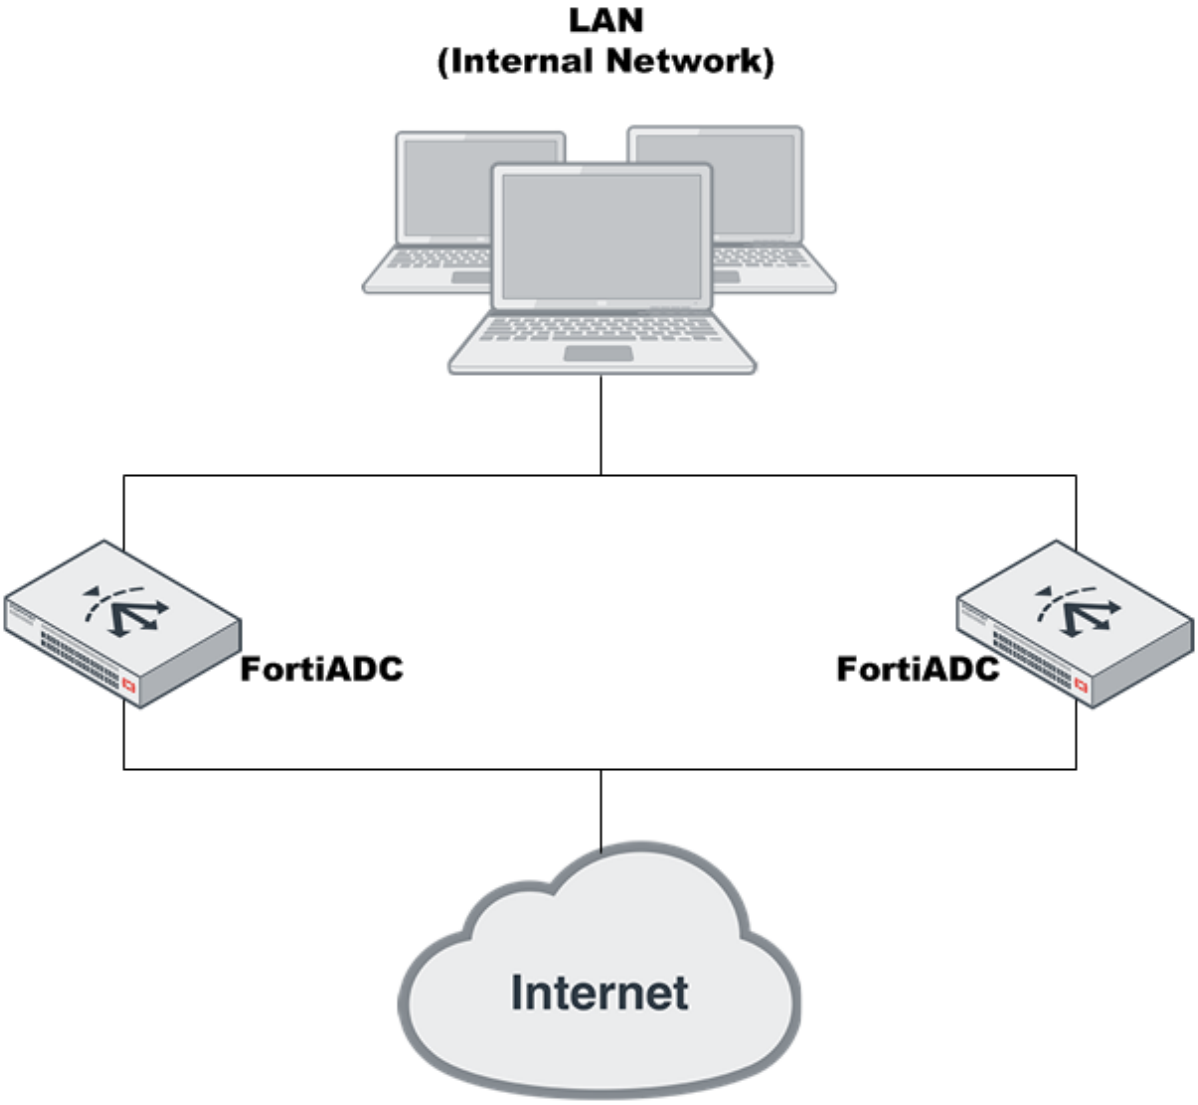
In a typical configuration, connected devices or servers are configured to use the floating IP address associated with the VRRP group as their default gateway. This floating IP is only active on the currently designated active VRRP group. During a failover event, the floating IP transitions to the newly active VRRP group, ensuring that the floating IP remains consistently available.

Refer to **Figure 1** above for an example of HA-VRRP mode. In this scenario, two VRRP groups are established: VRRP_Group1 and VRRP_Group2. FortiADC1 is configured as the active node for VRRP_Group1 and the passive node for VRRP_Group2. Conversely, FortiADC2 serves as the passive node for VRRP_Group1 and the active node for VRRP_Group2. The real servers are allocated to these two groups, with servers in VRRP_Group1 pointing their default gateway to the floating IP of VRRP_Group1, and servers in VRRP_Group2 pointing their default gateway to the floating IP of VRRP_Group2.

Under normal operational conditions, FortiADC1 manages the traffic for VRRP_Group1, while FortiADC2 handles traffic for VRRP_Group2. If a failure is detected in any monitored link or device, the HA peer will take over the traffic handling responsibilities to ensure uninterrupted service.

FortiADC units can be deployed as either primary or backup Virtual Router Redundancy Protocol (VRRP) routers and seamlessly integrate into existing VRRP configurations. In a VRRP setup, if a FortiADC unit designated as the primary fails, a backup unit automatically assumes its role, ensuring uninterrupted network traffic processing. Traffic directed to the failed primary unit is transparently redirected to the backup unit. Once the failed FortiADC unit is restored, it reverts to its primary role, resuming traffic processing for the network.

Figure 2: An Active-Active-VRRP cluster configuration using two FortiADC units

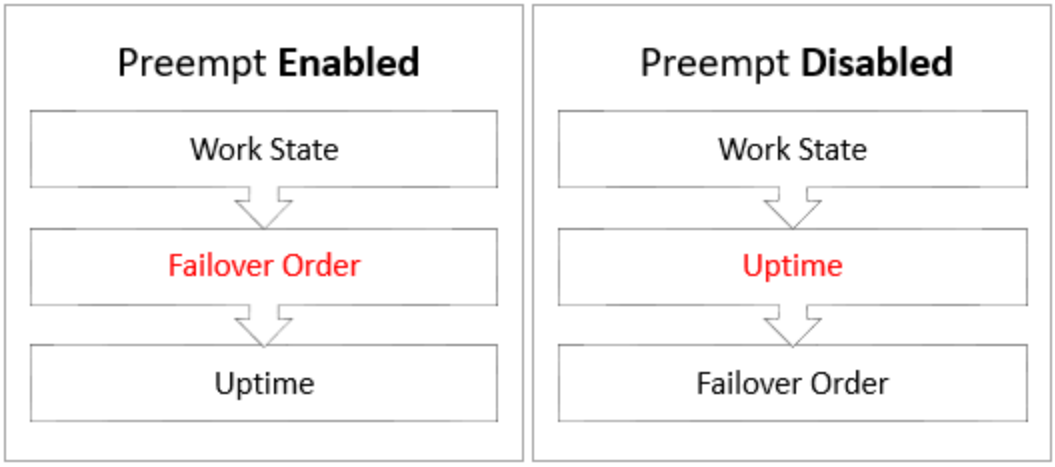


In an Active-Active-VRRP cluster, one node is designated as the primary node for a traffic group, while the remaining nodes serve as member nodes of the group. Traffic from upstream sources can be distributed across up to eight member nodes (one local and seven peers). The Active-Active-VRRP cluster also supports failover: if the primary node fails, the traffic group's responsibilities are automatically assumed by one of the backup nodes. This backup node will then send a gratuitous ARP message to adjacent devices, updating them to redirect traffic to its own MAC address across all network interfaces within the traffic group.

Active node election process in an Active-Active-VRRP cluster

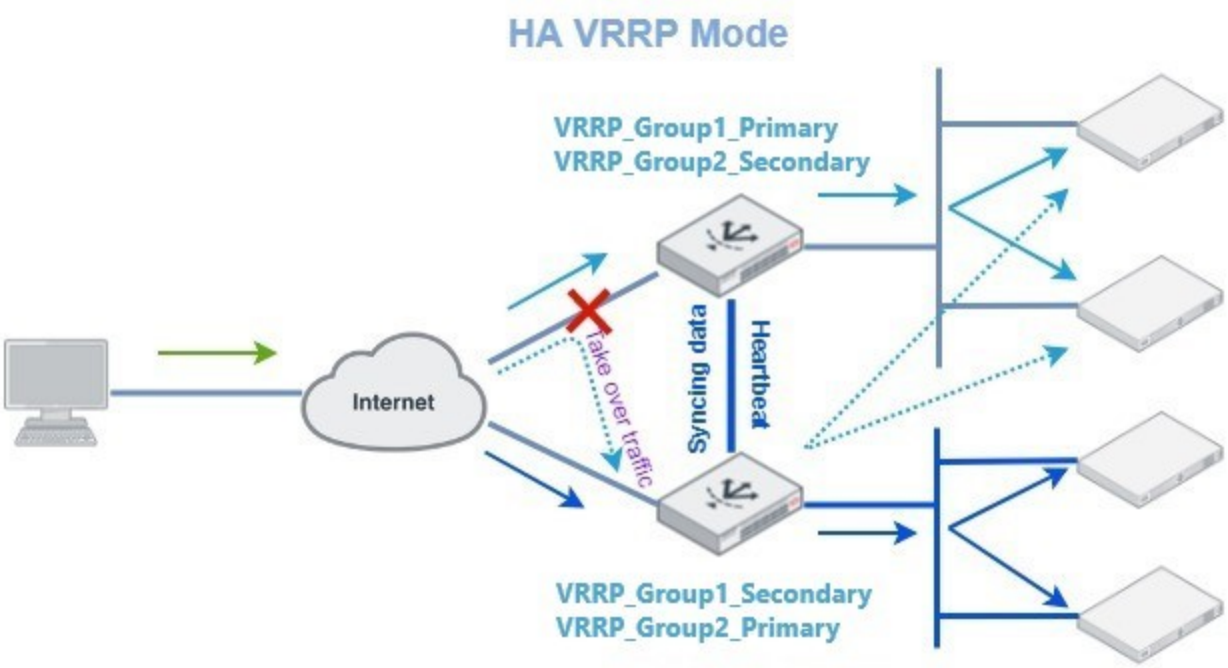
In an Active-Active-VRRP cluster, the Active and Passive nodes are elected by the traffic groups. The system determines the active node by evaluating and prioritizing specific operational conditions and configuration parameters. Depending on whether the Preempt setting has been enabled or disabled, the order each condition is considered differs slightly (illustrated in **Figure 3**).

Figure 3: Priority Criteria for Active Node Election



Criteria	Description
Work State	The "work state" refers to the operational status that is currently influenced solely by the remote IP check. If a device reports a remote IP check failure, it is considered down. In cases where one device shows a failure while another does not, the device without the failure is considered operational. If all devices report a remote IP check failure, the system proceeds to the next criterion.
Failover Order	Failover order is a configuration option within HA settings. It defines the sequence in which alternative devices should take over based on their Local Node ID.
Uptime	The "uptime" refers to the duration the device has been operational. The device with the longest uptime is given precedence. If multiple devices have identical uptime, the system proceeds to the next criterion.

HA-AA-VRRP configuration steps

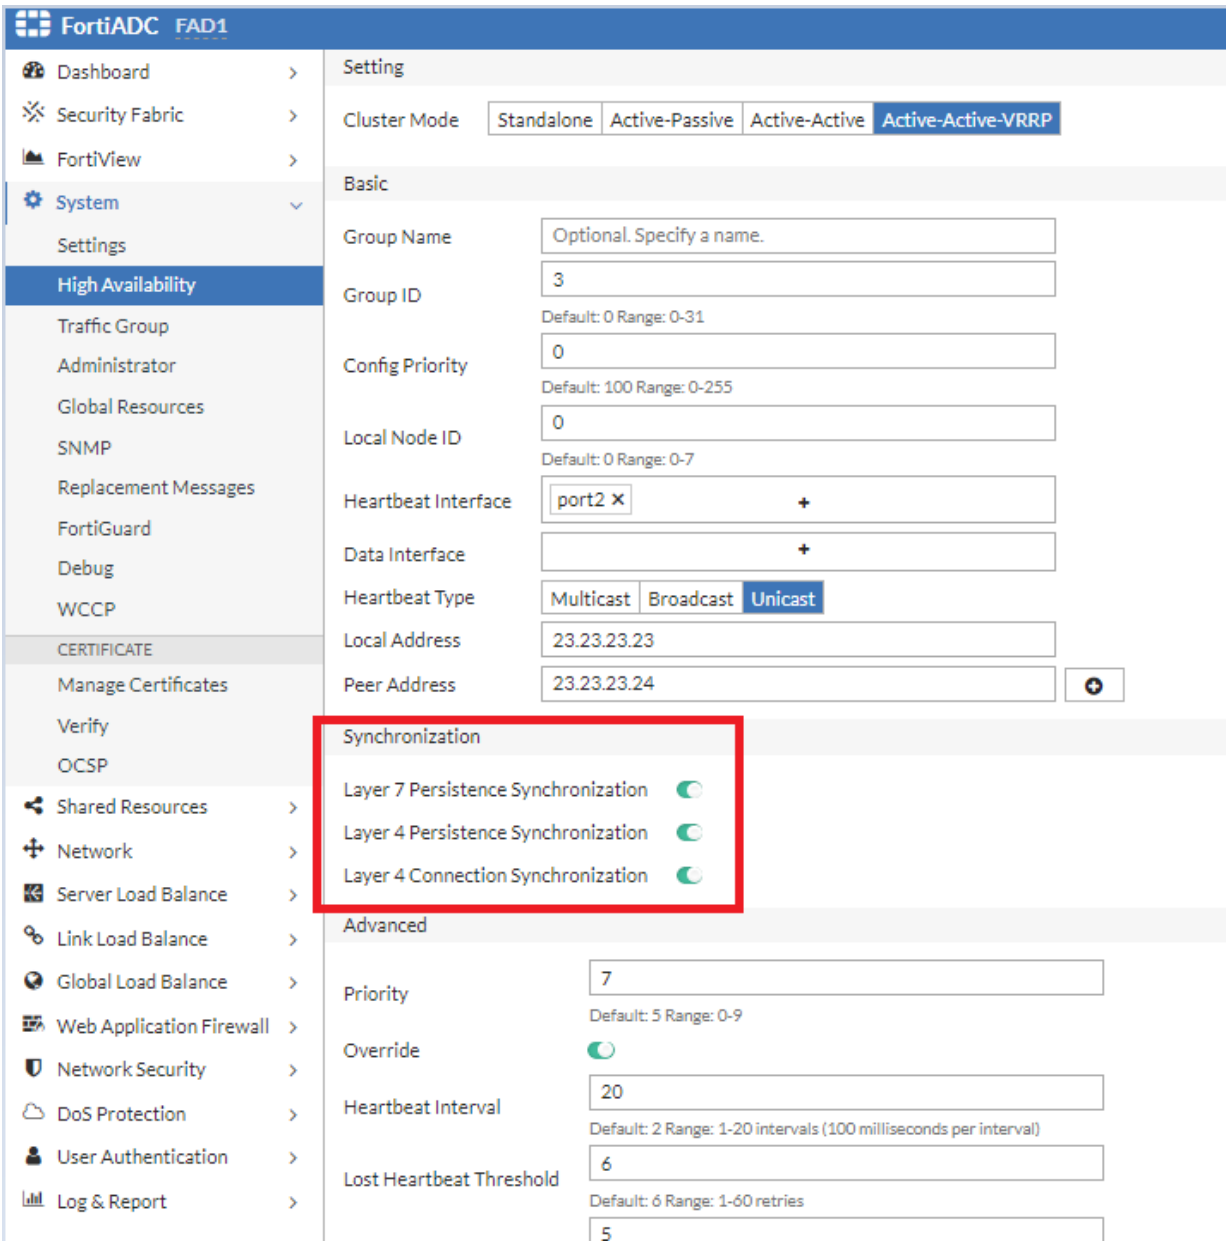


The following steps outline the deployment of an Active-Active-VRRP cluster in which each node is configured with two virtual servers: VS1, associated with the VRRP_Group1 traffic group, and VS2, associated with the VRRP_Group2 traffic group. In this setup, traffic directed to VS1 will be managed by FAD1, while traffic to VS2 will be managed by FAD2. In the event of a failure on one of the FortiADC units, the other device will automatically take over its traffic. Port2 is assigned to VRRP_Group1, and Port3 is assigned to VRRP_Group2.

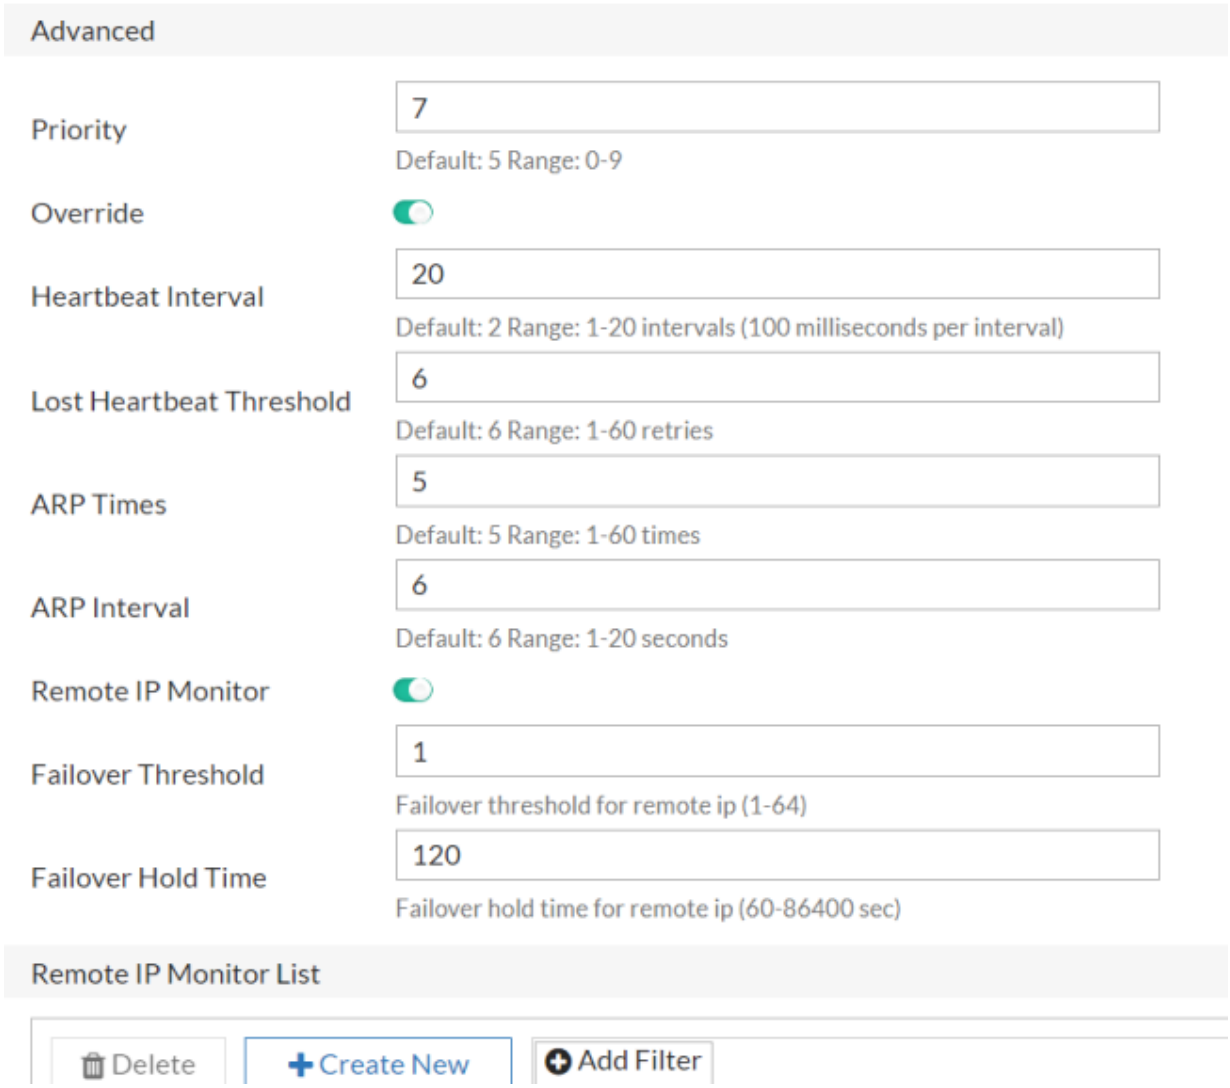
To ensure successful cluster negotiation, all HA devices must use identical heartbeat and data ports, share the same Group ID, and have unique Local Node IDs.

To deploy an Active-Active-VRRP 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 HA settings for each node.
   a. Log in to the appliance as the **admin** user.
   b. Complete the HA settings as described in Configuring HA settings.
   The example below shows the configuration for the Primary device FAD1.



Note: The configuration source is determined by the **Config Priority** set in the HA setting. When the configuration priority values of the nodes are different, the configuration of the device with the lower configuration priority will prevail. For the secondary node, ensure the Config Priority is set to a number higher than the device you want as the configuration source.

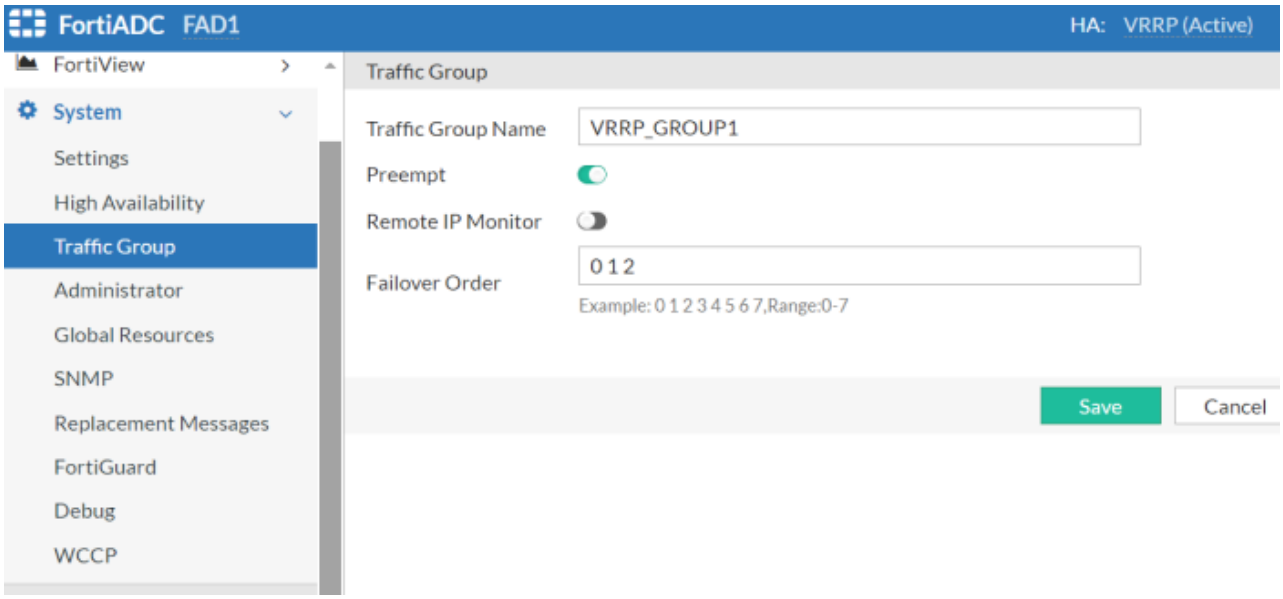


To configure in CLI:

```
config system ha
  set mode active-active-vrrp
  set hbdev port4 port5
  set group-id 15
  set local-node-id 0
  set group-name grp2
  set config-priority 20
  set override enable
  set l7-persistence-pickup enable
  set l4-persistence-pickup enable
  set l4-session-pickup enable
end
```

4. Configure the required VRRP traffic groups as described in Traffic Group on page 246.
   Once the devices in the cluster establishes the HA VRRP relationship, then configuration changes that occur on any

of the HA nodes can be synchronized to the other nodes. In this example, you can create the traffic groups on one of the nodes, FAD1.

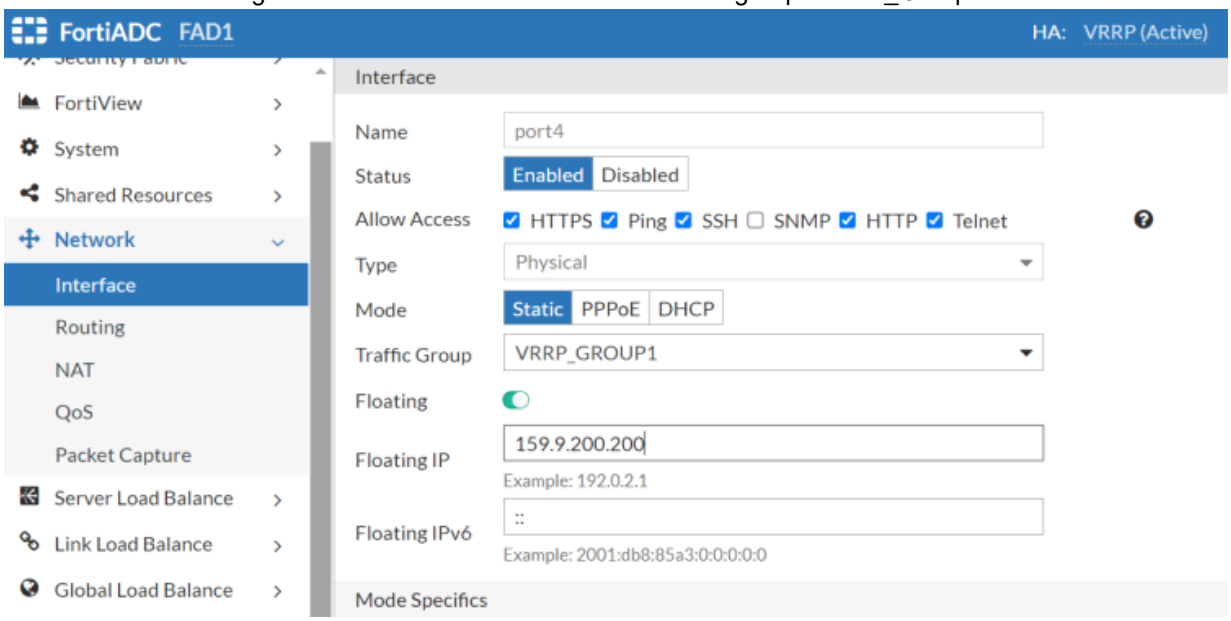


To configure in CLI:

```
config system traffic-group
  edit "VRRP_Group1"
    set failover-order 0 1
    set preempt enable
  next
  edit "VRRP_Group2"
    set failover-order 1 0
    set preempt enable
  next
end
```

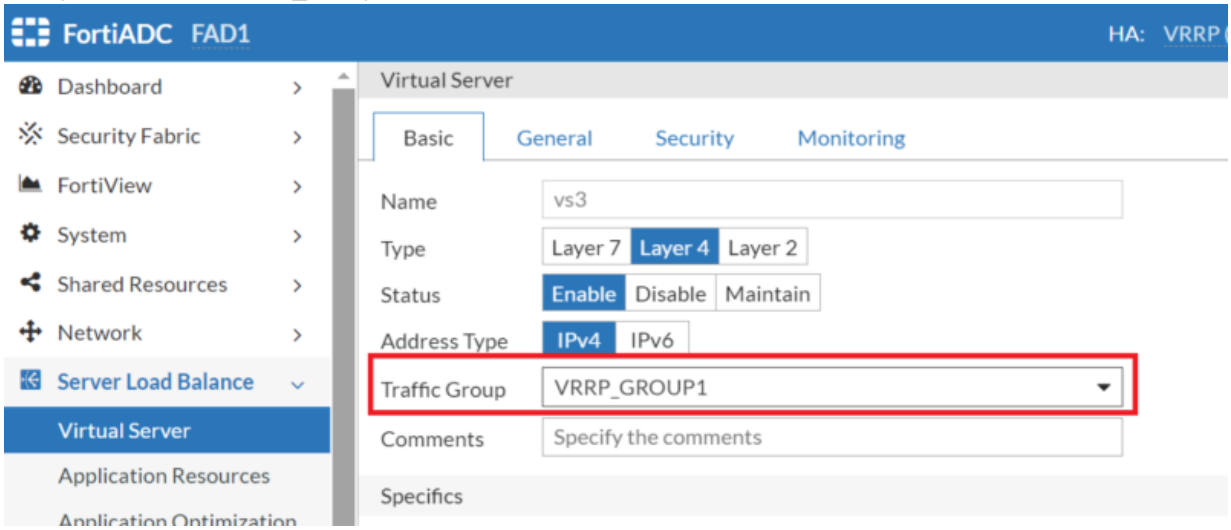
5. Assign the interface, virtual server, and other resources to the VRRP traffic group.
By default, all interfaces, virtual servers, and other resources are assigned to the "default" traffic group. It is recommended to assign these resources to a custom traffic group for better organization and management.

a. Edit the Interface configuration to associate with the custom traffic group "VRRP_Group1".



Note that the floating IP only functions within the Primary traffic group. In this example, port2 is part of VRRP_Group1, with FAD1 currently serving as the Primary for VRRP_Group1. As a result, the IP address '159.9.200.200' is currently active on FAD1. If FAD1 fails, FAD2 will assume the role of Primary for VRRP_Group1, and the IP address '159.9.200.200' will then become active on FAD2.

b. Edit the Virtual Server configurations to associate with their respective customer traffic groups: VS1 to VRRP_Group1 and VS2 to VRRP_Group2.



After saving the HA configuration changes, cluster members will initiate the join or rejoin process within the cluster. Upon saving the configuration on the primary node, it will automatically disseminate the updated configuration to the secondary node through the HA synchronization mechanism.

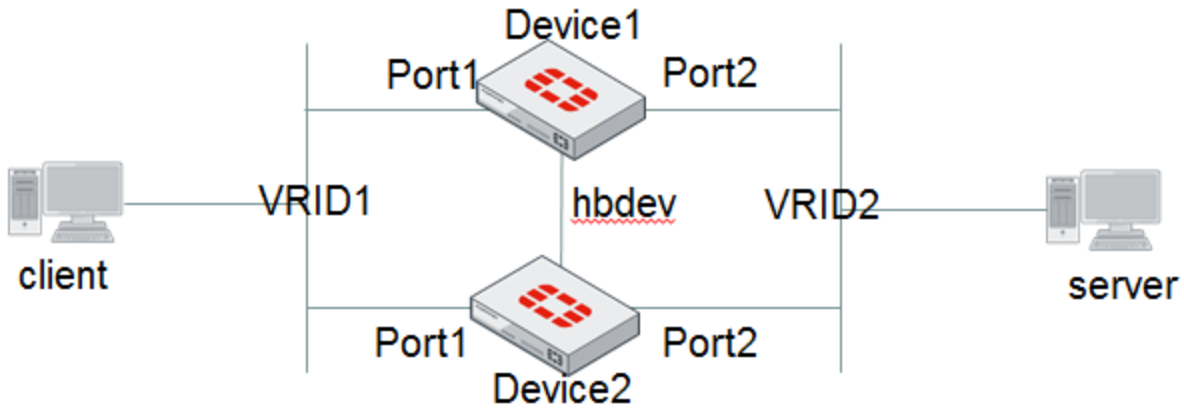
Traffic Group

From the **System > Traffic Group** sub-menu, you can configure traffic groups that are used in HA Active-Active-VRRP.

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 (primary or secondary) operating state and perform a VRRP fail-over in case the primary node fails. The different VRIDs have no relationship with each other.

In Traffic group on page 246, both VRID1 and VRID2 use Device1 as the primary. 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 SNAT pool, etc. to a traffic-group. With this configuration, it will trigger the whole traffic group to switch over when a resource fails.

The traffic group is designed to work with the HA active-active-VRRP mode. Normally, the number of traffic groups should be the same as the number of devices in an HA Active-Active-VRRP mode. In each traffic group, you should configure a different HA node as the primary device. For example, you have HA node A and node B. It's suggested to configure two traffic groups, where traffic group A uses node A as the primary, and node B as the secondary, while traffic group B uses node B as the primary, and node A as the secondary. With this configuration, all the nodes are actively processing traffic, and whichever node fails, its traffic and all related resources such as the floating IP address and virtual server can be taken over by a new primary.

Using traffic group with the HA active-active-VRRP mode can also achieve active-passive HA deployment. FortiADC comes with a predefined traffic group named "default". You can configure the resources such as the floating IP address and virtual server for the default traffic group, then specify the primary node and secondary node in the traffic group, so that when the primary node fails, the resources can be taken over by the new primary.

Please note that traffic group should be associated with a network interface. The floating IP address of the interface can be failed over to the new primary, but the IP address of the interface does not transfer among the HA nodes in this group,

because the interface IP address is not synchronized among HA nodes in active-active-VRRP mode and it always attaches to the physical device who owns the interface.

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 247.
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. For example, if node A was the primary and the traffic was taken over by node B during failover. Once node A comes back, it will again take over the primary role for this traffic group.
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. The number should be the "Local Node ID" in HA configuration.

Administrator

The **System > Administrator** sub-menu includes the following topics:

- Managing administrator users on page 248
- Creating administrator users on page 249
- Creating REST API administrator users on page 251

- Configuring access profiles on page 253
- Enabling password policies on page 255

Managing administrator users

This topic includes the following information:

- Administrator user overview
- REST API administrator user overview
- Create administrator users
- Create REST API 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.

REST API administrator user overview

As a REST API administrator user, you can generate a generalized authorization token to access the supported FortiADC REST APIs of other systems, such as the FNDN.

REST API administrator user accounts share much of the same features of normal administrator accounts; the primary difference is that REST API administrators cannot login through the GUI or CLI and must use their assigned API token to interact with FortiADC. REST API administrators can be created via GUI where they will automatically be assigned an API token. Users can generate access tokens that do not timeout or expire once they are assigned to the REST API administrator. These access tokens cannot be deleted, but can be regenerated. Once regenerated, the previous token will no longer be valid. When a REST API administrator is deleted, any existing associated tokens will be revoked.

REST API administrators can send requests to FortiADC with their API token as a header field.

An example is shown below:

```
curl 'https://XX.XX.XX.XX/api/XXXX/' -H 'APITOKEN: ede4b6632a464a469b85abedd7b5cc91'
```

Basic steps

1. Create REST API administrator user accounts with permissions provisioned by the profiles.
2. Save the automatically generated **API Key** for use later. This is your API authorization token and is only shown once after being generated, so ensure the API key is saved to a secure location.
3. Use the API key to access REST API resources.

Creating 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, LDAP or TACACS+ authentication, you must have already have created the RADIUS server, LDAP server or TACACS+ server configuration.
- You must have Read-Write permission for System settings.

To create an administrator user account:

1. Go to **System > Administrator**.
2. Click the **Admin** tab.
3. Click **Create New > Administrator** to display the configuration editor.
4. Complete the configuration as described in Administrator user configuration on page 249.
5. Click Save.

Administrator user configuration

Settings	Guidelines
Name	Name of the administrator account, such as `admin1` or `admin@example.com`. Do not use spaces or special characters except the 'at' symbol (`@`). The maximum length is 35 characters. If you use LDAP, RADIUS or TACACS+, specify the LDAP, RADIUS or TACACS+ 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, RADIUS or TACACS+ server. After you initially save the configuration, you cannot edit the name.
Global Admin	Select whether or not to allow the administrator account to have Global access, which is required to access all virtual domains, configure Automation Stitches, and create system backup files. • No — This is the default option. The administrator account will only have access to the virtual domain specified in this configuration. Administrators with no Global Admin permission cannot configure Automation Stitches or create system backup files.

Settings	Guidelines
	• Yes — The administrator account will have access to all virtual domains. With Global Admin access, the administrator can configure Automation Stitches and create system backup files.
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 *not* 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. • TACACS+ — Use a TACACS+ authentication server. Select the TACACS+ server configuration.
Password	The **Password** is available if **Authentication Type** is **Local**. 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	The **Confirm Password** is available if **Authentication Type** is **Local**. Re-enter the same password.

Two-factor Authentication	The **Two-factor Authentication** is available if **Authentication Type** is **Local**.
	Options:
	• None
	• FortiToken Cloud
	  • Email address—Set the email address registered with FortiToken Cloud
	  • Country dial code—Set country dial code of mobile phone number
	  • Phone number—Set mobile phone number registered with FortiToken Cloud
	Note: FortiADC does not support FortiToken Cloud functionality in HA condition.

Wildcard	The **Wildcard** option is available if **Authentication Type** is **RADIUS**, **LDAP** or **TACACS+**. Enable the wildcard option to allow multiple remote admin accounts to match one local admin account. This way, multiple RADIUS, LDAP or TACACS+ admin accounts can use one FortiADC admin account.
Restrict to trusted hosts	Enable/disable to restrict logins to trusted hosts only.
Trusted Hosts	The **Trusted Hosts** option is available if **Restrict to trusted hosts** is **enabled**.

Settings	Guidelines

	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 *not* 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 *one* computer, enter only its IP address and 32- or 128-bit netmask:
	`192.0.2.1/32`
	`2001:0db8:85a3::8a2e:0370:7334/128`
	To allow login attempts from any IP address (not recommended), enter:
	`0.0.0.0/0`
	Caution: If you restrict trusted hosts, do so for *all* 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 *one* 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 *after* 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 *only* this administrator will log in.

Creating REST API administrator users

A REST API administrator is required to generate an authorization token prior to sending requests for supported FortiADC REST APIs. You can create a REST API administrator account through the GUI and an authorization token, the API key, will be automatically generated and assigned to the user.



Although users can use an API request to create a REST API administrator account, the resulting token would not be properly assigned to the user. Without an assigned user this authorization token would be invalid and would not be able to access the supported FortiADC REST APIs.

Before you begin:

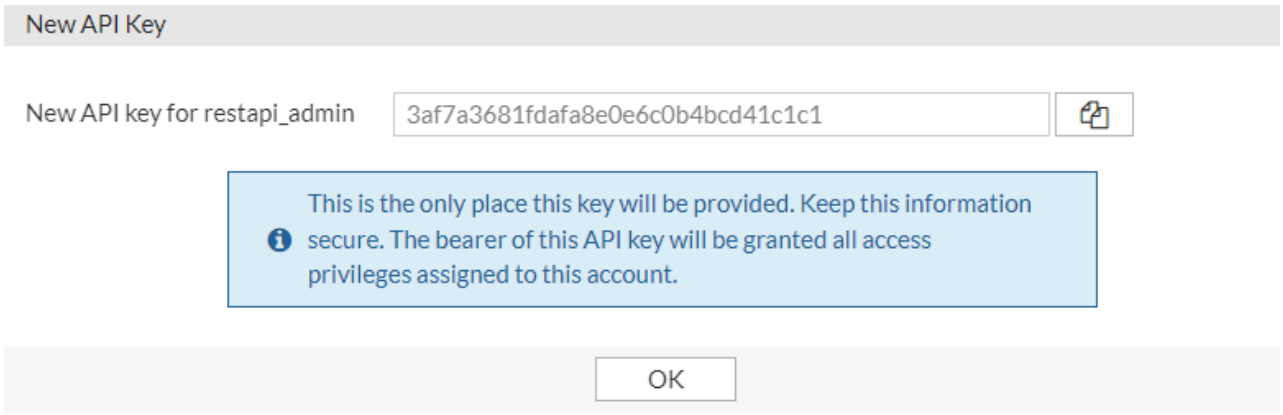
- You must have Read-Write permission for System settings.

To create a REST API administrator in the GUI:

1. Go to **System > Administrator**.
2. Select the **Admin** tab.
3. Click **Create New > REST API Admin** to display the configuration editor.
4. Configure the following settings:

Setting	Description
Name	Enter the login name of the REST API administrator account. The maximum length is 35 characters. Do not use spaces or special characters except the 'at' symbol (@). Using special characters like <, >, (,), #, ", or ' in the administrator account name can result in a cross-site scripting (XSS) vulnerability. After you initially save the configuration, you cannot edit the name.
Comments	Optionally, enter comments about the administrator account.
Global Admin	Select either of the following global admin access options: • No — The account can access the virtual domain specified in this configuration only. This is the default option. • Yes — The account can access all virtual domains.
Administrator profile	The **Administrator profile** option appears if **Global Admin** is **No**. Select a user-defined or predefined profile to use for the new administrator. The predefined profile named **super_admin_prof** is a special access profile used by the **admin** account. However, selecting this access profile will *not* 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.
CORS Allow Origin Toggle	Enable/disable for Cross-Origin Resource Sharing (CORS) for browsers.
CORS Allow Origin	The **CORS Allow Origin** option appears if **CORS Allow Origin Toggle** is enabled. Specify the URL that can access the REST API.
Restricted to trusted hosts	Enable/disable to use Trusted Hosts to allow specific IP addresses to log in to the REST API.
Trusted Hosts	The **Trusted Hosts** option appears if **Restricted to trusted hosts** is enabled. Specify the trusted host IP address and and netmask allowed to log in to the REST API. 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.

5. Click Save.
The **New API Key** pane opens.



The API key is the REST API authorization token that is used in REST API calls.

6. Copy the API key to a secure location.
   This API key will not be displayed anywhere else after you close the pane. If this one is lost or compromised, you can regenerate a new key by editing the REST API Admin user. Once regenerated, the previous token will no longer be valid.
7. Click **OK**.

Configuring 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 253 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

Web UI Menus	CLI Commands
	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 `config` command, there is an equivalent `get`/`show` command. The `config` commands require write permission. The `get`/`show` commands require read permission.	

Before you begin:

- You must have Read-Write permission for System settings.

To configure administrator profiles:

1. Click System > Administrator.
2. Click the **Access Profile** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in Configuring access profiles on page 253.
5. 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.

Settings	Guidelines
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.

Enabling 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:

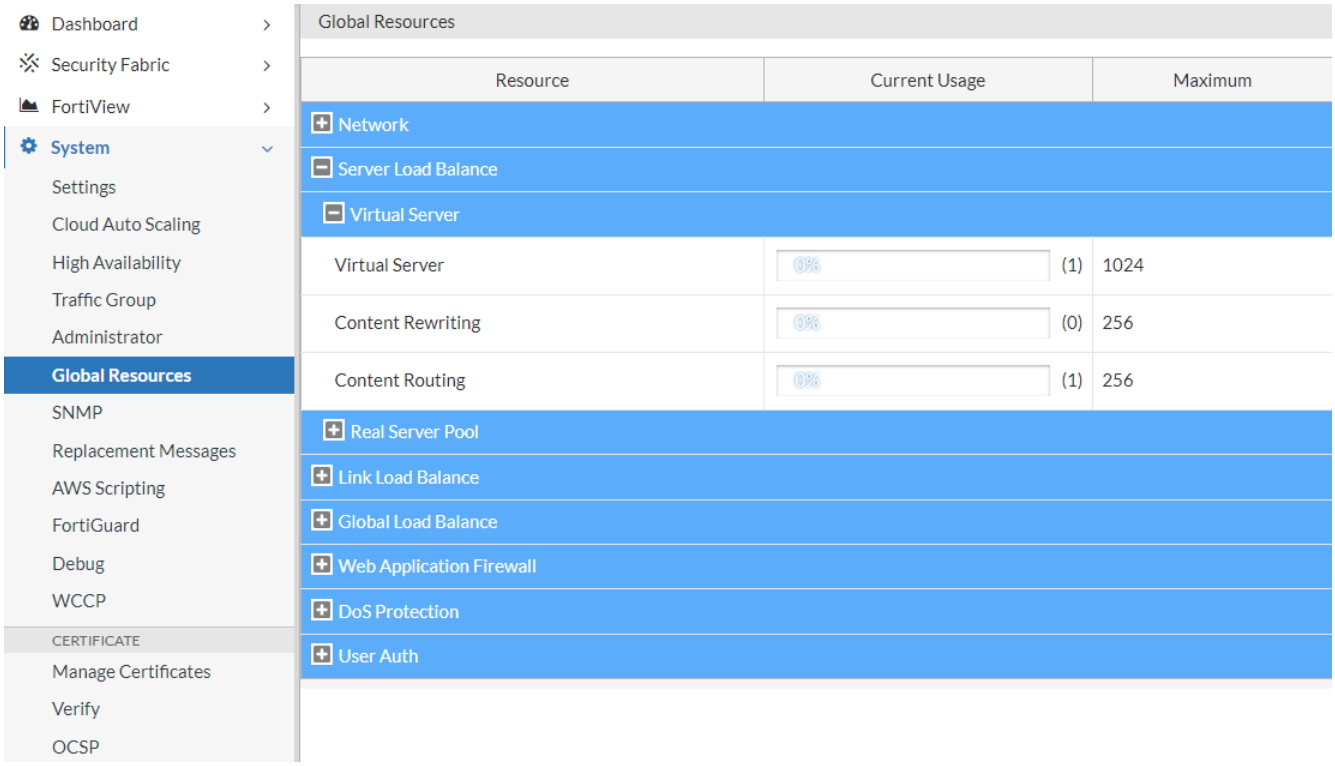
1. Go to **System > Administrator**.
2. Click the **Password Policy** tab.
3. Complete the configuration as described in Password policy configuration.
4. Click **Save**.

Password policy configuration

Settings	Guidelines
Password Policy	Disabled 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.

Global Resources

From the **System > Global Resources** sub-menu, the global system resource usage is displayed, including current and maximum usage per resource.



The Global Resources page is organized by modules, and you can click **+** to view the resource usage data listed under the module. Global Resources is available in both the VDOM settings and Global settings, however, different resource usage data is displayed. The Global Resources usage data displayed from the VDOM settings display the sum of the usage data from each VDOM.



The Global Resources page is not supported when FortiADC is in SSLi mode.

Global Resources from VDOM settings:

Module	Resource usage displayed
Network	
Routing	Static Policy
Server Load Balance	
Virtual Server	Virtual Server Content Rewriting Content Routing
Real Server Pool	Real Server Pool Real Server
Link Load Balance	
Link Group	Link Group Gateway
Global Load Balance	
FQDN	Host Virtual Server Pool
Global Object	Server Link
Web Application Firewall	WVS Task Scan Profile Web Anti-Defacement
DoS Protection	DoS Protection Profile
User Auth	Local User OAuth Proxy
User Group	User Group
Remote Server	NTLM Server
SAML	SAML Service Provider

Global Resources from Global settings:

Module	Resource usage displayed
Network	
Interface	Interface
User Auth	
Remote Server	LDAP Server RADIUS Server TACACS+ Server

WCCP

From the **System > WCCP** sub-menu, you can configure the WCCP (Web Cache Communication Protocol).

WCCP is a Cisco-developed content-routing protocol that provides a mechanism to redirect traffic flows in real-time. FortiADC supports Version 2 (WCCPv2).

With WCCP, 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. Configure the following settings:

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.

4. Click **Save**.
5. Enable WCCP in an interface.
   a. Go to **Networking > Interface**. Select an interface and open the dialogue.
   b. Under Mode Specifics, enable **WCCP** (default is disable).
   c. Click **Save**.
6. Enable WCCP in a virtual server for monitoring.



Only Layer 7 Virtual Servers are supported.

   a. Go to **Server Load Balance > Virtual Server**.
   b. Select a virtual server. Click the **Monitoring** tab.
   c. Enable **WCCP** (default is disable).
   d. Click **Save**.

SNMP

The **System > SNMP** sub-menu includes the following topics:

- Configuring SNMP on page 260
- Downloading SNMP MIBs on page 261
- Configuring SNMP v1/v2 on page 262
- Configuring SNMP v3 on page 263

Configuring SNMP

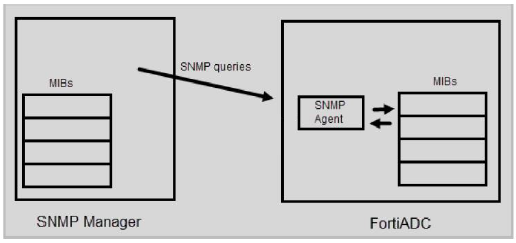
Many organizations use *SNMP* (simple network management protocol) to track the health of their systems. FortiADC supports SNMP v1, v2c, and v3.

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 260 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 Downloading SNMP MIBs on page 261.
- 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 261.
4. Save the configuration.

SNMP settings

Settings	Guidelines
SNMP Agent	Disabled by default. Enable to activate the SNMP agent so that the system can d receive SNMP queries.
Description	A description or comment about the system, such as `dont-reboot`. 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 (`jdoe`). The contact information can be up to 35 characters long, and can contain only letters (a-z, A-Z), numbers, hyphens (-) and underscores (_).
Location	Physical location of the appliance, such as `floor2`. 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 Downloading SNMP MIBs on page 261.

Downloading SNMP MIBs

FortiADC allows you to download full FortiADC and Fortinet Core MIB (Management Information Base) files, which provides more options for server load balance, global server load balance, link load balance, and firewall with SNMP traps.

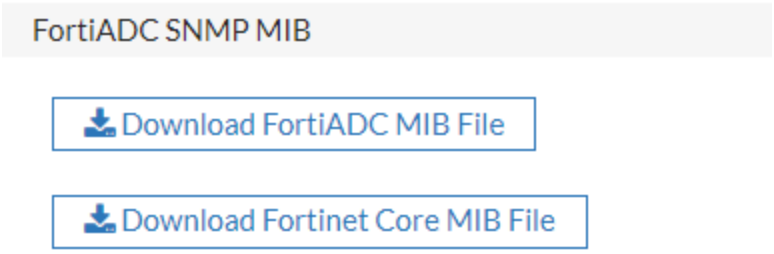


FortiADC 7.4.3 introduces the new SNMP MIB ifXTable with 64-bit counters to allow users to query 10 G statistics.

To download an SNMP MIB file:

1. Go to **System > SNMP**.
   The configuration page displays the **System Information** tab.

2. In the **FortiADC SNMP MIB** section, initiate the file download by clicking the **Download FortiADC MIB File** or **Download Fortinet Core MIB File**.



3. Follow the instructions onscreen to complete the download.

FortiADC MIBs

The following lists the MIBs used with FortiADC.

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).

Configuring 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 262.
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 `management`.

Settings	Guidelines
	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 `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.

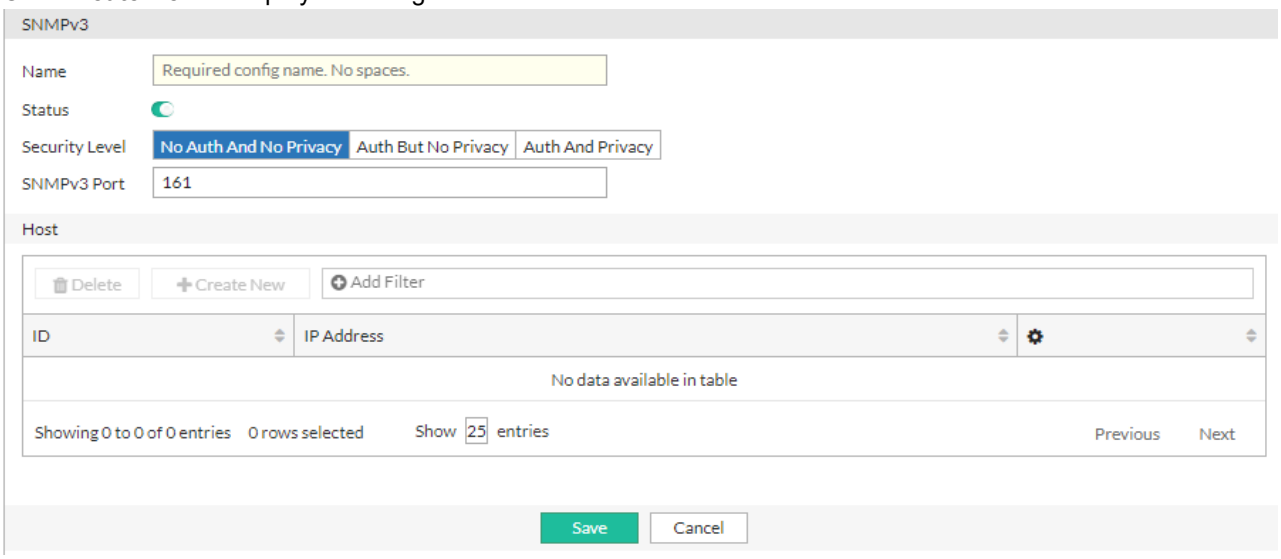
To test traps, cause one of the events that should trigger a trap.

Configuring SNMP v3

To configure SNMP v3:

1. Go to **System > SNMP**.
2. Click the **SNMPv3** tab.

3. Click **Create New** to display the configuration editor.

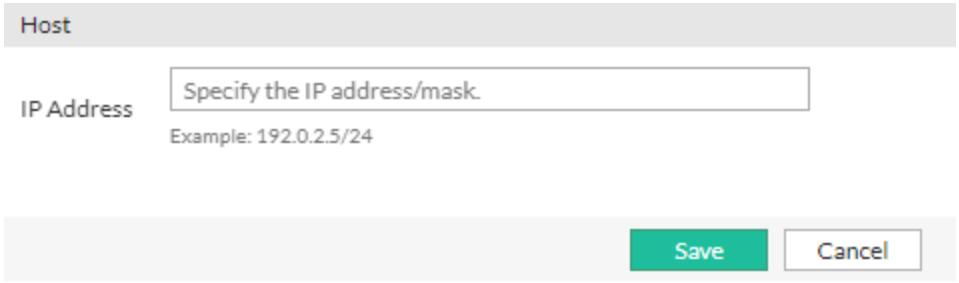


4. Configure the following SNMP v3 settings:

Settings	Guidelines
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.
Auth Algorithm	The **Auth Algorithm** option is available if the **Security Level** is **Auth But No Privacy** or **Auth And Privacy**. Select an Authentication Algorithm from the following options: • SHA1 • MD5 • SHA224 • SHA256 • SHA384 • SHA512
Auth Password	The **Auth Password** option is available if the **Security Level** is **Auth But No Privacy** or **Auth And Privacy**. Specify an Authentication Password.
Private Algorithm	The **Private Algorithm** option is available if the **Security Level** is **Auth And Privacy**. Select a Private Algorithm from the following options:

Settings	Guidelines
	• AES • DES • AES256 • AES256CISCO
Private Password	The **Private Password** option is available if the **Security Level** is **Auth And Privacy**. 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.

5. Click **Save**.
   After the SNMP v3 configuration is saved, the Host section becomes available to configure.
6. Under the **Host** section, click **Create New** to display the configuration editor.



7. Configure the following Host settings:

Settings	Guidelines
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.

8. Save the Host list configuration and then save the SNMP v3 configuration to commit the Host list changes.



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.

Replacement Messages

From the **System > Replacement Messages** sub-menu, you can configure the pre-login disclaimer message that shows before the user reaches the login page.



The Pre-login disclaimer message configuration is not supported when FortiADC is in SSLi mode.

To configure the pre-login disclaimer in the GUI:

1. Go to **System > Replacement Messages**.
2. In the Pre-login Disclaimer Message tab, edit the message.
3. Click **Save**.

Enable the pre-login banner

The pre-login banner needs to be enabled to show the pre-login disclaimer messages.

To enable the pre-login banner in the GUI:

1. Go to **System > Settings**.
2. In the **Basic** tab, enable **Pre Login Banner**.
3. Click **Save**.

To enable the pre-login banner in the CLI:

```
config system global
   set pre-login-banner enable
end
```

FortiGuard

The **System > FortiGuard** sub-menu includes the following topics:

- Connecting to FortiGuard services on page 267
- Configuring FortiGuard service settings on page 270

Connecting to FortiGuard services

After you have subscribed to FortiGuard services, configure your FortiADC to connect to the Internet so that it can reach the world-wide Fortinet Distribution Network (FDN) in order to:

- Verify its FortiGuard service licenses
- Download up-to-date signatures, IP lists, stolen account credentials, and engine packages

FortiADC appliances can often connect using the default settings. However, due to potential differences in routing and firewalls, you should confirm this by verifying connectivity.



You must first register the FortiADC appliance with Fortinet Customer Service & Support (https://support.fortinet.com/) to receive service from the FDN. The FortiADC appliance must also have a valid Fortinet Technical Support contract that includes service subscriptions and be able to connect to the FDN. For port numbers to use to validate the license and update connections, see Appendix B: Port Numbers.

Connecting your FortiADC to the FDN will enable FortiGuard to periodically update 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. However, if you want to perform a manual update, you must download the update file from the FortiGuard website.

To determine your FortiGuard license status

1. If your FortiADC appliance must connect to the Internet through an explicit (non-transparent) web proxy, configure the proxy connection (see Accessing FortiGuard via a proxy on page 269).
   If FortiADC is deployed in a closed network, you can also use FortiManager as a proxy and connect FortiADC with it to validate the license. FortiManager fully supports FortiGuard security service updates and license validation for FortiADC.
2. Go to **Dashboard > Main**.
3. In the **Licenses** widget, check the status icon for each service package.
   Valid — At the last attempt, the FortiADC appliance was able to successfully contact the FDN and validate its FortiGuard license. Continue with Scheduling automatic signature updates.
   Expired — At the last attempt, the license was **either** expired or FortiADC was unable to determine license status due to network connection errors with the FDN. See the following for how to verify the connection status.

   Your FortiADC appliance cannot detect the latest vulnerabilities and compliance violations unless it is licensed and has network connectivity to download current definitions from the FortiGuard service.

   If the connection **did not** succeed:
   - On FortiADC, verify the following settings:
     - time and time zone
     - DNS settings

- network interface up/down status and IP
- static routes

- On your computer, use `nslookup` to verify that FortiGuard domain names are resolving (license authentication queries are sent to `update.fortiguard.net`):

```
C:\Users\cschwartz>nslookup update.fortiguard.net
Server: google-public-dns-a.google.com
Address: 8.8.8.8

Non-authoritative answer:
Name: fds1.fortinet.com
Addresses: 209.66.81.150
209.66.81.151
208.91.112.66
Aliases: update.fortiguard.net
```

- Check the configuration of any NAT or firewall devices that exist between the FortiADC appliance and the FDN or FDS server override. On FortiADC, enter the execute ping and execute traceroute commands to verify that connectivity from FortiADC to the Internet and FortiGuard is possible:

```
FortiADC # exec traceroute update.fortiguard.net
traceroute to update.fortiguard.net (209.66.81.150), 32 hops max, 84 byte packets
1 192.0.2.2 0 ms 0 ms 0 ms
2 209.87.254.221 <static-209-87-254-221.storm.ca> 4 ms 2 ms 3 ms
3 209.87.239.161 <core-2-g0-3.storm.ca> 2 ms 3 ms 3 ms
4 67.69.228.161 3 ms 4 ms 3 ms
5 64.230.164.17 <core2-ottawa23_POS13-1-0.net.bell.ca> 3 ms 5 ms 3 ms
6 64.230.99.250 <tcore4-ottawa23_0-4-2-0.net.bell.ca> 16 ms 17 ms 15 ms
7 64.230.79.222 <tcore3-montreal01_pos0-14-0-0.net.bell.ca> 14 ms 14 ms 15 ms
8 64.230.187.238 <newcore2-newyork83_so6-0-0_0> 63 ms 15 ms 14 ms
9 64.230.187.42 <bxX5-newyork83_POS9-0-0.net.bell.ca> 21 ms 64.230.187.93 <BX5-
     NEWYORK83_POS12-0-0_core.net.bell.ca> 17 ms 16 ms
10 67.69.246.78 <Abovenet_NY.net.bell.ca> 28 ms 28 ms 28 ms
11 64.125.21.86 <xe-1-3-0.cr2.lga5.us.above.net> 29 ms 29 ms 30 ms
12 64.125.27.33 <xe-0-2-0.cr2.ord2.us.above.net> 31 ms 31 ms 33 ms
13 64.125.25.6 <xe-4-1-0.cr2.sjc2.us.above.net> 82 ms 82 ms 100 ms
14 64.125.26.202 <xe-1-1-0.er2.sjc2.us.above.net> 80 ms 79 ms 82 ms
15 209.66.64.93 <209.66.64.93.t01015-01.above.net> 80 ms 80 ms 79 ms
16 209.66.81.150 <209.66.81.150.available.above.net> 83 ms 82 ms 81 ms
```

License validation with FortiManager

If FortiADC is deployed in a closed network, you can validate your FortiADC-VM license through FortiManager because it has the built-in FDS (FortiGuard Distribution Servers) feature. This requires FortiManager to have Internet connection. To configure FortiADC-VM to validate its license using FortiManager, before you upload the license, enter the following command:

```
config system fortiguard
   set override-server-status enable
   set override-server-address <fortimanager_ip>:8890
end
```

where `<fortimanager_ip>` is the IP address of the FortiManager. (TCP port 8890 is the port where the built-in FDS feature listens for requests.)

For more information on the FortiManager built-in FDS feature, see the FortiManager Administration Guide.



FortiManager fully supports FortiGuard security service updates and license validation for FortiADC.

To verify FortiGuard update connectivity

1. If your FortiADC appliance must connect to the Internet (and therefore FDN) through an explicit (non-transparent) web proxy, first you must configure the proxy connection. For details, see Accessing FortiGuard via a proxy.
2. Go to **System > FortiGuard**.
   To access this part of the web UI, your administrator's account access profile must have **Read** and **Write** permission to items in the **System** category.
3. If you want your FortiADC appliance to connect to a specific FDS other than the default for its time zone, enable **Override Server** address and enter the IP address and port number of an FDS in the format `<FDS_ipv4>:<port_int>, such as 10.0.0.1:443`, or enter the domain name of an FDS.
4. Click **Save**.
5. Click **Update FortiGuard service definitions**.
   The FortiADC appliance tests the connection to the FDN and the server you specified to override the default FDN server. The time required varies by the speed of the FortiADC appliance's network connection, and by the number of timeouts that occur before the connection attempt is successful or the FortiADC appliance determines that it cannot connect. If you have enabled logging via **Log Settings > Event**, test results will be indicated in **Event Log**. If the connection test is successful, you would see this log message:

   ```
   VM License validated
   ```

   For more troubleshooting information, enter the following commands:

   ```
   diagnose debug enable
   diagnose debug module updated all
   ```

   These commands display cause additional information in your CLI console. For example:

   ```
   FortiADC # [update]: Poll timeout.
   FortiADC # *ATTENTION*: license registration status changed to 'VALID',please logout and
        re-login
   ```

   For example, poll (license and update request) timeouts can be caused by incorrectly configured static routes and DNS settings, links with high packet loss, and other basic connectivity issues. Unless you override the behavior with a specific FDS address (enable and configure Override Server address), FortiADC will connect to the FDN by communicating with the server closest according to the configured time zone. Timeouts can also be caused by configuring an incorrect time zone.

Accessing FortiGuard via a proxy

You can access FortiGuard via a web proxy server. Using the CLI, you can configure FortiADC to connect through an explicit (non-transparent) web proxy server to the FortiGuard Distribution Network (FDN) for signature updates. FortiADC connects to the proxy using the HTTP `CONNECT` method as described in RFC 2616 (http://tools.ietf.org/rfc/rfc2616.txt).

CLI Syntax

```
config system fortiguard
   set tunneling-status enable
   set tunneling-address 0.0.0.0
   set tunneling-password mypassword
   set tunneling-port 8080
```

```
   set tunneling-username FortADC
end
```

For details, see the FortiADC CLI Reference: https://docs.fortinet.com/product/fortiadc/.

Configuring FortiGuard service settings

FortiGuard periodically updates the WAF Signature Database, IP Reputation Database, and Geo IP Database.

From **System > FortiGuard**, you can configure FortiGuard settings on your FortiADC appliance through the FortiGuard Distribution Network (FDN).

Here, you can configure FortiADC to request for FortiGuard service updates from the FDN by Scheduling automatic signature updates on page 271 and/or Manually initiating update requests on page 273.

Before you begin:

You must have **Read** and **Write** permission for **System** settings.

Licenses

Under the **Licenses** section, you can check your FortiGuard license status and upgrade the license as needed.

Support Contract

Under the **Support Contract** section, you can review the following contract information and directly login to the Fortinet Service & Support website.

Support Type	Description
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 and updates

Under the **FortiGuard Services** section, you can review the list of your FortiGuard service entitlement and the status of each service.

From here, you can also manually update each service by uploading the update packages individually. You can obtain each update package from the FortiGuard website.

Alternatively, you can configure FortiADC to request for FortiGuard service updates from the FDN by doing either or both of the following:

- Scheduling automatic signature updates on page 271
- Manually initiating update requests on page 273

Service	Description
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.
Credential Stuffing Defense	Shows the version of the Credential Stuffing Defense file on your FortiADC unit.
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.
Intrusion Prevention	Shows the version of the Regular IPS Database, Extended IPS Database, and IPS Engine on your FortiADC unit. To manually update the file, click **Update** to display controls that enable you to select and upload the latest Intrusion Prevention file.
Antivirus	Shows the version of the Antivirus Regular Virus Database, Extended Virus Database, Extreme Virus Database, and AV Engine on your FortiADC unit. To manually update the file, click **Update** to display controls that enable you to select and upload the Antivirus files.
Threat Analytics	Shows the license status of the AI Threat Analytics on your FortiADC unit.
Data Loss Prevention (DLP)	Shows the version of the FortiGuard Data Loss Prevention (DLP) service on your FortiADC unit. To manually update the file, click **Update** to display controls that enable you to select and upload the latest DLP Signature file.
Advanced Bot Protection	Shows the license status of the Advanced Bot Protection on your FortiADC unit.

Scheduling automatic signature updates

You can configure the FortiADC appliance to periodically poll for FortiGuard service updates from the FDN, and automatically download and apply updates if they are available. For example, you may want to schedule update requests every night at 2 AM local time when traffic volume is light. You can also use the command `config system fortiguard` to upgrade from the Anycast server. For more information, see `set anycast {enable|disable}` in `config system fortiguard` in the FortiADC CLI Reference (https://docs.fortinet.com/product/fortiadc/).



You can manually upload update packages, or initiate an update request as an alternative or in conjunction with scheduled updates. For additional/alternative update methods, see Manually initiating update requests on page 273.

To configure automatic updates

1. Verify that the FortiADC appliance has a valid license and can connect to the FDN, or (if destination NAT is used, for example) the IP address that you are using to override the default IPs for FDN servers. For details, see Connecting to FortiGuard services on page 267 to determine your FortiGuard license status and to verify the FortiGuard update connectivity.
2. Go to **System > FortiGuard**.
   The page informs you if you are not registered or if registration has expired. If your registration is active, continue scheduling updates; otherwise, click **Register** or **Renew**.
3. Configure the following settings:

Setting	Guideline
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. Note: This feature provides another option for your FortiADCto 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 DNS	Click the button to enable or disable DNS via web proxy tunneling for FDN.
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.

4. Click **Save**.

Results of the update activity appear in **Log & Report > Event log** if you have enabled logging via **Log Settings > Event**.

When the FortiADC appliance requests an update, the event is recorded in **Log & Report > Event log**.

Example log messages include:

```
Update result: fcni=yes fdni=yes fsci=yes IP Reputation(4.00709) Geo IP(2.00094) Regular
     Virus Database(89.00510) Extended Virus Database(88.09720) Extreme Virus Database
     (88.09670) AV Engine(6.00162) from 173.243.140.6:443
```

Once the attack signature update is complete, FortiADC immediately begins to use them. No reboot is required.

Manually initiating update requests

If an important update has been released but there is too much time remaining until your appliance's next scheduled update poll, you can manually trigger the FortiADC appliance to connect to the FDN or FDS server override to request available updates for its FortiGuard service packages.



You can manually initiate updates as an alternative or in addition to other update methods. For details, see Scheduling automatic signature updates on page 271Scheduling automatic signature updates on page 271

To manually request updates

1. Before manually initiating an update, first verify that the FortiADC appliance has a valid license and can connect to the FDN or override server. For details, see Connecting to FortiGuard services on page 267 to determine your FortiGuard license status and to verify the FortiGuard update connectivity.
2. Go to **System > FortiGuard**.
3. Click **Update FortiGuard Service Definitions**.
   The web UI displays a message similar to the following:
   Update database successful, status refreshed.

Results of the update activity appear in **Log & Report > Event log** if you have enabled logging via **Log Settings > Event**.

When the FortiADC appliance requests an update, the event is recorded in **Log & Report > Event log**.

Example log messages include:

```
Update result: fcni=yes fdni=yes fsci=yes IP Reputation(4.00709) Geo IP(2.00094) Regular
     Virus Database(89.00510) Extended Virus Database(88.09720) Extreme Virus Database
     (88.09670) AV Engine(6.00162) from 173.243.140.6:443
```

Once the attack signature update is complete, FortiADC immediately begins to use them. No reboot is required.

Web Filter

Under the **Web Filter** section, you can configure your FortiGuard web filter settings.

Setting	Guideline
Cache Status	Click the button to enable or disable caching of the categorical lists of websites.

Setting	Guideline
	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.

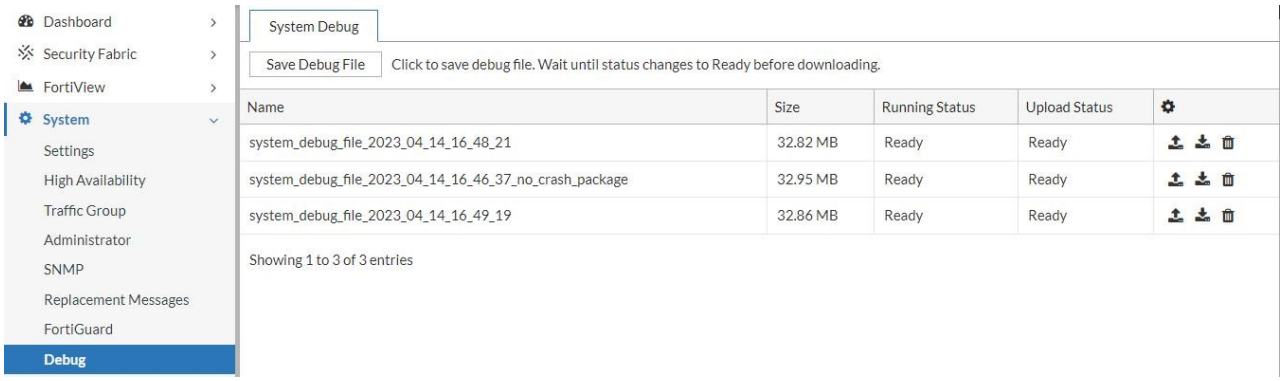
Debug

From the **System > Debug** sub-menu, you can use the Save Debug File feature to create an archive of various configuration files, logs and other details used to help in diagnosing any issues that may arise. The file is saved locally (maximum of 3 files) and can be downloaded to the logged-in admin PC or uploaded to an FTP server.

A key component of the system debug file is the crash log. However, downloading the entirety of the crash log may overburden the browser memory. For this reason, the **Save Debug File** feature offers the option to download a text file containing only the file names from the crash log as an alternative.

To create a system debug file:

1. Go to **System > Debug**.
2. Click **Save Debug File**.
3. Set the **Enable Crash Package** options:
   - Enable — The debug file will contain all core files and crash files. This is the default option.
   - Disable — The debug file will contain all core files and a text file containing only the file names of the crash files.
4. Click **OK**.
   The save debug will run, which may take several seconds. Check the **Running Status** for when the file is ready to download.
5. When the file is ready, as indicated in the **Running Status** column, you can download it via the GUI or upload it to an FTP server.



Column	Description
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. **Ready** — The file is ready to download via the GUI or upload to FTP.
Upload Status	The upload FTP status. **Running** — The File is uploading. Error messages will appear in the Upload Status column. **Ready** — The Upload is completed.

6. Optionally, you may terminate a debug file save request while its **Running Status** is still in **Running**.

Certificate

The **System > Certificate** sub-menu allows you to perform the following tasks relating to certificate management and validation:

- Manage Certificates on page 275
- Verify on page 292
- OCSP on page 297

Manage Certificates

Under the **System Certificate** sub-menu, the **Manage Certificates** page contains the features that allow you to perform the following certificate management tasks:

- Generating a local certificate on page 278
- Importing a local certificate on page 281
- Creating a local certificate group on page 289
- Importing intermediate CAs on page 290
- Creating an intermediate CA group on page 291

Managing and validating 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

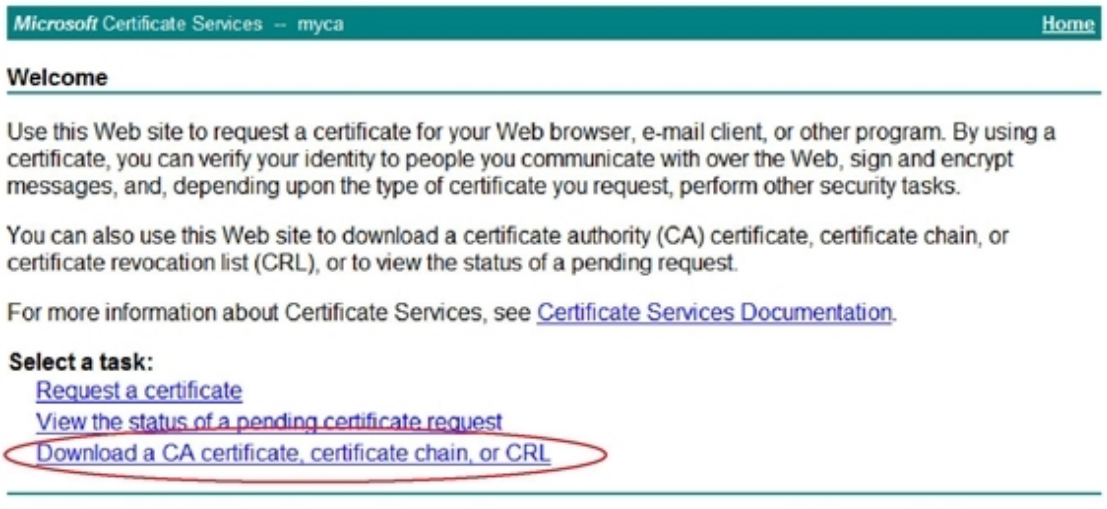
You must download the certificates from your backend servers so that you can import them into the FortiADC system.

This example shows how to download a CA certificate from Microsoft Windows 2003.

To download a CA certificate from Microsoft Windows 2003 Server:

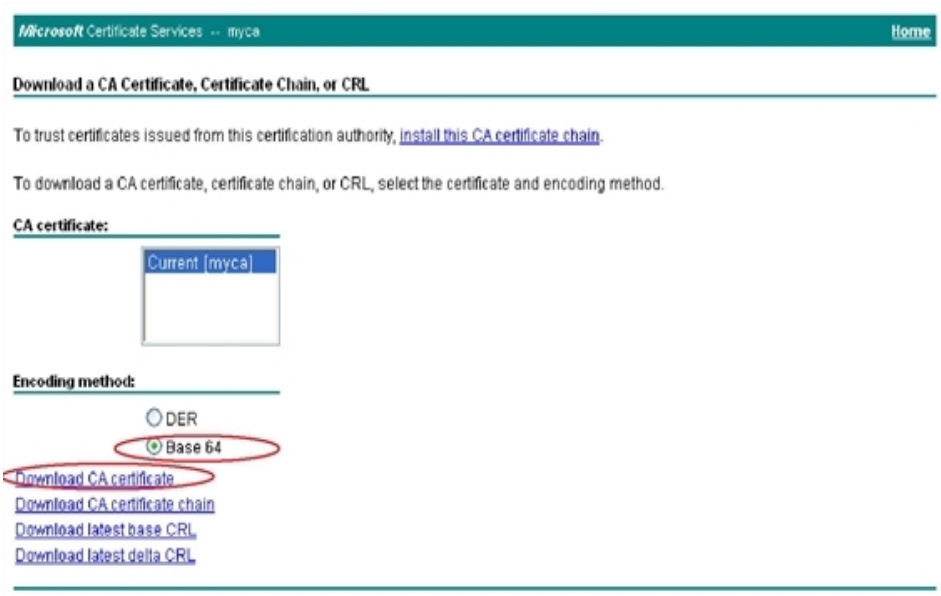
1. Go to `https://<ca-server_ipv4>/certsrv/`.
   where `<ca-server_ipv4>` is the IP address of your CA server.
2. Log in as Administrator. Other accounts may not have sufficient privileges.
   The Microsoft Certificate Services home page appears. Welcome page on page 277 is an example of this page.

Welcome page



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 278 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



Generating a local certificate

In order for FortiADC to authenticate client certificates, you can either generate a certificate signing request (CSR) or upload trusted CA certificates to FortiADC.

Many commercial certificate authorities (CAs) provide websites where you can generate your own 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 > Manage Certificates**.
2. Click the **Local Certificate** tab.

3. Click **Generate** to display the configuration editor.
4. Configure the following settings.

Setting	Description
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 `Subject:` line.
Subject Information	
ID Type	Select the type of identifier to use in the certificate to identify the virtual server: • Host IP—The *static* 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 `www.example.com`. 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 (http://) 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 `10.0.0.1`.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 `www.example.com`. The domain name *must* 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 `Host name mismatch` 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 `admin@example.com`. 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.

Setting	Description
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 `admin@example.com`.
Key Information	
Key Type	Select either of the following: • RSA • ECDSA
Key Size/ Curve Name	For RSA key, select one of the following key sizes: • **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.**

5. Click **Save**.
   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.

Importing a local certificate

You can import (upload) local certificates and their private key files into the FortiADC system to allow FortiADC to authenticate client certificates.

The following types of X.509 server certificates and private keys are supported:

- Base64-encoded
- PKCS #12 RSA-encrypted

After you have downloaded the local certificate and private key files, you can import them into the FortiADC system.

Alternatively, you can select the automated certificate type to use the ACME service to get the SSL/TLS certificates from Let's Encrypt or other ACME providers. Certificates imported through Let's Encrypt have a ninety-day lifetime (which may differ from other ACME providers). These certificates must be renewed prior to expiration. FortiADC supports the TLS-ALPN-01 and DNS-01 challenge types. The TLS-ALPN-01 challenge supports automatic certificate renewal. The DNS-01 challenge requires manual certificate renewal, however, only the DNS-01 challenge can issue certificates containing wildcard domain names.

Before you begin:

- You must have Read-Write permission for System settings.

To import a local certificate through file upload or using the ACME protocol:

1. Go to **System > Manage Certificates**.
2. Click the **Local Certificate** tab.
3. Click **Import** to display the configuration editor.
4. Select the local certificate **Type** from the drop-down menu.
   - **Certificate** — Use this option only if you have a certificate and its key in separate files.
   - **PKCS12 Certificate** — Use this option only if you have a PKCS #12 password-encrypted certificate with its key in the same file.
   - **Local CSR Certificate** — Use this option only if you have a CA-signed certificate that was originated from a CSR generated in FortiADC. See Importing a local certificate on page 281.
     Note: Ensure that the load-balancer (FortiADC appliance) you use to import a local certificate is the same appliance where the CSR was generated as that is where the key matching the certificate resides. The import operation will fail without the matching key on the same hardware system.
   - **Automated** — Use this option if you want to use the ACME protocol to get the certificates from Let's Encrypt or other ACME providers.
5. Configure the following settings based on the local certificate **Type**.

Setting	Description
Certificate	
Certificate Name	Specify the certificate name that can be referenced by other parts of the configuration, such as `www_example_com`. 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.
Input Type	Select either of the following:

Setting	Description
	• Upload • Manual Input
Certificate File	The **Certificate File** option appears if the **Input Type** is **Upload**. Browse for and upload the certificate file that you want to use.
Key File	The **Key File** option appears if the **Input Type** is **Upload**. Browse for and upload the corresponding key file.
Certificate	The **Certificate File** option appears if the **Input Type** is **Manual**. Paste the contents of the certificate file into the text box.
Key	The **Certificate File** option appears if the **Input Type** is **Manual**. Paste the contents of the key file into the text box.
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.
PKCS12 Certificate	
Certificate Name	Specify the certificate name that can be referenced by other parts of the configuration, such as `www_example_com`. 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.
Local CSR Certificate	
Certificate File	Browse for and upload the certificate file that you want to use.
Automated	
Certificate Name	Specify the certificate name that can be referenced by other parts of the configuration, such as `www_example_com`. The maximum length is 35 characters. Do not use spaces or special characters. **Note**: If the **Challenge Type** is **TLS-ALPN-01**, the **Certificate Name** must match the name of the "placeholder" certificate that is linked to the HTTPS virtual server. For details, see Fulfilling the ACME TLS-ALPN-01 challenge with a placeholder certificate on page 287.
Domain Name	Specify the web server domain(s) to be protected by the certificate. When inputting multiple domains, separate each domain using a comma with no additional spaces. For example: `example1.com,example2.com,example3.com`

Setting	Description
	Note: If the **Challenge Type** is **TLS-ALPN-01**, all domains should be resolved to the IP address set by the VS. For details, see Fulfilling the ACME TLS-ALPN-01 challenge with a placeholder certificate on page 287.
Email	Enter the email address that will receive notifications regarding the status of the certificate. Depending on which ACME service provider you use, you may receive notification for when the certificate request has been approved through the Certificated Services or when the certificate is due to expire.
Key Type	Select either of the following: • RSA • ECDSA **Note**: If the **Challenge Type** is **TLS-ALPN-01**, and you choose to fulfill the TLS-ALPN-01 challenge using a placeholder certificate, the **Key Type** must match the key type of the "placeholder" certificate that is linked to the HTTPS virtual server. For details, see Fulfilling the ACME TLS-ALPN-01 challenge with a placeholder certificate on page 287.
Key Size	The **Key Size** option appears if the **Key Type** is **RSA**. Select one of the following key sizes: • 2048 bit • 3072 bit • 4096 bit
Curve Name	The **Key Size** option appears if the **Key Type** is **ECDSA**. Select one of the following curve names: • prime256v1 • secp384r1 • secp521r1
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.
CA Group	Specify the name of the CA Group. FortiADC will use the CA certificate in the CA Group to verify the certificate sent by the ACME provider. From the drop-down, you may select previously configured CA Group or select **Create New** to create and configure a CA Group directly.
ACME Service	Select either of the following: • Let's Encrypt — use the Let's Encrypt certificate authority (https://letsencrypt.org/) as the ACME provider. • Other — use an ACME provider that is not Let's Encrypt, such as Buypass AS (https://www.buypass.com/).
ACME Server URL	The **ACME Server URL** option appears if the **ACME Service** is **Other**.

Setting	Description
	Specify the URL of the ACME server. The ACME request URL must begin with "`https://`". After you have obtained the ACME certificate from your chosen ACME service provider, you will need to provide the ACME server URL to connect to FortiADC. This will enable FortiADC to act as the ACME client to send the ACME request and receive the ACME certificate/key. **Note**: The ACME server URL is unique to the ACME service provider. Please refer to the documentation from your ACME provider for further information.
Challenge Type	The ACME server requires validation that you control the domain names in the certificate using "challenges" as defined by the ACME standard. FortiADC supports the TLS-ALPN-01 and DNS-01 challenge types. Select either of the following challenge types: • TLS-ALPN-01 — The TLS-ALPN-01 supports automatic certificate renewal. However, this method cannot be used to validate wildcard domains. In FortiADC, to fulfill the TLS-ALPN-01 challenge, the ACME server validates control of the domain name by connecting to the Virtual Server at one of the addresses resolved for the domain name. This is achieved by linking a certificate to an HTTPS virtual server to allow the ACME server resolving domain to point to its IP. Then FortiADC generates a temporary certificate to fulfill the validation. FortiADC offers two options to fulfill this requirement: 1) allow FortiADC to manage this process on the backend by selecting a **Local Certificate Group**, or 2) prepare the placeholder certificate **before** the certificate import configurations (for details, see Fulfilling the ACME TLS-ALPN-01 challenge with a placeholder certificate on page 287). • DNS-01 — The DNS-01 challenge can be used to issue certificates containing wildcard domain names. To use this challenge type, you will need to take steps to fulfill the challenge after completing the certificate import configurations (for details, see Fulfilling the ACME DNS-01 challenge on page 285). Certificates imported using the DNS-01 challenge need to be manually renewed.
Renew Window	The **Renew Window** option appears if the **Challenge Type** is **TLS-ALPN-01**. Specify a renew window (in minutes) to automatically renew the certificate before it expires. (Range: 0-43200 minutes). Setting the renew window to 0 will disable the automatic certificate renewal.
Local Certificate Group	The **Local Certificate Group** option appears if the **Challenge Type** is **TLS-ALPN-01**. Select a local certificate group to allow FortiADC to manage the placeholder certificate generation and removal process required for the TLS-ALPN-01 protocol on the backend. This is optional. Alternatively, you can prepare a placeholder certificate to fulfill the TLS-ALPN-01 challenge instead. For details, see Fulfilling the ACME TLS-ALPN-01 challenge with a placeholder certificate on page 287.
Challenge Wait Time	The **Challenge Wait Time** option appears if the **Challenge Type** is **DNS-01**.

Setting	Description
	Specify the ACME DNS-01 challenge wait time in minutes. (Range: 1-1440 minutes). The ACME DNS-01 challenge wait time refers to the amount of time you will have to fulfill the DNS-01 challenge. A longer challenge wait time is recommended to ensure enough time is allotted to perform the required Public DNS configuration changes and for the changes to take effect. For more information, see Fulfilling the ACME DNS-01 challenge on page 285.

6. Click **Save**.

Fulfilling the ACME DNS-01 challenge

The DNS-01 challenge asks you to prove that you control the DNS for your domain name by putting a specific value in a TXT record under that domain name.

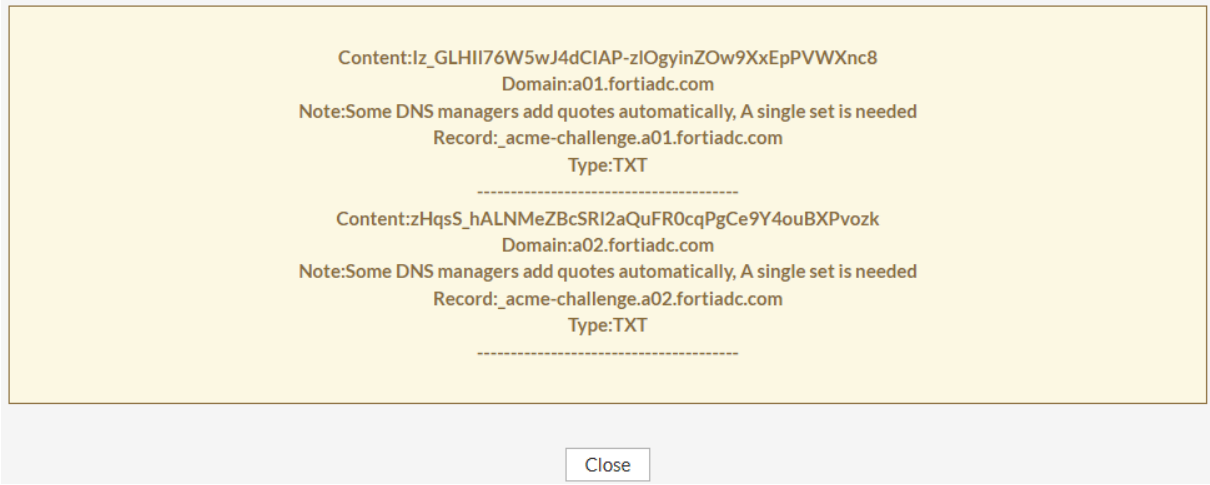
After you have saved your automated local certificate configuration, the ACME DNS challenge information is generated. With this information, you will configure your Public DNS Service to create the TXT record.



Certificates generated by the ACME DNS-01 challenge cannot be renewed automatically. Please manually renew the certificate before it expires.

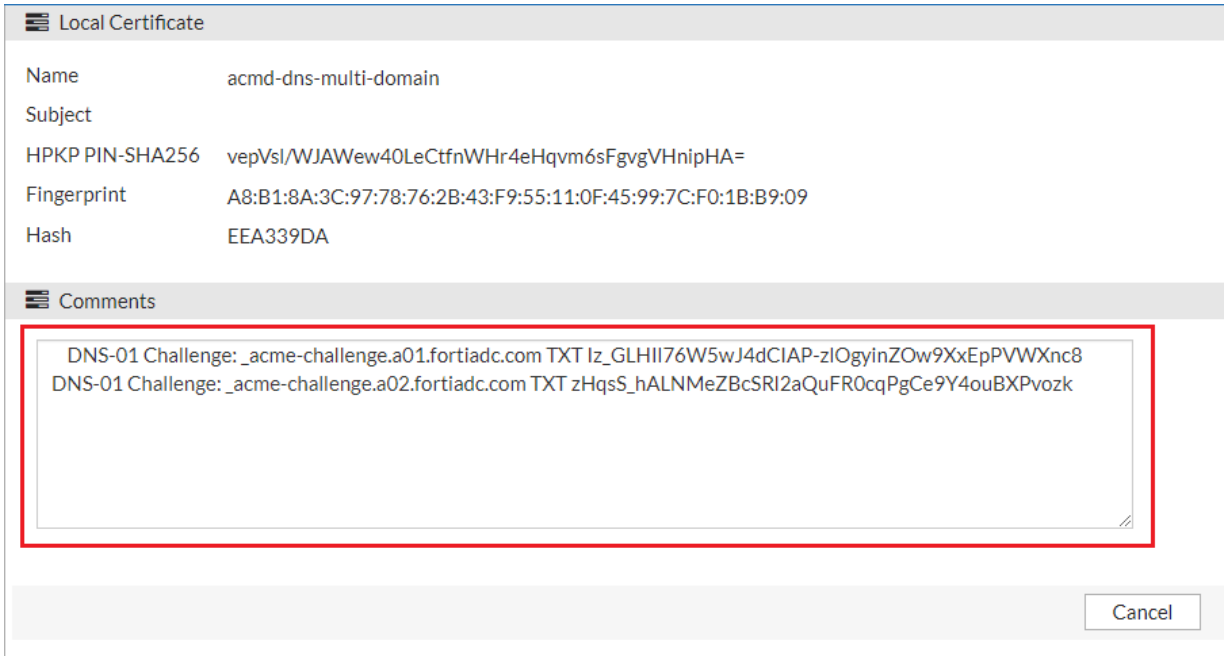
To add the record the DNS challenge information to the Public DNS Service:

1. Obtain the ACME DNS challenge information using either of the following methods.
   - After you save your automated local certificate configuration, you will be shown the challenge information for each domain. Save this information for use later.



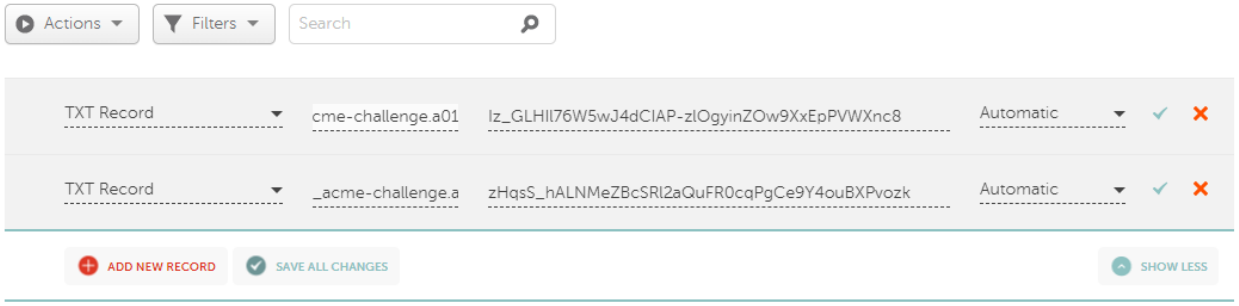
   Note: If using multiple domains, it is recommended you do not exceed 10 domains, as the excessive number of records cannot be displayed in the **Local Certificate** page. Which means that the DNS challenge information will only be available in this pop-up dialog, immediately after saving the certificate configuration.

- In the **Local Certificate** page, locate the local certificate record and click the (View icon) to see the details.



Note: If the DNS challenge information exceeds the 520 character limit, then no records will be displayed in the Comments box. If there are no records in the Comments box, you can view it in the event log.

2. Login to your DNS service provider and go to your DNS Domain management page.
3. Add a record for each domain and input the challenge information into the corresponding fields.



4. Save the changes.
   The DNS configuration changes may take several minutes to take effect.

The ACME provider will then query the DNS system for that record to find a match. If there is a match, the ACME certificate passes validation (certificate status will progress from Pending → OK). However, if the record is not found within the specified challenge wait time then the certificate validation fails (certificate status is Fail).

If the certificate validation fails, then you will need to delete the record and import a new automated local certificate to try again.



It is recommended to set a longer challenge wait time to allow enough time for the DNS configuration changes to take effect. If the DNS configuration changes has not taken effect at the time the ACME provider queries the DNS system for the TXT record, then the validation will fail. Various factors may influence the speed of the DNS (such as the DNS service provider, network speed, network traffic), so the DNS configuration changes may take as long as 20 minutes to take effect.

Fulfilling the ACME TLS-ALPN-01 challenge with a placeholder certificate

The TLS-ALPN-01 challenge is performed entirely at the TLS layer (port 443), and requires the FortiADC virtual server to validate control of the domain name by connecting to a TLS server at one of the addresses resolved for the domain name and verifying that a certificate with specific content is presented.

FortiADC offers two options to fulfill the ACME TLS-ALPN-01 challenge: 1) select a certificate group from the **Local Certificate Group** field to allow FortiADC to manage the placeholder certificate generation and removal process on the backend, or alternatively 2) you can prepare the placeholder certificate prior to configuring the automated certificate using the TLS-ALPN-01 challenge.

In this section, we will cover how to fulfill the ACME TLS-ALPN-01 challenge by using a placeholder certificate.

In FortiADC, to fulfill the TLS-ALPN-01 challenge, the ACME server validates control of the domain name by connecting to the Virtual Server at one of the addresses resolved for the domain name. This is achieved by linking a certificate to an HTTPS virtual server to allow the ACME server resolving domain to point to its IP. Then FortiADC generates a temporary certificate to fulfill the validation.

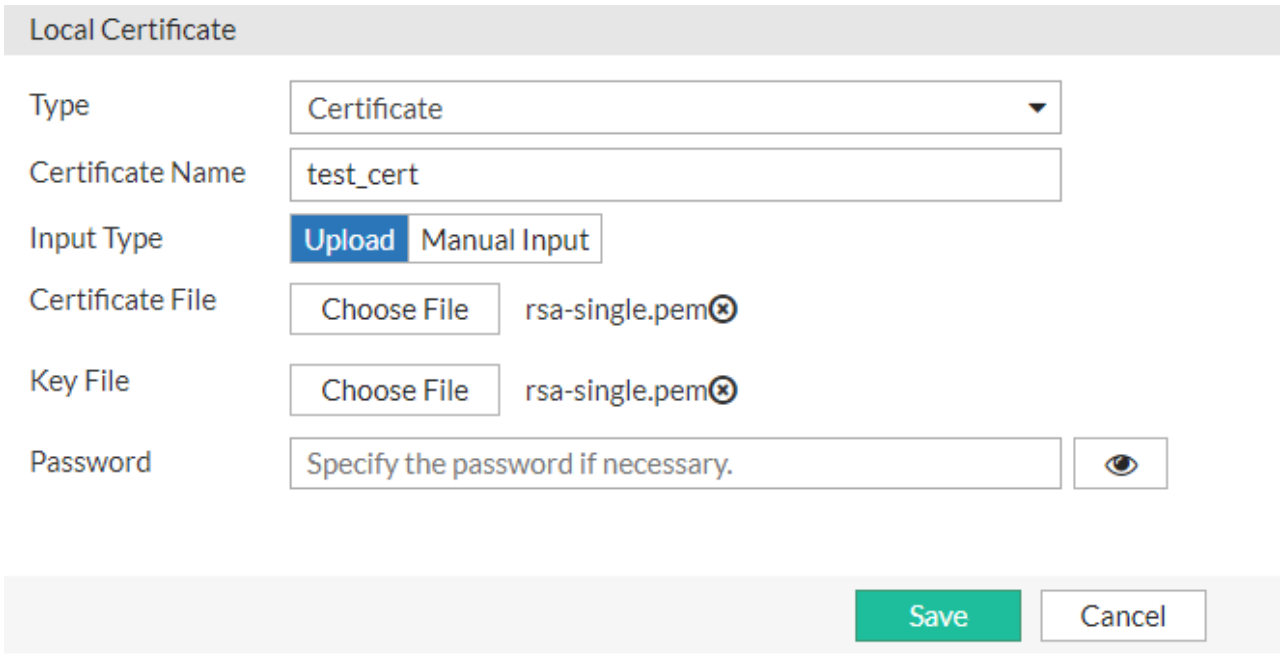
Prepare the following before you configure the automated certificate using the TLS-ALPN-01 challenge:

- A valid local certificate that functions as a placeholder
- An HTTPS virtual server to link the placeholder certificate

Once the placeholder certificate has been linked to the HTTPS virtual server, you will then use the placeholder certificate name and the domain name from the virtual server to import the automated certificate using the TLS-ALPN-01 challenge. This certificate then replaces the placeholder certificate so that it will be linked to the HTTPS virtual server to fulfill the TLS-ALPN-01 challenge.

To prepare the placeholder certificate and HTTPS virtual server for the ACME TLS-ALPN-01 challenge:

1. Generate or import a local certificate. This certificate must be valid (Status is OK). Ensure the Key Type of this placeholder certificate matches the key type of the automated certificate you intend to import. In the example below, the placeholder certificate is RSA, so the automated certificate you will be importing must also be RSA. Record the certificate name for use in later steps. For details, see Importing a local certificate on page 281 or Importing a local certificate on page 281.



Note: If importing a local certificate, you should only import the following certificate types: Certificate, PKCS12 Certificate and Local CSR Certificate. As the placeholder certificate must be valid, it is not recommended to use an Automated certificate type for this purpose since this type of certificate cannot be valid until the ACME challenge is fulfilled.

2. Create a local certificate group and add the placeholder certificate you have created previously under this certificate group. Select the placeholder certificate from the **Local Certificate** drop-down and leave all other parameters as default. Record the certificate group name for use in later steps. For details, see Creating a local certificate group on page 289.
3. Create a Client SSL profile and add the certificate group you have created previously as the **Local Certificate Group**. Record the Client SSL profile name for use in later steps. For details, see Configuring Client SSL profiles on page 512.
4. Create an HTTPS virtual server. Apply the Client SSL profile you have created previously. For details, see Configuring virtual servers on page 384.
   The Address of this HTTPS virtual server must be associated to a domain to ensure it can be reached by the ACME provider. It is recommended that this domain be registered at a DNS service provider so you can set the domain to point to a specific IP address. Record the domain for use in later steps.
5. Import the automated certificate using the TLS-ALPN-01 challenge type.
   Input the information for the following settings according to the guidelines below.

Setting	Guideline
Certificate Name	The name must match the name of the placeholder certificate. Once this automated certificate configuration is completed, it will replace the placeholder certificate.
Domain Name	Input the domain(s) of the HTTPS virtual server that has been linked to the placeholder certificate. The ACME provider will reach the domain(s) that points to the HTTPS virtual server IP address.
Key Type	The Key Type must match the placeholder certificate.

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.
- You must 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 > 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. Under the **Group Member** section, click **Create New** to display the configuration editor.
6. Complete the configuration as described in Local certificate group configuration on page 289.
7. 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 297.
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.

Settings	Guidelines
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 ECSDA certificates, RSA is the only choice. So, having both an RSA certificate and an ECSDA 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.
- You must 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 > 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 291.

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 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:

1. Go to **System > Manage Certificates**.
2. Click the **Intermediate CA Group** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in Intermediate CA group configuration on page 291.
5. 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	

Settings	Guidelines
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,

Verify

Under the **System Certificate** sub-menu, the **Verify** page contains the features that allow you to validate client certificates and real server certificates from within the FortiADC system:

- Validating certificates on page 292
- Importing CRLs on page 294
- Creating a CA group on page 295
- Importing CAs on page 296

Validating certificates

Certificate verification rules specify the CA certificates to use when validating client certificates, and they specify a certificate revocation list (CRL) and/or online certificate status protocol (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.

For a client certificate to be valid, it must meet the following criteria:

- Must not be expired or not yet valid
- Must not be revoked by either CRL or, if enabled, OCSP
- Must be signed by a certificate authority (CA) whose certificate you have imported into the FortiADC appliance

Before you begin:

- You must have Read-Write permission for System settings.
- You must have already created a 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.



For the same certificate object you can configure multiple CRL files.

To configure a certificate verification object:

1. Go to **System > Verify**.
   The configuration page displays the **Verify** tab.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in Certificate verify configuration on page 293.
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 `customize-error-ignore`. The option is disabled by default. If it's enabled, you are required to select the `ca-ignore-errors` and `cert-ignore-errors`, as described below.
`ca-ignore-errors`	Note: CLI only. When `customize-error-ignore` 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 `customize-error-ignore` is disabled (by default), the CLI shows the following: `ca-ignore-errors: UNABLE_TO_GET_ISSUER_CERT UNABLE_TO_GET_CRL CERT_UNTRUSTED`
`cert-ignore-errors`	Note: CLI only. When `customize-error-ignore` is enabled, the following options become available for you to choose from: • `UNABLE_TO_GET_ISSUER_CERT` • `UNABLE_TO_GET_CRL`

Settings	Guidelines
	• 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 `customize-error-ignore` is disabled (by default), the CLI shows the following: `cert-ignore-errors: UNABLE_TO_GET_CRL`
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 certificates 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.
- You must 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 > 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 295.
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.

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.
- You must have already added the CAs to the CA certificate store.

To create a CA group:

1. Go to **System > 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 296.

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.

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 certficate 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.
- You must 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 > 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 297.
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.

OCSP

Under the **System Certificate** sub-menu, the **OCSP** page contains the features that support OCSP (Online Certificate Status Protocol) on FortiADC:

- OCSP stapling on page 297
- Adding OCSPs on page 299
- Importing OCSP signing certificates on page 303

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 Generating a local certificate.

- 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 > OCSP**.
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 298.
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. 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.

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.
- You must know the URL of an OCSP server.
- You must have downloaded the certificate and key files, and have access to them for upload.
- You must have already imported the OCSP signing certificates into FortiADC. See Importing OCSP signing certificates on page 303 and Creating a CA group on page 295.

To add an OCSP verify object:

1. Go to **System > OCSP**.
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 299.
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.

Settings	Guidelines
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: **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.
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.
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.

Settings	Guidelines
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 the Max age check.
Advanced	
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 next update 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 next update field. 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 OCSP server sets the Next Update field in OCSP response.
Caching	Enable or disable OCSP caching. This is enabled by default. **Note**: For a detailed discussion about the function of OCSP caching, see OCSP caching.
Caching This Update 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 This Update Extra Maxage both exist. If one of them is -1, the other one will be used.
Caching Next Update 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 in the OCSP response. 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 Next Update 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. This is disabled by default.
Tunneling Status	Enable or disable tunneling. This is disabled by default. Once enabled, all required settings for the tunneling function becomes configurable.

Settings	Guidelines
	Tunneling, also known as port forwarding, is a method of securely transmitting private data over a public network. It works by encapsulating the private data within public network protocols, ensuring that routing nodes in the public network remain unaware that the transmission belongs to a private network.
Tunneling IP	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.

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 Next Update 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 "Caching This Update Extra Maxage" and the "Caching Next Update Ahead Time" to manipulate the caching behavior.

Configure OCSP caching from the Console

```
config system certificate ocsp
   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.
- You must have the remote certificates downloaded onto you local machine so that you can upload it to FortiADC.

To import an OCSP-signing certificate:

1. Go to **System > OCSP**.
2. Click the **OCSP Signing Certificates** tab.
3. Click **Import** to display the configuration editor.
4. Complete the configuration as described in Importing an OCSP signing certificate on page 303.
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.

Network

The **Network** menu contain features and configurations that allow you to manage your network on FortiADC.

This section is organized into the following sub-menu topics:

- Interface on page 304
- Routing on page 322
- NAT on page 345
- QoS on page 350
- Packet Capture on page 353

Interface

The **Network > Interface** sub-menu allows you to perform the following tasks relating to the network interface:

- Configuring network interfaces on page 307
- Configuring the management interface on page 313
- Linking VDOMs for inter-VDOM routing on page 317
- Configuring virtual overlay networks on page 315
- Transparent mode on page 322

Interface Overview

This section covers the following topics:

- Physical interface
- VLAN interface
- Aggregate interface
- Loopback interface
- Softswitch

Physical interfaces

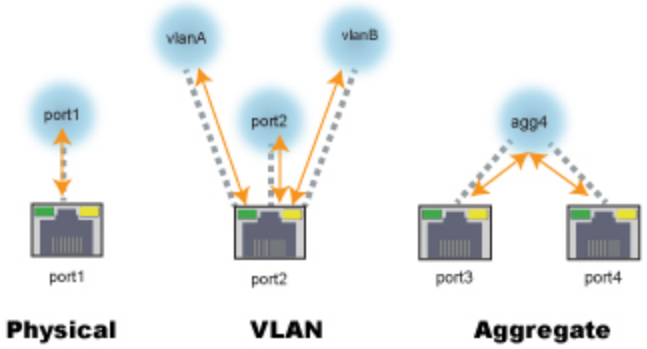
Each physical network port (or vNIC on FortiADC-VM) has a network interface that directly corresponds to it—that is, a “physical network interface.”

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 305 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 305 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
...		
Connectivity layers that will be considered when distributing frames among the aggregated physical ports: •**Layer 2** •**Layer 2-3** •**Layer 3-4**		

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 configure the network interface by editing the configuration, however, 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:

1. Go to **Network > Interface**.
2. 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.
3. Complete the configuration as described in Network interface configuration on page 307.
4. Save the configuration.

Network interface configuration

Settings	Guidelines
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.
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.

Settings	Guidelines
	• 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 the management interface on page 313.
Virtual Domain	If applicable, select the virtual domain to which the configuration applies.
Type	Select from the following: • VLAN • Aggregate • Loopback • Softswitch **Note:** 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: • VLAN—A logical interface you create to VLAN subinterfaces on a single physical interface. • Aggregate—A logical interface you create to support the aggregation of multiple physical interfaces.
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. • DHCP—Use DHCP to automatically assign IP addresses and other communication parameters, including subnet masks, default gateway addresses, and DNS servers to the host.
Traffic Group	Select either of the following: • Default • Create New Available only if **Static** is selected for **Mode**.

Settings	Guidelines
Floating	Enable/Disable floating IP. Available only if **Static** is selected for **Mode**.
Floating IP	Enter the floating IP. Available only if **Floating** is enabled. **Note**: Ensure the Floating IP is different from the Interface IP, otherwise network issues will occur due to the interface/port conflict.
Type Specifics	
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.
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
Softswitch	
Member	Select the interfaces that are included in the softswitch.
Mode Specifics	
Static	

Settings	Guidelines
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.
WCCP	Enable/disable WCCP to redirect traffic flows in real-time.
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.
Trust IP Address	Enable/disable the Trust IPs Access Control (TIAC) feature to restrict access to management interfaces according to the Trust IP Address List. If the source IP is not on the Trust IP Address List, the device will refuse the client directly. To add IP addresses to the Trust IP Address List, enable the feature and save the configuration. After you have saved it the first time, you can edit it to add IPs to the list.
PPPoE	
WCCP	Enable/disable WCCP to redirect traffic flows in real-time.
Trust IP Address	Enable/disable the Trust IPs Access Control (TIAC) feature to restrict access to management interfaces according to the Trust IP Address List. If the source IP is not on the Trust IP Address List, the device will refuse the client directly. To add IP addresses to the Trust IP Address List, enable the feature and save the configuration. After you have saved it the first time, you can edit it to add IPs to the list.
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.
DHCP	
WCCP	Enable/disable WCCP to redirect traffic flows in real-time.

Settings	Guidelines
Trust IP Address	Enable/disable the Trust IPs Access Control (TIAC) feature to restrict access to management interfaces according to the Trust IP Address List. If the source IP is not on the Trust IP Address List, the device will refuse the client directly. To add IP addresses to the Trust IP Address List, enable the feature and save the configuration. After you have saved it the first time, you can edit it to add IPs to the list.
Retrieve Gateway	Use the default gateway retrieved from the DHCP server instead of the one configured in the FortiADC system settings.
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.
Trust IP Address 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.
Type	Select the IP address type: • IPv4/Netmask • IPv4 Address Range • IPv6/Netmask • IPv6 Address Range
IPv4/Netmask, IPv6/Netmask	Specify the IP address that can access the interface.
Address Range	Specify a range of IP addresses that can access the interface.
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 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 in 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 in 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 in the CLI:

```
config system interface
   edit <specified_name>
      set type vlan
      set vlanid <number>
      set interface <port_name>
      set ip <ip&netmask>
   end
```

To enable/disable the Trust IP Address status in the CLI:

```
config system interface
   edit <port_name>
      set trust-ip <enable | disable>
```

To configure the Trust IP Address List in the CLI:

```
config trust-ip-list
   edit <name>
      set type {ip-netmask | ip-range}
      set ip-network <ip&netmask>
      set start-ip <ip>
      set end-ip <ip>
   next
   edit <name>
      set type {ip6-netmask | ip6-range}
      set ip6-network <ip6&netmask>
      set start-ip6 <ip6>
      set end-ip6 <ip6>
   next
end
```

Configuring the 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 secondary 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 Configuring the management interface on page 313 at the end of this section.



- As the management interface is a global configuration, it 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 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 23 of the Getting Started on page 22 section of this Handbook.

To configure the management interface:

1. From FortiADC's global interface, go to **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 314.
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. *Note: It must be promiscuous mode to work.* Promiscuous mode is required because dedicated management interface is a virtual interface and does not share the physical port mac address.
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 Trust IP	Enable/disable the Trust IPs Access Control (TIAC) feature to restrict access to the management interface according to the Trust IP Address List. If the source IP is not on the Management Trust IP Address List, the device will refuse the client directly. To add IP addresses to the Management Trust IP Address List, enable the feature and save the configuration. After you have saved it the first time, you can edit it to add IPs to the list. **Note**: For HA configurations, the Management Trust IP Address list will not be synchronized to peer nodes.
Management Trust IP Address 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.
Type	Select the IP address type: • IPv4/Netmask • IPv4 Address Range
IPv4/Netmask	Specify the IP address that can access the interface.
Address Range	Specify a range of IP addresses that can access the interface.

"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
end
```

Configuring virtual overlay networks

A virtual overlay network uses tunneling protocols to extend isolated network segments between servers for multi-tenant data center networks. Overlay networking can be implemented using various networking protocols and standards. FortiADC supports the Virtual Extensive LAN (VXLAN) and Network Virtualization using Generic Routing Encapsulation (NVGRE) protocols.

VXLAN and NVGRE are officially documented in RFC 7348 and RFC 7637, respectively.

To configure a virtual overlay network:

1. From a virtual domain, go to **Network > Interface**.
2. Click the **Overlay Tunnel** tab.
3. Click **Create New** to display the configuration editor.
4. Configure the following Overlay Tunnel settings:

Setting	Guideline
Name	Specify a unique name for the overlay tunnel. Specify a unique name for the overlay tunnel. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No spaces. After you initially save the configuration, you cannot edit the name.
Mode	Select a virtual overlay networking protocol: • VXLAN • NVGRE VXLAN is the default option.

Setting	Guideline
VXLAN Type	The **VXLAN Type** option is available if the **Mode** is **VXLAN**. Select the VXLAN gateway type: • Linux VXLAN • Flannel VXLAN Linux VXLAN is the default option.
Interface	Specify the outing interface for VXLAN encapsulated traffic.
IP Version	Select an IPv4 address type: • IPv4 Unicast • IPv4 Multicast
Multicast TTL	The **Multicast TTL** option is available if the **IP Version** is **IPv4 Multicast**. Specify the multicast TTL. Valid values are from 0 (default) to 255, minimum value is 1.
Destination IP	Specify the destination IP address. **Note**: For IPv4 unicast, you specify multiple IP addresses; for IPv4 multicast, specify one IP address only.
Port	The **Port** option is available if the **Mode** is **VXLAN**. Specify the VXLAN destination port (number). The default is 4789. The valid range is 1–6553.
VNI	The **VNI** option is available if the **Mode** is **VXLAN**. Specify the VXLAN network ID. The valid range is 1–16777215.
VSID	The **VSID** option is available if the **Mode** is **NVGRE**. Specify the Virtual Subnet Identifier (VSID) for NVGRE.

5. Click **Save**.
   Once the Overlay Tunnel settings are saved, the Remote Host MAC Mapping will be available to configure.
6. In the **Remote Host MAC Mapping** section, click **Create New** to display the configuration editor.
7. Configure the following Remote Host MAC Mapping settings:

Setting	Guideline
Host MAC	Specify the MAC address of the remote host.
VTEP	Specify the IPv4 address of the virtual tunnel endpoint.

8. Click **Save**.
   The dialog closes once the Remote Host MAC Mapping configuration is saved.
9. Click **Save** to update the Overlay Tunnel configuration.

- Virtual overlay networking is supported on all FortiADC hardware platforms.
- For both VXLAN and NVGRE, the mapping can be configured either manually or dynamically (the system learns from the remote VTEP).
- For VXLAN, IPv4 multicast and IPv4 unicast IP addresses can be used as the destination IP of the VTEP.
- For NVGRE, only IPv4 unicast IP addresses can be used as the destination IP of the VTEP.
- Once an overlay tunnel is created, an interface with the same name as the tunnel name will be created automatically. It can be used to encapsulate or decapsulate VXLAN or NVGRE packets.

Linking VDOMs for inter-VDOM routing

VDOM links allow VDOMs to communicate internally without using additional physical interfaces.

Inter-VDOM routing is the communication between VDOMs. VDOM links are virtual interfaces that connect VDOMs. A VDOM link contains a pair of interfaces, each one connected to a VDOM and forming either end of the inter-VDOM connection.

When VDOMs are configured on your FortiADC unit, configuring inter-VDOM routing and VDOM links is like creating a VLAN interface. VDOM links can be managed in either the CLI or in the network interface list in the GUI.

Inter-VDOM routing is only available for these classic scenarios: static route, PBR, L4 SLB, L7 SLB and NAT. It is currently not supported in IPv6 related configurations.

To create and configure a VDOM-link pair in the GUI:

1. Go to **Network > Interface**.
2. Scroll to the **Vdom Link** section.
3. Click **Create New**.

4. Configure the following interface settings for each VDOM in the link:

Virtual Domain	Select the VDOM to link for inter-VDOM routing.
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. **Note:** The IP address cannot be 0.0.0.0/0.
Allow Access	Allow inbound service traffic. Select from the following options: • 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 “ping”). • SSH — Enables SSH connections to the CLI. We recommend this option instead of Telnet. • SNMP — Enables SNMP queries to this network interface. • 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. • 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.
Status	The Status 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.

5. Click **Save**.



By default, VDOM links created in the GUI are ethernet links. The link type can't be changed after it has been created. To set a VDOM link type to point-to-point (ppp), it needs to be created in the CLI.

To create a VDOM-link pair in the CLI:

```
config global
   config system vdom-link
      edit <vdom-link-name>
         set type {ethernet|ppp}
      next
   end
```

Using this command will automatically create a VDOM-link pair in the system interface. However, by default, these VDOM links will not be assigned an IP address or `allowaccess` options, so you would not be able to route traffic between the VDOM links until these settings are configured in the system interface.

To configure the interface settings for the VDOM-link pair in the CLI:

```
config system interface
```

```
   edit <vdom-link-name0>
      set type vdom-link
      set vdom <vdom-name>
      set ip <ip&netmask>
      set allowaccess {http https ping snmp ssh telnet}
   next
   edit <vdom-link-name1>
      set type vdom-link
      set vdom <vdom-name>
      set ip <ip&netmask>
      set allowaccess {http https ping snmp ssh telnet}
   next
end
```

To delete a VDOM link in the GUI:

1. Go to **Network > Interface**
2. Select a **VDOM Link** and click **Delete**.

To delete a VDOM link in the CLI:

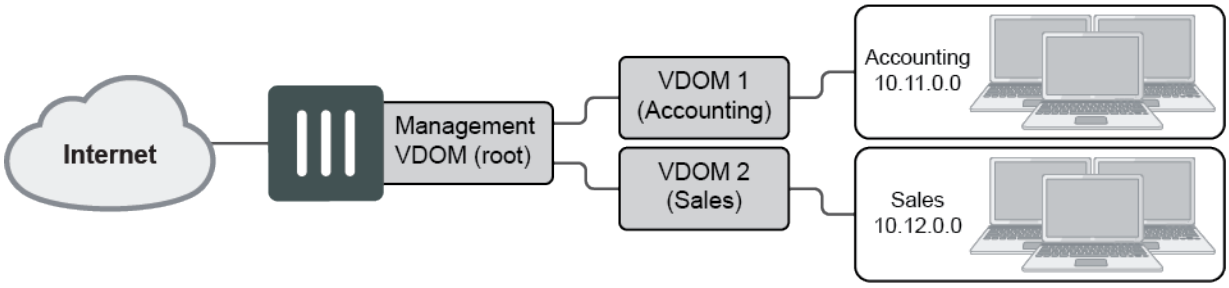
```
config global
   config system vdom-link
      delete <vdom-link-name>
   end
end
```

Test the configuration

When the inter-VDOM routing has been configured, test the configuration to confirm proper operation. Testing connectivity ensures that physical networking connections and FortiADC unit interface configurations are properly configured.

The easiest way to test connectivity is to use the `ping` and `traceroute` commands to confirm the connectivity of different routes on the network.

Example



This example shows how to configure a FortiADC unit to use inter-VDOM routing for a static route scenario.

There are two departments of a company, Accounting and Sales, that need to connect to one FortiADC unit, and the company uses a single ISP to connect to the Internet. To achieve this, two separate pairs of VDOM-links will need to be

created and configured to link the traffic between VDOM 1 (Accounting) and the Management VDOM (root), and between VDOM 2 (Sales) and the Management VDOM (root).

This example includes the following general steps. We recommend following the steps in the order below.

1) Create the VDOM-link pairs in the system interface

Create the two sets of VDOM-link pairs in the system interface: VDOM 1 to the Management VDOM (Accounting-root) and VDOM 2 to the Management VDOM (Sales-root).

To create the Accounting to root VDOM-link pair:

```
config global
   config system vdom-link
      edit Accounting-root
         set type ethernet
      next
   end
end
```

This will automatically create the pair of VDOMs named Accounting-root0 and Accounting-root1 in the system interface. These VDOMs are created with default interface settings so they will need to be edited to enable the routing between VDOM 1 (Accounting) and the Management VDOM (root).

To create the Sales to root VDOM-link pair:

```
config global
   config system vdom-link
      edit Sales-root
         set type ethernet
      next
   end
end
```

This will automatically create the pair of VDOMs named Sales-root0 and Sales-root1 in the system interface. These VDOMs are created with default interface settings so they will need to be edited to enable the routing between VDOM 2 (Sales) and the Management VDOM (root).

2) Configure the VDOM-link pairs in the system interface

In the system interface, edit and configure the two sets of VDOM-link pairs created through `config system vdom-link` to enable the inter-VDOM routing between VDOM 1 (Accounting) and the Management VDOM (root), and VDOM 2 (Sales) and the Management VDOM (root).

To configure the interface settings for the Accounting-root VDOM-link pair:

```
config system interface
   edit Accounting-root0
      set type vdom-link
      set vdom Accounting
      set ip 111.111.111.2/24
      set allowaccess https ping ssh
   next
   edit Accounting-root1
      set type vdom-link
```

```
      set vdom root
      set ip 111.111.111.1/24
      set allowaccess https ping ssh
   next
end
```

To configure the interface settings for the Sales-root VDOM-link pair:

```
config system interface
   edit Sales-root0
      set type vdom-link
      set vdom Sales
      set ip 122.122.122.2/24
      set allowaccess https ping ssh
   next
   edit Sales-root1
      set type vdom-link
      set vdom root
      set ip 122.122.122.1/24
      set allowaccess https ping ssh
   next
end
```

3) Configure the static routes

Configure the static route to send the traffic back to the Internal VDOM for each inter-VDOM link. For each VDOM-link pair, specify the destination default route to point to the Internal VDOM network IP address, and specify the gateway IP address of the next-hop router to point to the other end of the inter-VDOM link. Then set the other VDOM in the link pair to send back the traffic by setting the gateway IP address to point back to the other VDOM in the link.

Configuring the static route for the Accounting-root VDOM-link pair

For the Accounting-root VDOM-link pair, set the destination for the Management VDOM (root) as the Internal VDOM 1 (Accounting) network, and set the gateway routing IP address to point to VDOM 1 (Accounting). Then set the gateway routing IP address for VDOM 1 (Accounting) to point to the Management VDOM (root).

To route the Management VDOM (root):

```
config router static
   edit 1
      set destination 20.24.2.0/24
      set gateway 111.111.111.2
   next
end
```

To route the VDOM 1 (Accounting):

```
config router static
   edit 2
      set gateway 111.111.111.1
   next
end
```

Configuring the static route for the Sales-root VDOM-link pair

For the Sales-root VDOM-link pair, set the destination for the Management VDOM (root) as the Internal VDOM 2 (Sales) network, and set the gateway routing IP address to point to VDOM 2 (Sales). Then set the gateway routing IP address for VDOM 2 (Sales) to point to the Management VDOM (root).

To route the Management VDOM (root):

```
config router static
   edit 3
      set destination 20.24.3.0/24
      set gateway 122.122.122.2
   next
end
```

To route the VDOM 2 (Sales):

```
config router static
   edit 4
      set gateway 122.122.122.1
   next
end
```

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.

Routing

The **Network > Routing** sub-menu allows you to perform the following tasks relating to network routing:

- Configuring static routes on page 323
- Configuring policy routes on page 324
- Configuring OSPF routes on page 325
- Configuring ISP routes on page 333
- Configuring BGP routes on page 335
- Configuring Bidirectional Forwarding Detection (BFD) on page 339
- Access list vs. prefix list on page 340
- Configuring an IPv4 access list on page 341
- Configuring an IPv6 access list on page 342

- Configuring an IPv4 prefix list on page 343
- Configuring an IPv6 prefix list on page 344

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 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 may 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 **Network > 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 323.
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,

Settings	Guidelines
	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, does not need to 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 may 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 will look 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 **Network > 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 325.
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.

Configuring OSPF routes

OSPF (Open Shortest Path First) is described in RFC 2328, 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 supporting 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 **Network > 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 326.
5. Save the configuration.

OSPF configuration

Settings	Guidelines
Router ID	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	

Settings	Guidelines
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: • 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 accidently 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.
Type	Area type setting: •
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.

Settings	Guidelines
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.
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 accidently 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.

Settings	Guidelines
Key	A string of up to 16 characters to be hashed with the cryptographic MD5 hash function.

Configuring OSPFv3 routes

FortiADC supports OSPFv3 (Open Shortest Path First version 3), a link-state routing protocol for IPv6 networks that extends OSPFv2's core functionality, including shortest path calculation using the Dijkstra algorithm and maintaining a Link-State Database (LSDB) for network topology management.

This initial implementation enables FortiADC to advertise IPv6 virtual server IP addresses, as well as "connected" and "static" IPv6 routes, improving routing efficiency in IPv6 environments.

Before you begin:

- You must know how OSPFv3 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 OSPFv3:

1. Navigate to **Network > Routing**.
2. Click the **OSPFv3** tab.
3. Configure the Router settings.

Setting	Description
Router ID	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 OSPFv3 domain to the OSPFv3 speaker.
Redistribute Connected	Enable/disable to redistribute connected routes to OSPFv3. Redistributed routes are distributed into OSPFv3 as Type-5 External LSAs into links to areas. This is disabled by default.
Redistribute Static	Enable/disable to redistribute static routes to OSPFv3. Redistributed routes are distributed to OSPFv3 as Type-5 External LSAs into links to areas. This is disabled by default. When enabled, the static routing will be advertised to the OSPFv3 neighbors.

4. Configure the **Area**.
   a. Under the **Area** section, click **Create New** to display the configuration editor.
   The maximum number of supported Area configurations is 1.
   b. Configure the Area ID:



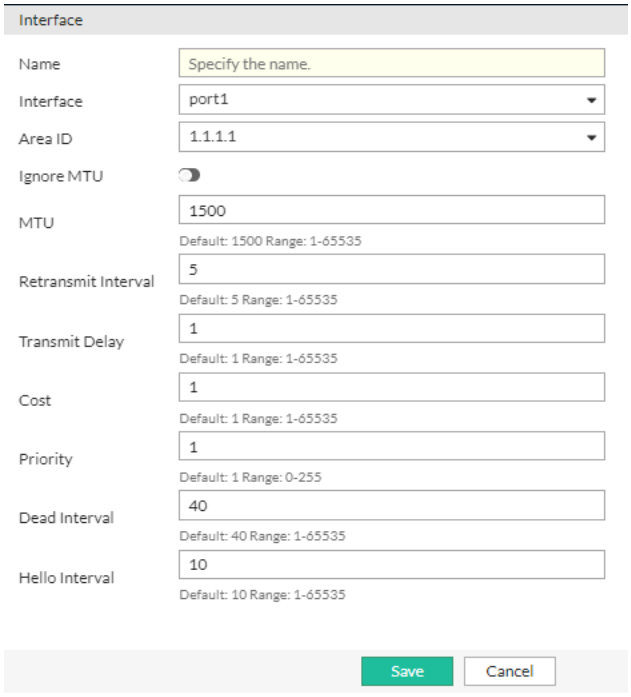
Setting	Description
Area	32-bit number that identifies the OSPFv3 area. An OSPFv3 area is a smaller part of the larger OSPFv3 network. Areas are used to limit the link-state updates that are sent out. The flooding used for these updates would overwhelm a large network, so it is divided into these smaller areas for manageability.

c. Click **Save** to commit the changes and exit the dialog.
The new entry will appear under the Area section.

5. Configure the **Interface**.
The IPv6 virtual server IP address will be advertised to OSPFv3 neighbors by putting the virtual server interface in the OSPFv3 interface.

a. Under the Interface section, click **Create New** to display the configuration editor.
The maximum number of supported Interface configurations is 128.

b. Configure the Interface settings:



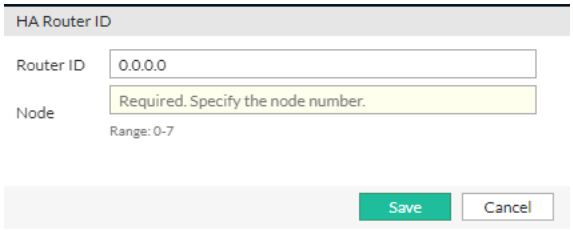
Setting	Description
Name	Configuration name. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No spaces. After you initially save the configuration, you cannot edit the name.

Setting	Description
Interface	Select the interface on which to enable OSPFv3.
Area ID	Select the Area ID.
Ignore MTU	Ignores MTU mismatches between OSPFv3 neighbors. This is disabled by default.
MTU	Specify the maximum packet size for transmission. The default value is 1500, and the valid range is 1-65535.
Retransmit Interval	Specify the interval between LSA retransmissions when acknowledgements are not received. The default value is 5 seconds, and the valid range is 1-65535.
Transmit Delay	Additional time added to the LSA age during transmission. The default value is 1 second, and the valid range is 1-65535.
Cost	Set the link cost for the specified interface, with lower metric values being preferred for path selection. This cost value is applied to the router-LSA's metric field and used in the SPF (Shortest Path First) calculation. The default value is 1, and the valid range is 1-65535.
Priority	The router's priority setting determines its eligibility for DR/BDR (Designated Router/Backup Designated Router) election, with higher values increasing the likelihood of becoming the DR. The router with the highest priority will have a greater chance of being elected as the DR, while setting the priority to 0 makes the router ineligible for DR election. The default value is 1, and the valid range is 0-255.
Dead Interval	The Dead Interval defines the time, in seconds, after which a neighbor is declared down if no Hello packet is received. This interval must be consistent across all routers on a shared network. The default value is 40 seconds, and the valid range is 1-65535.
Hello Interval	The Hello Interval specifies the time between Hello packets sent on a configured interface to maintain neighbor relationships, while the Dead Interval determines the time after which a neighbor is declared down if no Hello packet is received. Both intervals must be consistent across all routers on the same network to ensure proper OSPF operation. The default value is 10 seconds, and the valid range is 1-65535.

c. Click **Save** to commit the changes and exit the dialog.
The new entry will appear under the Interface section.

6. Configure the **HA Router ID**.
   a. Under the HA Router ID section, click **Create New** to display the configuration editor.
   The maximum number of supported HA Router ID configurations is 128.

b. Configure the HA Router ID settings:



Setting	Description
Router ID	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 OSPFv3 Router ID; when it is in HA mode, it uses the HA Router list ID. Specify a unique 32-bit router ID for the OSPFv3 process in dotted decimal notation. This router ID must be an IP address assigned to the router and must remain unique across the entire OSPFv3 domain to ensure proper identification of the OSPFv3 speaker. In an HA Active-Active-VRRP deployment, each Router ID must be configured independently for each node. In an HA Active-Passive deployment, the router ID defaults to 0.0.0.0 in the neighbor node.
Node	Specify the HA Node ID. The valid range is 0-7.

c. Click **Save** to commit the changes and exit the dialog.
The new entry will appear under the HA Router ID section.

7. Click **Save** again to commit the changes to each section of the OSPFv3 configuration.

Repeat the configuration steps for OSPFv3 on the neighbor node. OSPFv3 requires at least two OSPFv3-capable devices to form a neighbor relationship and exchange routing information.

Troubleshooting

The following CLI commands have been introduced to retrieve information about the OSPFv3 configuration.

Command	Guidelines
`get router ospf6`	Use this command to retrieve basic information of the OSPFv3 configuration.

`get router info6 ospf6 database`
`get router info6 ospf6 interface`
`get router info6 ospf6 neighbor`
`get router info6 ospf6 route`
`get router info6 ospf6 status`
	Use these commands to retrieve detailed information of the OSPFv3 configuration.

`diagnose debug module ospf6d`	Use this command to retrieve detailed debug information for the OSPFv3.

Configuring 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 **Network > 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 333.
5. Save the configuration.

ISP Route configuration

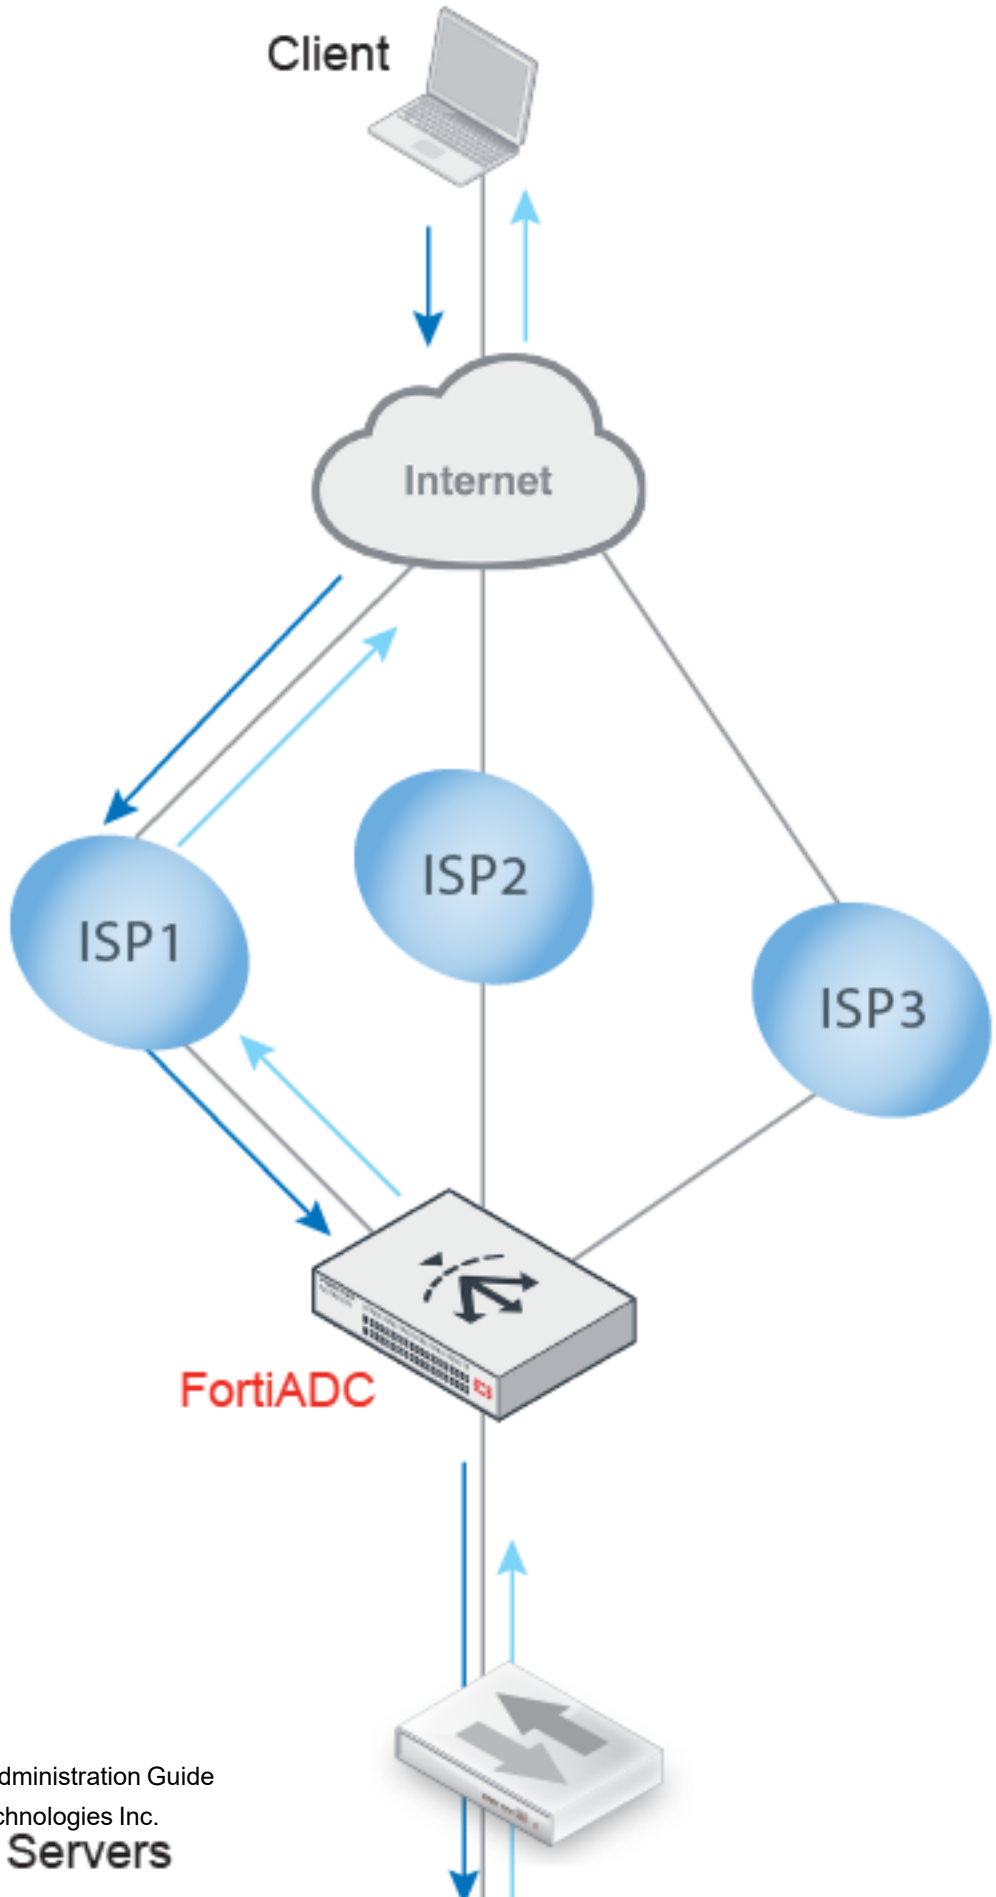
Settings	Guidelines
Destination	Select an ISP address book configuration object. **Note**: Two ISP routes cannot reference the same ISP address book. The ISP routing feature does not support multipath routing.
Gateway	IP address of the gateway router that can route packets to the destination IP address 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 333, 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


Client
Internet
ISP1
ISP2
ISP3
FortiADC
Servers

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
```

Configuring BGP routes

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 touting 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.



FortiADC is designed for BGP node and for BGP route injection (distribute VS public IP to BGP network). It's not recommend to deploy it as a core BGP routing.

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.

Before you begin:

- You must know how BGP has been implemented in your network, i.e., the configuration details of the implementation.
- You must have Read-Write permission for System settings.
- You must have configured all the required access (IPv6) lists and prefix (IPv6) lists. See Access list vs. prefix list.

To configure BGP routes:

1. Click **Network > Routing**.
2. Click the **BGP** tab.
3. Make the desired entries and/or selections as described in BGP configuration on page 337.
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 *keepalive* 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 *keepalive* 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
BFD	Enable to activate Bidirectional Forwarding Detection (BFD) on the BGP session. When BFD detects a path failure, a neighbor Down event is notified immediately to the BGP process, triggering a BGP neighbor status change. For details, see Configuring Bidirectional Forwarding Detection (BFD) on page 339.
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).

Configuring Bidirectional Forwarding Detection (BFD)

You can use Bidirectional Forwarding Detection (BFD) to quickly detect Border Gateway Protocol (BGP) session failures by enabling quicker rerouting of traffic in the event of a link or peer failure. After you configure the BFD object, you can enable BFD in the BGP Neighbor configuration. For more information about the BGP, see Configuring BGP routes on page 335.

The BFD protocol serves as an uncomplicated Hello mechanism designed to identify network failures. Within this framework, Hello packets are transmitted at predetermined, regular intervals. The protocol detects a neighbor failure when a routing device ceases to receive a reply within a specified interval. The failure detection timers employed by BFD possess shorter time limits compared to default failure detection mechanisms for BGP, thereby delivering quicker detection capabilities and a marked decrease in reconvergence time.

After BGP neighbors are configured to perform rapid state detection through BFD, BGP will notify the BFD process. Based on the BGP neighbor's IP address, the BFD process establishes neighbors and starts the negotiation process. The bidirectional detection starts once the BFD neighbor status is UP. The interval between BFD packets is configurable

within 200 milliseconds to 30 seconds. When BFD detects a path failure, a neighbor Down event is notified immediately to the BGP process, triggering a BGP neighbor status change.

Before you begin:

- You must know how BGP has been implemented in your network, i.e., the configuration details of the implementation.
- You must have Read-Write permission for System settings.

To configure BFD:

1. Go to **Network > Routing**.
2. Click the **BFD** tab.
3. Click **Create New** to display the configuration editor.
4. Configure the following Interface settings:

Setting	Description
Name	Specify a unique name for the BFD configuration object. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No space is allowed. **Note**: Once saved, the name of a BFD configuration cannot be changed
Interface	Specify the Interface to assign for BFD. The BFD Interface can refer to the Link Load Balance Ingress Interface specified in the Link Policy. Any Layer 3 interfaces that receive and send external packets can be assigned for BFD except for Loopback interfaces.
Desired Min Transmit Interval	Specify the desired minimum transmit interval for BFD liveness detection in milliseconds. Default: 750ms, Range: 200ms-30000ms. This refers to the interval that the FortiADC would like to use when transmitting BFD Control packets.
Detect Multiplier	Specify the detection time multiplier. Default: 3 Range: 1-20. The negotiated transmit interval, multiplied by this value, provides the Detection Time for the receiving system.
Required Min Receive Interval	Specify the required minimum receive interval for BFD liveness detection in milliseconds. Default: 500ms, Range: 200ms-30000ms. This refers to the minimum interval after which the local routing device must receive a reply from a neighbor with which it has established a BFD session.

5. Click **Save**.
   Once you have saved the BFD object, you can enable BFD in a BGP Neighbor configuration.

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 341 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 wan to include access lists in BGP routing configuration, we highly recommend that you have them configured ahead of time.

Prefix lists

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 IPv4 access list

FortiADC 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 IPv4 access list:

1. Go to **Network > Routing**.
2. Click the **Access List** tab.
3. Click **Create New** to display the configuration editor.

4. Configure the following Access List settings:

Setting	Guideline
Name	Enter a unique name for the new access list. The name can be up to 35 alphanumeric characters long, including . (period), : (colon), _ (underscore), and - (hyphen). No space is allowed.
Description	Optionally, enter a brief description of the access list. The description can be up to 1023 alphanumeric characters long, with no restriction on use of special characters. Space between characters is allowed.

5. Click **Save**.
Once the Access List configuration is saved, the Rule will be available to configure.
6. In the **Rule** section, click **Create New** to display the configuration editor.
7. Configure the following Access List Rule settings:

Setting	Guideline
Action	Select the action as **Permit** or **Deny**.
IPv4 Prefix	Enter the IPv4 address/subnet mask in the format of 0.0.0.0./0.

8. Click **Save**.
The dialog closes once the Rule configuration is saved.
9. Click **Save** to update the Access List configuration.

Configuring an IPv6 access list

FortiADC 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 IPv6 access list:

1. Go to **Network > Routing**.
2. Click the **Access List** tab.
3. Click **Create New** to display the configuration editor.
4. Configure the following Access IPv6 List settings:

Setting	Guideline
Name	Enter a unique name for the new access list. The name can be up to 35 alphanumeric characters long, including . (period), : (colon), _ (underscore), and - (hyphen). No space is allowed.
Description	Optionally, enter a brief description of the access list. The description can be up to 1023 alphanumeric characters long, with no restriction on use of special characters. Space between characters is allowed.

5. Click **Save**.
Once the Access IPv6 List configuration is saved, the Rule will be available to configure.
6. In the **Rule** section, click **Create New** to display the configuration editor.

7. Configure the following Access IPv6 List Rule settings:

Setting	Guideline
Action	Select the action as **Permit** or **Deny**.
IPv6 Prefix	Enter the IPv6 address/subnet mask in the format of 0.0.0.0./0.

8. Click **Save**.
   The dialog closes once the Rule configuration is saved.
9. Click **Save** to update the Access IPv6 List configuration.

Configuring an IPv4 prefix list

FortiADC supports IPv4 prefix lists over BGP routing. If you are configuring BGP routing using IPv4, you must configure prefix lists using the IPv4 protocol.

To configure an IPv4 prefix list:

1. Go to **Network > Routing**.
2. Click the **Prefix List** tab.
3. Click **Create New** to display the configuration editor.
4. Configure the following Prefix List settings:

Setting	Guideline
Name	Enter a unique name for the new access list. The name can be up to 35 alphanumeric characters long, including . (period), : (colon), _ (underscore), and - (hyphen). No space is allowed.
Description	Optionally, enter a brief description of the access list. The description can be up to 1023 alphanumeric characters long, with no restriction on use of special characters. Space between characters is allowed.

5. Click **Save**.
   Once the Prefix List configuration is saved, the Rule will be available to configure.
6. In the **Rule** section, click **Create New** to display the configuration editor.
7. Configure the following Prefix List Rule settings:

Setting	Guideline
Action	Select the action as **Permit** or **Deny**.
IPv4 Prefix	Enter the IPv4 address/subnet mask in the format of 0.0.0.0./0.
GE	Specify the **GE** (greater than and equal to) values. Set to 0 or a number larger than netmask.
LE	Specify the **LE** (less than and equal to) values. Set to 0 or a number no smaller than GE.

8. Click **Save**.

The dialog closes once the Rule configuration is saved.

9. Click **Save** to update the Prefix List configuration.

Configuring an IPv6 prefix list

FortiADC supports IPv6 prefix lists over BGP routing. If you are configuring BGP routing using IPv6, you must configure prefix lists using the IPv6 protocol.

To configure an IPv6 prefix list:

1. Go to **Network > Routing**.
2. Click the **Prefix List** tab.
3. Click **Create New** to display the configuration editor.
4. Configure the following Prefix IPv6 List settings:

Setting	Guideline
Name	Enter a unique name for the new access list. The name can be up to 35 alphanumeric characters long, including . (period) , : (colon), _ (underscore), and - (hyphen). No space is allowed.
Description	Optionally, enter a brief description of the access list. The description can be up to 1023 alphanumeric characters long, with no restriction on use of special characters. Space between characters is allowed.

5. Click **Save**.
   Once the Prefix IPv6 List configuration is saved, the Rule will be available to configure.
6. In the **Rule** section, click **Create New** to display the configuration editor.
7. Configure the following Prefix IPv6 List Rule settings:

Setting	Guideline
Action	Select the action as **Permit** or **Deny**.
IPv6 Prefix	Enter the IPv6 address/subnet mask in the format of 0.0.0.0./0.
GE	Specify the **GE** (greater than and equal to) values. Set to 0 or a number larger than netmask.
LE	Specify the **LE** (less than and equal to) values. Set to 0 or a number no smaller than GE.

8. Click **Save**.
   The dialog closes once the Rule configuration is saved.
9. Click **Save** to update the Prefix List configuration.

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.

Configuring source NAT

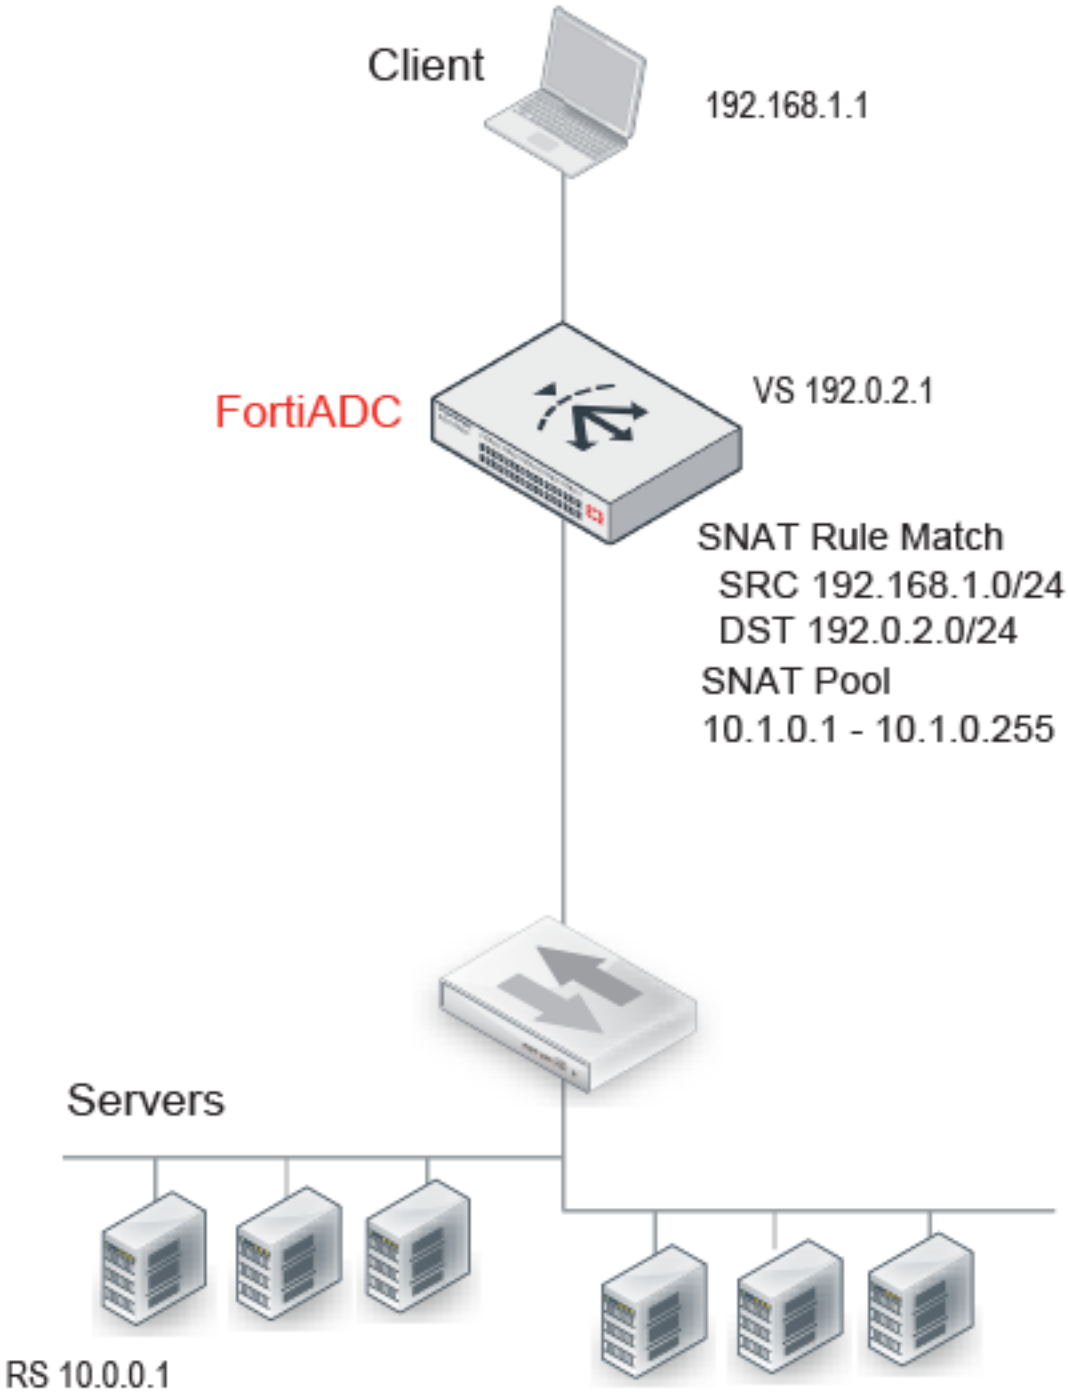
You can use source NAT (SNAT) when clients have IP addresses from private networks. This ensures you do not have multiple sessions from different clients with source IP 192.168.1.1, for example. Or, you can map all client traffic to a single source IP address because a source address from a private network is not meaningful to the FortiADC system or backend servers.

SNAT on page 345 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



SNAT

	Ingress	Egress
SRC	192.168.1.1	10.1.0.1
DST	192.0.2.1	192.0.2.1

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 **Network > 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 347.
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.

Configuring 1-to-1 NAT

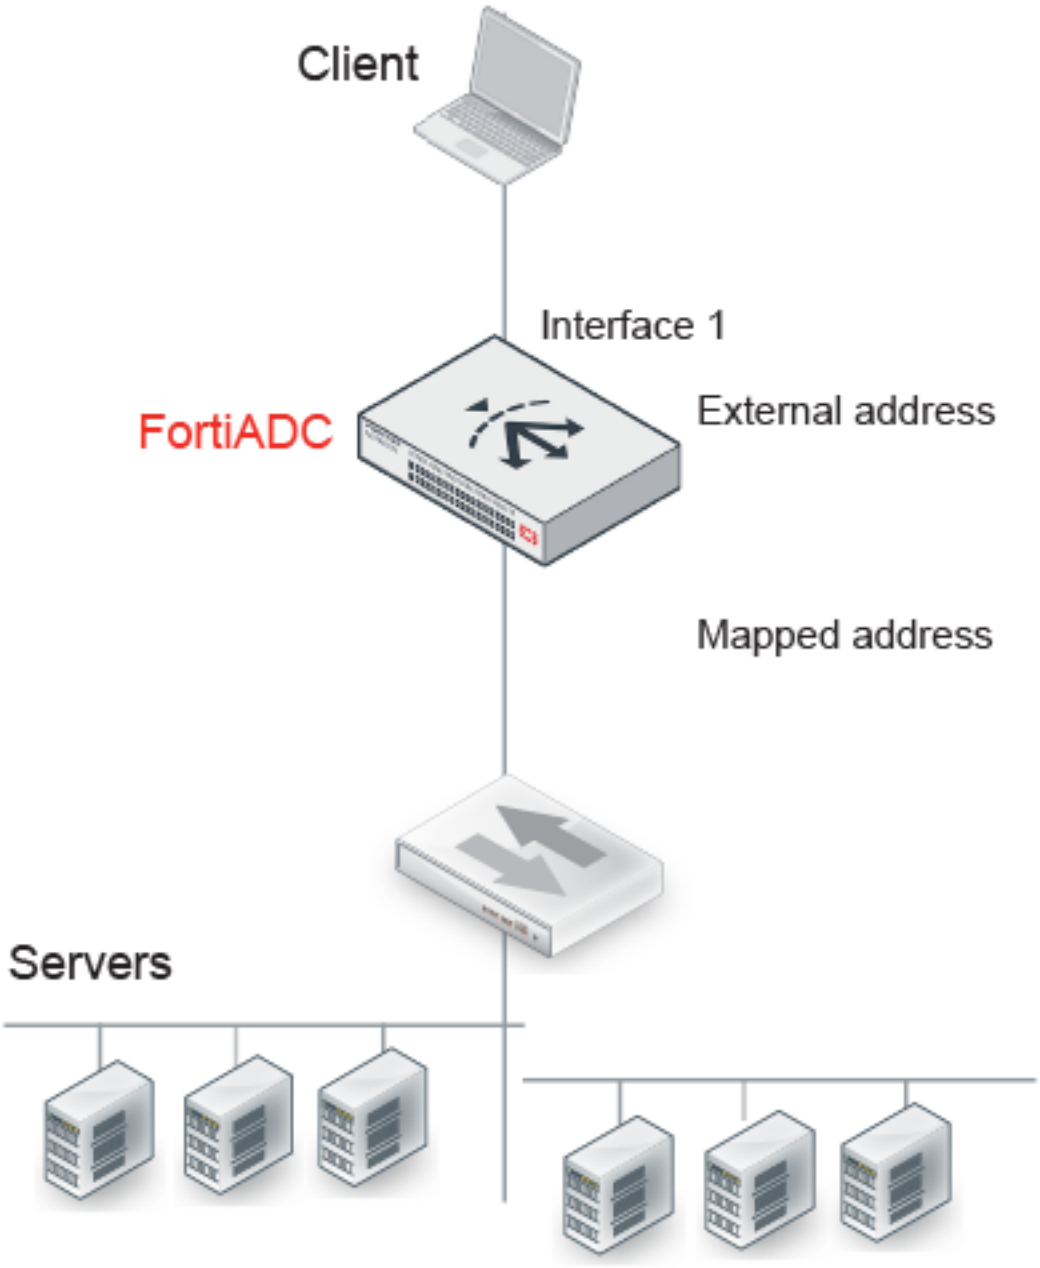
You can use 1-to-1 NAT when you want to publish public or “external” IP addresses for FortiADC resources but want the communication among servers on the internal network to be on a private or “internal” IP address range.

One-to-One NAT on page 348 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



1 to 1 NAT

	IP address	Port
External	192.0.2.0 - 192.0.2.9	80
Mapped	198.51.100.0 - 198.51.100.9	8080

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 **Network > 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 350.
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 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 **Network > QoS**.
   The configuration page displays the **QoS Filter** tab.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in QoS filter configuration on page 351.
4. Save the configuration.
5. 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.

Settings	Guidelines
	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 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 **Network > 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 352.
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 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 **Network > 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 353
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.

Packet Capture

From the **Network > Packet Capture** sub-menu, you can configure the tcpdump utility that 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.

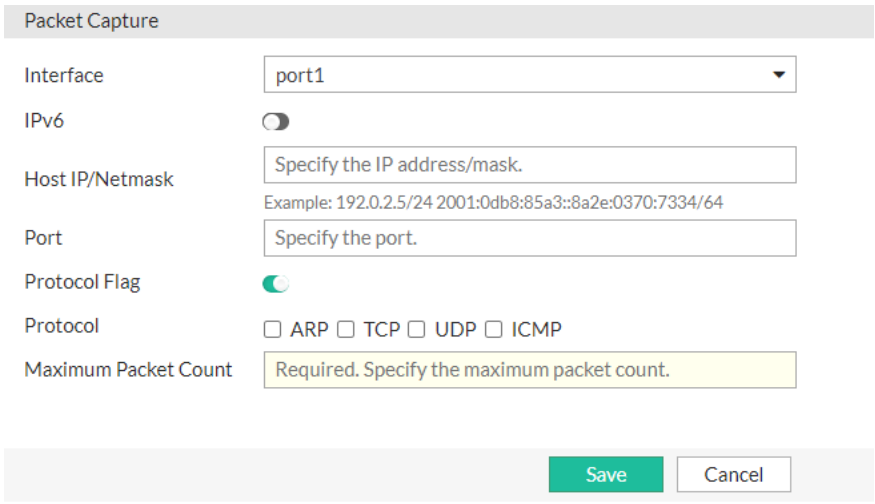
Before you begin:

- 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.

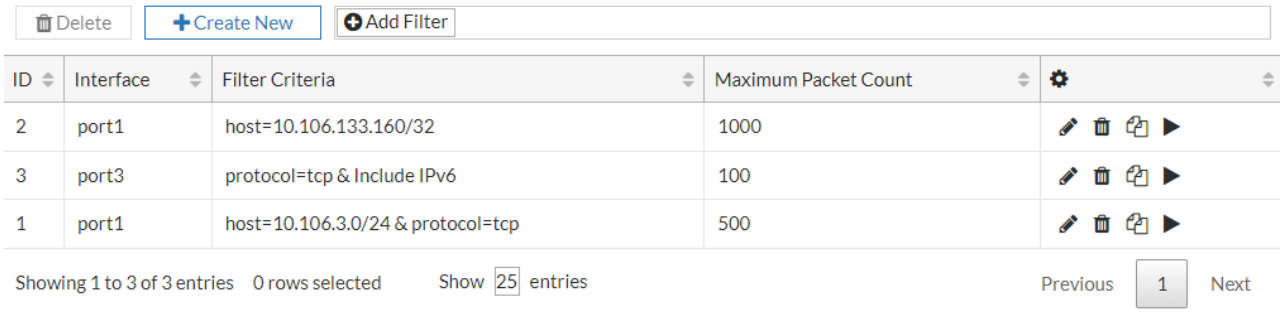
To use the web UI version of tcpdump:

1. Go to **Network > 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 354.

   Packet capture configuration page



Packet capture toolbar



Shared Resources

The **Shared Resources** menu contain features and configurations that allow you to manage resources that are used in various configurations and policies across the FortiADC system.

This section is organized into the following sub-menu topics:

- Health Check on page 355
- Schedule Group on page 368
- Address on page 368
- Service on page 375

Health Check

The **Shared Resources > Health Check** sub-menu includes the following topics:

- Configuring health checks on page 355
- Monitoring health check status on page 365
- Configuring health check scripts on page 366

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 356 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.



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.

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**.
   The configuration page displays the **Health Check** tab.
2. Click **Create New** to display the configuration editor.
3. Enter a name for the Health Check configuration and select the **Type**. The Type option determines what parameters

will need to be configured for that health check type.

4.

Parameter	Guideline
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 one of the following health check types: • ICMP • TCP Echo • TCP • HTTP • HTTPS • DNS • RADIUS • SMTP • POP3 • IMAP4 • RADIUS Accounting • FTP • TCP Half Open Connection • TCP SSL • SNMP • SSH • L2 Detection • UDP • SIP • SIP TCP • SNMP Custom • RTSP • MySQL • Diameter • Script • Oracle • LDAP • MSSQL • LDAPS

5. Configure the **General** settings that apply to all health check types.

Parameter	Guideline
Destination Address Type	• IPv4 • IPv6 • FQDN — destination FQDN type is only supported for LDAP and LDAPS health check types. **Note**:

Parameter	Guideline
	For the LDAP or LDAPS health check types, if Verify Host Certificate is enabled, the destination address type must match the CN in the LDAP/S server certificate as either IP address or FQDN. For example, if the CN in the LDAP/S server certificate is FQDN, then the destination address in the health check configuration must be FQDN as well.
Destination Address	The **Destination Address** option is available if **Destination Address Type** is **IPv4** or **IPv6**. 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.
FQDN	The **FQDN** option is available if **Destination Address Type** is **FQDN**. Specify the destination FQDN (Fully Qualified Domain Name).
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.

6. Configure the **Specifics** settings as required.

Setting	Guidelines
ICMP	
No specific options	Simple ping to test connectivity.
TCP Echo	
No specific options	Simple ping to test connectivity.
TCP / TCP Half Open Connection / UDP	

Setting	Guidelines
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-version	Specify the HTTP version
Additional-string	Attach some string to HTTP header content
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.

Setting	Guidelines
Receive String	A string expected in return when the HTTP GET request is successful.
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.

Setting	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

Setting	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.
SNMP Custom List	
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_UINTEGER
SNMP Counter	Specify the value for the evaluation. The range is 1-2147483647.
SNMP Compare	• Equal • Less • Greater **Less** is the default option.
Name	Specify the SNMP custom name.
Weight	Specify the SNMP custom weight.
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.

Setting	Guidelines
Username	Specify the database user name. (Optional)
Password	Specify the database password, if applicable.
MySQL Server Type	Select either of the following: • Primary (Default) • Secondary
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 filed 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

Setting	Guidelines
	• SID • Connect string
Service name	Use this to specify the service name.
SID	Use this to specify the SID.
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
LDAP	
Port	Port Listening port number of the backend server. Usually LDAP is 389.
Password	The corresponding password.
Attribute	Attributes for the LDAP health check object.
BaseDN	The distinguished name where a LDAP server will search from.
BindDN	The distinguished name used to bind to a LDAP server.
Filter	Criteria to use in selecting results.
MSSQL	
Port	Specify the listening port number of the MSSQL server. Valid values range from 0 to 65535.
Username	Specify the database user name. (Optional)
Password	Specify the database password, if applicable.
Database	Specify the name of the MSSQL database.
MSSQL Send String	Specify the MSSQL send string.
MSSQL Receive String	Specify the MSSQL receive string.
Row	The row in which the send string (command) takes effect.
Column	The column in which the send string (command) takes effect.
LDAPS	

Setting	Guidelines
Port	Port Listening port number of the backend server. The default port is 636 for LDAPS.
Password	The corresponding password.
Bind DN	The distinguished name where a LDAPS server will search from.
Base DN	The distinguished name used to bind to a LDAPS server.
Filter	Criteria to use in selecting results.
Attribute	Attributes for the LDAPS health check object.
Verify Host Certificate	Enable to verify the LDAPS server certificate. This is disabled by default.
CA	The **CA** option is available if **Verify Host Certificate** is **enabled**. Specify the CA certificate.

7. Click **Save**.
   After the Health Check configuration is saved, you can select it in the SLB server pool, LLB link group, or GLB server configuration.



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 servers in real time directly from your desktop, as described below.

Before you begin:

- You must have already created a Health Check configuration.

To configure a health check monitor:

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 366.
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 or FQDN of the remote server.
Health Check	Select the health check configuration.
Direct Route Mode	Enable/disable Direct Route Mode.
Real Server IP	Specify the IP address. Available only if **Direct Route Mode** is enabled.
Count	Specify the count. Range is 1 - 1000 (default = 1).
Interval	Specify the amount of time after the previous health check before this health check executes, in seconds (1 - 3600, default = 5).
Timeout	Specify the timeout period between health checks, in seconds (1 - 3600, default = 3).

Configuring health check scripts

You can upload or input scripting to run health check on FortiADC. When you set a new pool with the health check script, health check will run this script periodically. The health check will then use the result of the script as its result.

The following should be noted when using the Health Check Script feature:

- The supported scripting type is shell script.
- Some commands, such as curl, and ping are supported.
- You should use a new root file-system to avoid the script accessing your current file-system.
- You should limit the quantity of the script and not allow too many scripts to run at the same time.

Before you begin:

- You should have knowledge of some basic concepts of shell scripting, such as variables, if/else statements, and while loop.

To configure a health check script:

1. Go to **Shared Resources > Health Check**.
2. Click the **Health Check Script** tab.
3. Configure the health check script, as described below.

Settings	Guidelines
Name	Enter a unique name for the health check script. After you initially save the configuration, you cannot edit the name.
Input	Type or paste the script.

4. Click **Save**.
5. Go back to the Health Check tab. Create a new health check configuration and select one of the health check scripts you created.

Predefined Health Check Scripts

The following table describes the predefined health check scripts. You can get started with these or create custom scripts.

Predefined script	Usage
CURL_HTTP_CODE	Uses curl to check the status of server
ICMP	Uses ping to check the status of server
PORT_STATUS	Uses nc to check status of a port on server.

Health Check Script programs and commands

The following table describes the predefined commands.

Predefined commands	Usage
$NODE_IP	The IP address of real server.
$NODE_PORT	The port of the real server.
$UP:	Sets the status of health check to UP.
$DOWN:	Sets the status of health check to DOWN.

The following programs and commands can be used in the health check script:

- hostname
- ipcalc
- stty
- usleep
- busybox
- curl
- ed
- nmeter
- xargs
- tail
- basename
- dirname
- expr
- killall
- mkfif
- pwdx
- shred
- split
- traceroute6
- cat
- date
- egrep
- ifconfig
- iplink
- ln
- mktemp
- nameif
- ping6
- rm
- sleep
- true
- base64
- chmod
- dns domain name
- fgrep
- ip
- ls
- mount
- netstat
- printenv
- rmdir
- stat
- umount
- bash
- cp
- echo
- grep
- ipaddr
- kill
- mkdir
- mpstat
- pidof
- ps
- sed
- uniq
- uptime
- link
- mknod
- mv
- ping
- pwd
- slattach
- touch
- vconfig
- adduser
- arping
- fakeidentd
- inetd
- killall5
- nbd-client
- powertop
- pmap
- pstree
- iproute
- telnet
- seq
- sort
- sum
- test
- tr
- truncate
- unix2dos
- wc
- yes
- cmp
- dos2unix
- find
- lsof
- nc
- nslookup
- pgrep
- printf
- time
- traceroute
- tty
- unlink
- whoami
- awk
- diff
- du
- flock
- md5sum
- netcat
- passwd
- pkill
- pscan
- realpath
- smemcap
- ssl_client

Schedule Group

From the **Shared Resources > Schedule Group** sub-menu, you can 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 368.
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.

Address

The **Shared Resources > Address** sub-menu includes the following topics:

- Configuring IPv4 address groups on page 369
- Creating IPv4 address objects on page 370
- Configuring IPv6 address groups on page 370
- Creating IPv6 address objects on page 371
- Managing ISP address books on page 372

Configuring IPv4 address groups

You can configure IPv4 address group objects when you have more than one IPv4 address object you want to specify in rules that match source or destination addresses. For example, if you subscribe user 1 and user 2 to a group of links, then you can create rules that match the user 1 OR user 2 address space and load balance the set of gateways assigned to them.

The following policies use IPv4 address groups:

- Link load balancing policies

Basic Steps

1. Create IPv4 address objects.
2. Configure IPv4 address group objects. You can add up to 256 members in a group.
3. Select the IPv4 address groups when you configure your policies.

Before you begin:

- You must have Read-Write permission for System settings.

To configure an IPv4 address group:

1. Go to **Shared Resources > Address**.
   The configuration page displays the **Address Group** tab.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in Address Group configuration on page 369.
4. 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 IPv4 address objects

You can create IPv4 address objects to specify matching source and destination addresses in policies.

The following policies use IPv4 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 IPv4 address objects.
2. Select them when you configure address groups or policies.

Before you begin:

- You must have Read-Write permission for System settings.

To create an IPv4 address object:

1. Go to **Shared Resources > Address**.
2. Click the **Address** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in Address object configuration on page 370.
5. 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	Specify a subnet using the IPv4 address/mask notation.
Address Range	Specify the start and end of an address range.

Configuring IPv6 address groups

You can configure IPv6 address group objects when you have more than one IPv6 address object you want to specify in rules that match source or destination addresses. For example, if you subscribe user 1 and user 2 to a group of links, then you can create rules that match the user 1 OR user 2 address space and load balance the set of gateways assigned to them.

The following policies use IPv6 address groups:

- Link load balancing policies

Basic Steps

1. Create IPv6 address objects.
2. Configure IPv6 address group objects. You can add up to 256 members in a group.
3. Select the IPv6 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. Go to **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 371.
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 can create IPv6 address objects to specify matching source and destination addresses in policies.

The following policies use IPv6 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 IPv6 address objects.
2. Select them when you configure address groups or policies.

Before you begin:

- You must have Read-Write permission for System settings.

To create an IPv6 address object:

1. Go to **Shared Resources > Address**.
2. Click the **IPv6 Address** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in IPv6 Address object configuration on page 372.
5. 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
IPv6/Netmask	Specify a subnet using the IPv6 address/mask notation.
Address Range	Specify the start and end of an address range.

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 373 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.
- 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.

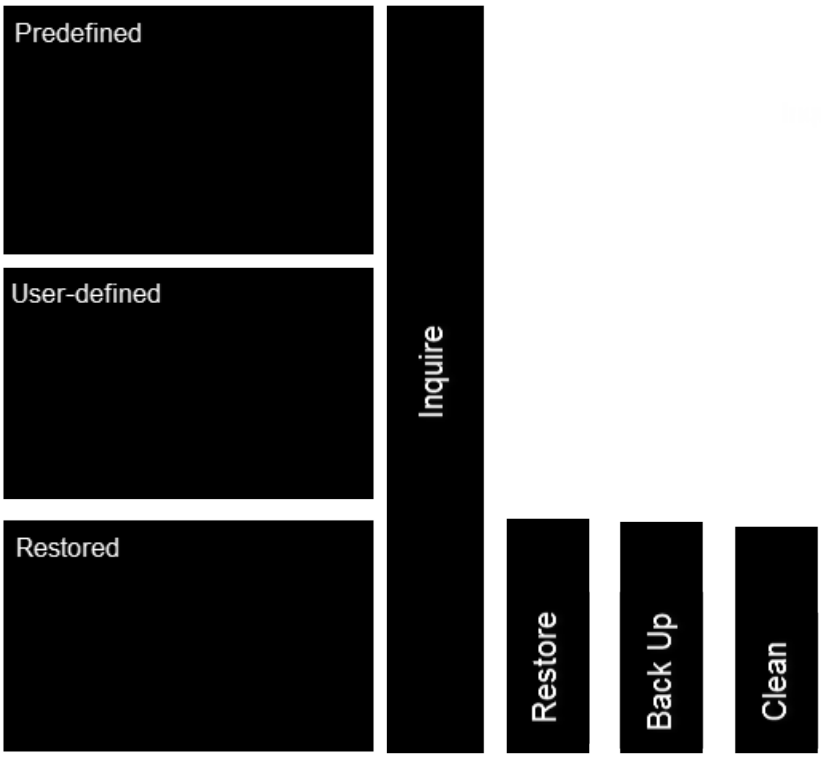
You can also create new user-defined entries for other ISPs.

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. This operation will back up the current restored ISP address books, however, it does not back up the predefined addresses or user-configured entries.

You use the **Clean** utility to erase entries that were imported from the text file. This operation will erase the current restored ISP address books, however, it 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.

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 374.
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.

Settings	Guidelines
Province	Select the associated province location. The configuration supports the following selections: Anhui Beijing Chongqing Fujian Gansu Guangdong Guangxi Guizhou Hainan Hebei Heilongjiang Henan Hubei Hunan Jiangsu Jiangxi Jilin Liaoning Neimenggu Ningxia Qinghai Shandong Shanghai Shanxi (Taiyuan) Shanxi (Xian) Sichuan Tianjin Xianggang Xinjiang Xizang Yunnan Zhejiang Unknown

Service

The **Shared Resources > Service** sub-menu includes the following topics:

- Creating service groups on page 375
- Creating service objects on page 376

Creating service groups

You can configure service group objects when you have more than one service you want to specify in a rule that matches services. 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**.
   The configuration page displays the **Service Group** tab.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in Service Group configuration on page 376.
4. 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.

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.

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 377.
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
Minimum Destination Port	1
Maximum Destination Port	65535

Server Load Balance

As an advanced application delivery controller (ADC), FortiADC supports load balancing for your application servers to enhance application availability and performance. After you have deployed the FortiADC, the traffic can be routed to the FortiADC *virtual server* and distributed among the destination *real servers*. To optimize server response times and server capacity, FortiADC can assign traffic based on various load-balancing algorithms and health checks to ensure application availability.

FortiADC's **Server Load Balance** functionality is supported by the following core modules that work in tandem to load-balance your application servers:

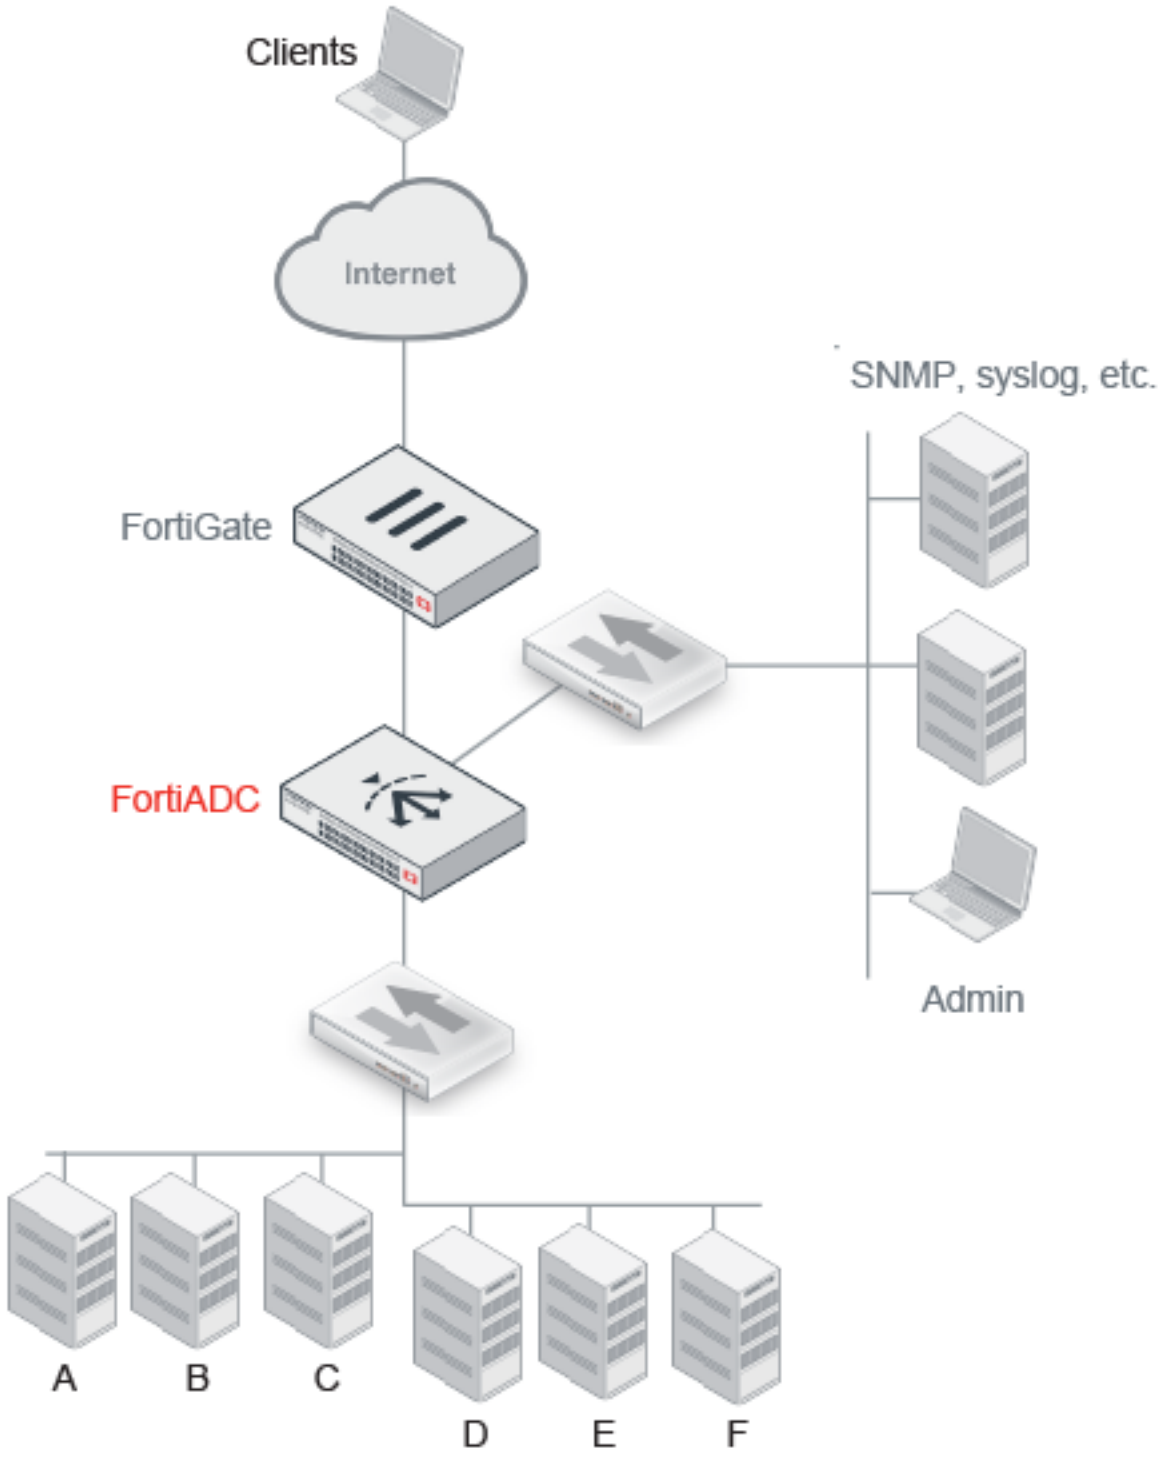
- Virtual Server on page 383 — The Virtual Server module contains configurations to define the application delivery control, wherein three classes are supported: Layer 7, Layer 4, and Layer 2.
- Application Resources on page 450 — The Application Resources module contains configurations to define the protocols and load-balancing methods that will be used to distribute traffic.
- Application Optimization on page 542 — The Application Optimization module contains tools to further optimize application response times and server capacity, such as defining compression and caching rules.
- Real Server Pool on page 551 — The Real Server Pool module contains configurations to define the destination real servers to direct the traffic.

Additional functionality such as HTTP scripting and SSL Forward Proxy is offered for advanced customization of your load-balancing needs.

- Scripting on page 569 — Lua scripting is supported for HTTP/HTTPS applications to perform actions that are not supported by the current built-in feature set.
- SSL-FP Resources on page 594 — FortiADC offers resources to support SSL Forward Proxy functionality, such as SSL decryption exceptions using L2 Exception Lists and Web Filtering, and Certificate Caching to enhance performance.

Basic network topology on page 378 illustrates an example of a basic load balancing deployment. The FortiADC appliance is deployed in front of a server farm, where the network interfaces are connected to three subnets: 1) a subnet for management traffic; 2) a subnet that hosts the real servers A, B, and C; and 3) a different subnet that hosts the real servers D, e, and F. In this topology, the FortiADC system performs health checks on the real servers and distributes traffic to them according to the system logic and user-defined settings.

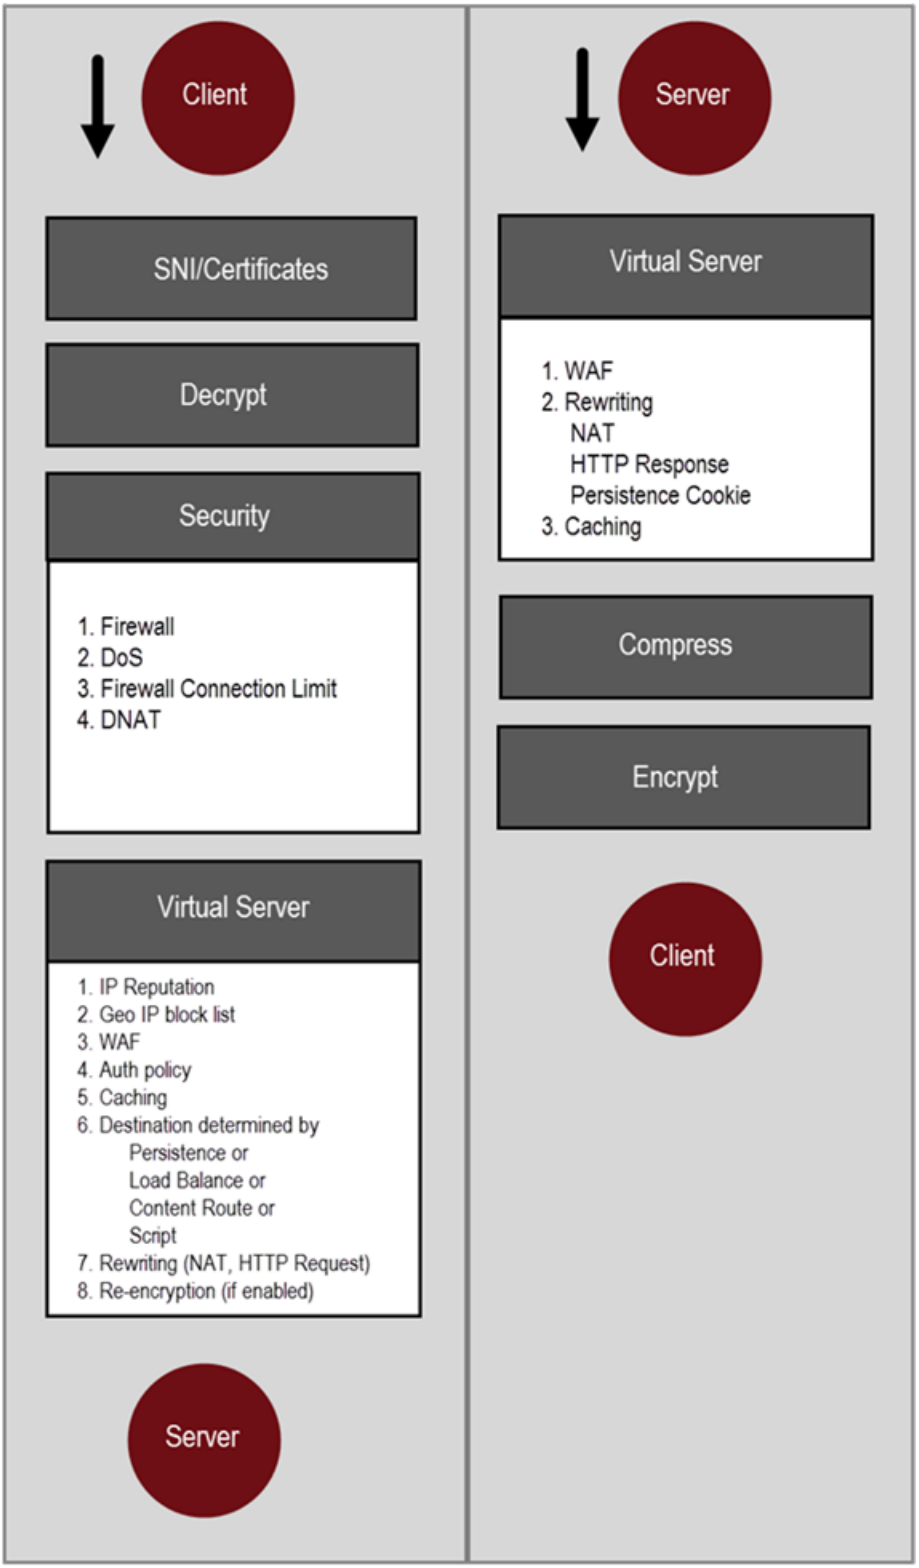
Basic network topology



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.

The FortiADC processing for application using HTTP on page 380 illustrates the order in which the FortiADC features process Client-to-Server and Server-to-Client traffic in an HTTP application server.

FortiADC processing for application using HTTP


Client
SNI/Certificates
Decrypt
Security
1. Firewall
2. DoS
3. Firewall Connection Limit
4. DNAT
Virtual Server
1. IP Reputation
2. Geo IP block list
3. WAF
4. Auth policy
5. Caching
6. Destination determined by
Persistence or
Load Balance or
Content Route or
Script
7. Rewriting (NAT, HTTP Request)
8. Re-encryption (if enabled)
Server
Server
Virtual Server
1. WAF
2. Rewriting
NAT
HTTP Response
Persistence Cookie
3. Caching
Compress
Encrypt
Client

In the Client-to-Server direction:

1. If the SNI or SSL decryption is applicable, the system acts on those exchanges.
2. The security module rules filter the traffic. If the traffic is not dropped, it will continue to the virtual server module.
3. The virtual server security features are applied. If the traffic is not dropped, it will continue for further processing until it is forwarded to the server.
   a. If a caching rule applies, the FortiADC cache serves the content and the request will not be forwarded to a backend server.
   b. If the system selects a destination server based on a persistence rule, content route, or script, then the load-balancing rules will not be applied.
   c. 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:

1. The server traffic reaches the virtual server where the WAF HTTP response, rewriting, persistence, and caching rules are applied.
2. If applicable, the FortiADC compresses and encrypts the server response traffic and forwards the packets to the client.

Server load balancing configuration overview

The FortiADC server load-balancing configuration framework is designed to be modular to support the granularity of the FortiADC application delivery control rules. This allows you to tailor configurations based on traffic type by specifying different options and rules for each type.

FortiADC provides a wide range of configuration options that allow you to customize your server load-balancing to your desired use-case. Many predefined options and configurations are also available to help you get started on load-balancing your application server. Be aware that some configurations and options are minimally required to run server load-balancing.

The steps below describes the basic configuration workflow for server load-balancing:

1. Configure health check rules and real server SSL profiles. For details, see Configuring health checks on page 355 and Configuring real server SSL profiles on page 562.
   This step is optional if you intend to use predefined health check rules and predefined real server SSL profiles. However, if you do intend to use custom health check rules and real server SSL profiles, ensure to configure them before you configure the pools of real servers.
2. Configure real server pools. For details, see Using real server pools on page 552.
   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. These configurations are all contained under the Application Resources module. For details, see Application Resources on page 450.
   This step is optional if you intend to use predefined persistence rules, profiles, and methods.
4. Configure the virtual server. For details, see Configuring virtual servers on page 384.
   This step is **required**. In your virtual server configuration, you define the load-balancing rules by specifying the class of application delivery control (Layer 7, Layer 4, or Layer 2), and the persistence rules, profiles, and load-balancing

methods to use to assign traffic. You can reference predefined or custom defined objects configured in previous steps.

Example workflow

For a members-only HTTPS web server farm, you may 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

Virtual Server

The Virtual Server plays a vital role in the FortiADC Server Load Balance functionality by managing and distributing incoming network traffic to the appropriate backend servers and resources.

Within an ADC, the virtual server acts as a logical entity that receives and processes incoming clients. In FortiADC, the virtual server is configured with specific settings to define how to load-balance different types of traffic.

At the core of the Virtual Server module is the **Virtual Server** configuration that serves as a container for all configuration objects that defines the application delivery control:

- **Application Profile** — Defines the settings for specific protocols and how the FortiADC server should handle its traffic. For details, see Configuring Application profiles on page 451.
- **Session Persistence** — Persistence rules are used to identify traffic that should not be load-balanced, but instead be forwarded to the same backend server based on specific session attributes. For details, see Configuring persistence rules on page 527.
- **Load Balance Method** — Specifies the load-balancing algorithm to use for distributing traffic. For details, see Configuring load-balancing (LB) methods on page 526.
- **Content Routing** — Defines how Layer 7 traffic should be routed based on the TCP/IP or HTTP header content. For details, see Configuring content routes on page 412.
- **Clone Pool** — Specifies the set of destinations of monitor servers that will be sent a copy of the traffic to detect suspicious patterns before allowing traffic to reach the real server. For details, see Using clone pools on page 430.
- **Security Profiles** — Specifies the applicable security policies to enforce, such as policies from the Web Application Firewall, Anti Virus, DoS Protection, and Intrusion Prevention. For details, see Web Application Firewall on page 660, AntiVirus on page 841, DoS Protection Profile on page 867, and Intrusion Prevention on page 836.

Module configurations

The **Server Load Balance > Virtual Server** sub-menu contains the following configurations:

- Virtual Server — see Configuring virtual servers on page 384.
- Content Rewriting — see Using content rewriting rules on page 397.
- Content Routing — see Configuring content routes on page 412.
- NAT Source Pool — see Using source pools on page 416.
- Schedule Pool — see Using schedule pools on page 429.
- Clone Pool — see Using clone pools on page 430.

Configuring virtual servers

The Virtual Server configuration defines what type of incoming client traffic will be received by FortiADC and how that traffic will be processed. The configuration framework is designed to be modular to support the granularity of the FortiADC application delivery control rules. This allows you to tailor configurations to your desired use-case.

FortiADC supports the following three classes of application delivery control:

- Layer 7 — Load-balancing, persistence, and routing are based on Layer 7 (Application Layer) objects, such as HTTP headers and cookies. As such, Web Application Firewall policies can be applied to Layer 7 Virtual Servers. Furthermore, additional functionality exclusive to HTTP/HTTPS can also be applied, including Application Optimization and Scripting.
- Layer 4 — Load-balancing, persistence, and network address translation are based on Layer 4 (Transport Layer) objects, such as source and destination IP addresses.
- Layer 2 — FortiADC offers the Layer 2 option to load-balance connections among multiple next-hop gateways when the request's destination IP is unknown.

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.

Virtual Server configuration modes

FortiADC offers two Virtual Server configuration modes:

- **Basic Mode** — The configuration process is simplified to only include the minimally required parameters to configure a virtual server: Application type, IP address, Port, Interface, and Real Server Pool. FortiADC will then configure "advanced parameters" using default values. To configure a virtual server in Basic Mode, see Basic Mode virtual server configuration on page 385.

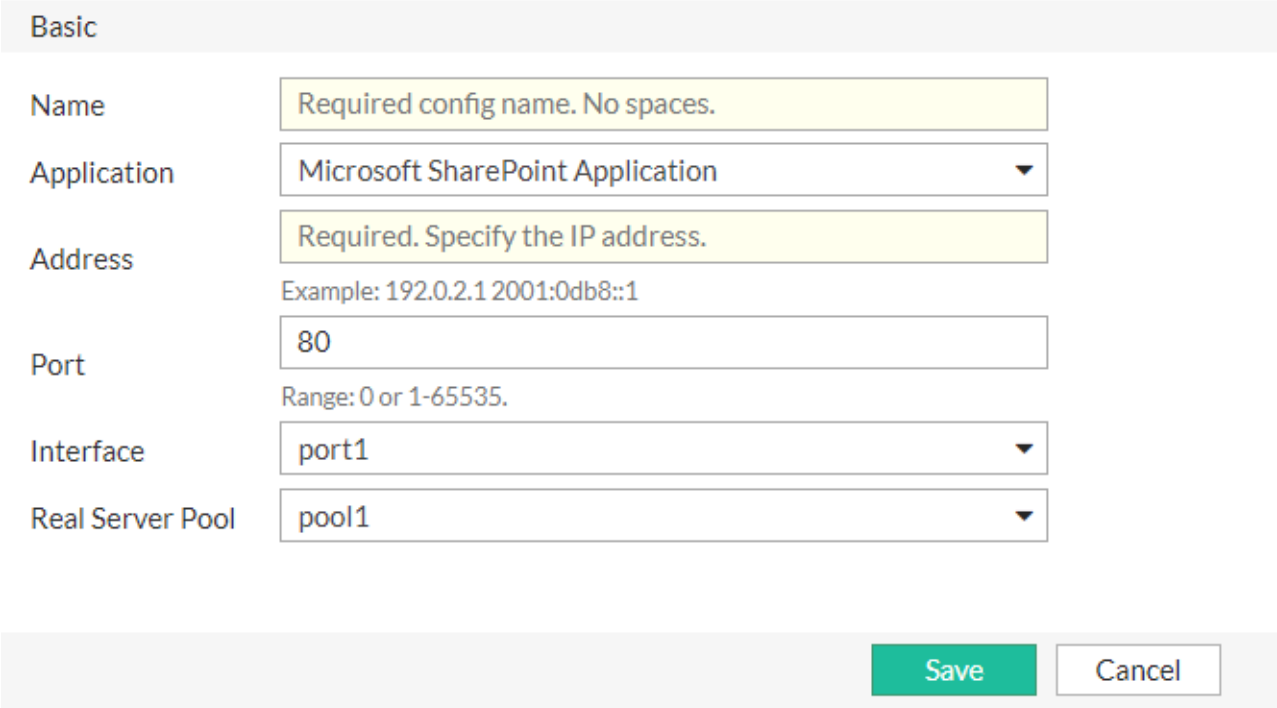
- **Advanced Mode** — In the Advanced Mode configuration, users are offered the option to configure all parameters, both basic and advanced. This involves multiple tabs and sections, with many parameters dependent on the specific virtual server type and Application profile. This configuration mode is ideal for "power users" who are knowledgeable and have an advanced understanding of load-balancing and virtual servers. To configure a virtual server in Advanced Mode, see Advanced Mode virtual server configuration on page 387.

Basic Mode virtual server configuration

The Basic virtual server configuration mode simplifies the configuration process by requiring only the basic parameters to be configured and applying default values for advanced parameters.

To configure a virtual server using Basic Mode:

1. Go to **Server Load Balance > Virtual Server**.
   The configuration page displays the **Virtual Server** tab.
2. Click **Create New > Basic Mode** to open the Basic Mode configuration editor.



3. Configure the following Basic settings:

Setting	Description
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”. Once saved, the name of a virtual server configuration cannot be changed.
Application	Select an application from the following: • Microsoft SharePoint Application • Microsoft Exchange Server Application

Setting	Description
	• 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 • ISO8583 • L7 TCP • L7 UDP
Address	Specify the IP address provisioned for the virtual server. **Note:** The **Address** option is *not available* if **Application** is **IP**.
Port	Specify the port, ports, or a range of ports. 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. The default port number is 80. Valid values are from 0 to 65535. Port 0 applies to Layer-4 virtual servers only. **Note:** The **Port** option is *not available* if **Application** is **IP**.
Interface	Select a network interface from the list menu, or specify a new one.
Real Server Pool	Specify the real server pool. If a real server pool is already configured, it will be available to select from the drop-down menu, otherwise, you can create a new real server pool.
SSL	The **SSL** option is available if **Application** is **HTTP(S)**. Enable/disable SSL. This is disabled by default. Once SSL is enabled, the **Client SSL Profile** option becomes available and you must select a profile.
Client SSL Profile	The **Client SSL Profile** option is available if **Application** is **HTTP2 H2**, **HTTP (S)** with SSL enabled, **TCPS**, or **SMTP**. Select a Client SSL Profile.
Protocol	The **Protocol** option is available if **Application** is **IP**.

Setting	Description
	Specify the protocol. You can enter up to eight numeric values or value ranges corresponding to the protocols you would like to use, separated by space.
Domain Name	The **Domain Name** option is available if **Application** is **SMTP**. Specify the FQDN.

4. Click **Save**.
   FortiADC creates the virtual server configuration and automatically configures the advanced parameters associated with the selected Application.

Advanced Mode virtual server configuration

The Advanced Mode virtual configuration offers more granularity with many more parameters than in Basic Mode. The available parameters are primarily dependent on the **Type** and **Application Profile**.

The settings for Advanced Mode configuration are separated into tabs to configure specific virtual server functionality.

Four main tabs that are available by default:

- **Basic** — Includes basic settings that specify the name, type, status, and address type. Depending on the virtual server type selected, you can configure Specifics such as the schedule pool, content routing, content writing, packet forwarding method, or NAT source pool.
- **General** — Includes the core parameters that determines the connection, application profile, persistence, load-balancing method, and more parameters that define how to handle the traffic.
- **Security** — Policies from various security modules can be applied, depending on whether the policy is compatible with the virtual server type.
- **Monitoring** — Various logs can be enabled to monitor the virtual server activity.

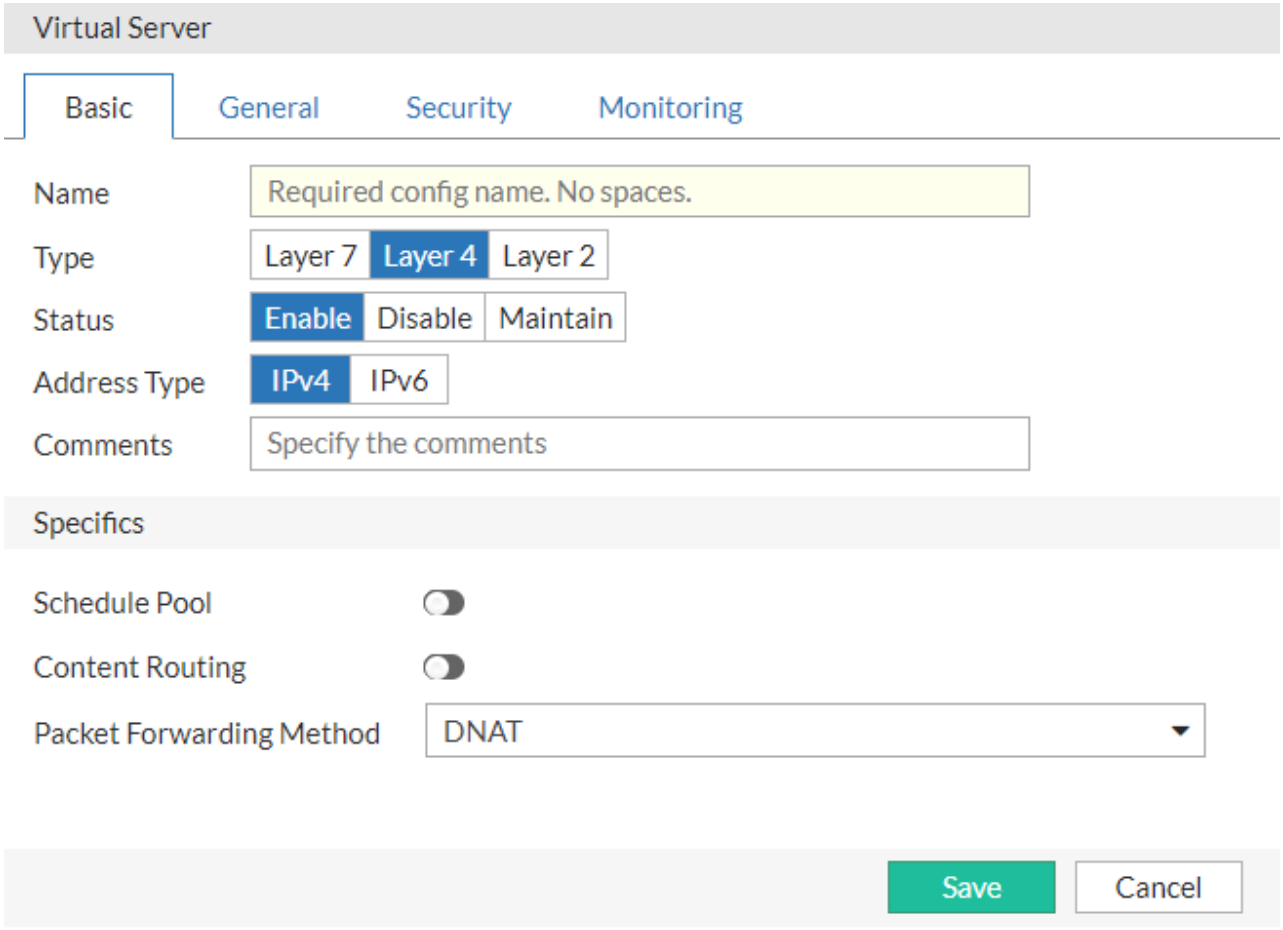
Two additional tabs become available when Layer 7 Type is selected:

- **Application Optimization** — Available only for Layer 7 HTTP and HTTPS application profiles. In this tab, you can select a user-defined Page Speed profile or create a new one.
- **SSL Traffic Mirror** — Available only for Layer 7 HTTPS and TCPS application profiles. In this tab, you can enable SSL Traffic Mirror to send a copy of traffic decrypted by SSL inspection to the specified outgoing interfaces.

To configure a virtual server using the Advanced Mode:

1. Go to **Server Load Balance > Virtual Server**.
   The configuration page displays the **Virtual Server** tab.

2. Click **Create New > Advanced Mode** to open the Advanced Mode configuration editor.



3. Configure and save the settings in the **Basic** tab.

Setting	Description
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.
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. Depending on your Type selection, the Layer 7, Layer 4, or Layer 2 **Specifics** configuration section will appear.
Status	• Enable — The virtual server can receive new sessions.

Setting	Description
	• 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.
Address Type	• IPv4 • IPv6 **Note**: IPv6 is not supported for FTP, HTTP Turbo, RDP, or SIP profiles.
Comment	A string used to describe the purpose of the configuration

a. If the **Type** is **Layer 7**, configure the following Layer 7 **Specifics** settings:

Setting	Description
Schedule Pool	Enable/disable the Schedule Pool. This is disabled by default. **Note**: If Schedule Pool is enabled, Content Routing becomes unavailable.
Schedule Pool List	The **Schedule Pool List** option appears if **Schedule Pool** is enabled. Select the schedule pool(s) and arrange them in a desired order.
Content Routing	Enable/disable the Content Routing. This is disabled by default. 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	The **Content Routing List** option appears if **Content Routing** is enabled. Select the content-routing rules and 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	Enable/disable the Content Rewriting. This is disabled by default. **Note:** • This option applies to Layer-7 only. • This option does NOT apply to SIP profiles.
Content Rewriting List	The **Content Rewriting List** option appears if **Content Rewriting** is enabled. Select the content rewriting rules and 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.
NAT Source Pool List	Select one or more source pool configuration objects and arrange them in a desired order. See Using source pools. **Note**:

Setting	Description

	By default, the same IP pool cannot be set up in different virtual servers. However, you can enable IP address sharing through the CLI to allow the source pool to be set up in different virtual servers.
	To enable IP address sharing:
	`config system global`
	`  set share-ip-address enable`
	`end`

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). 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.

b. If the **Type** is **Layer 4**, configure the following Layer 4 **Specifics** settings:

Setting	Description
Schedule Pool	Enable/disable the Schedule Pool. This is disabled by default. **Note**: If Schedule Pool is enabled, Content Routing becomes unavailable.
Schedule Pool List	The **Schedule Pool List** option appears if **Schedule Pool** is enabled. Select the schedule pool(s) and arrange them in a desired order.
Content Routing	Enable/disable the Content Routing. This is disabled by default. 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	The **Content Routing List** option appears if **Content Routing** is enabled. Select the content-routing rules and 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.
Packet Forwarding Method	**Note:** This setting applies to Layer-4 virtual servers only.

Setting	Description
	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 Tpye 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.
	Note:
	By default, the same IP pool cannot be set up in different virtual servers. However, you can enable IP address sharing through the CLI to allow the source pool to be set up in different virtual servers.
	To enable IP address sharing:
	`config system global`
	`  set share-ip-address enable`
	`end`

c. If the **Type** is **Layer 2**, configure the following Layer 2 **Specifics** settings:

Setting	Description
Schedule Pool	Enable/disable the Schedule Pool. This is disabled by default. **Note**: If Schedule Pool is enabled, Content Routing becomes unavailable.
Schedule Pool List	The **Schedule Pool List** option appears if **Schedule Pool** is enabled. Select the schedule pool(s) and arrange them in a desired order.
Content Routing	Enable/disable the Content Routing. This is disabled by default. Note:

Setting	Description
	• 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	The **Content Routing List** option appears if **Content Routing** is enabled. Select the content-routing rules and 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.

4. Configure and save the settings in the **General** tab.

Setting	Description
Configuration	
Use Azure LB Backend IP	The **Use Azure LB Backend IP** option is only available if you have deployed FortiADC instances in Azure via the ARM template for VRRP HA. For details, see Azure LB Backend on page 166 and the FortiADC Document on Deploying High Availability on Azure. Enable or disable the virtual server to reference Azure LB Backend IP as the virtual server IP.
Azure LB Backend IP	The **Azure LB Backend IP** option is only available if you have deployed FortiADC instances in Azure via the ARM template for VRRP HA. For details, see Azure LB Backend on page 166 and the FortiADC Document on Deploying High Availability on Azure. Select the Azure LB Backend IP Configuration used as the virtual server IP.
Address	Enter the IP address provisioned for 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,

Setting	Description
	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. **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 100,000,000. 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	This option is available if **Layer 4** is selected as the **Type** in the **Basic** settings. With Layer 4 profiles 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.
Resources	
Profile	Select a predefined or user-defined profile configuration object. See Configuring Application profiles. **Note**: Only TCP, UDP and IP profiles are available for Layer 2 VS Type with IPv6 Address.
Client SSL Profile	**Note**: This setting applies to **HTTPS**, **TCPS**, HTTP2 H2, **SMTP**, and **FTPS** 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. **Note**: If a ZTNA Profile is referenced in the VS, ensure the client SSL profile has enabled client certificate verification for the corresponding EMS CA certificate object. See Configuring Client SSL profiles on page 512.
Persistence	Select a predefined or user-defined persistence configuration object. See Configuring persistence rules. **Note:** The persistence rule with **Match Across Virtual Servers** enabled works only with L4 virtual servers or the L7 virtual server whose profile is LB_PROF_RADIUS.

Setting	Description
Method	Select a predefined or user-defined method configuration object. See Configuring load-balancing (LB) methods on page 526.
Real Server Pool	Select a real server pool configuration object. See Configuring real server pools. **Note**: For Layer 2 VS Type, the available real server pools are dependent on the Address (IPv4 or IPv6) selected in Basic settings.
Clone Pool	Select a configuration object. See Configuring a clone pool on page 431. **Note**: This option is not available if the VS Type is **Layer 2** and Address is **IPv6**.
Auth Policy	This option is available if **Layer 7** is selected as the **Type** in the **Basic** settings. Select an authentication policy configuration object. HTTP/HTTPS only. See Authentication Policy.
Scripting	This option is available only if Scripting is enabled AND if **Layer 7** is selected as the **Type** in the **Basic** settings. Select the scripting object(s) and arrange them in a desired order. **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.
AD FS Published Service	This option is available if **Layer 7** is selected as the **Type** in the **Basic** settings. Select an AD FS configuration object. HTTPS only. See AD FS Proxy on page 917.
L2 Exception List	This option is available if **Layer 2** is selected as the **Type** in the **Basic** settings AND an HTTP/HTTPS server load-balancing profile is selected. Select an exception configuration object. See Configuring an L2 Exception List.
HTTP Redirect to HTTPS	This option becomes available when an HTTPS server load-balancing profile is selected. Enable/disable HTTP redirect to HTTPS. This option is disabled by default. 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**. You can either accept the default port (80), or specify up to eight ports or ranges of ports of your preference.
Error Page (This section is only available for Layer 7 virtual servers)	
Error Page	Select an error page configuration object. See Configuring error pages. **Note**: Not supported for SIP profiles.

Setting	Description
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. Maximum 1023 bytes. **Note**: Not supported for SIP profiles.
FortiGSLB (This section is not available for Layer 2 virtual servers)	
Public IP Type	Set the Public IP type for the virtual server as either IPv4 or IPv6.
Public IPv4/IPv6	Enter the virtual server public IP address.
One Click GSLB Server	Enable/disable the FortiGSLB One Click GSLB server.
Host Name	The **Host Name** option is available if **One Click GSLB Server** is enabled. Enter the hostname part of the FQDN, such as `www`. **Note:** You can specify the @ symbol to denote the zone root. The value substitute for @ is the preceding $ORIGIN directive.
Domain Name	The **Domain Name** option is available if **One Click GSLB Server** is enabled. The domain name must end with a period. For example, `example.com.`

5. Configure and save the settings in the **Security** tab. The security settings available are dependent on the virtual server type selected in **Basic** settings and the server load-balancing profile in **General** settings.

Setting	Description
WAF Profile	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. AV profile can support HTTP/HTTPS/SMTP. See Creating an AV profile on page 842.
DoS Protection Profile	Select a DoS protection profile configuration object or create a new one. See DoS Protection Profile on page 867.
Captcha Profile	Select a Captcha configuration object. See Configuring Captcha on page 541.
ZTNA Profile	**Note**: This setting applies to **Layer 7 HTTPS** and **TCPS** applications only. Select a ZTNA Profile object. See Configuring a ZTNA Profile on page 865

6. Configure and save the settings in the **SSL Traffic Mirror** tab.
Note: The SSL Traffic Mirror settings are only accessible if Layer 7 is selected as the Type in the Basic settings AND an HTTPS or TCPS server load-balancing profile is selected in General settings.

Setting	Description
SSL Traffic Mirror	Enable/disable SSL Traffic Mirror.
Mirror To	The **Mirror To** field appears when **SSL Traffic Mirror** is enabled. Select the ports from the list of Available Items.

7. Configure and save the settings in the **Application Optimization** tab.
Note: The Application Optimization settings are only accessible if Layer 7 is selected as the Type in the Basic settings AND an HTTP or HTTPS server load-balancing profile is selected in General settings.

Setting	Description
Page Speed	Select a page speed optimization profile.

8. Configure and save the settings in the **Monitoring** tab.

Setting	Description
Traffic Log	Enable/disable 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.
FortiView	Enable/disable to view this virtual server from FortiView.
WCCP	The WCCP option is only available for Layer 7 virtual servers. Enable/disable Web Cache Communications Protocol.

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:0

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

HTTP header rewriting on page 397 summarizes the HTTP header fields that can be rewritten.

HTTP header rewriting

Direction	HTTP Header
HTTP Request	• Host • Referer

Direction	HTTP Header
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 399.
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 — `https://example.com/content/index.html` • Relative URL — `content/index.html` 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: `https://example.com/content/index.html` **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. • 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: `https://example.com/content/index.html`

Settings	Guidelines
	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.

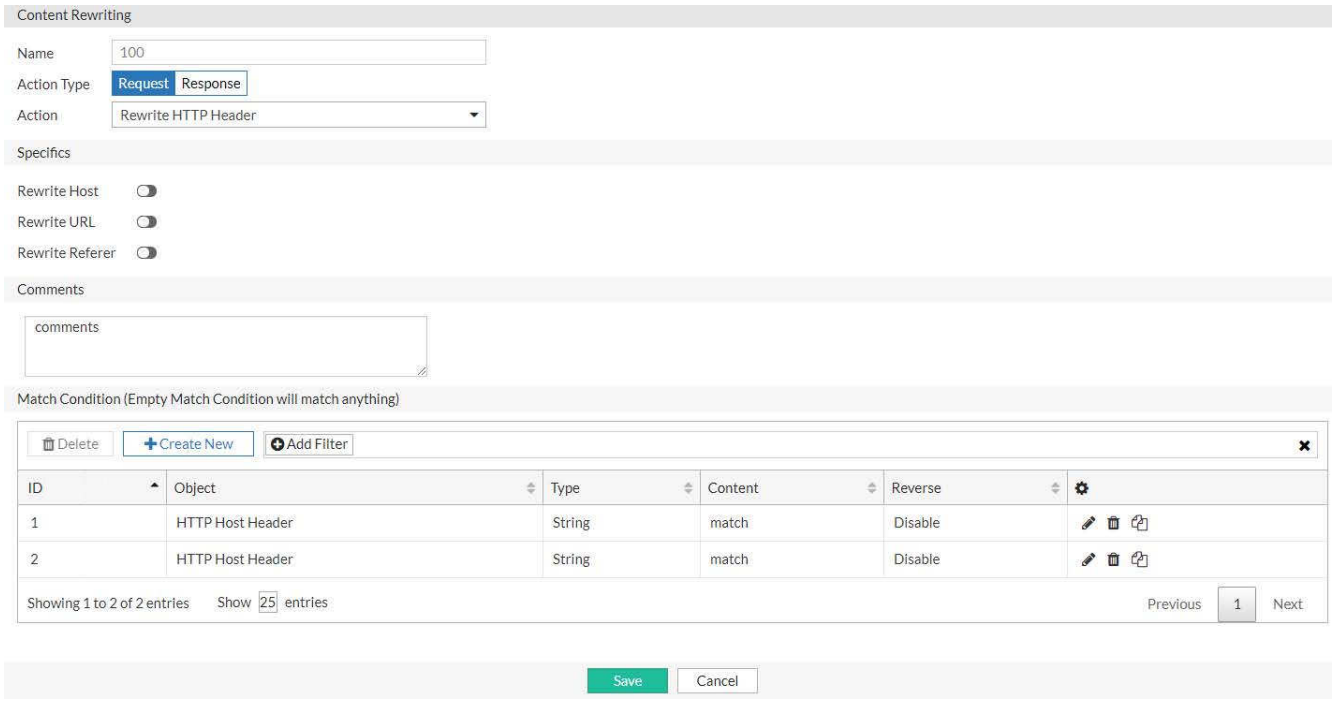
Example: Redirecting HTTP to HTTPS

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 400 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`.

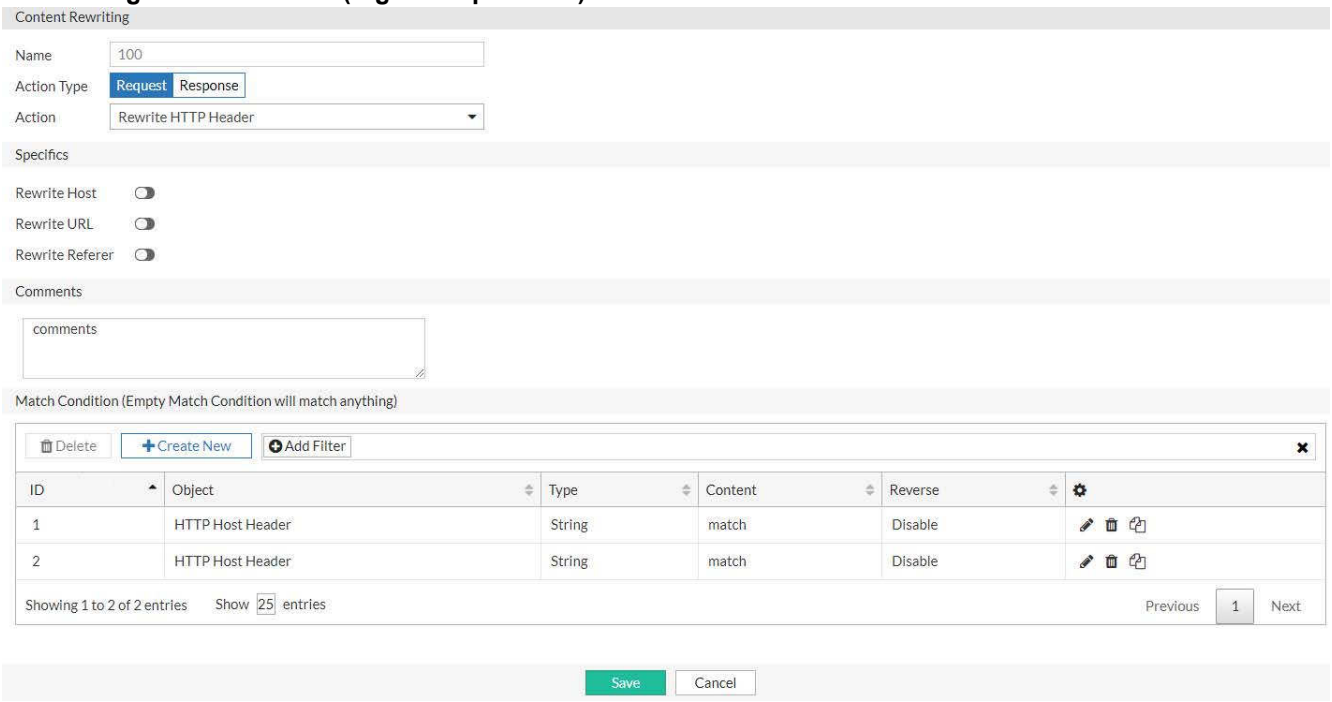
Redirecting HTTP to HTTPS (literal string)



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 402 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)



Common PCRE syntax elements on page 402 describes commonly used PCRE syntax elements.

PCRE examples submitted to the FortiGate Cookbook on page 405 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.	Let's say the regular expressions in a condition table have the following capture groups: (a)(b)(c(d))(e)

Pattern	Usage	Example
	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 <= *n* <= 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.	This syntax results in back-reference variables with the following values: `$0` — a `$1` — b `$2` — cd `$3` — d `$4` — e
\	Escape character. Except, if it is followed by an alphanumeric character, the alphanumeric character is *not* matched literally as usual. Instead, it is interpreted as a regular expression token. For example, \w matches a word, as defined by the locale. Except, if it is followed by regular expression special character: *.|^$?+\(){}[]\ 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.	**Text:** /url?parameter=value **Regular expression:** \?param **Matches:** ?param
.	Matches any single character *except* \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 *not* 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) *unless* followed by a question mark (?), which makes it optional.	**Text:** www.example.com **Regular expression:** w+ **Matches:** www

Pattern	Usage	Example
	Does not match if there is not at least 1 instance.	Would also match “w”, “ww”, “wwww”, 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 *many* times as possible (also called “greedy” matching). *? — Match as *few* 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 “wwww” in wwww.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 *not* include those next characters in the returned match string (if any). 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 *not* when it is part of “catch”.	**Text:** /url?parameter=valuepack **Regular expression:** p(?=arameter) **Matches:** p, but only in “parameter, *not* in “pack”, which does not end with “arameter”.
^	Matches either: the *position* of the beginning of a line (or, in multiline mode, the first line), *not* 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, *or* when you want to match anything that is *not* a specific character.	**Text:** /url?parameter=value **Regular expression:** ^/url **Matches:** /url, but *only* if it is at the beginning of the path string. It will *not* match “/url” in subdirectories. **Text:** /url?parameter=value **Regular expression:** [^u] **Matches:** /rl?parameter=value

Pattern	Usage	Example
$	Matches the *position* of the end of a line (or, in multiline mode, the entire string), *not* 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, `[@-B]` 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 *not* 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 “wwww”, 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.
|	Matches *either* the character/capture group before *or* 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.

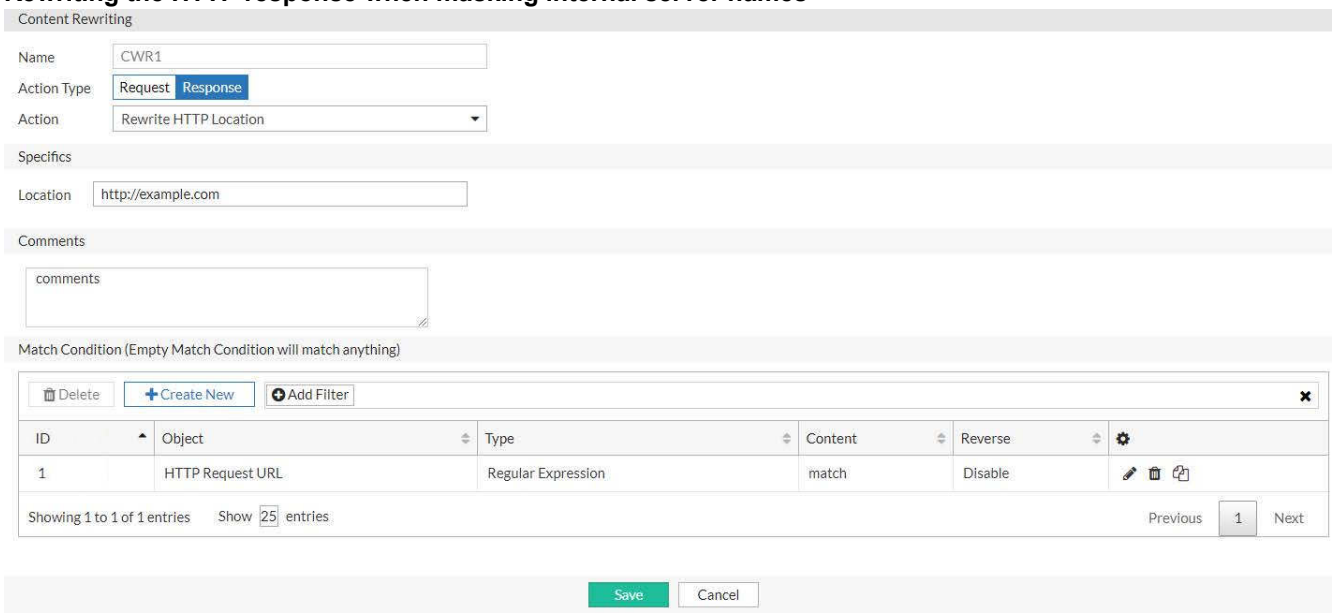
Regular Expression	Usage
\b10\.1\.1\.1\b	A specific IPv4 address.
\b(25[0-5]|2[0-4][0-9]|[01]?[0-9][0-9]?) \.(25[0-5]|2[0-4][0-9]|[01]?[0-9][0-9]?) \.(25[0-5]|2[0-4][0-9]|[01]?[0-9][0-9]?) \.(25[0-5]|2[0-4][0-9]|[01]?[0-9][0-9]?) \b	Any IPv4 address.
(?i)\b.*\.(a(c|d|e(ro)?|f|g|i|m|n|o|q|r|s (ia)?|t|y|w|x|z) |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|du|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|vy) |m(a|c|d|e|g|h|il|k|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|el|f|g|h|j|k|l|m|n|o|p|r(avel)?|t|v|w|z) |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	Any domain name.
(?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 407 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



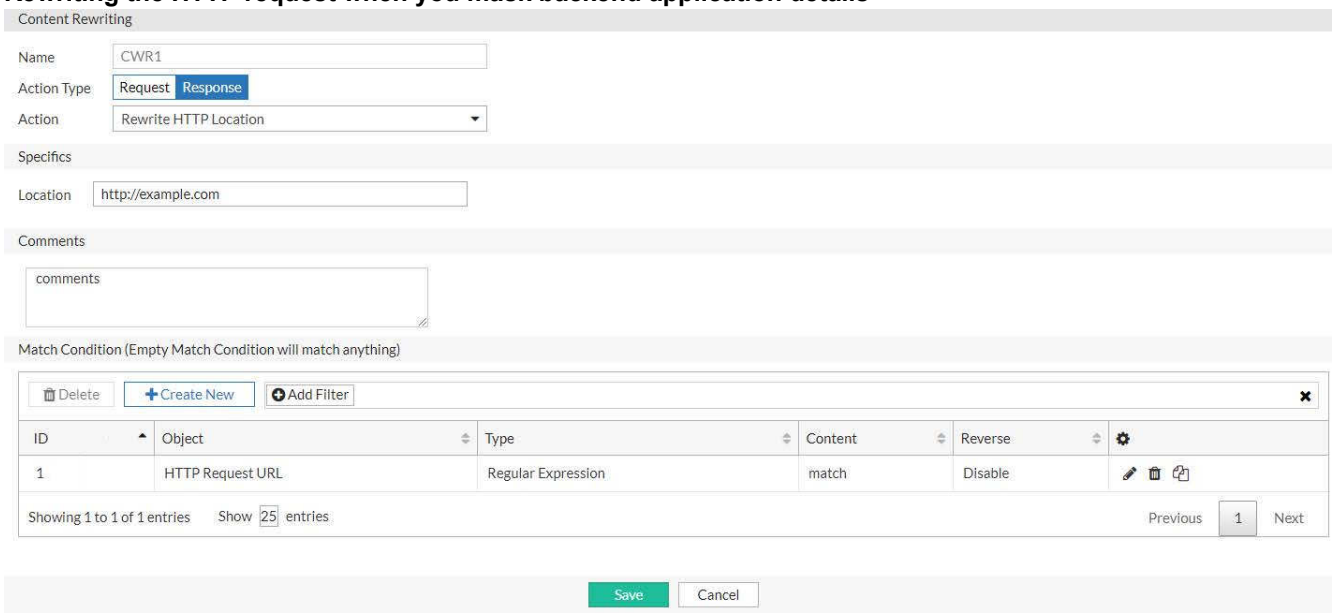
Example: Rewriting the HTTP request and response to mask application details

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.

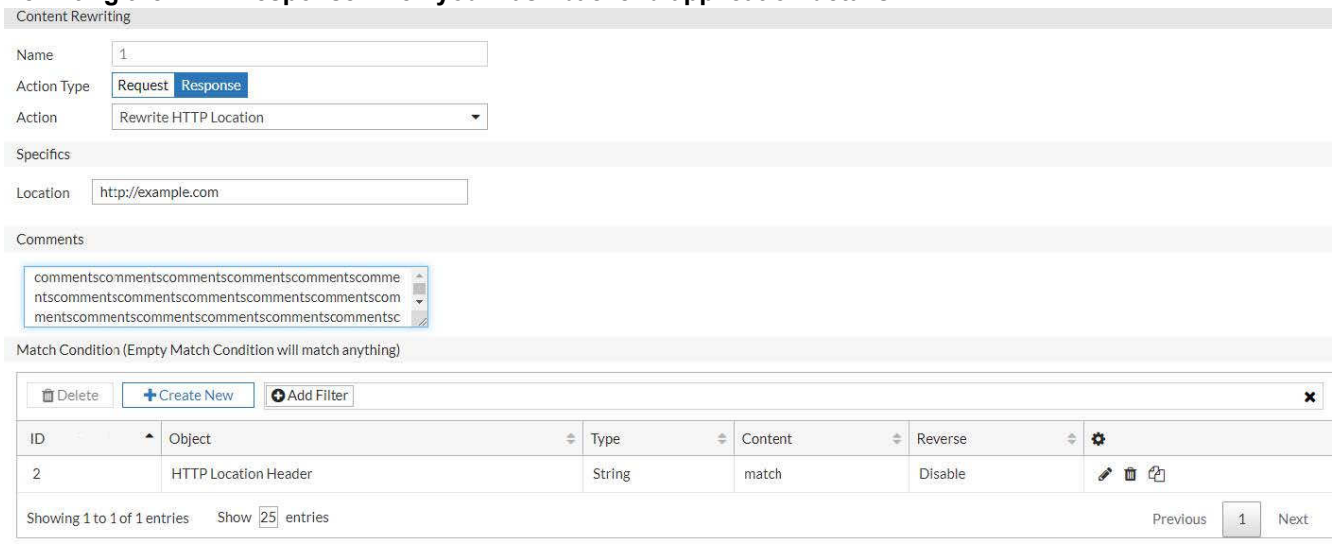
Rewriting the HTTP request when you mask backend application details on page 408 shows an HTTP request rule. In the match condition table, the rule is set to match traffic that has the Host header domain `example.com` and the relative URL `/blog` in the HTTP request URL. The rule action rewrites the request URL to the internal URL `http://www.example.com/wordpress/?feed=rss2`.

Rewriting the HTTP request when you mask backend application details



Rewriting the HTTP response when you mask backend application details on page 408 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

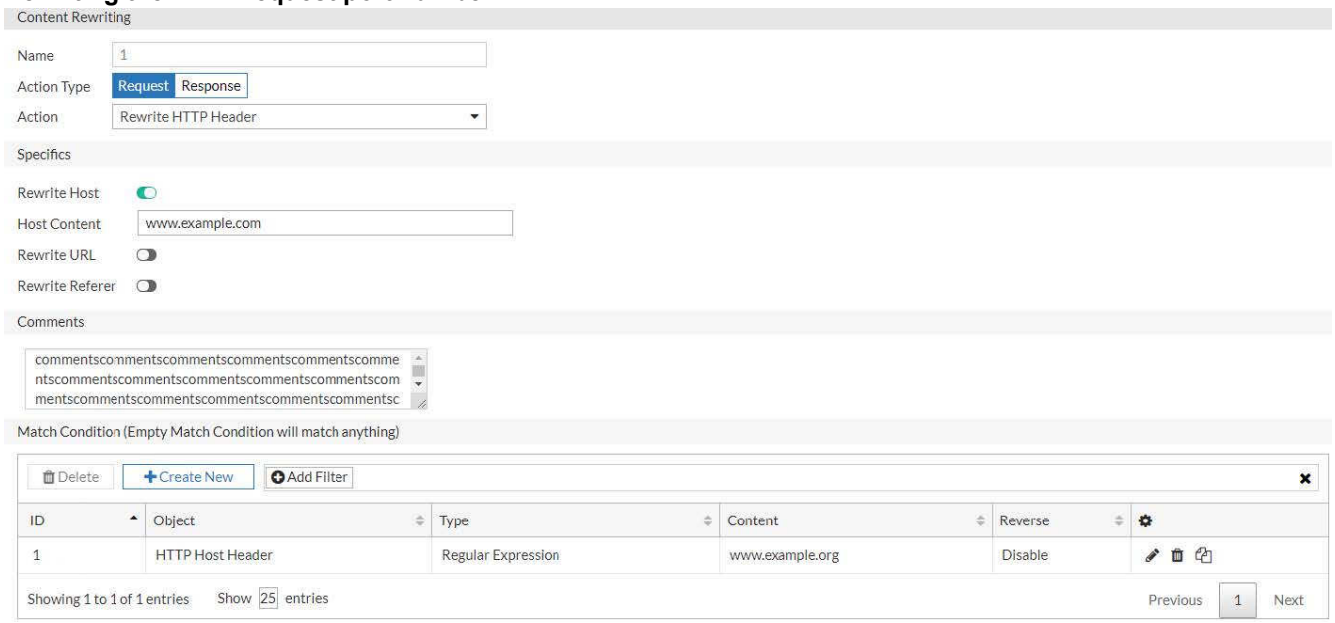


Example: Rewriting the HTTP request to harmonize port numbers

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 409 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



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

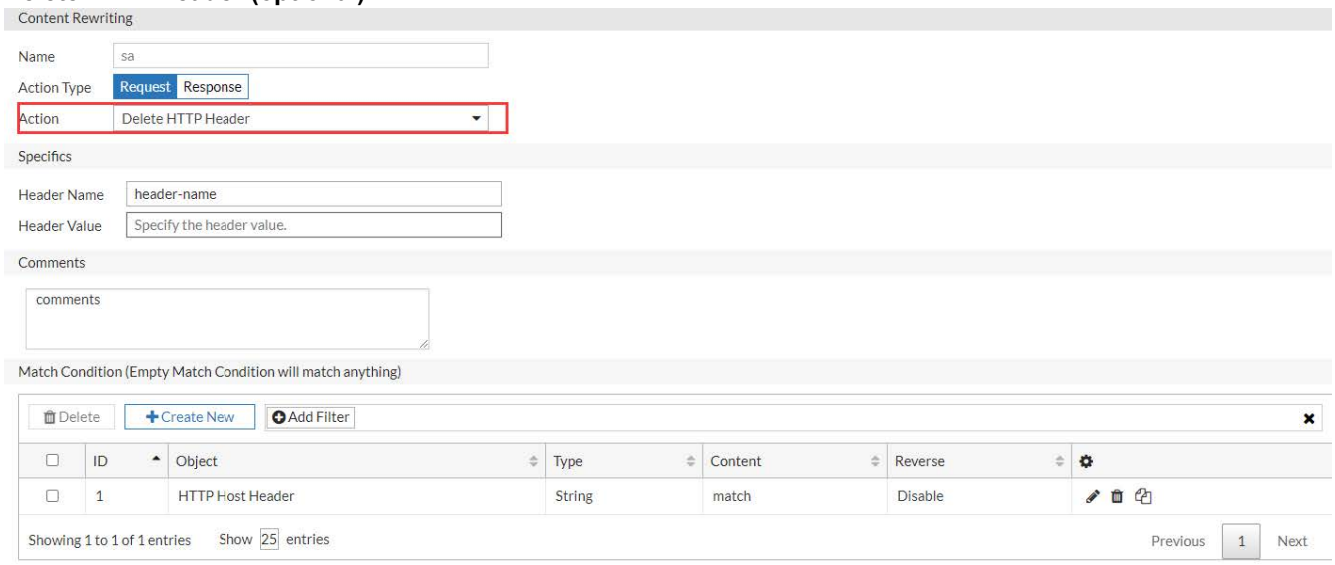
Support for HSTS and HPKP can be implemented for both SSL offloading and forward proxy.

SSL offloading

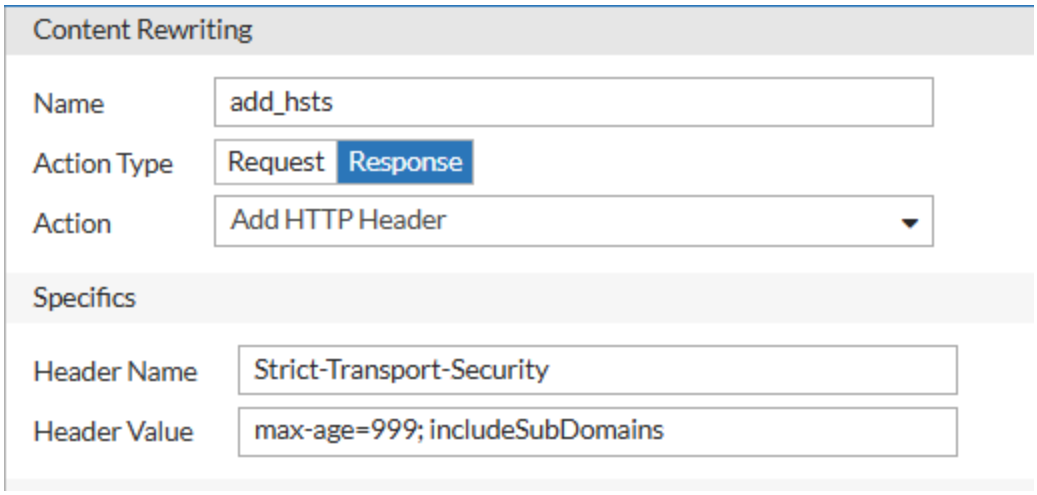
On the **Server Load Balance > Virtual Server > Content Rewriting** page, do the following:

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 411.
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 412.

Delete HTTP header (optional)



Add HTTP header



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. Enter a name for the Content Routing configuration and select the Type:

Setting	Description
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	• Layer 4 • Layer 7

The content routing settings applicable to the selected **Type** displays.

5. If the Type is Layer 7, configure the following settings.
   a. Configure the General settings:

Setting	Description
Schedule Pool	Enable to select a schedule pool for content routing. This is disabled by default.
Schedule Pool List	The **Schedule Pool List** option is available if **Schedule Pool** is **enabled**. Select the schedule pool objects.
Real Server Pool	The **Real Server Pool** option is available if **Schedule Pool** is **disabled**. Select the real server pool configuration.
Persistence	The **Persistence** option is configurable if **Inherit** (Persistence) is **disabled**. Select a session persistence type if not inheriting persistence from the virtual server.
Inherit (Persistence)	Enable to use the persistence object specified in the virtual server configuration. This is enabled by default.
Method	The **Method** option is configurable if **Inherit** (Method) is **disabled**. Select a load balancing method type if not inheriting method from the virtual server.
Inherit (Method)	Enable to use the method specified in the virtual server configuration.

   b. Click **Save** to commit the General settings to the Content Routing configuration.
   Once the configuration is saved, the Match Condition (Empty Match Condition will match anything) section becomes configurable.
   c. Under the Match Condition (Empty Match Condition will match anything) section, click **Create New** to display the configuration editor.

d. Configure the Match Condition (Empty Match Condition will match anything) settings:

Setting	Description
Object	Select content matching conditions based on the following parameters: • HTTP Host Header • HTTP Referer Header • HTTP Request URL • Source IP Address • SNI **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	The **Type** option is **not available** if **Object** is **Source IP Address**. • 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.
Negative	Enable to rule matches if traffic does not match the expression. This is disabled by default.
Ignore Case	The **Ignore Case** option is **not available** if **Object** is **Source IP Address**. Enable to allow users to let the match be case sensitive. This is enabled by default.

e. Click **Save** to commit the Match Condition configuration and exit from the editor.

6. If the Type is Layer 4, configure the following settings.

a. Configure the Specifics settings:

Setting	Description
Source (IPv4)	Address/mask notation to match the source IPv4 address in the packet header.
Source (IPv6)	Address/mask notation to match the source IPv6 address in the packet header.

b. Configure the General settings:

Setting	Description
Schedule Pool	Enable to select a schedule pool for content routing. This is disabled by default.
Schedule Pool List	The **Schedule Pool List** option is available if **Schedule Pool** is **enabled**. Select the schedule pool objects.
Real Server Pool	The **Real Server Pool** option is available if **Schedule Pool** is **disabled**. Select the real server pool configuration.

Setting	Description
Persistence	The **Persistence** option is configurable if **Inherit** (Persistence) is **disabled**. Select a session persistence type if not inheriting persistence from the virtual server.
Inherit (Persistence)	Enable to use the persistence object specified in the virtual server configuration. This is enabled by default.
Method	The **Method** option is configurable if **Inherit** (Method) is **disabled**. Select a load balancing method type if not inheriting method from the virtual server.
Inherit (Method)	Enable to use the method specified in the virtual server configuration.
Packet Forward Method	Select the packet forwarding method: • Inherit • Full NAT
Source Pool List	The **Source Pool List** option is available if **Packet Forward Method** is **Full NAT**. Select the source pool objects for packet forwarding.

7. Save the configuration.
   Once saved, the new Content Routing configuration will be listed in the Content Routing page.

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).



There are no validation checks for duplicate addresses in the NAT source pool for HA synchronization, SNAT, 1-to-1 NAT, and VIP, as the ha-mgmt-ip is not synchronized between the HA nodes. For these configurations, ensure the starting and ending IPs in the address range of the NAT source pool are not duplicates.

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 417.
5. Save the configuration.

Source pool configuration

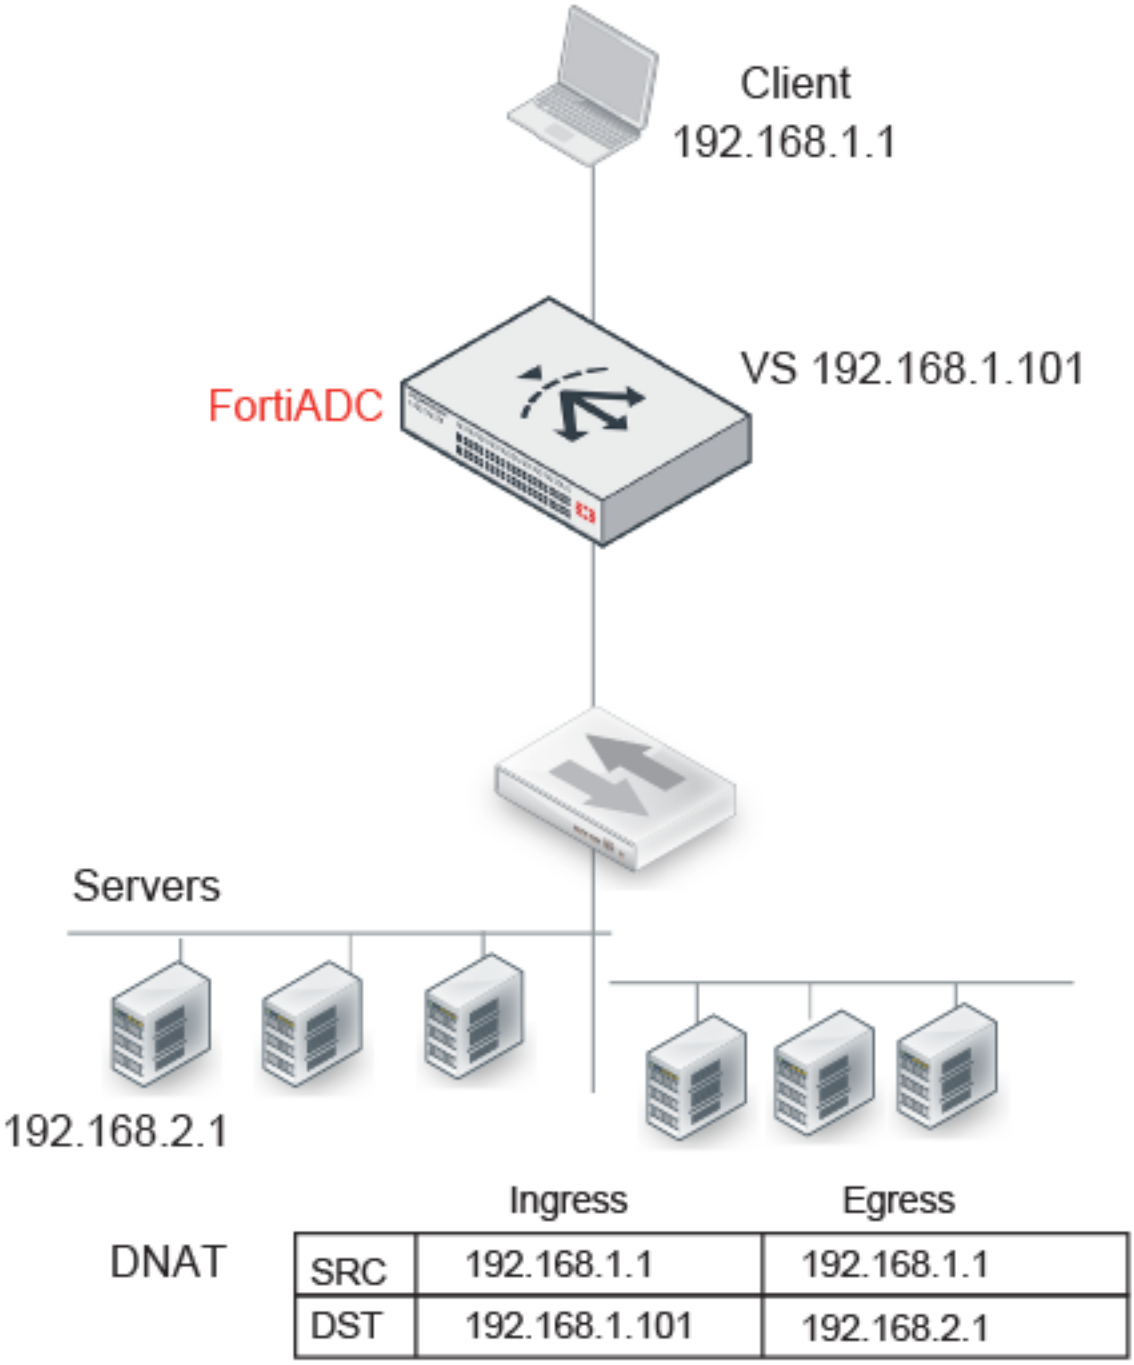
Settings	Guidelines
Name	Configuration name. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No spaces. You reference this name in the virtual server configuration. **Note**: After you initially save the configuration, you cannot edit the name.
Interface	Interface to receive responses from the 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 418 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



DNAT	Ingress	Egress
SRC	192.168.1.1	192.168.1.1
DST	192.168.1.101	192.168.2.1

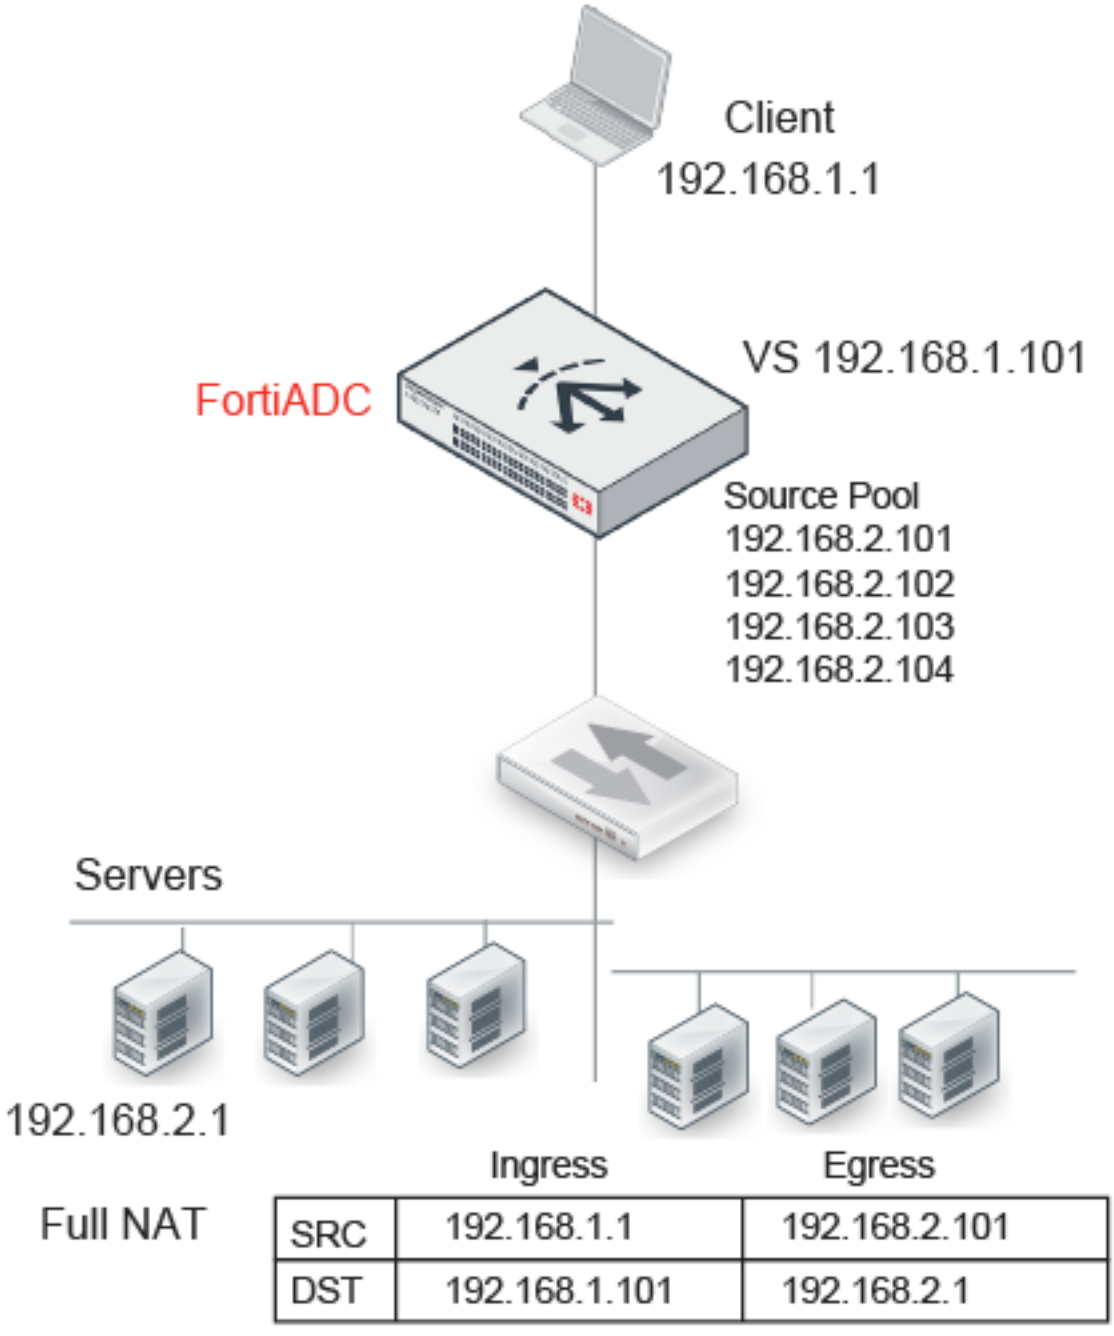
Example: full NAT

Full NAT on page 420 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



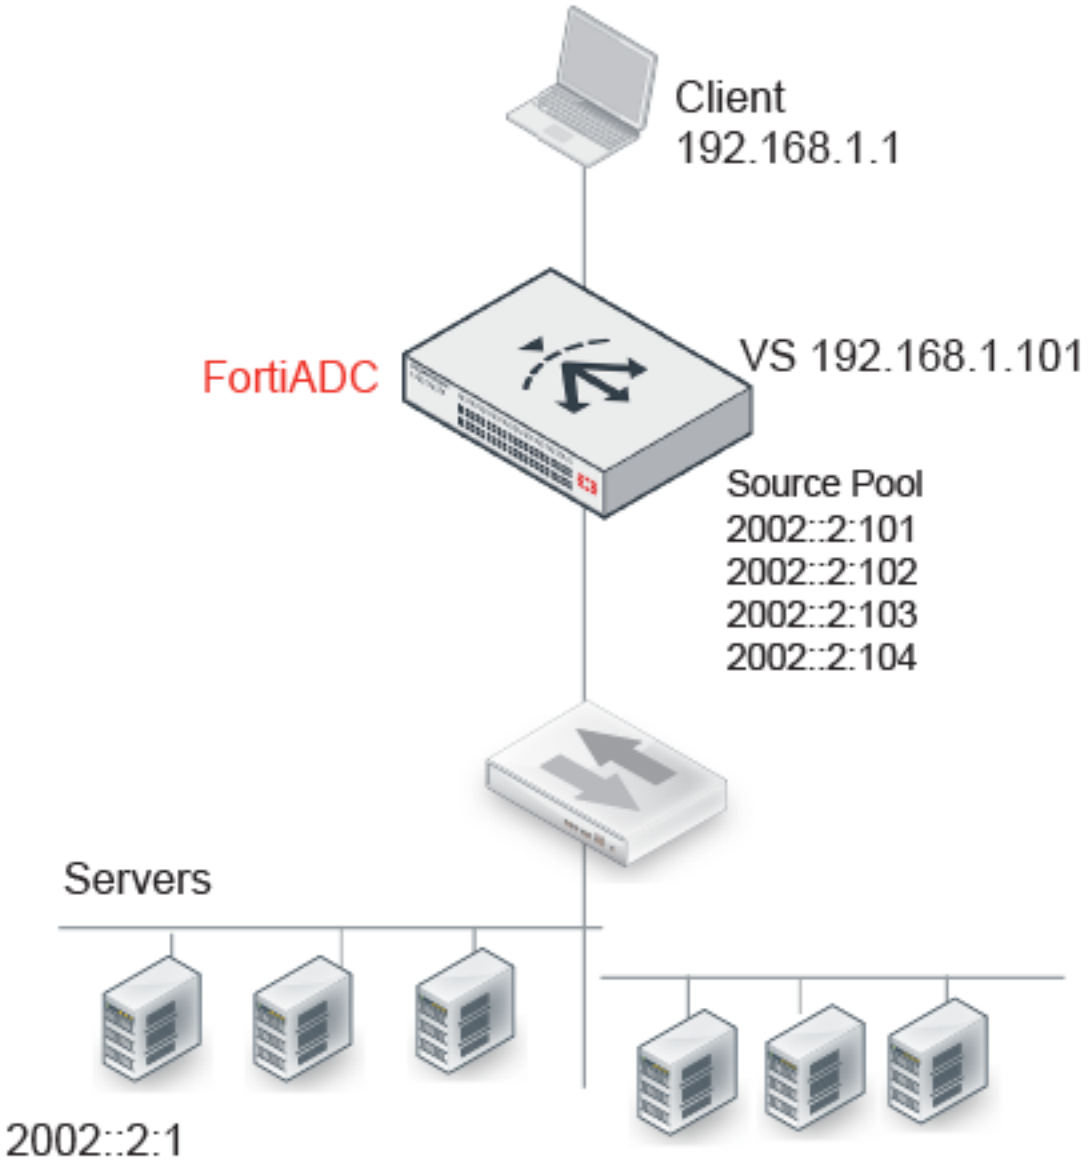
Full NAT		Ingress	Egress
	SRC	192.168.1.1	192.168.2.101
	DST	192.168.1.101	192.168.2.1

Example: NAT46 (Layer 4 virtual servers)

NAT46 (Layer 4 virtual servers) on page 421 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)



NAT46

	Ingress	Egress
SRC	192.168.1.1	2002::2:101
DST	192.168.1.101	2002::2:1

Limitations: NAT46 (Layer 4 virtual servers)

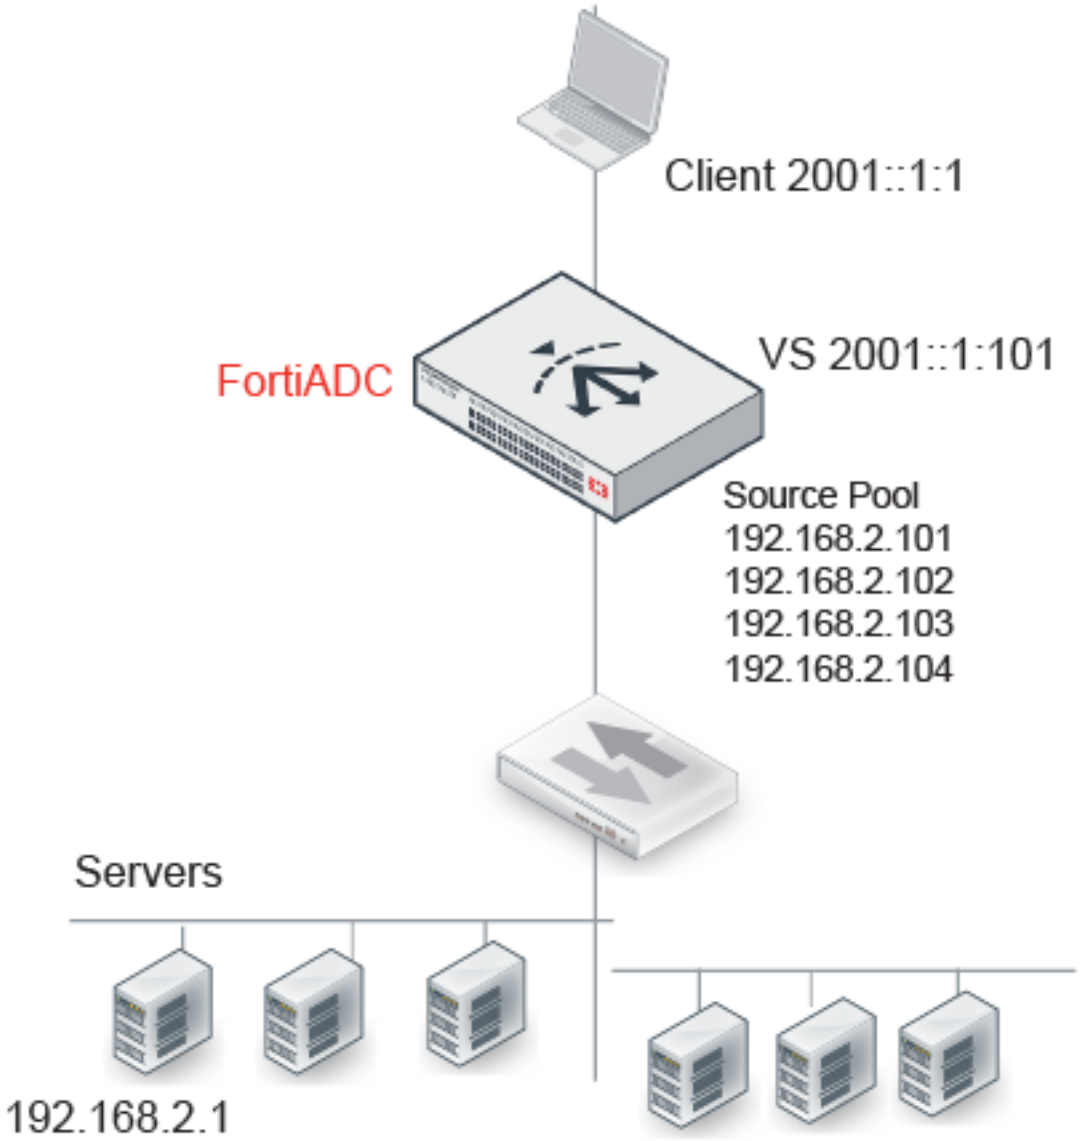
Features	Notes
Profile	Not Supported: FTP
ICMP	ICMP traffic is dropped.

Example: NAT64 (Layer 4 virtual servers)

NAT64 (Layer 4 virtual servers) on page 423 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)



NAT64

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

Limitations: NAT64 (Layer 4 virtual servers)

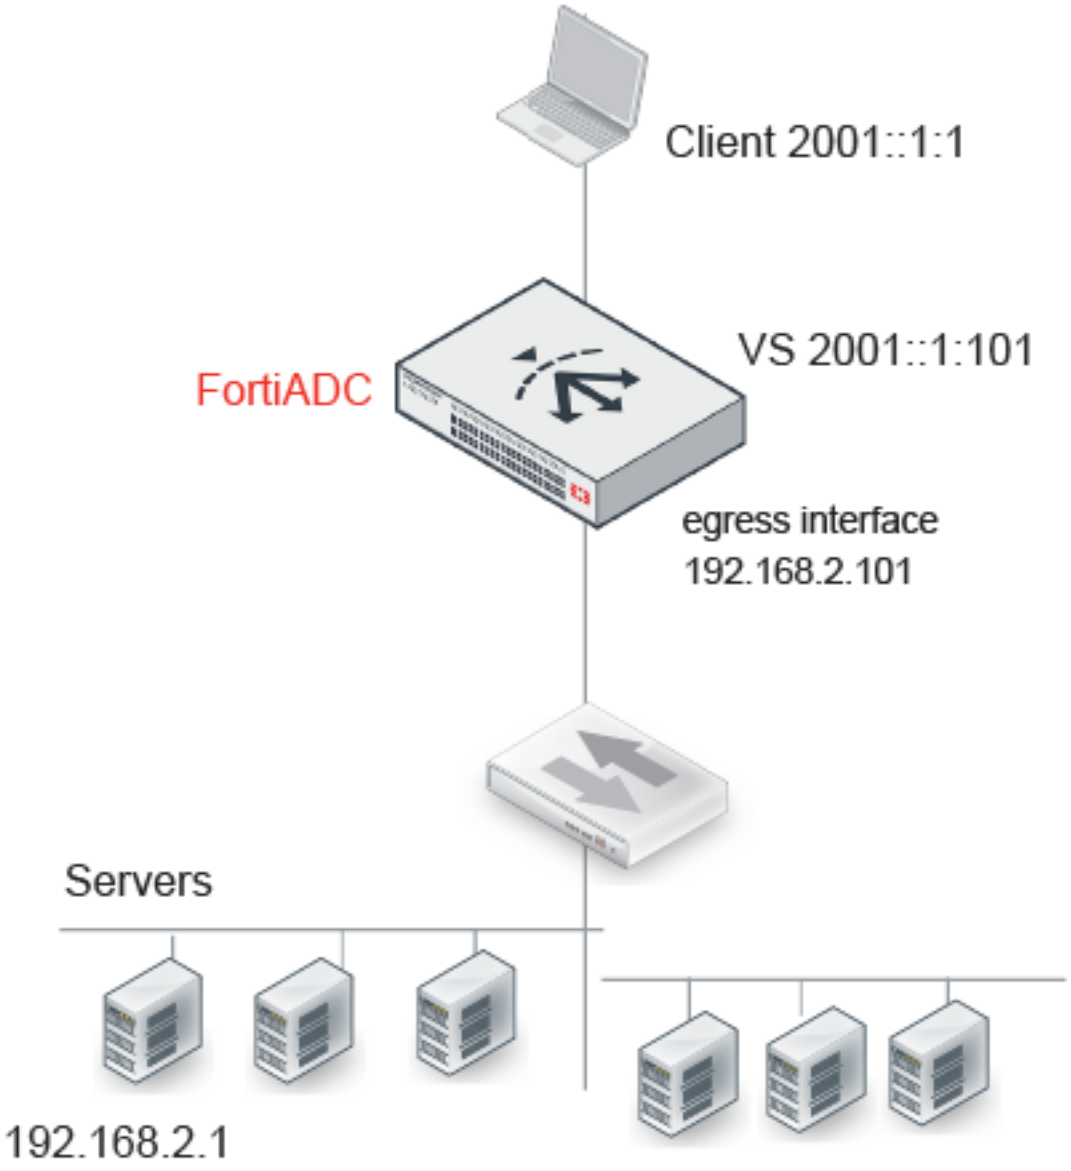
Features	Notes
Profiles	Not Supported: FTP
ICMP	ICMP traffic is dropped.
Security	Not Supported: IP Reputation, DoS protection, Security logs and reports

Example: NAT46 (Layer 7 virtual servers)

NAT46 (Layer 7 virtual servers) on page 425 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)



NAT64

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

Limitations: NAT46 (Layer 7 virtual servers)

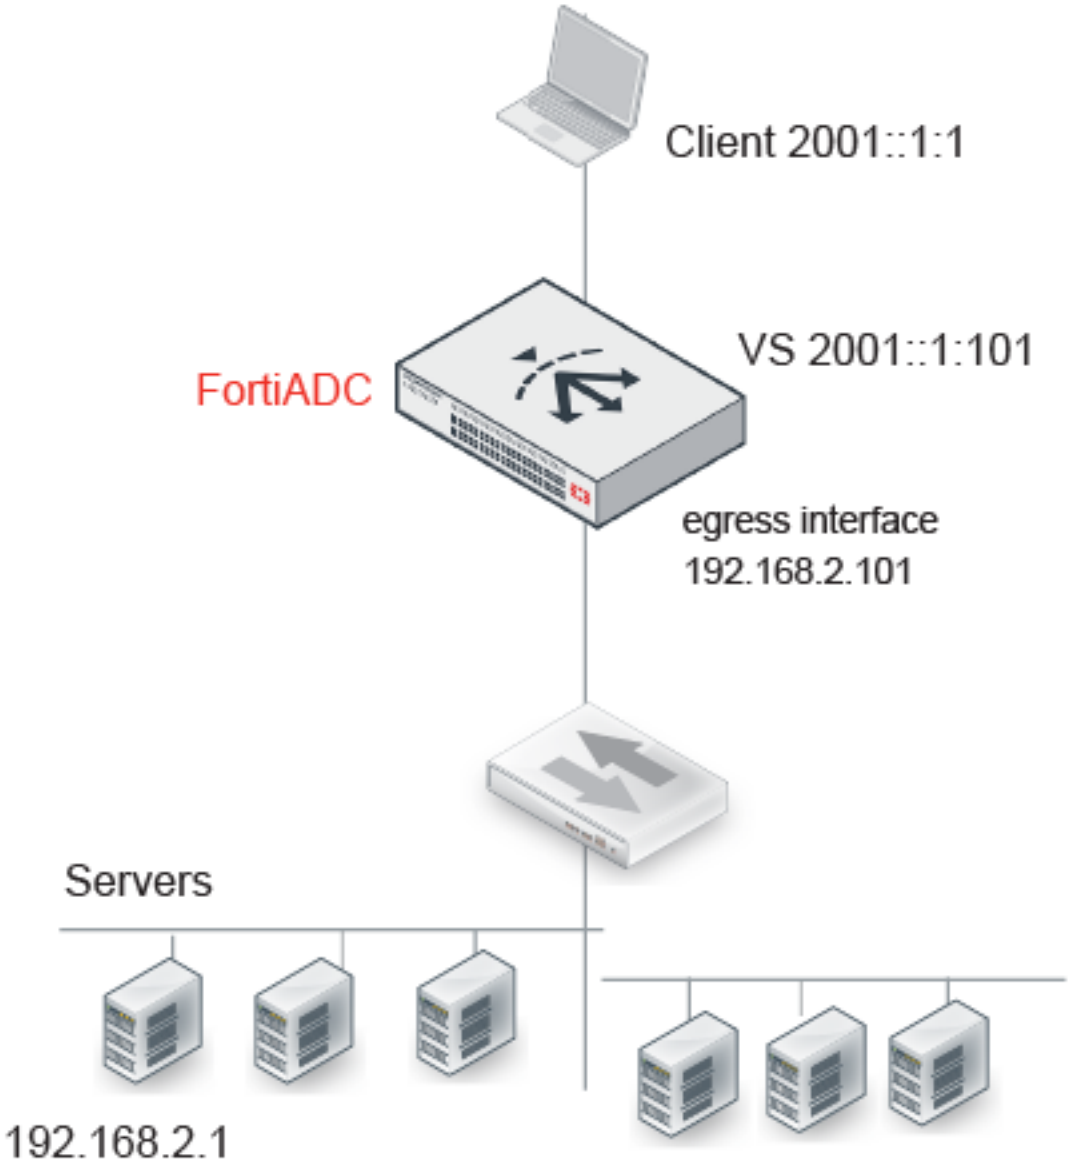
Feature	Note
Profiles	Not Supported: RADIUS, HTTP Turbo
Profile options	Not supported: Source Address (Using the original source IP address for the connection to the real server is contrary to the purpose of NAT.)
Virtual server options	Not supported: Connection Rate Limit
Real server pool options	Not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

Example: NAT64 (Layer 7 virtual servers)

NAT64 (Layer 7 virtual servers) on page 427 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)



NAT64

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

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 becomes 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. Consequently, when members of a pool become inactive, the system will react in the same way as it does when they fail their health check and the Health Check Down Option is set to "Reject" — 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 should 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 > Schedule Pool).
5. Configure content routing rules (Server Load Balance > 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. Go to **Server Load Balance > Virtual Server.**
2. Click 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 consists of a set of destination monitor servers. FortiADC protects real-server pools by duplicating incoming traffic and sending a copy to the clone pool, which retains it for monitoring purposes. The clone pool is assigned to a virtual server, and within this pool are a set of monitor servers, some of which may be IDS (Intrusion Detection System) servers. These IDS servers analyze the traffic to detect suspicious patterns but do not perform firewall functions, such as blocking traffic. Instead, the IDS servers may, for example, send notifications like an email when potential threats are detected.

Important: The clone pool receives the same traffic as the real-server pool, ensuring it mirrors the traffic for monitoring purposes.

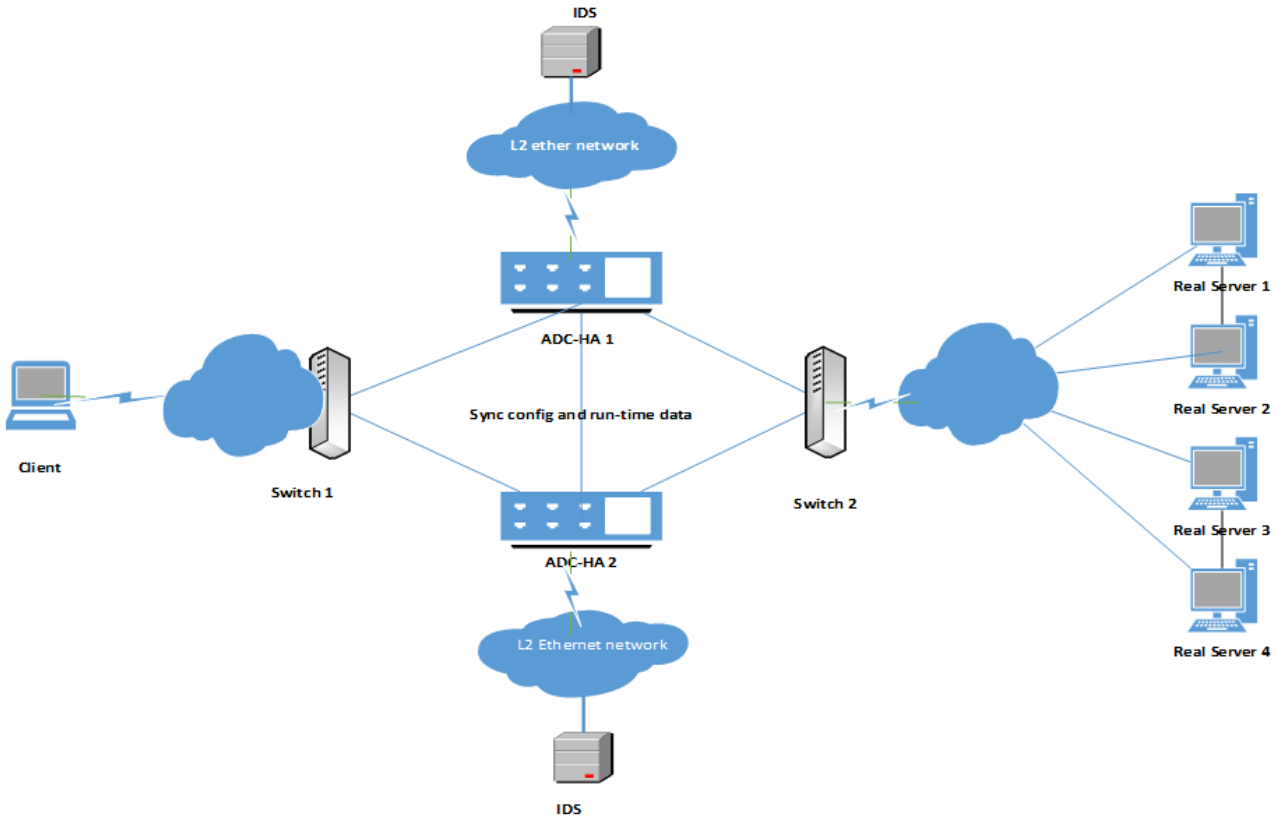
To configure a clone pool, you first create a pool of IDS or sniffer devices and then assign the pool as a clone pool to a virtual server. The clone pool feature is the recommended method for copying production traffic to IDS systems or sniffer devices. Note that when you create the clone pool, the service 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 431 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**.
2. Click the **Clone Pool** tab.
3. Click **Create New** to display the configuration editor.
4. Enter a name for the configuration and click **Save**.
   Once the clone pool configuration is saved, the **Pool Member** section becomes configurable.
5. Under the Pool Member section, click **Create New** to create a member inside your clone pool. You can create up to four members.
6. 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 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.

Entry/Selection	Description
	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.

Load Balancing for Layer 4 Virtual Servers

Overview

In FortiADC, Layer 4 load balancing provides a high-performance method for distributing network traffic across multiple backend servers using transport layer information such as IP addresses and TCP/UDP ports. This approach is crucial for environments requiring speed and low latency.

To deploy Layer 4 load balancing in FortiADC, administrators set up virtual server instances that act as the front-end for client requests. These virtual servers manage incoming traffic according to predefined policies and health checks, ensuring an even distribution of requests and avoiding overloading any single server.

Layer 4 virtual servers support five specific packet forwarding methods each offering unique benefits to meet various network requirements:

- Direct Routing on page 434
- DNAT on page 436
- Full NAT on page 438
- Tunneling on page 439
- NAT46 on page 440

These packet forwarding methods allow FortiADC to provide the flexibility needed to tailor load balancing configurations to specific requirements, optimizing traffic management and enhancing both performance and security.

This topic will cover how to configure Layer 4 virtual servers in FortiADC, including selecting the appropriate packet forwarding method and setting up load balancing policies to ensure a scalable, high-performance application delivery environment.

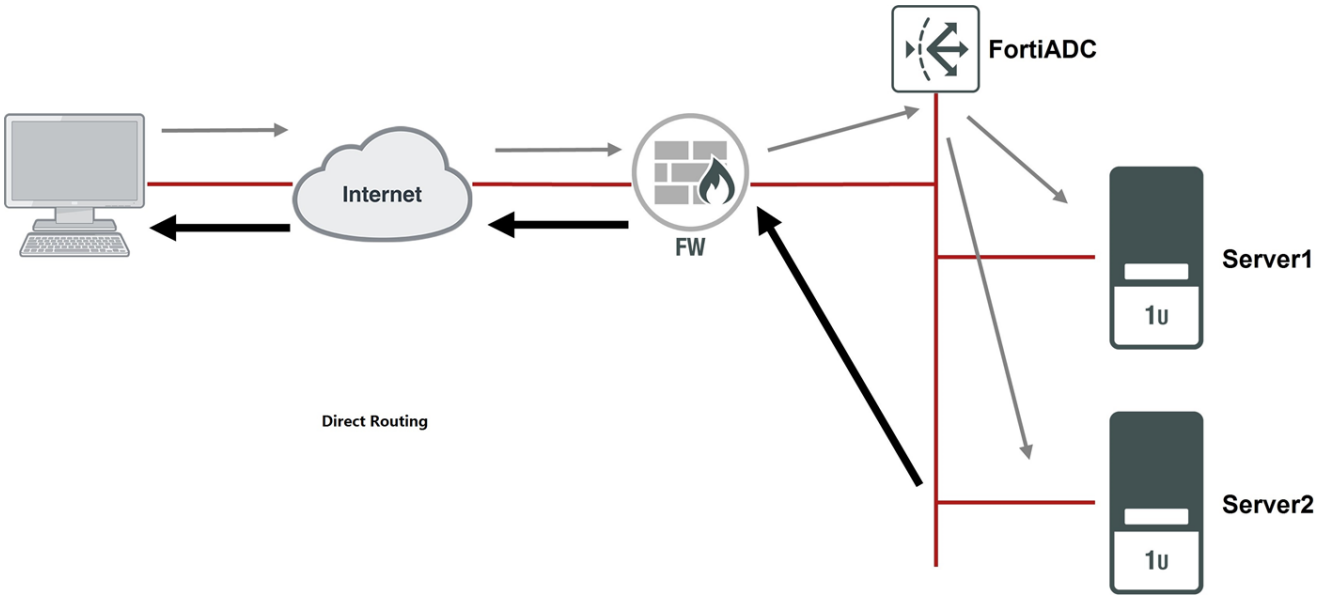
For detailed example of a Layer 4 server load-balancing deployment, see Example: Configuring FortiADC Layer 4 Load Balancing with FullNAT Mode, TCP Profile, and Weighted Round Robin Algorithm on page 441.

Direct Routing

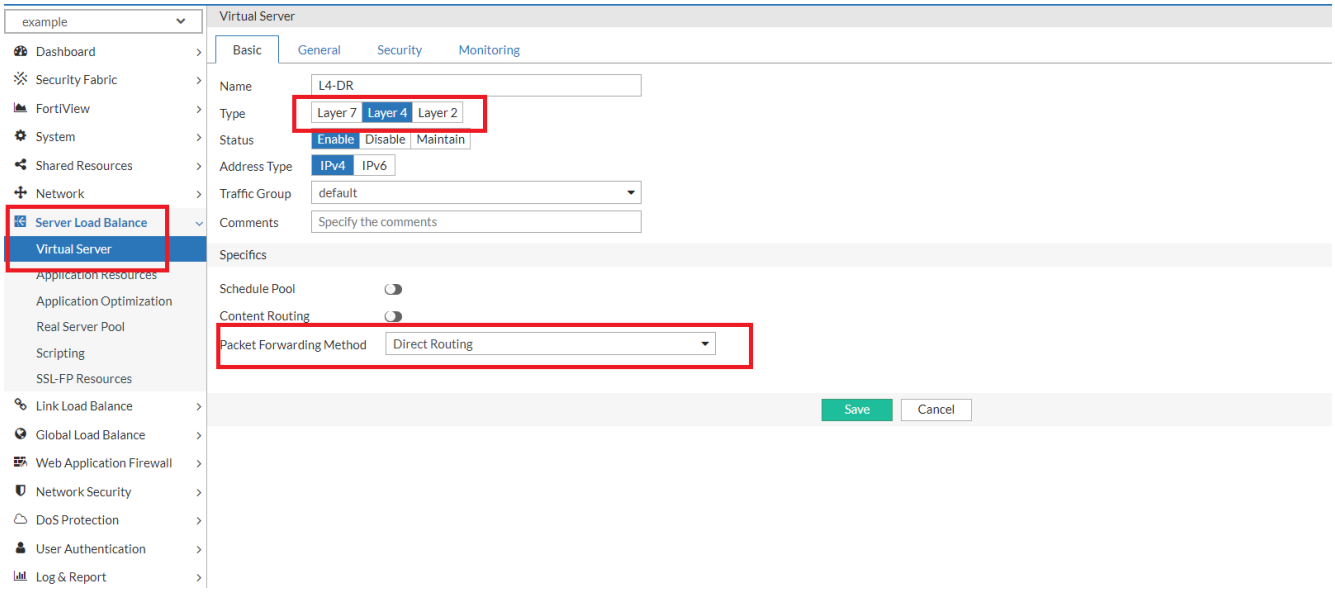
Direct Routing forwards packets directly to backend servers without modifying their IP addresses, preserving the original client IP information. This method enhances performance by minimizing latency and is useful for applications that require client IP visibility.

In Direct Routing mode, FortiADC forwards packets to backend servers by altering the destination MAC address to that of the selected real server. This method operates transparently, allowing the real server to see the original source IP address of the client. As a result, the Direct Routing mode preserves the client's IP address, making it ideal for scenarios where the backend servers need to maintain visibility into the originating client.

Topology



GUI



When using Direct Routing mode, the real server expects to handle traffic directed to both its own IP address and the virtual server IP (VIP) address. As such, both the real server and the load-balanced application must be configured to respond to requests sent to these addresses.

- **Subnet Configuration** — FortiADC must have an interface in the same subnet as the real servers to ensure the Layer 2 connectivity necessary for Direct Routing mode to function correctly.
- **VIP Placement** — The virtual IP (VIP) can be located on the same subnet as the real servers or on a different subnet, provided that FortiADC has an interface in the subnet where the VIP resides.
- **Port Translation Limitation** — Direct Routing mode does not support port translation; the port used by the real server must match the port of the virtual IP (VIP).

Example: Configure Direct Routing Mode for a Windows Server

To enable Direct Routing mode on a Windows Server, you need to configure a loopback adapter and assign the virtual server IP (VIP) to this adapter. This configuration ensures proper traffic routing and IP handling for Layer 4 load balancing. Follow the steps detailed below.

Install the Loopback Adapter:

1. Open a command prompt with administrative privileges.
   a. Click **Start**, type **cmd** into the search box.
   b. Right-click **cmd-exe**, and select **Run as administrator**.
2. In the command prompt, execute **hdwwiz.exe** and press **Enter** to launch the Hardware Wizard.
3. Proceed with the installation.
   a. Click **Next** to start the hardware installation wizard.
   b. Select **Install the hardware that I manually select from a list (Advanced)** and click **Next**.
4. Select the **Network adapters** from the list and click **Next**.
5. Install the Loopback Adapter.
   a. From the list of manufacturers, select **Microsoft**.
   b. Select **Microsoft KM-TEST Loopback Adapter** (suitable for Windows 10) and click **Next**.
6. Confirm and complete the installation.
   a. Click **Next** to proceed with the installation of the loopback adapter.
   b. Click **Finish** once the installation is complete.

Configure the Loopback Adapter:

1. Assign the virtual server IP (VIP) to the loopback adapter.
   a. Open **Network and Sharing Center** and navigate to **Change adapter settings**.
   b. Locate the newly added loopback adapter, right-click it, and select **Properties**.
   c. Set the VIP as the static IP address for the loopback adapter.
2. Rename network interfaces.
   a. Rename the interface connected to FortiADC to **nic_to_adc**.
   b. Rename the newly added loopback adapter to **loopback**.

Configure Interface Settings:

1. Enable weak host receive and send settings to support Direct Routing.
   a. Open a command prompt with administrative privileges.
   b. Execute the following commands to configure network interface settings.

```
netsh interface ipv4 set interface "nic_to_adc" weakhostreceive=enabled
netsh interface ipv4 set interface "loopback" weakhostreceive=enabled
netsh interface ipv4 set interface "loopback" weakhostsend=enabled
```

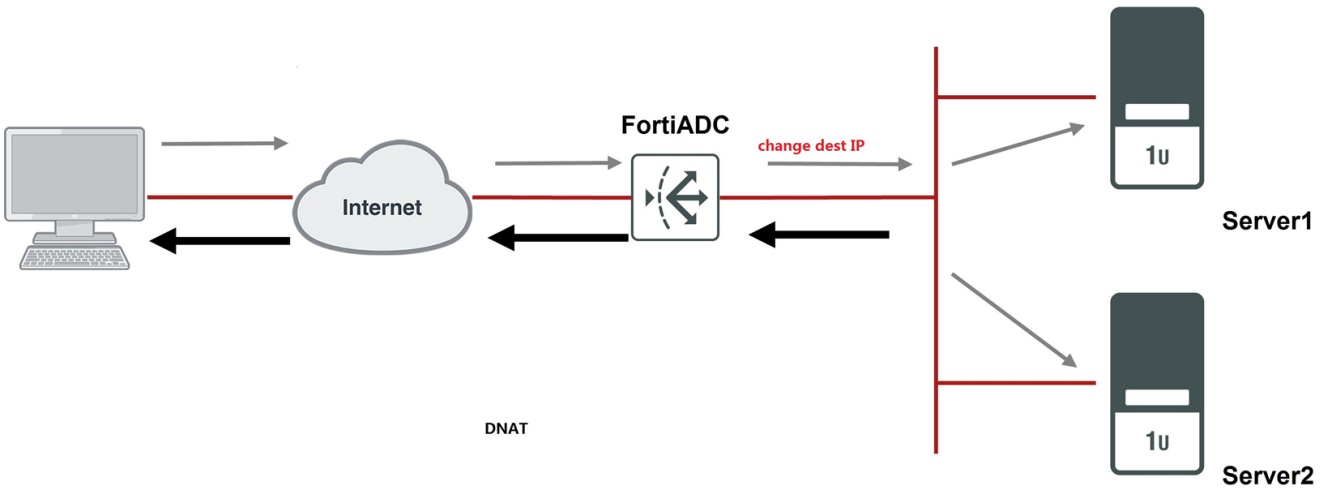
DNAT

DNAT (Destination Network Address Translation) modifies the destination IP address of incoming packets to direct traffic to backend servers, effectively hiding the server's internal IP addresses from external clients, thus enhancing

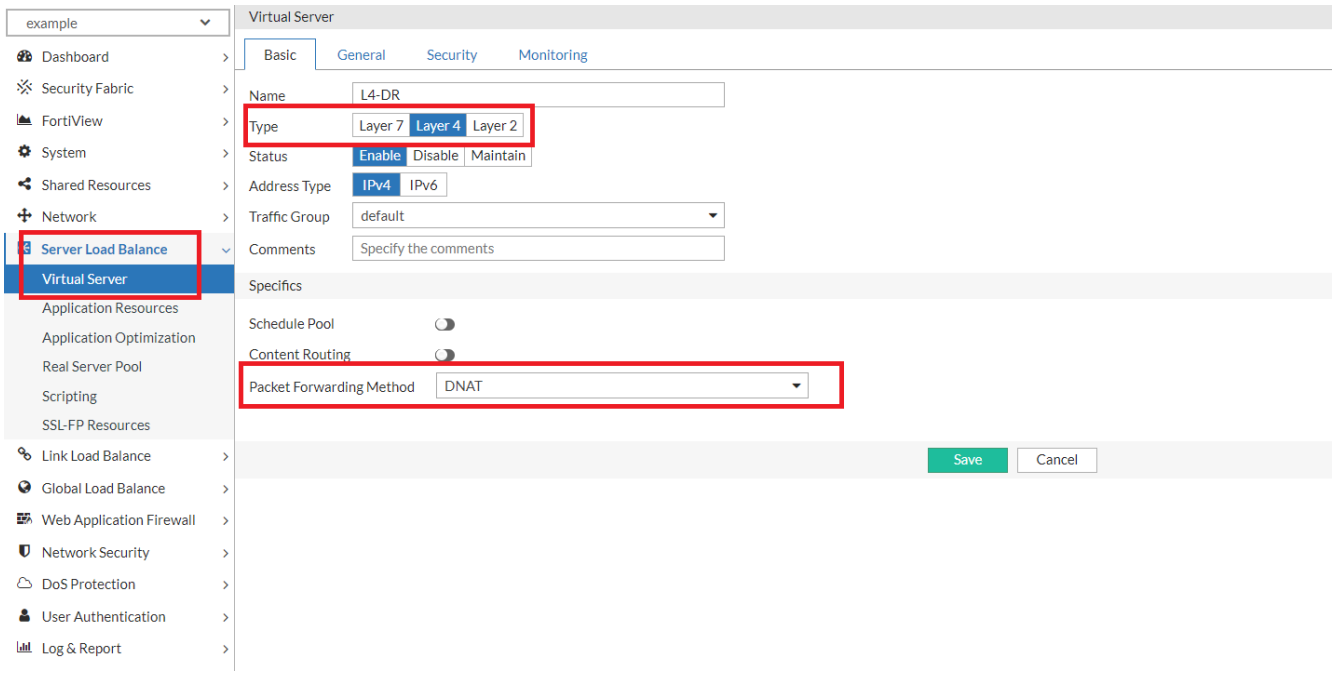
security.

Typically, Layer 4 DNAT in FortiADC requires two interfaces: one for connecting to the clients and another for connecting to the real servers. When a packet passes through the FortiADC virtual server, its destination IP address is modified to direct traffic to the appropriate backend server.

Topology



GUI



When using DNAT as the packet forwarding method in FortiADC, it is essential to ensure that response traffic from the backend servers returns through the FortiADC device. By setting the default gateway of each server to FortiADC's IP address on the same subnet or VLAN, you establish a direct path for all outbound traffic. This configuration allows FortiADC to perform the necessary address translation and maintain session consistency.

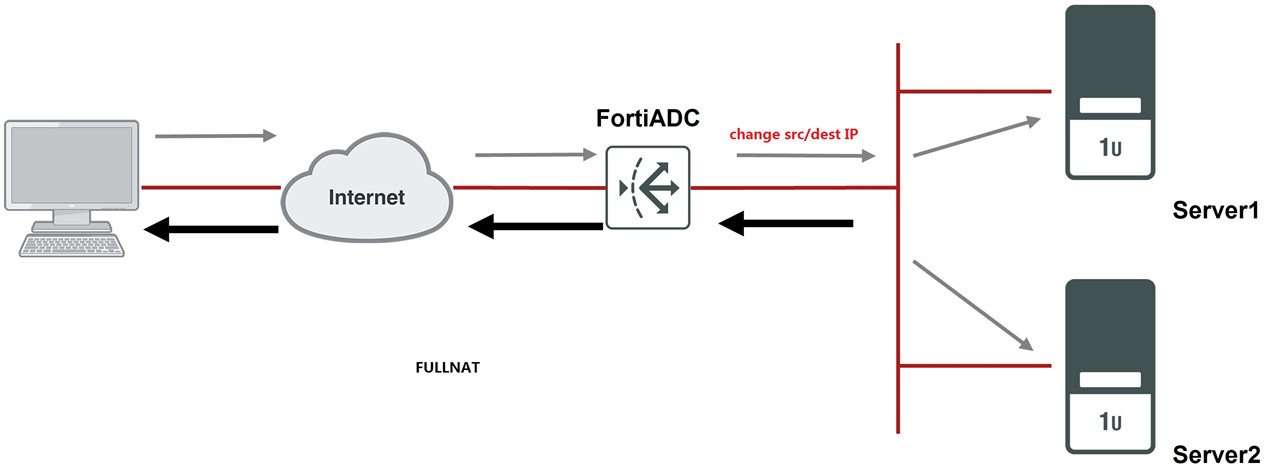
Alternatively, using static routes on the servers to direct responses to FortiADC's IP address ensures that return packets are sent back through the load balancer. This approach is particularly useful in complex network environments where multiple gateways might exist. By routing responses through FortiADC, the device can correctly reverse the DNAT operation, preserving the client-server communication flow and maintaining proper traffic management and security policies.

Full NAT

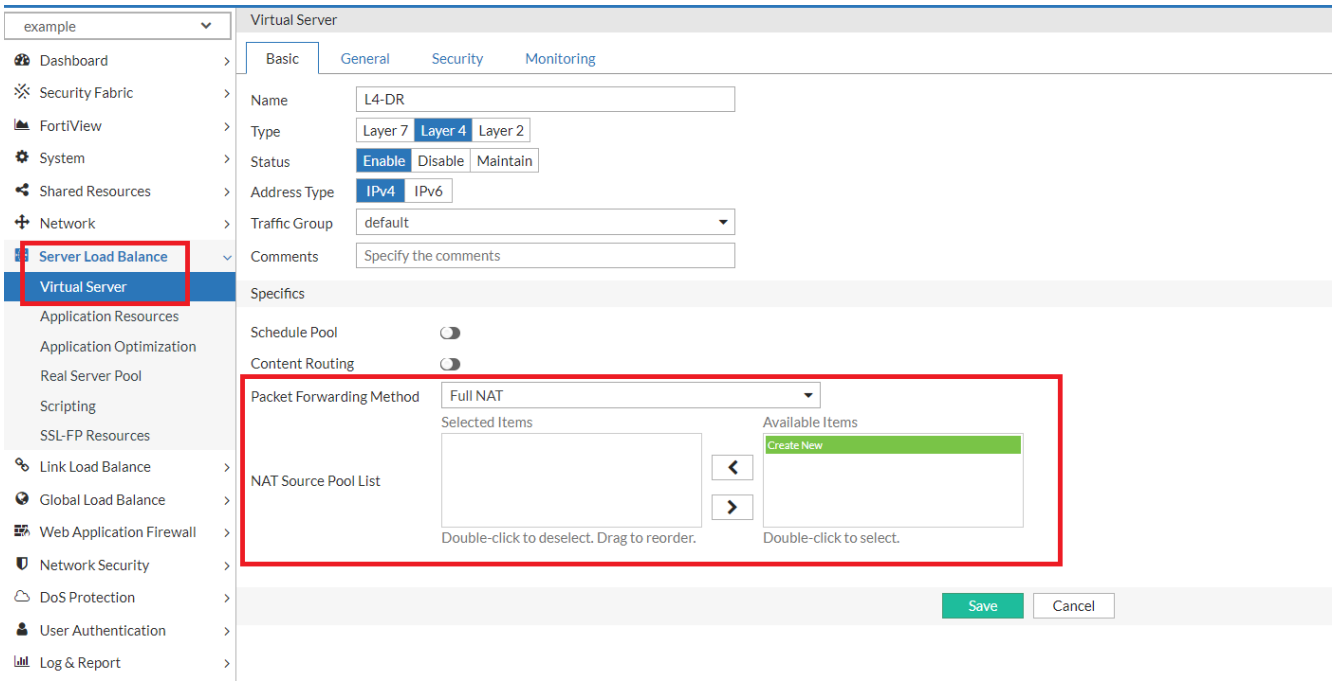
Full NAT translates both the source and destination IP addresses of packets. This method is ideal for complex network environments where complete address translation is necessary for communication between clients and servers across different subnets.

In FortiADC, a Layer 4 Full NAT virtual server modifies both the source and destination IP addresses of packets before forwarding them to the real servers. Users can define a custom range of IP addresses for the NAT source pool and select this range in the Pool List. Typically, the NAT source pool's address range is configured within the same network subnet as the real servers.

Topology



GUI

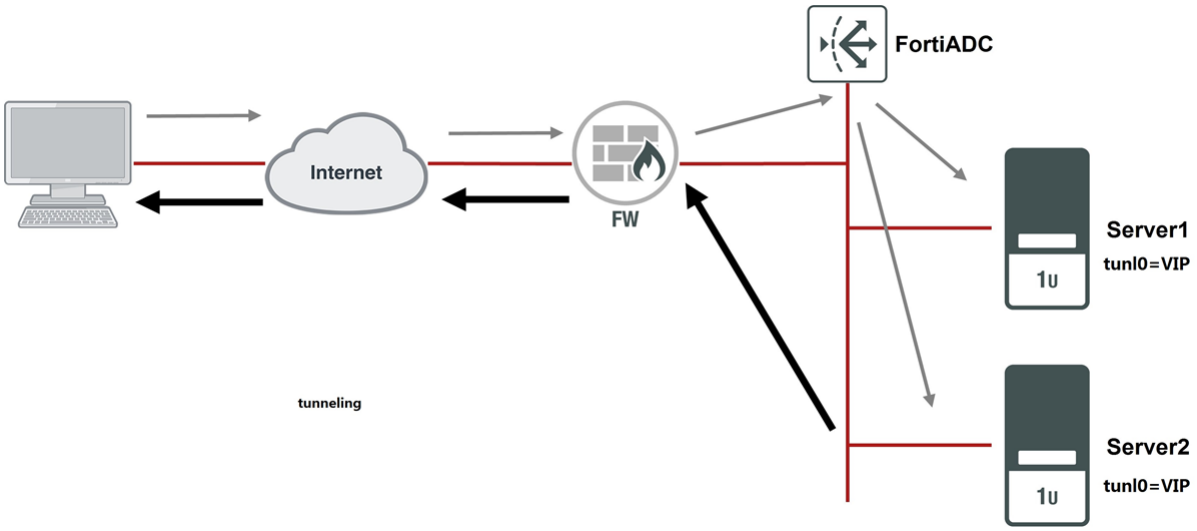


Tunneling

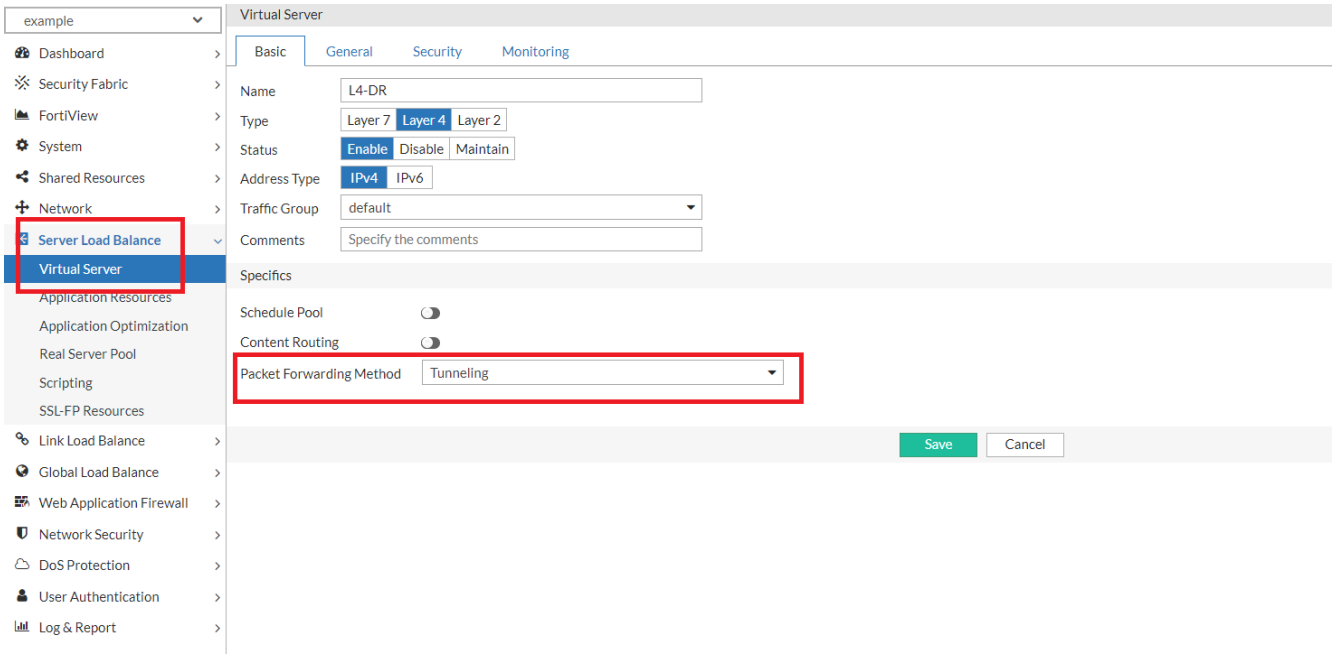
Tunneling encapsulates the original packets within a new packet before forwarding them to backend servers. This method supports scenarios where traffic needs to traverse different networks or requires additional encapsulation for security or protocol compatibility.

In Tunneling mode, a virtual server in FortiADC operates similarly to Direct Routing but adds an additional layer of encapsulation. The FortiADC virtual server encapsulates the original packet, containing the client IP and virtual server IP, within an IP-in-IP (IPIP) packet. The outer IPIP packet uses the FortiADC's IP address as the source and the real server's IP address as the destination. This encapsulated packet is then placed in an output chain and routed to the real server. Upon receiving the packet on a "tunl0" device, the real server decapsulates the IPIP packet, extracting the original packet with the client IP and virtual server IP. The real server then processes this packet and sends the response back to the client.

Topology



GUI

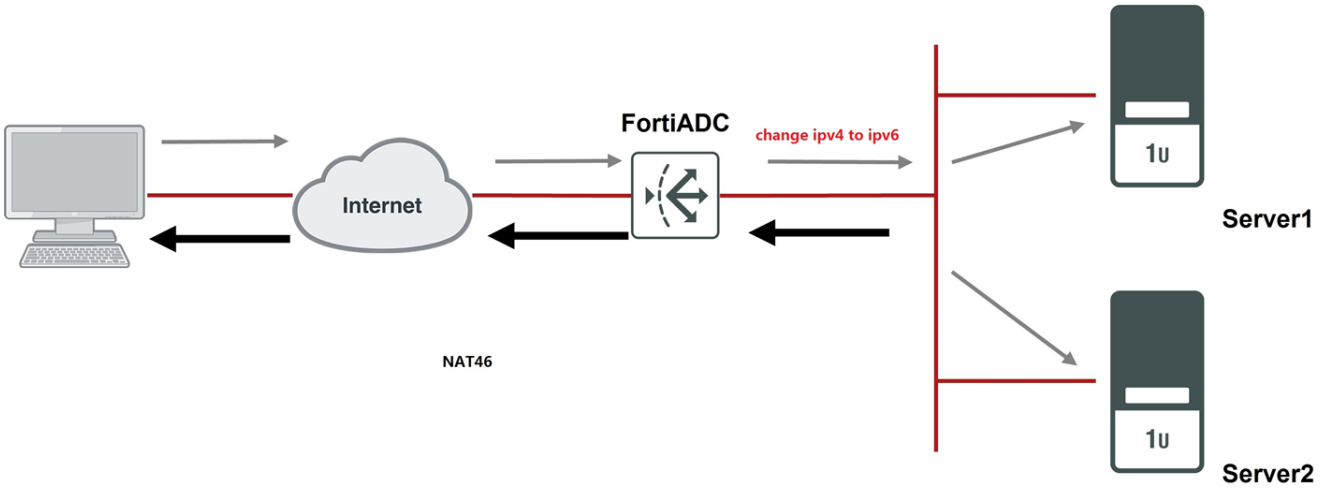


NAT46

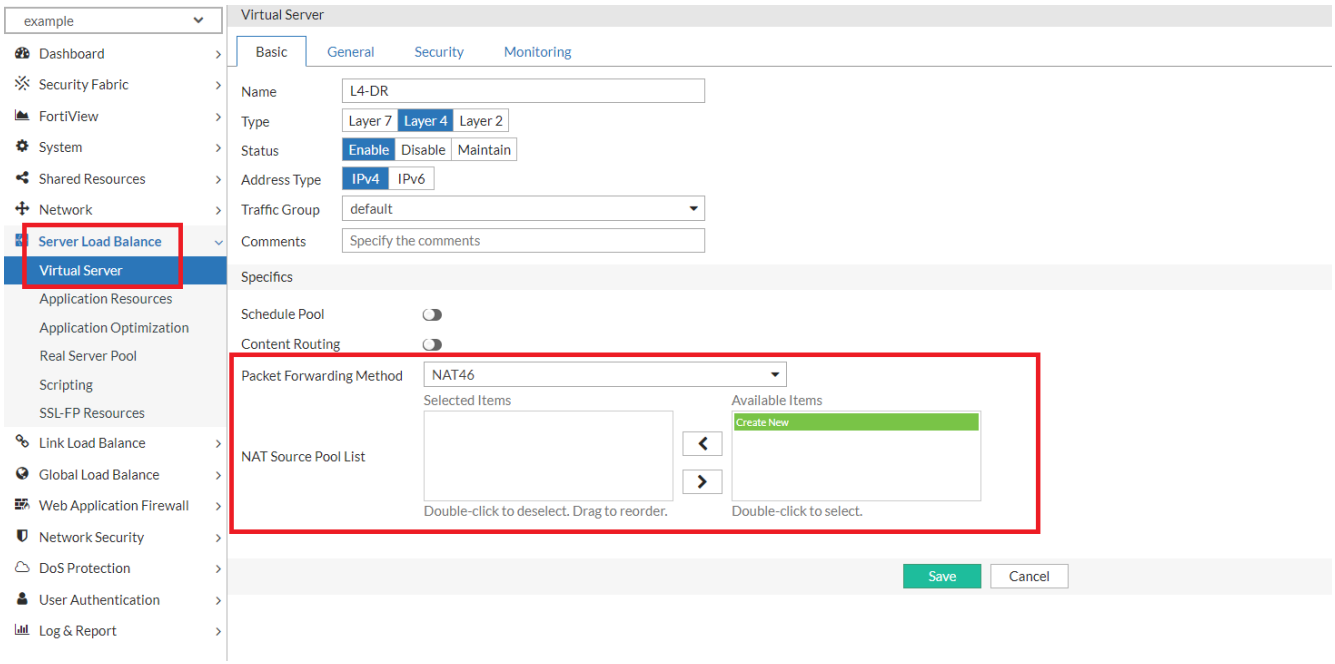
NAT46 (Network Address Translation for IPv4-to-IPv6) handles the translation between IPv4 and IPv6 traffic, allowing seamless communication between networks that use different IP versions. This method is crucial for environments transitioning to IPv6 or interacting with systems that use mixed IP versions.

In NAT46 mode, the virtual server converts the source address of packets from IPv4 to IPv6 using the IP addresses specified in the NAT source pool. The translated packets are then forwarded to the IPv6 real servers for processing.

Topology

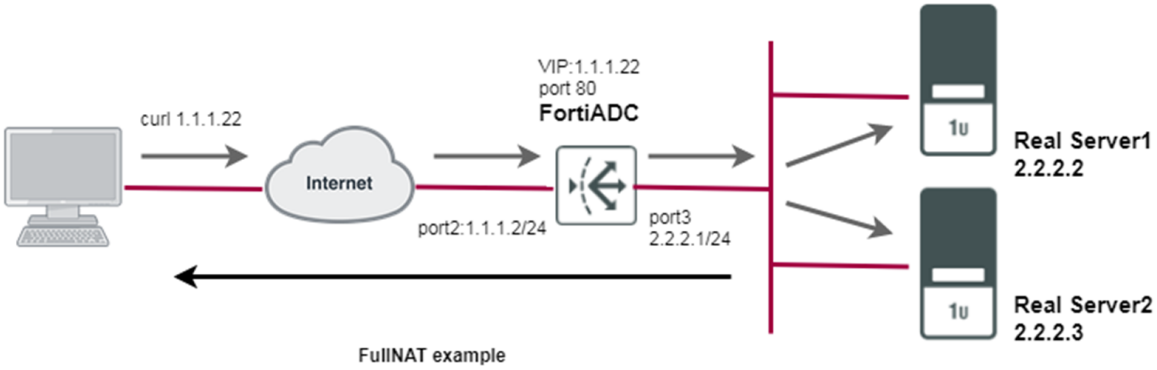


GUI



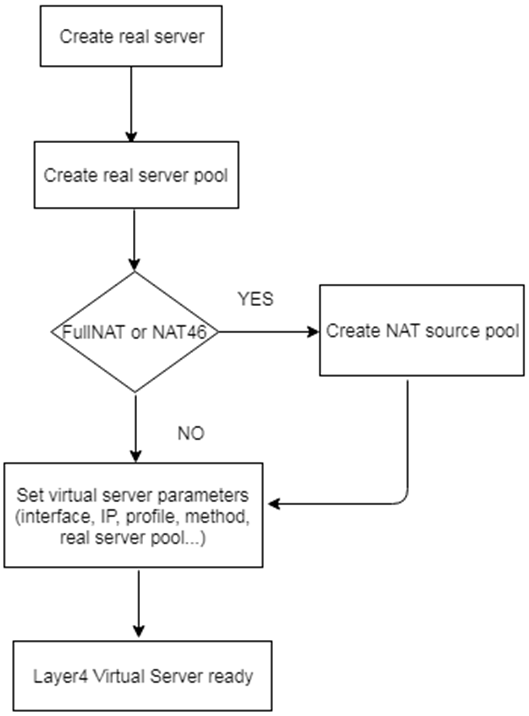
Example: Configuring FortiADC Layer 4 Load Balancing with FullNAT Mode, TCP Profile, and Weighted Round Robin Algorithm

This section demonstrates the configuration of a FortiADC appliance for Layer 4 load balancing using FullNAT mode. The setup employs a TCP profile to manage connection parameters and utilizes the Weighted Round Robin algorithm to efficiently distribute traffic across backend servers. This configuration exemplifies a robust approach to optimizing traffic flow and enhancing the reliability of networked services.



FullNAT example

Deployment workflow



basic Layer4 VS steps

The workflow above outlines the essential steps for deploying a Layer 4 virtual server for load balancing:

Step 1: Create the Real Server on page 443 — Define the backend servers that will receive traffic from the virtual server.

Step 2: Create the Real Server Pool on page 443 — Group the real servers into a pool that the virtual server will use for distributing traffic.

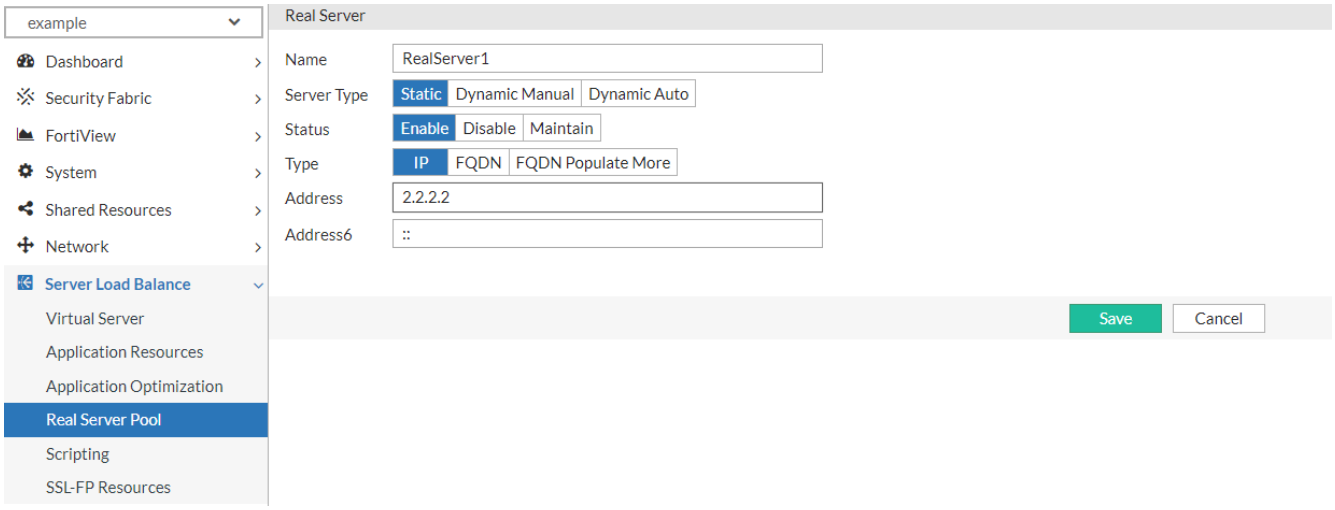
Step 3: Create a NAT Source Pool (if needed for FullNAT or NAT46) on page 445 — If using FullNAT or NAT46 as the packet forwarding method, define a NAT source pool to handle source IP address translation.

Step 4: Configure the Virtual Server Parameters on page 445 — Set up the virtual server by specifying its interface, IP address, profile, load-balancing method, and associated real server pool.

Step 5: Deploy the Layer 4 Virtual Server on page 447 — With all components configured, the Layer 4 virtual server is ready for use.

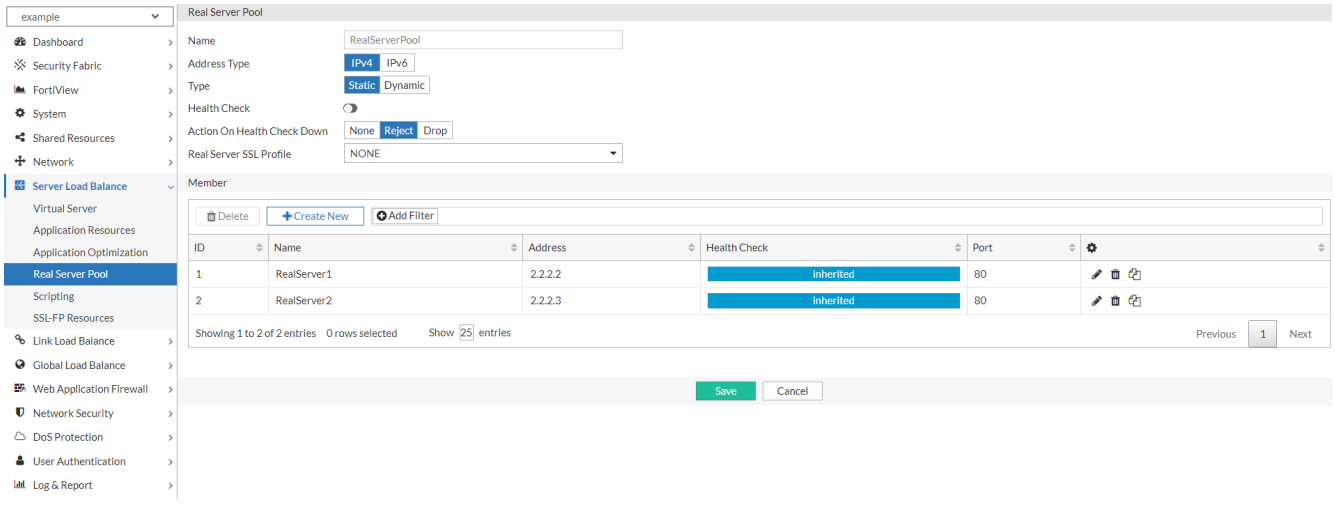
Step 1: Create the Real Server

In this step, you define the backend servers that will handle traffic from the virtual server. This involves specifying each server's IP address, hostname, and port number, as well as setting up health checks to monitor server availability. Proper configuration ensures that only healthy servers receive traffic, maintaining service reliability.



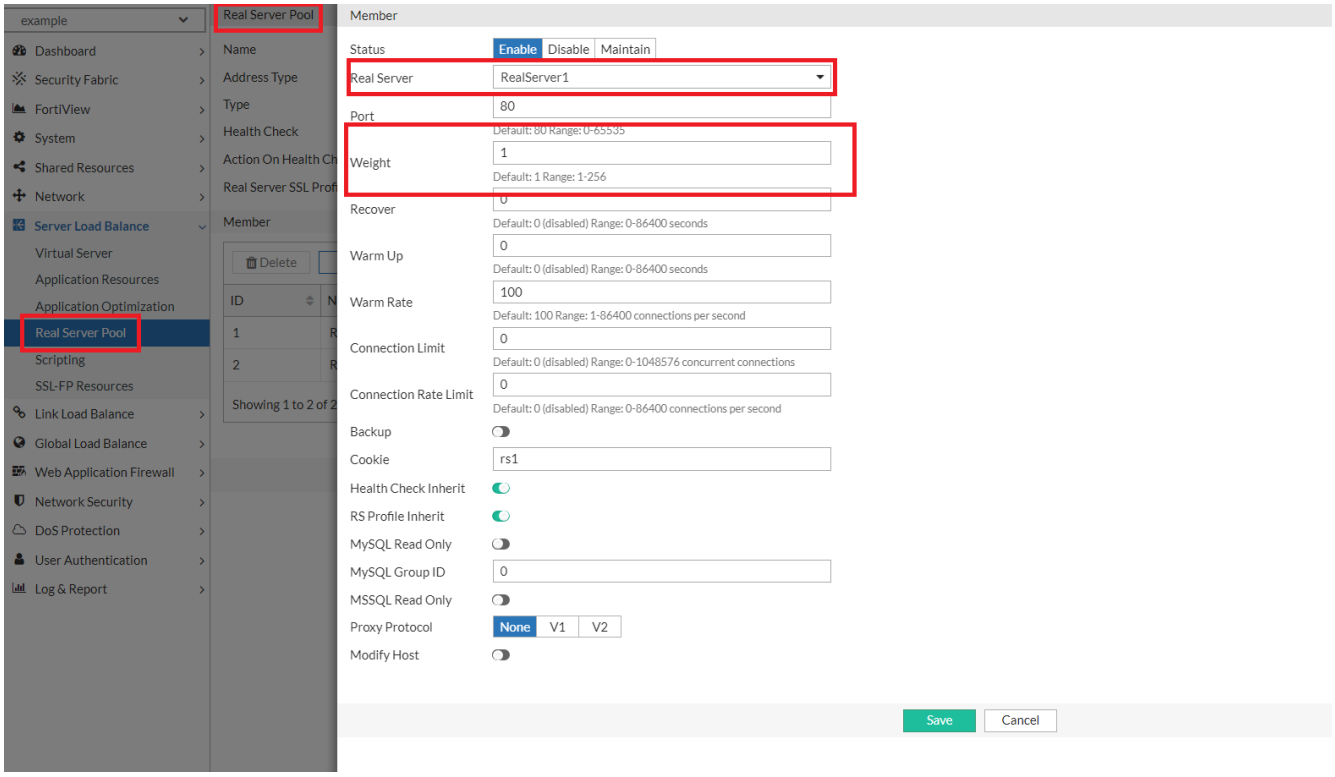
Step 2: Create the Real Server Pool

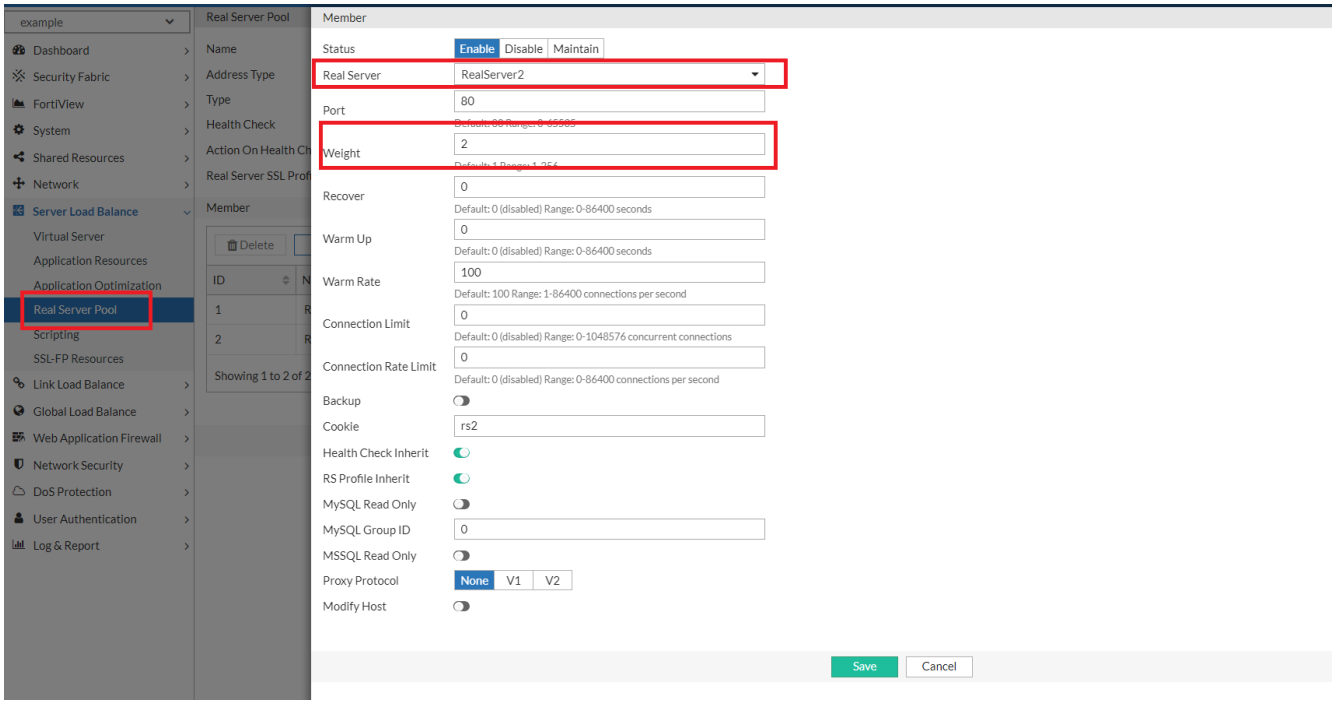
Next, you group the defined real servers into a pool that the virtual server will use to distribute traffic. This pool enables efficient traffic management by applying a load-balancing algorithm, such as round-robin or least connections. Optional persistence settings can be added to keep client sessions consistent, ensuring a smooth user experience.



Optionally, you can set the real server's weight.
To assign different weights to the real servers, adjust the weight settings in the Real Server Pool. For example, if

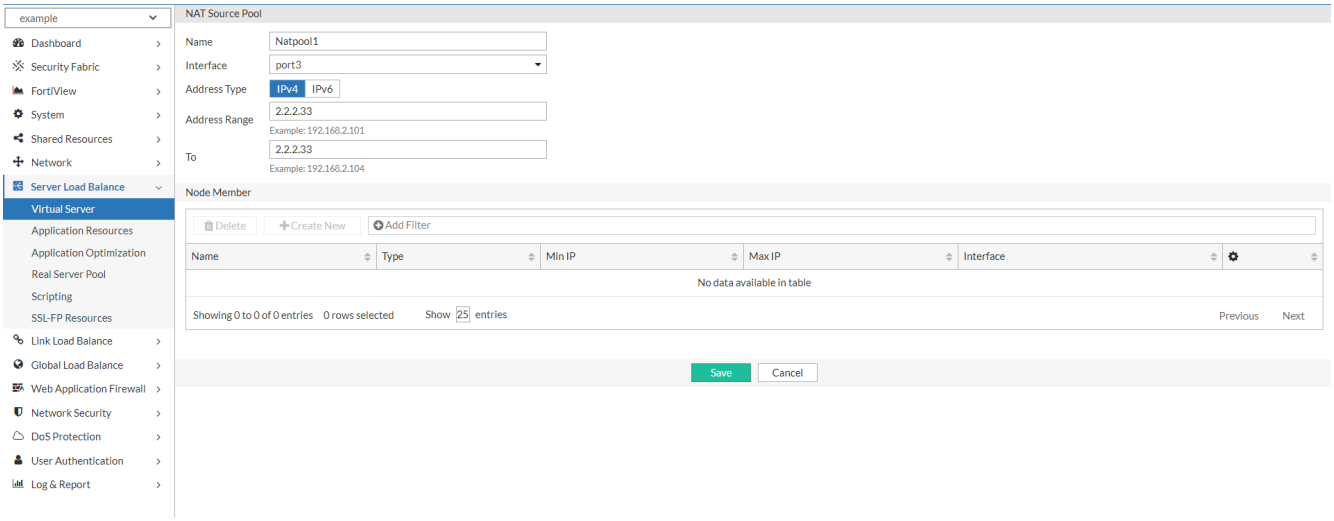
RealServer1 is assigned a weight of 1 and RealServer2 a weight of 2, the distribution of connections will reflect these ratios. In a scenario where there are 30 total connections, RealServer1 will handle 10 connections, while RealServer2 will manage 20 connections. This weighted distribution allows for more granular control over how traffic is allocated among your servers, ensuring that resources are utilized according to their capacity.





Step 3: Create a NAT Source Pool (if needed for FullNAT or NAT46)

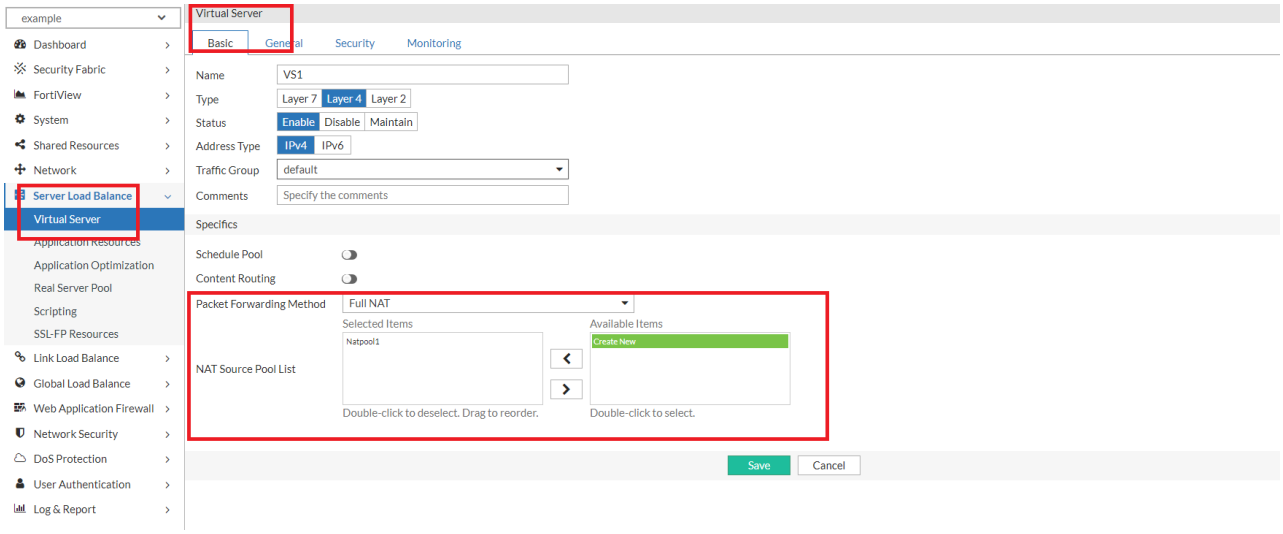
If your setup requires FullNAT or NAT46, you need to configure a NAT source pool. This pool handles source IP address translation, allowing packets to be properly routed back to clients. Defining how IP addresses are assigned and monitored is essential for maintaining consistent communication across different networks or IP versions.



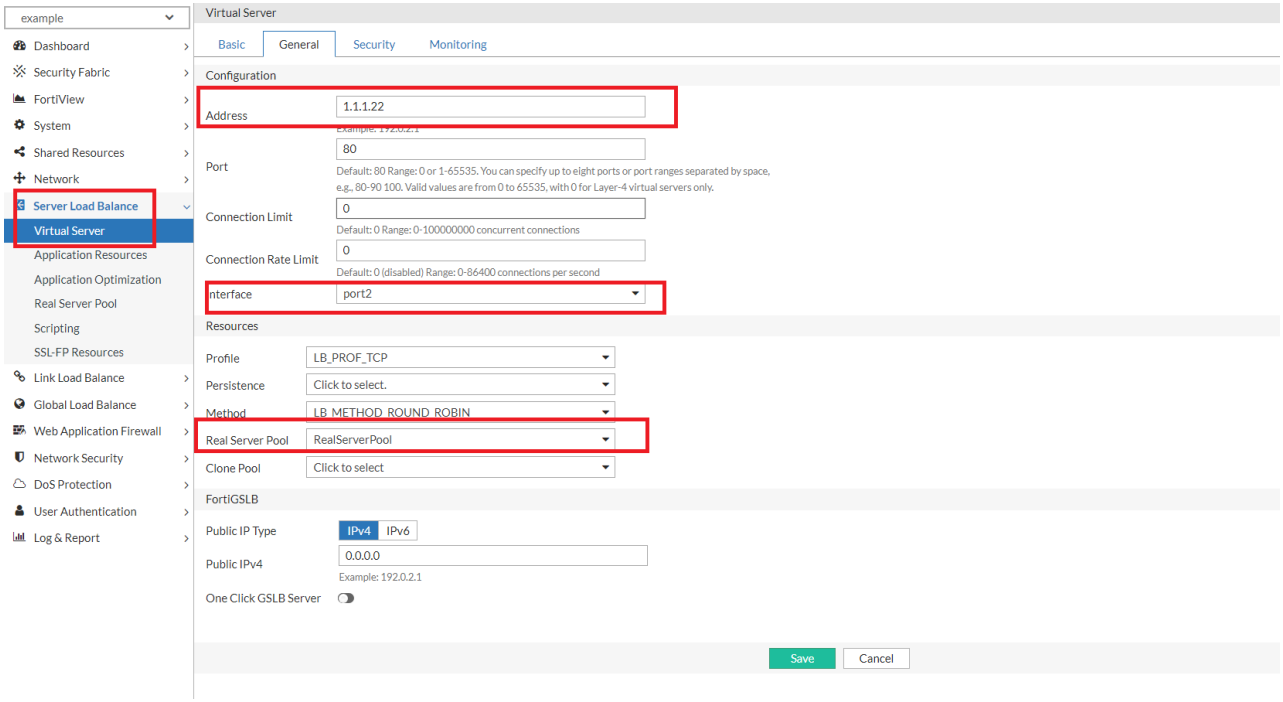
Step 4: Configure the Virtual Server Parameters

In this phase, you configure the virtual server by assigning it an IP address, selecting a protocol, and linking it to the real server pool. Additional settings, like security profiles and load-balancing methods, are applied to ensure the virtual server distributes traffic efficiently and securely to the backend servers.

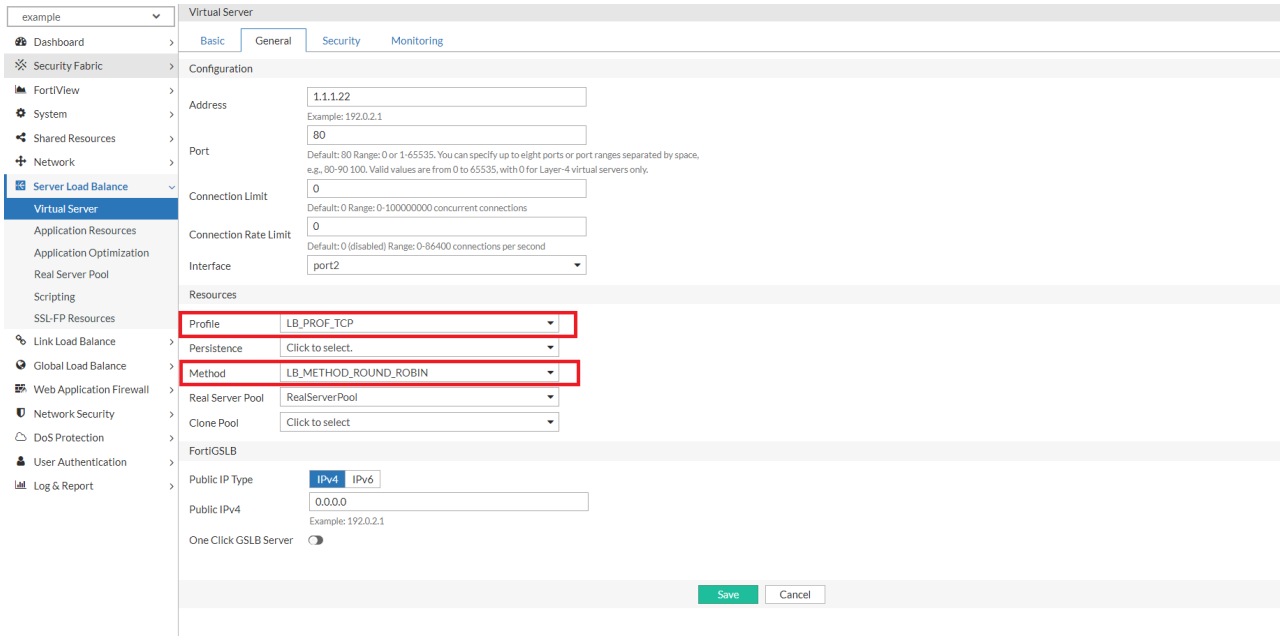
1. Create a FullNAT virtual server.



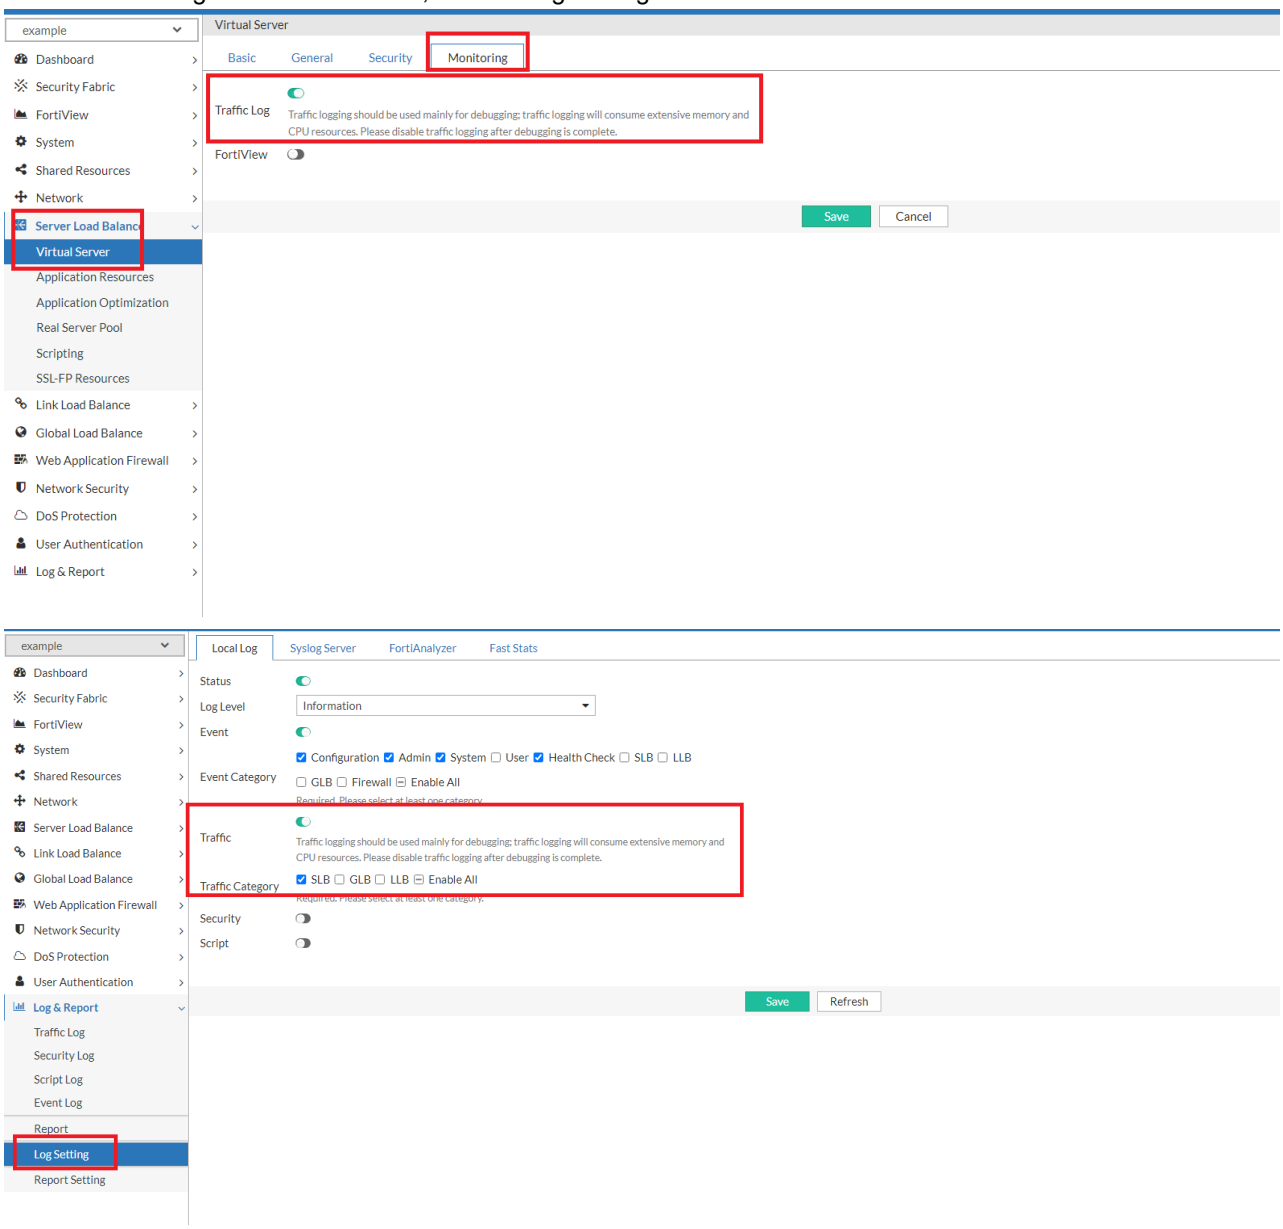
2. Select the real server pool previously created.



3. Set the **Profile** as **LB_PROF_TCP** (or any Application Profile object with the TCP type) and the **Method** as **LB_METHOD_ROUND_ROBIN**.



Enable Traffic Log in the virtual server, and the Log Settings for SLB.



Step 5: Deploy the Layer 4 Virtual Server

Finally, with all components in place, you deploy the Layer 4 virtual server. This step involves applying the configuration to the load balancer and conducting tests to ensure traffic is correctly routed to the backend servers. Ongoing monitoring is essential to maintain optimal performance and quickly address any issues that arise.

1. Generate and test traffic from the client by utilizing command-line tools such as "curl" or "wget". These tools allow you to perform HTTP requests, enabling you to validate connectivity, measure response times, and ensure the correct distribution of traffic across the configured servers.

```
[root@nxcentos ~]# wget http://1.1.1.22/big.iso
--2024-07-10 10:02:10--  http://1.1.1.22/big.iso
Connecting to 1.1.1.22:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 869269504 (829M) [application/octet-stream]
Saving to: 'big.iso.25'

 0% [                                                                    ] 1,376,256   20.0KB/s  eta 11h 45m
```

```
[root@nxcentos ~]# curl 1.1.1.22

2.2.2.3
```

2. To verify the effectiveness of the Weighted Round Robin (WRR) method, send three requests to the virtual server. You should observe that two responses are returned from RealServer2 and one from RealServer1, reflecting the configured weight distribution. This confirms that the WRR algorithm is correctly balancing the traffic according to the assigned weights.

```
[root@nxcentos ~]# curl 1.1.1.22

2.2.2.2

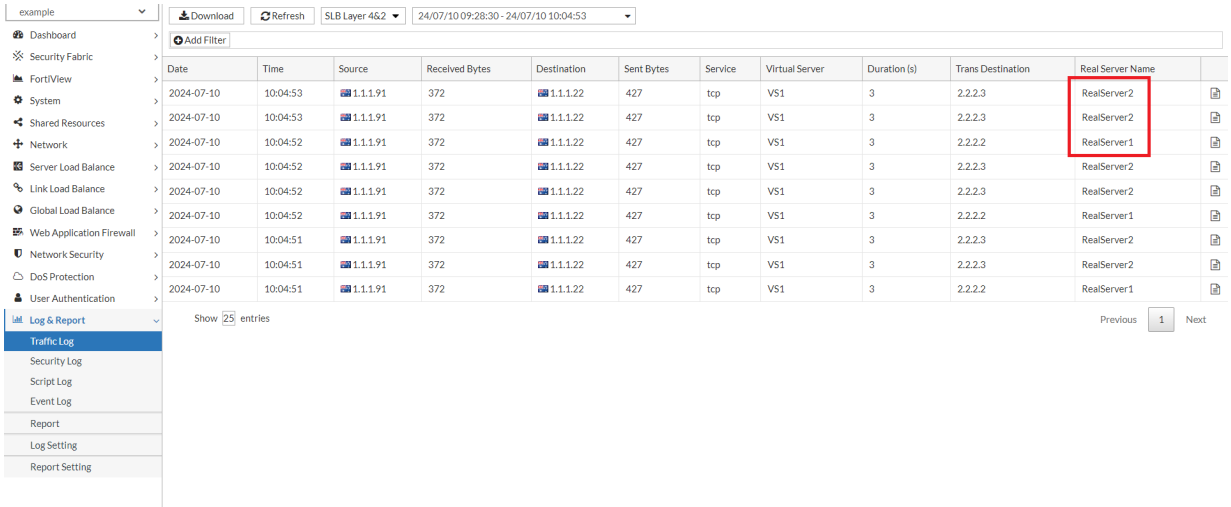

[root@nxcentos ~]# curl 1.1.1.22

2.2.2.3


[root@nxcentos ~]# curl 1.1.1.22

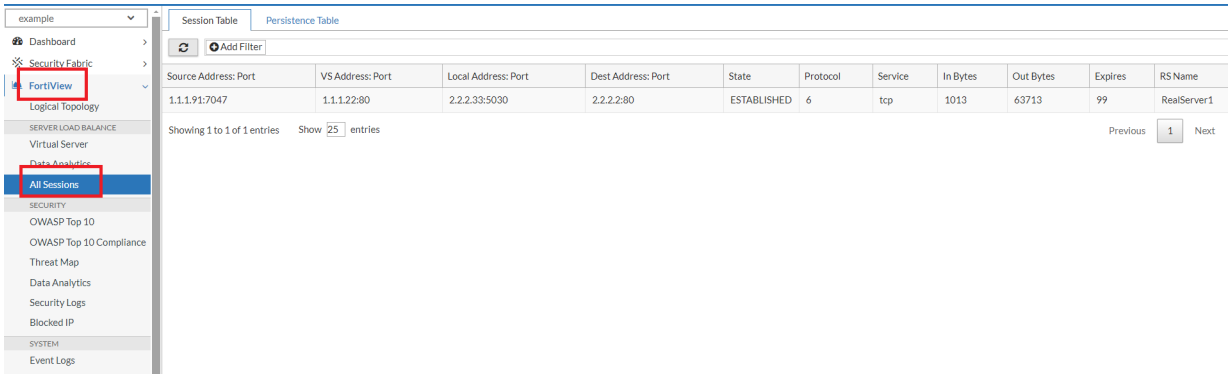
2.2.2.3
```

3. Review the traffic logs and session table data to verify the accuracy of traffic distribution and session persistence. This will help ensure that the configured load-balancing settings are functioning as expected and provide insights into connection patterns and server utilization.
   a. Examine the traffic logs to verify that the Weighted Round Robin (WRR) method is functioning as expected. The logs should reflect the proportional distribution of traffic in line with the configured server weights, confirming the correct operation of the WRR algorithm.
   The traffic logs confirm that traffic is being directed to the real servers according to their assigned weights, with a distribution ratio of 1:2 between RS1 and RS2.



b. Check the session table to verify that active connections are distributed according to the configured load-balancing settings. This provides real-time confirmation that the Weighted Round Robin (WRR) algorithm is correctly allocating sessions across the real servers.

FortiView:



CLI:

```
FortiADC (example) # diagnose server-load-balance session list
client-ip/port virtual-server-ip/port local-ip/port real-server-ip/port protocol service state in-bytes out-bytes expire virtual-server-name real-server-name
1.1.1.91 7047 1.1.1.22 80 2.2.2.33 5030 2.2.2.2 80 6 tcp ESTABLISHED 25613 2643905 99 VS1 RealServer1
```

Application Resources

The Application Resources module contains configuration objects that define the Virtual Server settings for application delivery control, such as the protocol settings, session persistence, and load-balancing methods to use to distribute traffic.

The **Server Load Balance > Application Resources** sub-menu contains the following configurations:

- **Application Profile** — The Application Profile configuration object plays a crucial role by defining how the FortiADC virtual server should handle traffic for specific protocols. Each Application Profile type is based on specific protocols that can only operate at certain OSI layers, which dictates the virtual server type, load-balance method, and persistence methods that can be used.
  For details, see Configuring Application profiles on page 451.
- **Client SSL** — The Client SSL profile is used to manage the SSL session between the client and the proxy, which allows the FortiADC to accept and terminate client requests sent via the SSL protocol. For details, see Configuring Client SSL profiles on page 512.
- **HTTP2 Profile** — For HTTP/S Application Profiles, you can enable it to support HTTP/2 by specifying an HTTP2 Profile. The HTTP/2 protocol uses binary encoding rather than text encoding, making it more efficient, and allowing it to support functionality such as server push, multiplexing of requests, and header compression.
  For details, see Configuring HTTP2 profiles on page 518.
- **HTTP3 Profile** — For HTTP/S Application Profiles, you can enable it to support HTTP/3 by specifying an HTTP3 Profile. HTTP/3 is the latest version of the HTTP protocol and unlike previous versions which relied on TCP to handle streams in the HTTP layer, HTTP/3 uses QUIC (Quick UDP Internet Connections), a multiplexed transport protocol built on UDP. The HTTP/3 protocol has a lower latency as a result of using QUIC, allowing it to have a quicker handshake for establishing a secure session compared to HTTP/2 which achieves this using TCP and TLS.
  For details, see Configuring HTTP3 profiles on page 521.
- **LB Method** — The LB Method configuration object defines the load-balancing algorithm, types, or techniques to use for distributing traffic, which include commonly used methods such as Round-robin, IP hash, and Least connection.
  For details, see Configuring load-balancing (LB) methods on page 526.
- **Persistence** — The Persistence configuration object allows you to define persistence rules to identify traffic that should not be load-balanced, but instead be forwarded to the same backend server based on specific session attributes. The persistence rule is then evaluated before load-balancing rules, so that when packets received by FortiADC matches an entry in the persistence session, the packets are then forwarded to the server that established the connection and load-balancing rules will not be applicable.
  For details, see Configuring persistence rules on page 527.
- **Error Page** — Applicable to HTTP/S load-balancing, the Error Page configuration allows you to customize an HTML error page to respond to clients attempting HTTP/S connections when backend real servers are unavailable.
  For details, see Configuring error pages on page 535.
- **Decompression** — Applicable to HTTP/S load-balancing, Decompression rules allow FortiADC to decompress the HTTP response traffic that was previously compressed to allow modules to perform inspection or modification before passing to the backend server.
  For details, see Configuring decompression rules on page 536.
- **Captcha** — FortiADC allows administrators to validate incoming users with CAPTCHAs to determine whether a client is a regular user or an attacker. The Captcha configuration allows you to customize a CAPTCHA page to use your own verifications in the event a WAF or DoS attack is triggered.
  For details, see Configuring Captcha on page 541.

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. Each Application Profile Type is compatible with specific virtual server types, load-balancing methods, and content routing types. To help you get started, FortiADC offers predefined profile configurations you can apply directly in the virtual server configuration, or use as a template to create a user-defined profile.

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.

Configuration Overview

The Application Profile configuration workflow is entirely dependent on the specific protocol. Each protocol, or Profile Type, requires its own set of parameters to be configured. The steps below outlines the basic workflow to configure a custom Application Profile.

1. Go to **Server Load Balance > Application Resources**.
   The configuration page displays the **Application Profile** tab.
2. Click **Create New** to display the configuration editor.
3. In the **Name** field, specify a unique name for the custom Application Profile configuration object. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No space is allowed.
   Once saved, the name of a Application Profile configuration cannot be changed.
4. From the **Type** field, select a protocol from the drop-down menu. This selection determines the parameters to configure.
   To view the configuration parameters for each protocol type, see Application Profile Types on page 451 where you can navigate to the desired Profile Type.
5. Click **Save** to save the configuration.



All values in the predefined profiles are view-only, and cannot be modified. However, 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.

Application Profile Types

Each Application Profile Type requires specific parameters to be configured. FortiADC supports the following table lists of Application Profile Types and their usage and compatibilities. Use the links under the **Profile Type** column to navigate to the configuration parameters applicable to each Application Profile Type and view its available predefined Profile configuration.

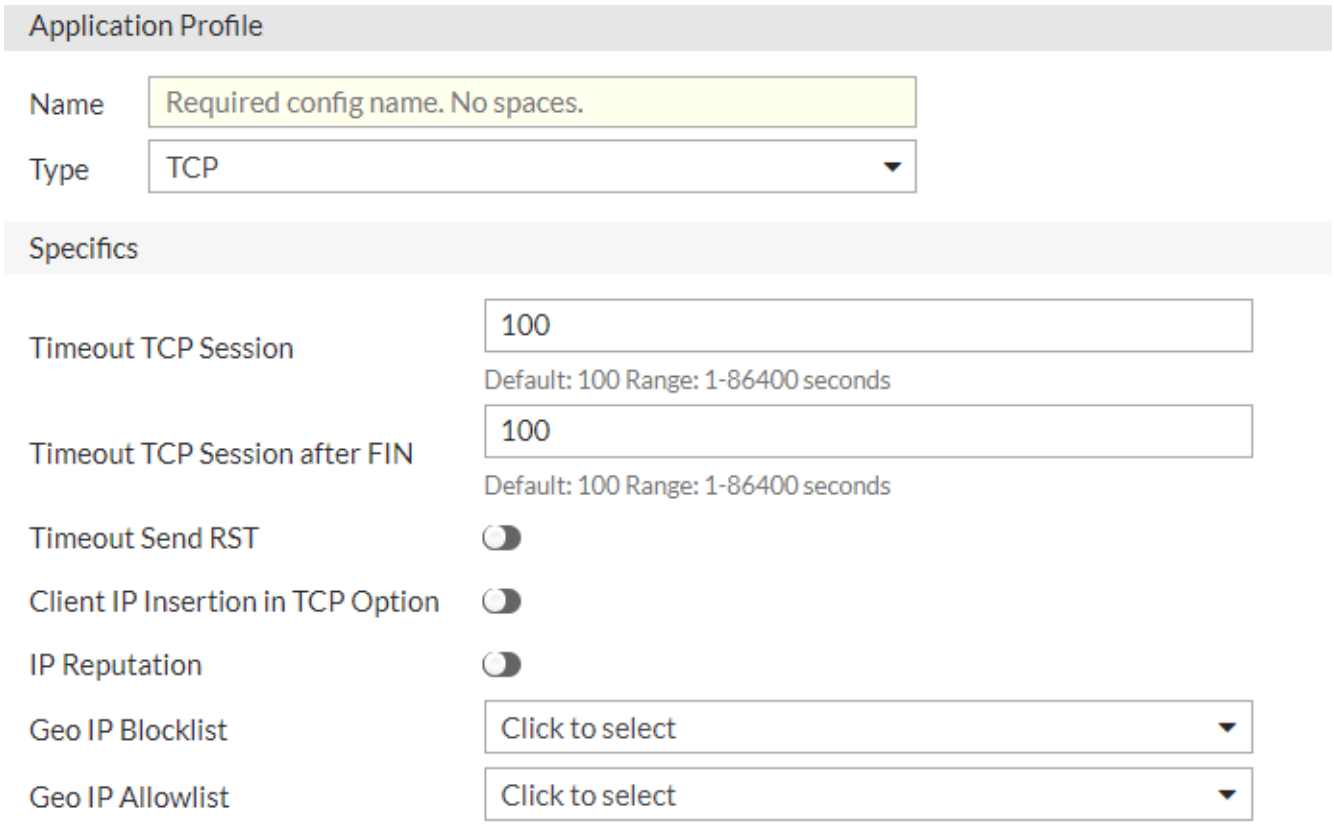
Profile Type	Usage	VS Type	LB Methods	Persistence
TCP on page 455	Use for other TCP protocols.	Layer 4	Round Robin, Least Connection, Fastest Response, Dynamic Load	Source Address, Source Address Hash, Source Address-Port Hash
		Layer 2	Round Robin, Least Connection, Fastest Response, Destination IP Hash, Dynamic Load	
UDP on page 457	Use with UDP servers.	Layer 4	Round Robin, Least Connection, Fastest Response, Dynamic Load	Source Address, Source Address Hash, Source Address-Port Hash
		Layer 2	Round Robin, Least Connection, Fastest Response, Dynamic Load, Destination IP Hash	
HTTP on page 459	Use for standard, unsecured web server traffic.	Layer 7	Round Robin, Least Connection, URI Hash, Full URI Hash, Host Hash, Host Domain Hash, Dynamic Load	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
		Layer 2	Round Robin, Least Connection, URI Hash, Full URI Hash, Host Hash, Host Domain Hash, Dynamic Load, Destination IP Hash	
FTP on page 465	Use with FTP servers.	Layer 7	Round Robin, Least Connection	Source Address, Source Address Hash
		Layer 4	Round Robin, Least Connection, Fastest Response, Dynamic Load	Source Address, Source Address Hash, Source Address-Port Hash
RADIUS on page 467	Use with RADIUS servers.	Layer 7	Round Robin	RADIUS attribute
TCPS on page 468	Use for secured TCP when offloading TLS/SSL from the backend servers. You must import the backend server certificates into FortiADC and select them in the TCPS profile.	Layer 7	Round Robin, Least Connection, Dynamic Load	Source Address, Source Address Hash, Source Address-Port Hash, SSL Session ID
		Layer 2	Round Robin, Least Connection, Destination IP Hash, Dynamic Load	

Profile Type	Usage	VS Type	LB Methods	Persistence
HTTPS on page 470	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	Round Robin, Least Connection, URI Hash, Full URI Hash, Host Hash, Host Domain Hash, Dynamic Load	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, SSL Session ID
		Layer 2	Round Robin, Least Connection, URI Hash, Full URI Hash, Host Hash, Host Domain Hash, Dynamic Load, Destination IP Hash	
HTTP Turbo on page 477	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 Connection, Fastest Response	Source Address
SIP on page 478	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
RDP on page 482	Use with Windows Terminal Service (Remote Desktop Protocol).	Layer 7	Round Robin, Least Connection	Source Address, Source Address Hash, Source Address-Port Hash, RDP Cookie
DNS on page 484	Use with DNS servers.	Layer 7	Round Robin, Least Connection	Not supported yet.

Profile Type	Usage	VS Type	LB Methods	Persistence
IP on page 485	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, Dynamic Load	Source Address, Source Address Hash, Source Address-Port Hash
SMTP on page 486	Use with SMTP servers.	Layer 7	Round Robin, Least Connection	Source Address, Source Address Hash
MySQL on page 488	MySQL network protocol stack (such as MySQL-Proxy) which parses and builds MySQL protocol packets.	Layer 7	Round Robin, Least Connection	N/A
RTSP on page 488	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
RTMP on page 489	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
Diameter on page 490	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)
ISO8583 on page 491	Use with ISO8583 servers.	Layer 7	Round Robin	N/A
MSSQL on page 493	MSSQL network protocol stack, which parses and builds MSSQL protocol packets.	Layer 7	Least Connection	N/A
EXPLICIT_HTTP on page 494	A simple explicit/forward HTTP proxy mode.	Layer 7	N/A	N/A

Profile Type	Usage	VS Type	LB Methods	Persistence
	In this mode, you do not need to add backend real server pools. The destination IP address of the downstream is specified by the URL or Host field of the client request.			
L7 TCP on page 497	Use for other TCP protocols.	Layer 7	Round Robin, Least Connection	Source Address, Source Address Hash
L7 UDP on page 498	Use with UDP servers.	Layer 7	Round Robin, Least Connection	Source Address, Source Address Hash

TCP



Parameter	Description
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.

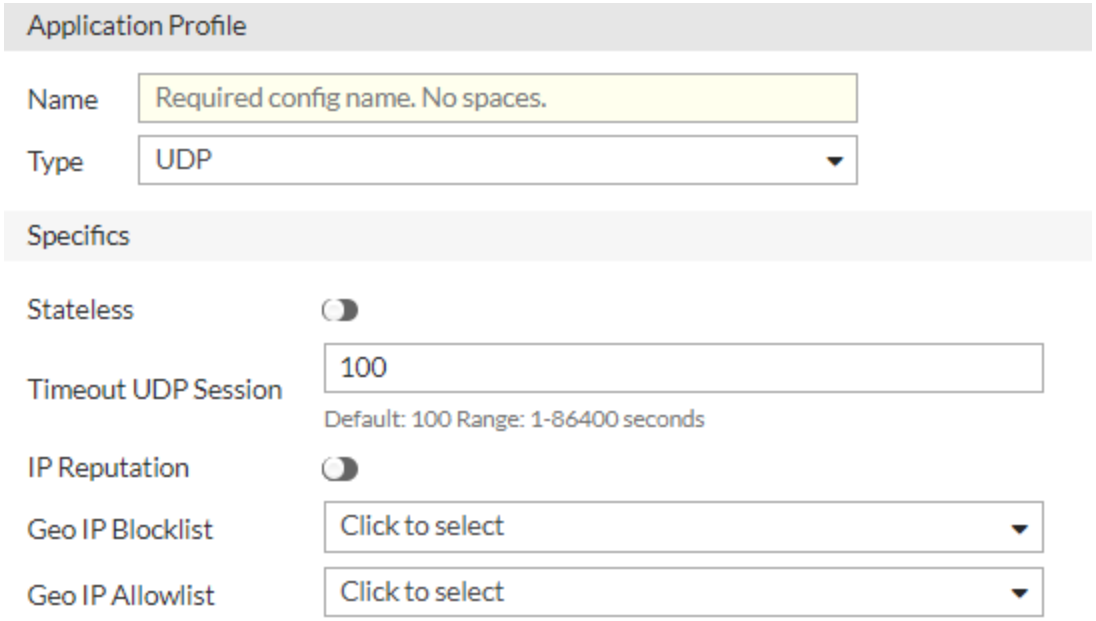
Parameter	Description
Timeout TCP Session after FIN	Client-side connection timeout. The default is 100 seconds. The valid range is 1 to 86,400.
Timeout Send RST	Enable to send TCP RST to the client and real server when the TCP session expires. This is disabled by default. **Note**: This function is supported for both IPv4 and IPv6 in L4 and L2 virtual servers. For L4 virtual servers, Timeout Sent RST is supported for DNAT/FullNAT/NAT46/NAT64 packet forwarding methods.
Client IP Insertion in TCP Option	Enable to insert the Client IP in the TCP option. This is disabled by default. This function is applicable in Layer 4 TCP virtual servers using DNAT or Full NAT as the packet forwarding method. By inserting the Client IP in the TCP option it allows the real server to easily retrieve the client address even after the Client IP has been modified during packet forwarding. **Note**: FortiADC cannot insert the Client IP in the TCP option under the following conditions: • If the TCP option number already exists. • If the length of the TCP header will exceed 60 bytes once the client address is inserted.
Client IP Insertion in TCP Option Number	The **Client IP Insertion in TCP Option Number** option is available if **Client IP Insertion in TCP Option** is enabled. Specify the TCP Option Number which indicates the TCP option kind to be collected. The default is 28. The valid range is 2 to 255.
IP Reputation	Enable to apply the FortiGuard IP reputation service. See Managing IP Reputation policy settings.
Geo IP Blocklist	Select a Geo IP block list configuration object. See Using the Geo IP block list.
Geo IP Allowlist	Select an allowlist configuration object. See Using the Geo IP allowlist.

Predefined TCP Profile: LB_PROF_TCP

Parameter	Default value
Timeout TCP Session	100
Timeout TCP Session after FIN	100
Timeout Send RST	Disabled
Client IP Insertion in TCP Option	Disabled
IP Reputation	Disabled
Geo IP Blocklist	None
Geo IP Allowlist	None

Navigate back to Application Profile Types on page 451.

UDP



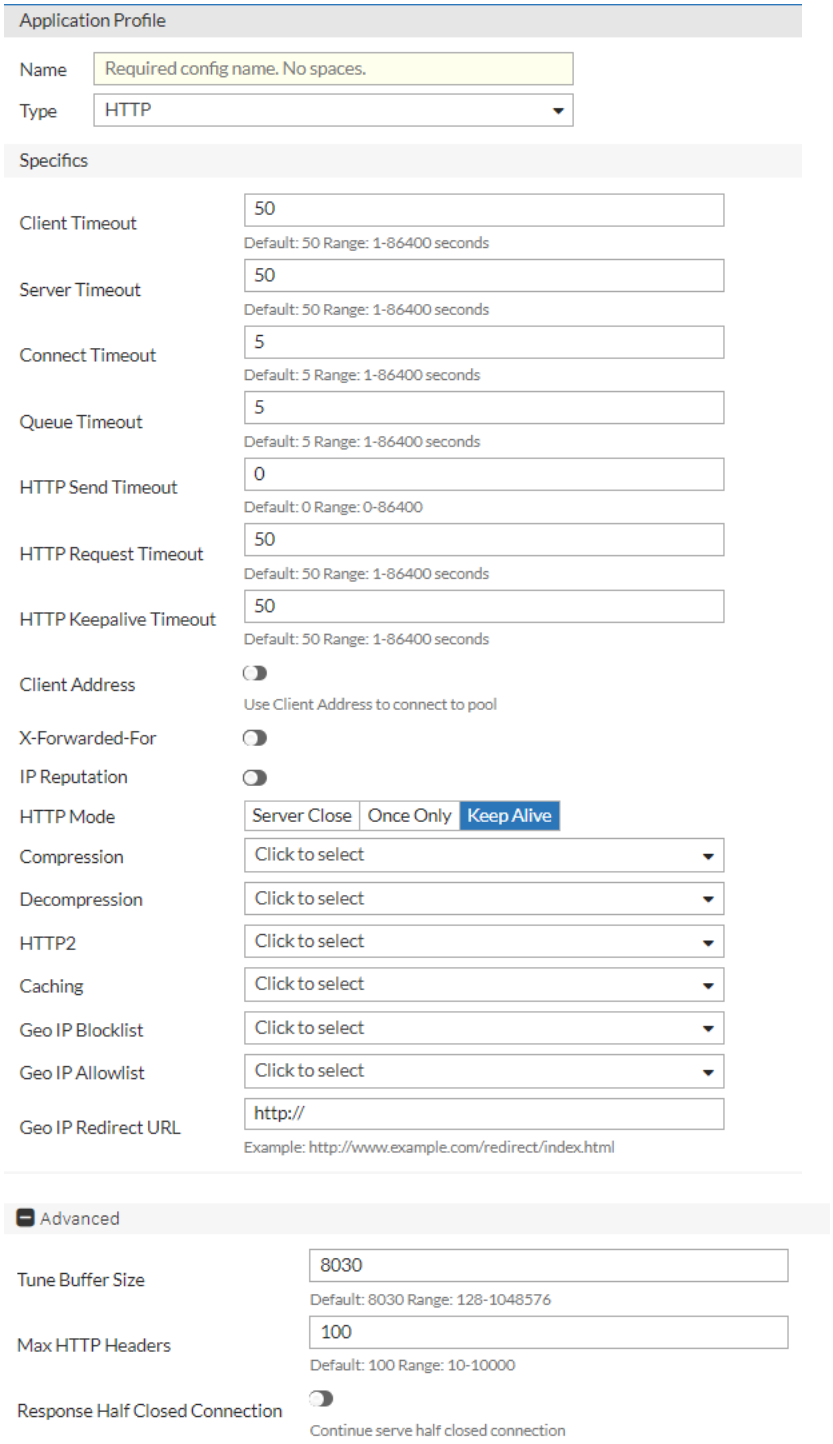
Parameter	Description
Stateless	Enable to apply UDP stateless function. A stateless virtual server accepts traffic that matches the virtual server address and forwards packets to the real servers without checking for pre-existing connections in the connection table. Any new connection entry is immediately discarded. This configuration is ideal for high-throughput, one-way UDP traffic processing. **Examples of Stateless Virtual Servers**: • Load balancing syslog traffic to a pool of syslog servers (commonly for FortiAnalyzer devices) • Load balancing DNS queries to a pool of DNS servers that use Direct Server Return (DSR), where the FortiADC is not the gateway for DNS response delivery.
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 Blocklist	Select a Geo IP block list configuration object. See Using the Geo IP block list.
Geo IP Allowlist	Select an allowlist configuration object. See Using the Geo IP allowlist.

Predefined UDP Profile: LB_PROF_UDP

Parameter	Default value
Timeout UDP Session	100
IP Reputation	Disabled
Stateless	Disabled
Geo IP Blocklist	None
Geo IP Allowlist	None

Navigate back to Application Profile Types on page 451.

HTTP


Application Profile
Name
Required config name. No spaces.
Type
HTTP
Specifics
Client Timeout
50
Default: 50 Range: 1-86400 seconds
Server Timeout
50
Default: 50 Range: 1-86400 seconds
Connect Timeout
5
Default: 5 Range: 1-86400 seconds
Queue Timeout
5
Default: 5 Range: 1-86400 seconds
HTTP Send Timeout
0
Default: 0 Range: 0-86400
HTTP Request Timeout
50
Default: 50 Range: 1-86400 seconds
HTTP Keepalive Timeout
50
Default: 50 Range: 1-86400 seconds
Client Address
Use Client Address to connect to pool
X-Forwarded-For
IP Reputation
HTTP Mode
Server Close
Once Only
Keep Alive
Compression
Click to select
Decompression
Click to select
HTTP2
Click to select
Caching
Click to select
Geo IP Blocklist
Click to select
Geo IP Allowlist
Click to select
Geo IP Redirect URL
http://
Example: http://www.example.com/redirect/index.html
Advanced
Tune Buffer Size
8030
Default: 8030 Range: 128-1048576
Max HTTP Headers
100
Default: 100 Range: 10-10000
Response Half Closed Connection
Continue serve half closed connection

Parameter	Description
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. If you apply an **HTTP2 profile** with **Backend HTTP2** enabled, the **Client Timeout** specifies the maximum time a client has to complete key actions, such as responding to a server request (including receiving the full response) or disconnecting from the server after completing a transaction.
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. If you apply an **HTTP2 profile** with **Backend HTTP2** enabled, the **HTTP Request Timeout** specifies the maximum time allowed for a client to send a complete HTTP request to the server.
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.
Client Address	Use the original client IP address as the source address when connecting to the real server.

Parameter	Description
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.
IP Reputation Redirect URL	Type a URL including the FQDN/IP and path, if any, to which a client will be redirected if the request violates the IP reputation policy.
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.
Decompression	Select a decompression configuration object. See Configuring decompression rules on page 536.
HTTP2	Select an HTTP2 Profile configuration object. See Configuring HTTP2 profiles on page 518. **Note**: When an HTTP2 profile is applied with **Backend HTTP2** enabled, the definitions of **Client Timeout** and **HTTP Request Timeout** will differ from their standard interpretation. Please refer to the respective descriptions of **Client Timeout** and **HTTP Request Timeout** for further details.
Caching	Select a caching configuration object. See Using caching features.
Geo IP Blocklist	Select a Geo IP block list configuration object. See Using the Geo IP block list.
Geo IP Allowlist	Select an allowlist configuration object. See Using the Geo IP allowlist.
Geo IP Redirect URL	For HTTP, if you have configured a Geo IP redirect action, specify a redirect URL.

Parameter	Description
Advanced Settings	
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. • 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.
Response Half Closed Connection	Enable to continue to serve half-closed connections.

Predefined HTTP Profiles:

LB_PROF_HTTP

Parameter	Default value
Client Timeout	50
Server Timeout	50
Connect Timeout	5
Queue Timeout	5
HTTP Request Timeout	50
HTTP Keepalive Timeout	50
Client Address	Disabled
X-Forwarded-For	Disabled
X-Forwarded-For Header	Blank
IP Reputation	Disabled
HTTP Mode	Keep Alive
Compression	None
Decompression	None
Caching	None

Parameter	Default value
Geo IP Blocklist	None
Geo IP Allowlist	None
Geo IP Redirect URL	http://
HTTP Send Timeout	5
HTTP2	None

LB_PROF_HTTP_SERVERCLOSE

Parameter	Default value
Client Timeout	50
Server Timeout	50
Connect Timeout	5
Queue Timeout	5
HTTP Request Timeout	50
HTTP Keepalive Timeout	50
Buffer Pool	Enabled
Client 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 Blocklist	None
Geo IP Allowlist	None
Geo IP Redirect URL	http://
HTTP Send Timeout	0
HTTP2	None

LB_PROF_HTTP2_H2C

Parameter	Default value
Client Timeout	50
Server Timeout	50
Connect Timeout	5
Queue Timeout	5
HTTP Send Timeout	0
HTTP Request Timeout	50
HTTP Keepalive Timeout	50
Client Address	Disabled
X-Forwarded-For	Disabled
X-Forwarded-For Header	None
IP Reputation	Disabled
HTTP Mode	Keep Alive
Compression	None
Decompression	None
HTTP2	LB_HTTP2_PROFILE_DEFAULT
Caching	None
Geo IP Blocklist	None
Geo IP Allowlist	None
Geo IP Redirect URL	http://
Tune Buffer Size	17418
Max HTTP Headers	200
Response Half Closed Connection	Disabled

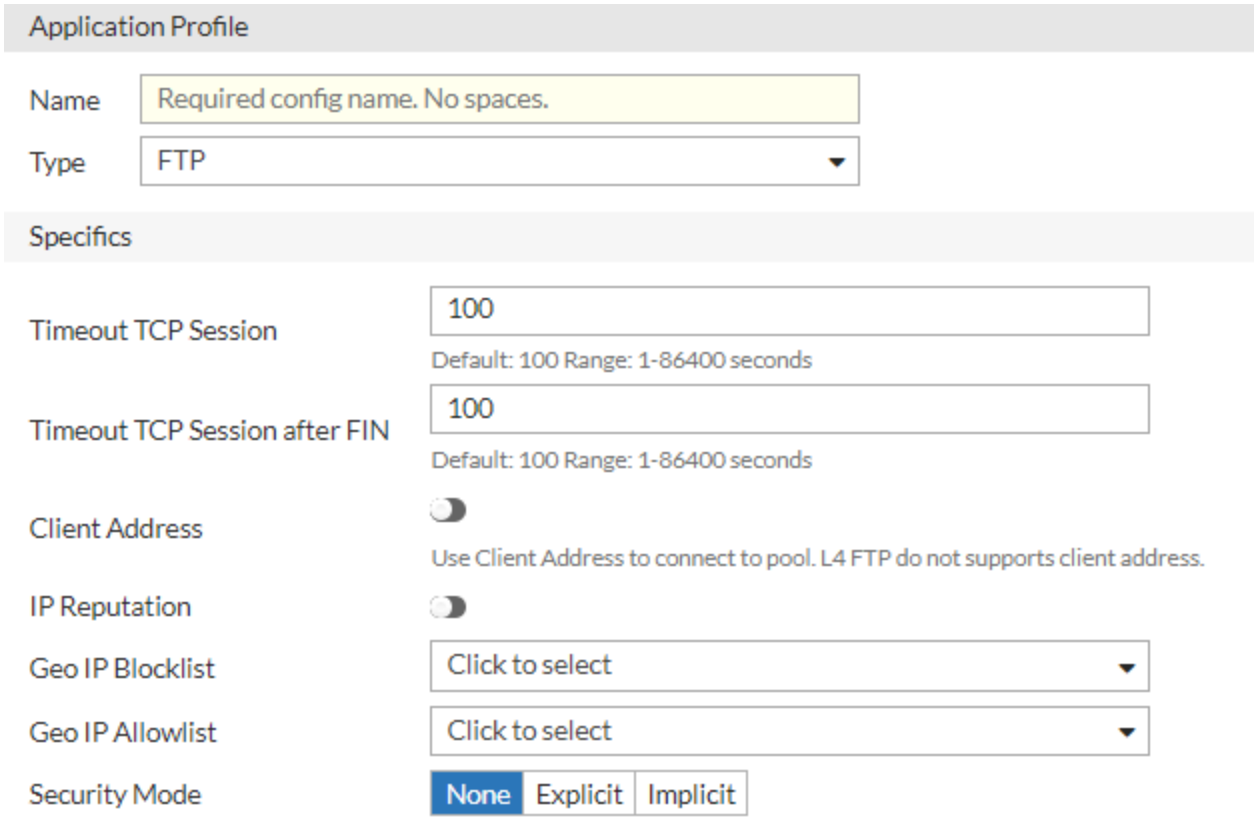
LB_PROF_HTTP2_END2END_H2C

Parameter	Default value
Client Timeout	50
Server Timeout	50
Connect Timeout	5
Queue Timeout	5
HTTP Request Timeout	50

Parameter	Default value
HTTP Keepalive Timeout	50
X-Forwarded-For	Disabled
HTTP Mode	Keep Alive
HTTP2	LB_HTTP2_PROFILE_END2END_H2
Tune Buffer Size	16384
Max HTTP Headers	200

Navigate back to Application Profile Types on page 451.

FTP



Parameter	Description
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.

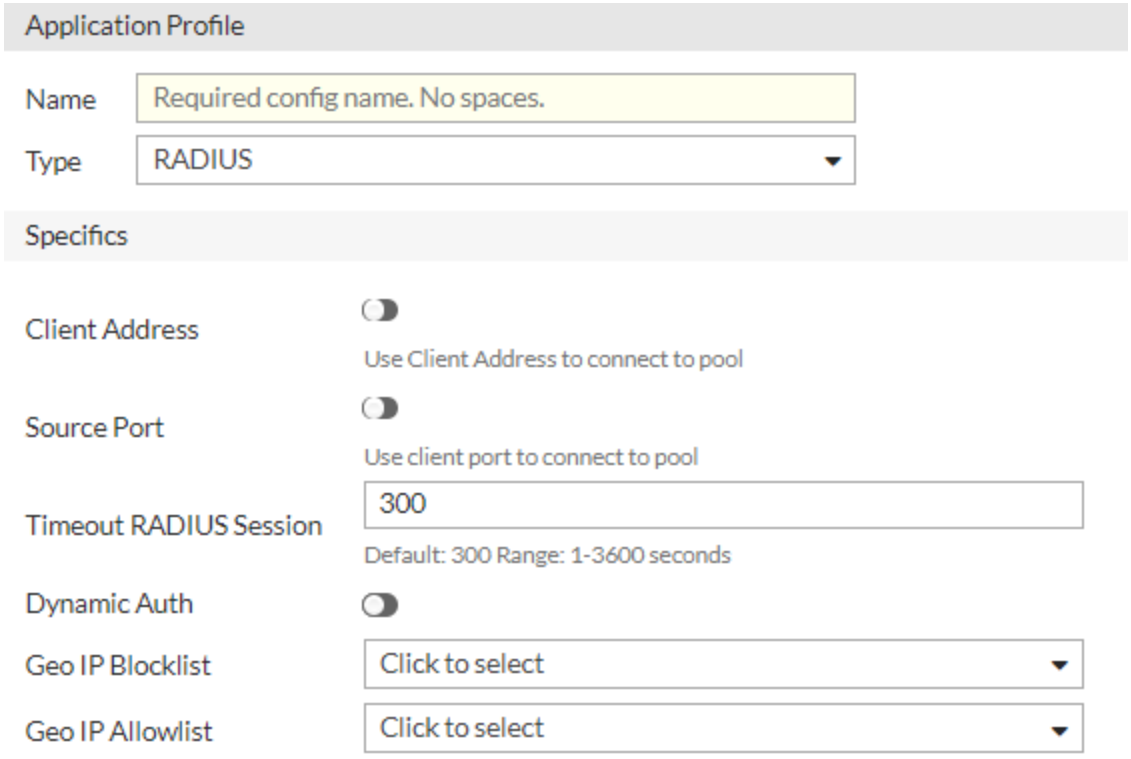
Parameter	Description
Timeout TCP Session after FIN	Client-side connection timeout. The default is 100 seconds. The valid range is 1 to 86,400.
Client Address	Use the original client IP address as the source address when connecting to the real server.
IP Reputation	Enable to apply the FortiGuard IP reputation service. See Managing IP Reputation policy settings.
Geo IP Blocklist	Select a Geo IP block list configuration object. See Using the Geo IP block list.
Geo IP Allowlist	Select an allowlist configuration object. See Using the Geo IP allowlist.
Security Mode	Select either of the following: • None • Explicit • Implicit

Predefined FTP Profile: LB_PROF_FTP

Parameter	Default value
Timeout TCP Session	100
Timeout TCP Session after FIN	100
IP Reputation	Disabled
Geo IP Blocklist	None
Geo IP Allowlist	None
Client Address	Off
Security Mode	None

Navigate back to Application Profile Types on page 451.

RADIUS



Parameter	Description
Client Address	Use the original client IP address as the source address when connecting to the real server.
Source Port	Use the original client port as the source port when connecting to the real server.
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.
Geo IP Blocklist	Select a Geo IP block list configuration object. See Using the Geo IP block list.
Geo IP Allowlist	Select an allowlist configuration object. See Using the Geo IP allowlist.

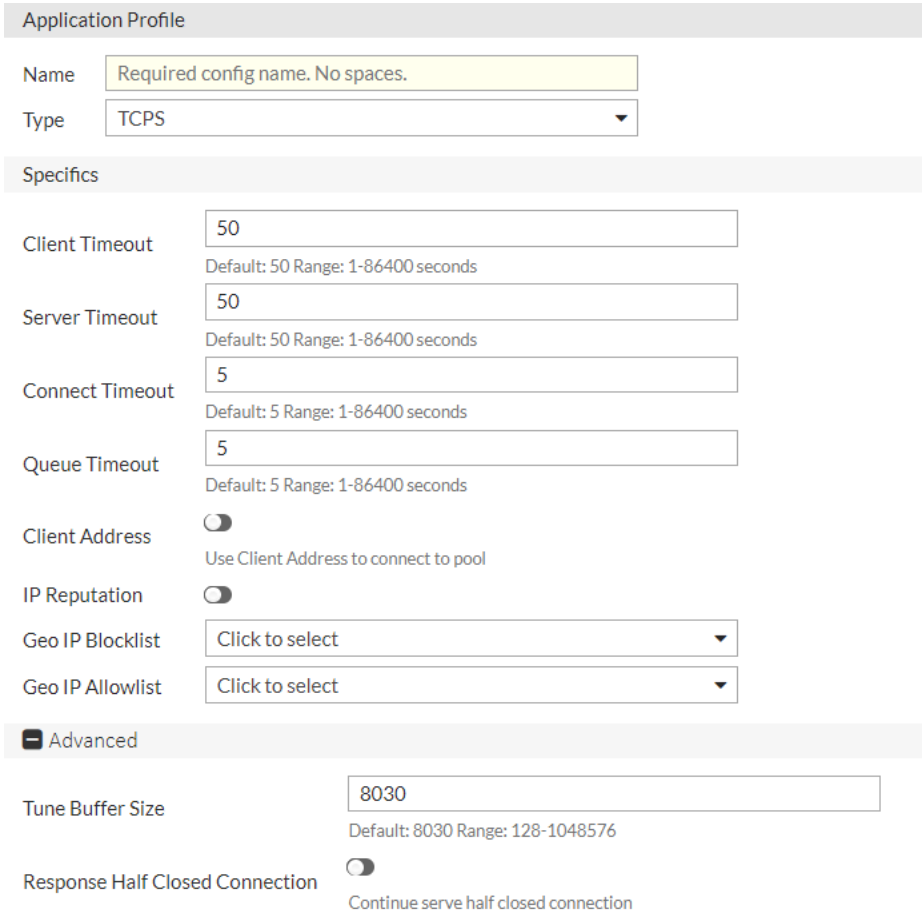
Predefined RADIUS Profile: LB_PROF_RADIUS

Parameter	Default value
Client Address	Off
Source Port	Off

Parameter	Default value
Dynamic Auth	Disabled
RADIUS Session	300
Geo IP Blocklist	None
Geo IP Allowlist	None

Navigate back to Application Profile Types on page 451.

TCPS



Parameter	Description
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.

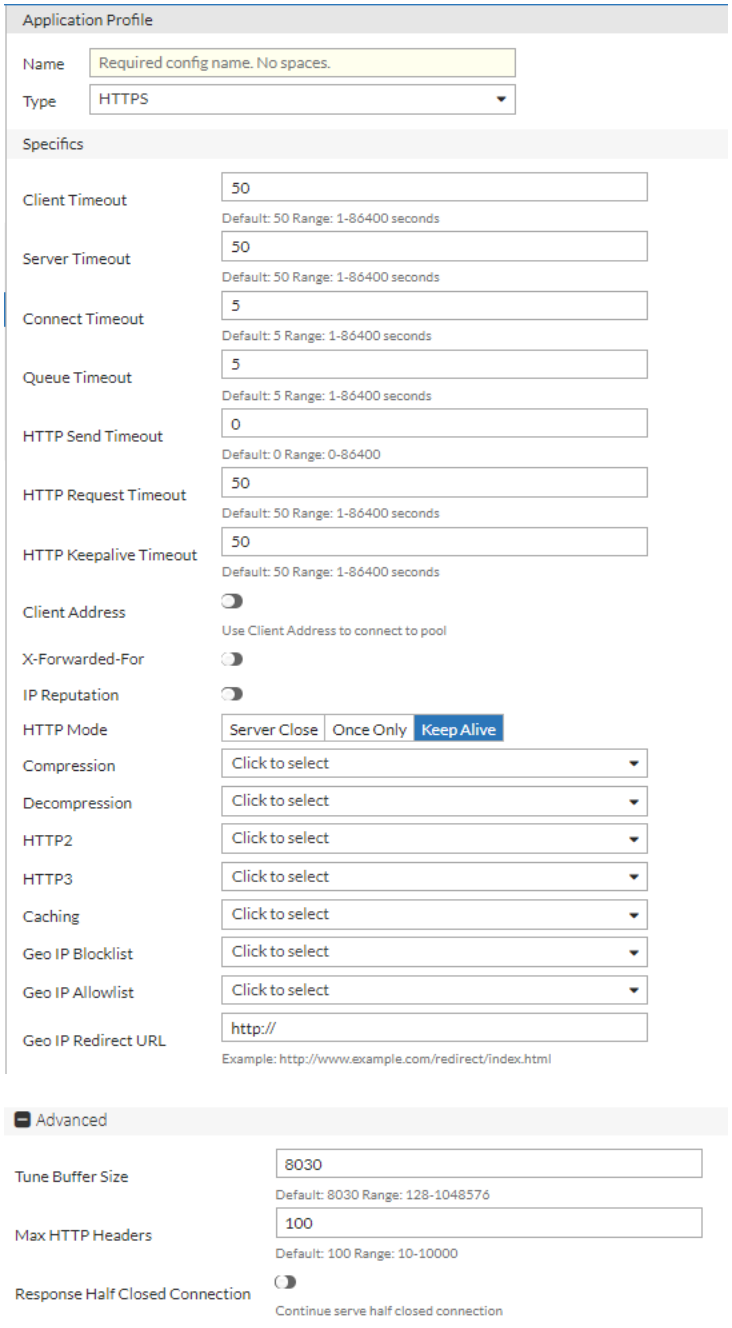
Parameter	Description
Connect Timeout	Multiplexed server-side TCP connection timeout. Usually less than the client-side timeout. The default is 5 seconds. The valid range is 1 to 3,600.
Queue Timeout	Specifies how long connection requests to a backend server remain in a queue if the server has reached 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.
Client 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.
Geo IP Blocklist	Select a Geo IP block list configuration object. See Using the Geo IP block list.
Geo IP Allowlist	Select an allowlist configuration object. See Using the Geo IP allowlist.
Advanced Settings	
Tune Buffer Size	Adjust the value of the TCPS 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.
Response Half Closed Connection	Enable to continue to serve half-closed connections.

Predefined TCPS Profile: LB_PROF_TCPS

Parameter	Default value
Client Timeout	50
Server Timeout	50
Connect Timeout	5
Queue Timeout	5
Client Address	Disabled
IP Reputation	Disabled
Geo IP Blocklist	None

Navigate back to Application Profile Types on page 451.

HTTPS



Parameter	Description
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. The default is 50 seconds. The valid range is from 1 to 3,600.

Parameter	Description
	If you apply an **HTTP3 profile** or an **HTTP2 profile** with **Backend HTTP2** enabled, the **Client Timeout** specifies the maximum time a client has to complete key actions, such as responding to a server request (including receiving the full response) or disconnecting from the server after completing a transaction.
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. The default is 50 seconds. The valid range is 1 to 3,600. If you apply an **HTTP3 profile** or an **HTTP2 profile** with **Backend HTTP2** enabled, the **HTTP Request Timeout** specifies the maximum time allowed for a client to send a complete HTTP request to the server.
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.
Client Address	Use the original client IP address as the source address when connecting to the real server. **Note**: IPv6 address types cannot be used as the client address to connect to the backend real servers in HTTPS virtual servers.
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.

Parameter	Description
	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.
IP Reputation Redirect URL	Type a URL including the FQDN/IP and path, if any, to which a client will be redirected if the request violates the IP reputation policy.
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.
Decompression	Select a decompression configuration object. See Configuring decompression rules on page 536.
HTTP2	Select an HTTP2 Profile configuration object. See Configuring HTTP2 profiles on page 518. **Note**: When an **HTTP2 profile** is applied with **Backend HTTP2** enabled, the definitions of **Client Timeout** and **HTTP Request Timeout** will differ from their standard interpretation. Please refer to the respective descriptions of **Client Timeout** and **HTTP Request Timeout** for further details.
HTTP3	Select an HTTP3 Profile configuration object. See Configuring HTTP3 profiles on page 521. **Note**: When an **HTTP3 profile** is applied, the definitions of **Client Timeout** and **HTTP Request Timeout** will differ from their standard interpretation. Please refer to the respective descriptions of **Client Timeout** and **HTTP Request Timeout** for further details.
Caching	Select a caching configuration object. See Using caching features.

Parameter	Description
Geo IP Blocklist	Select a Geo IP block list configuration object. See Using the Geo IP block list.
Geo IP Allowlist	Select an allowlist configuration object. See Using the Geo IP allowlist.
Geo IP Redirect URL	For HTTP, if you have configured a Geo IP redirect action, specify a redirect URL.
Advanced Settings	
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. • 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.
Response Half Closed Connection	Enable to continue to serve half-closed connections.

Predefined HTTPS Profiles:

LB_PROF_HTTPS

Parameter	Default value
Client Timeout	50
Server Timeout	50
Connect Timeout	5
Queue Timeout	5
HTTP Request Timeout	50
HTTP Keepalive Timeout	50
Client Address	Disabled
X-Forwarded-For	Disabled
X-Forwarded-For Header	None
IP Reputation	Disabled
HTTP Mode	Keep Alive

Parameter	Default value
Compression	None
Decompression	None
Caching	None
Geo IP Blocklist	None
Geo IP Allowlist	None
Geo IP Redirect URL	http://
HTTP Send Timeout	0
HTTP2	None

LB_PROF_HTTPS_SERVERCLOSE

Parameter	Default value
Client Timeout	50
Server Timeout	50
Connect Timeout	5
Queue Timeout	5
HTTP Request Timeout	50
HTTP Keepalive Timeout	50
Client Address	Disabled
X-Forwarded-For	Disabled
X-Forwarded-For Header	None
IP Reputation	Disabled
HTTP Mode	Server Close
Compression	None
Decompression	None
Caching	None
Geo IP Blocklist	None
Geo IP Allowlist	None
Geo IP Redirect URL	http://
HTTP Send Timeout	0
HTTP2	None

LB_PROF_HTTP2_H2

Parameter	Default value
Client Timeout	50
Server Timeout	50
Connect Timeout	5
Queue Timeout	5
HTTP Send Timeout	0
HTTP Request Timeout	50
HTTP Keepalive Timeout	50
Client Address	Disabled
X-Forwarded-For	Disabled
X-Forwarded-For Header	None
IP Reputation	Disabled
HTTP Mode	Keep Alive
Compression	None
Decompression	None
HTTP2	LB_HTTP2_PROFILE_DEFAULT
Caching	None
Geo IP Blocklist	None
Geo IP Allowlist	None
Geo IP Redirect URL	http://
Tune Buffer Size	17418
Max HTTP Headers	200
Response Half Closed Connection	Disabled

LB_PROF_HTTP2_END2END_H2

Parameter	Default value
Client Timeout	50
Server Timeout	50
Connect Timeout	5
Queue Timeout	5
HTTP Request Timeout	50

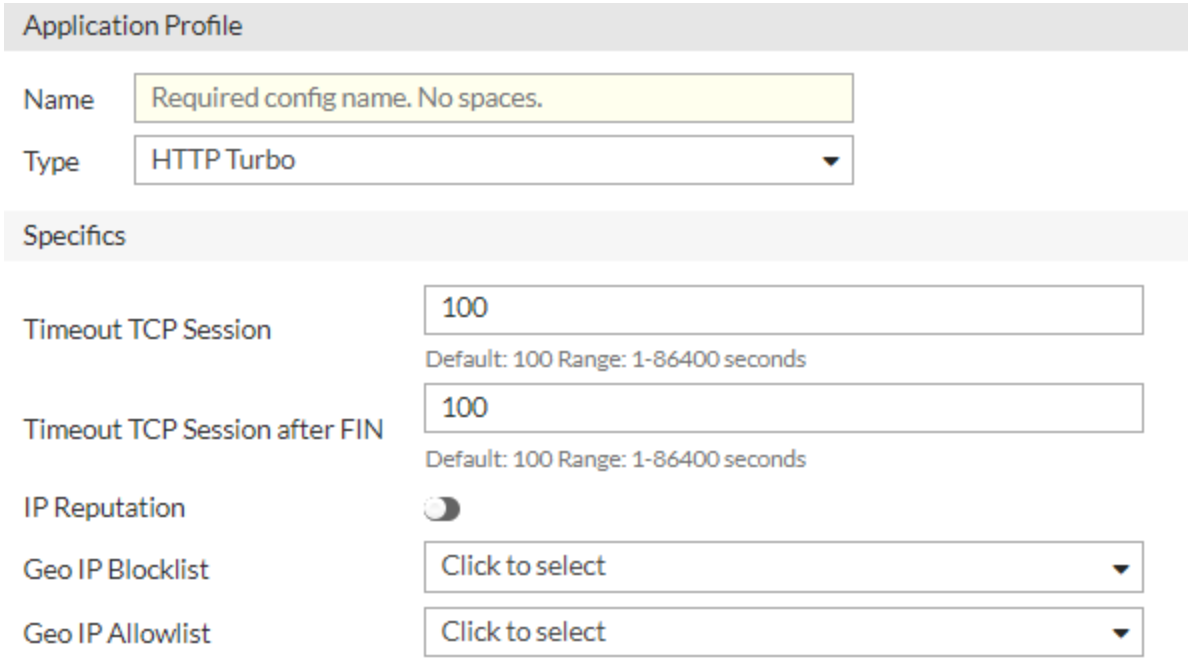
Parameter	Default value
HTTP Keepalive Timeout	50
Client Address	Disabled
X-Forwarded-For	Disabled
HTTP Mode	Keep Alive
HTTP2	LB_HTTP2_PROFILE_END2END_H2
HTTP3	None
Tune Buffer Size	16384
Max HTTP Headers	200

LB_PROF_HTTP3

Parameter	Default value
Client Timeout	50
Server Timeout	50
Connect Timeout	5
Queue Timeout	5
HTTP Request Timeout	50
HTTP Keepalive Timeout	50
Client Address	Disabled
X-Forwarded-For	Disabled
X-Forwarded-For Header	None
HTTP Mode	Keep Alive
HTTP3	LB_HTTP3_PROFILE_DEFAULT
Tune Buffer Size	32768
Max HTTP Headers	200

Navigate back to Application Profile Types on page 451.

HTTP Turbo



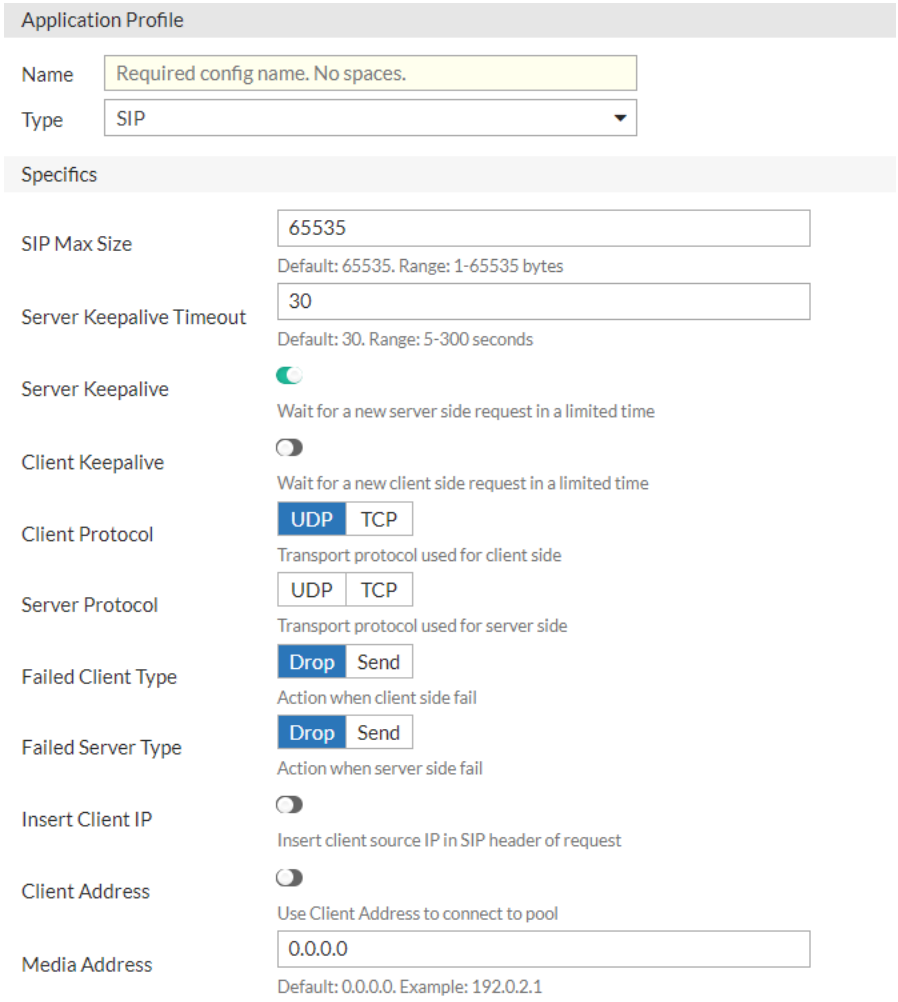
Parameter	Description
Timeout TCP Session	Client-side timeout for connections where the client has not sent a FIN signal, but the connection has been idle. The default is 100 seconds. The valid range is 1 to 86,400.
Timeout TCP Session after FIN	Client-side connection timeout. The default is 100 seconds. The valid range is from 1 to 86,400.
IP Reputation	Enable to apply the FortiGuard IP reputation service.
Geo IP Blocklist	Select a Geo IP block list configuration object. See Using the Geo IP block list.
Geo IP Allowlist	Select an allowlist configuration object. See Using the Geo IP allowlist.

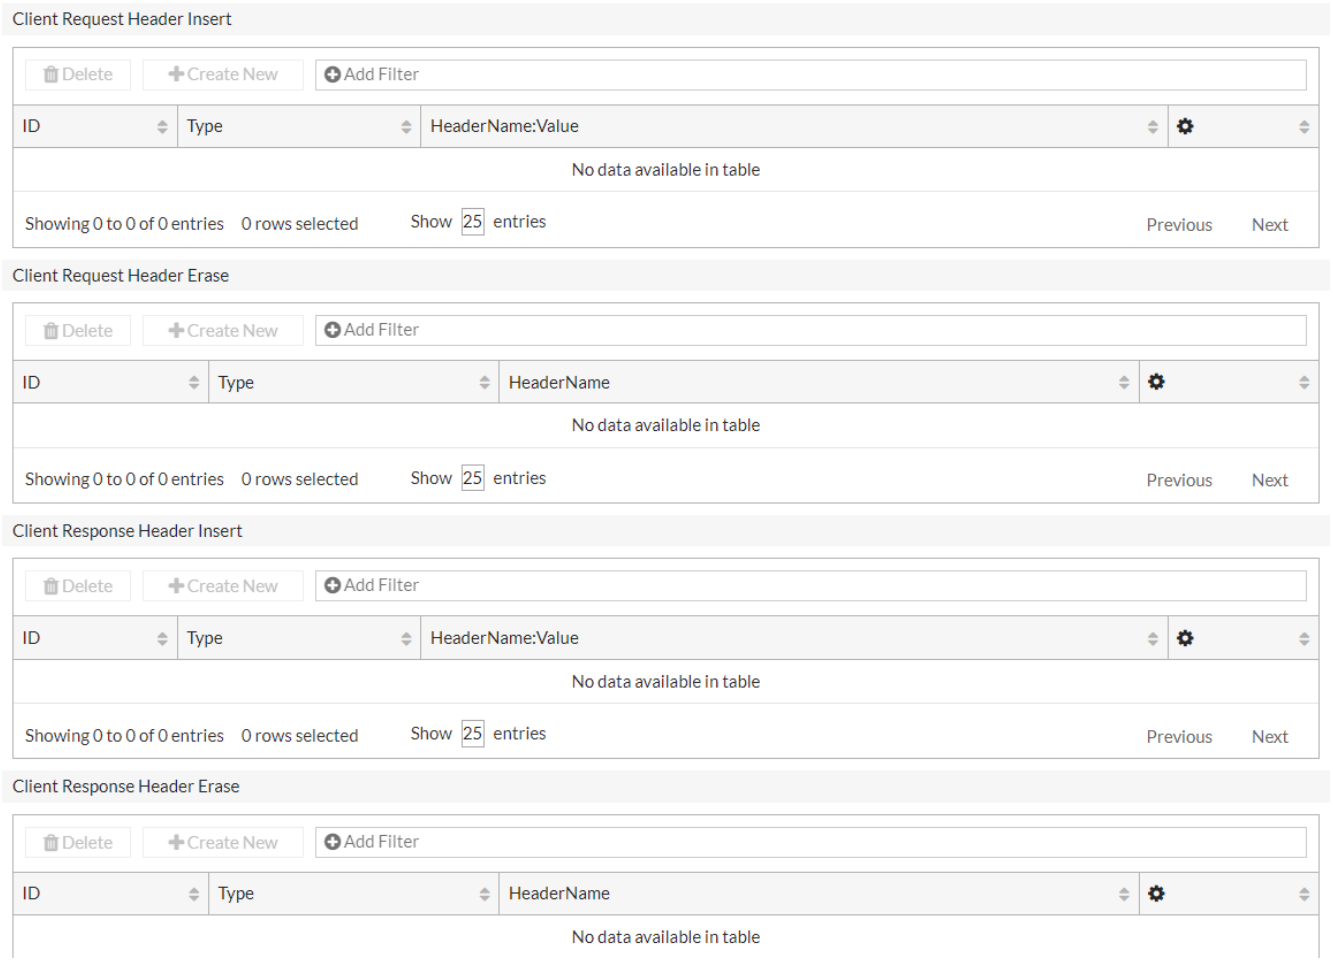
Predefined HTTP Turbo Profile: LB_PROF_TURBOHTTP

Parameter	Default value
Timeout TCP Session	100
Timeout TCP Session after FIN	100
IP Reputation	Disabled
Geo IP Blocklist	None
Geo IP Allowlist	None

Navigate back to Application Profile Types on page 451.

SIP


Application Profile
Name
Required config name. No spaces.
Type
SIP
Specifics
SIP Max Size
65535
Default: 65535. Range: 1-65535 bytes
Server Keepalive Timeout
30
Default: 30. Range: 5-300 seconds
Server Keepalive
Wait for a new server side request in a limited time
Client Keepalive
Wait for a new client side request in a limited time
Client Protocol
UDP
TCP
Transport protocol used for client side
Server Protocol
UDP
TCP
Transport protocol used for server side
Failed Client Type
Drop
Send
Action when client side fail
Failed Server Type
Drop
Send
Action when server side fail
Insert Client IP
Insert client source IP in SIP header of request
Client Address
Use Client Address to connect to pool
Media Address
0.0.0.0
Default: 0.0.0.0. Example: 192.0.2.1



Parameter	Description
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.
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.

Parameter	Description
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 Address	Use the original client IP address as the source address in the connection to the real server.
Media Address	Change the media address of SIP payload to specified address. 0.0.0.0 is default.
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.
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.

Parameter	Description
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.

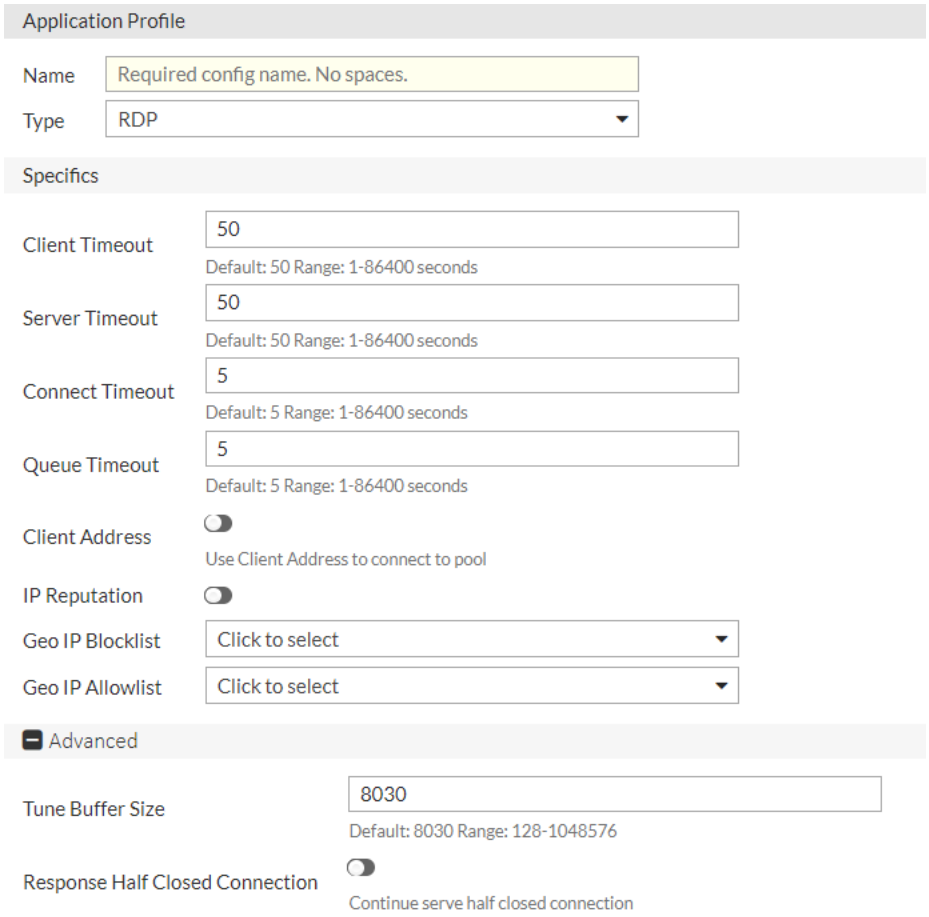
Predefined SIP Profile: LB_PROF_SIP

Parameter	Default value
SIP Max Size	65535
Server Keepalive Timeout	Server Keepalive Timeout
Server Keepalive	Enabled
Client Keepalive	Disabled
Client Protocol	UDP
Server Protocol	None
Failed Client Type	Drop

Parameter	Default value
Failed Server Type	Drop
Insert Client IP	Disabled
Geo IP Blocklist	None
Geo IP Allowlist	None
Client Address	Off
Media Address	0.0.0.0

Navigate back to Application Profile Types on page 451.

RDP



Parameter	Description
Client Timeout	Client-side TCP connection timeout. The default is 50 seconds. The valid range is 1 to 3,600.

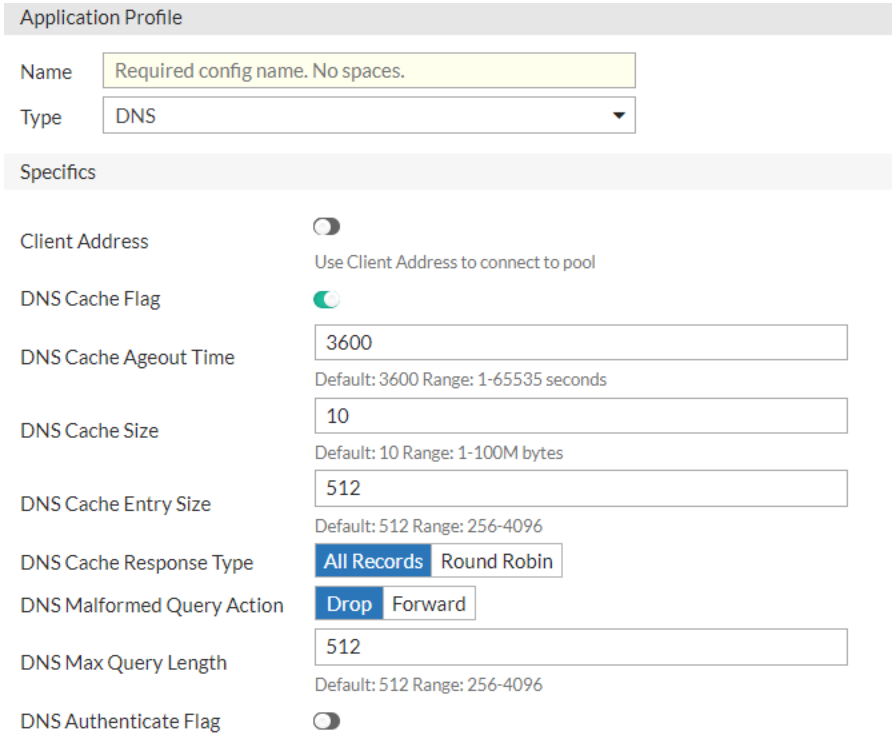
Parameter	Description
Server Timeout	Server-side IP session timeout. The default is 50 seconds. The valid range is 1 to 3,600.
Connect Timeout	Multiplexed server-side TCP connection timeout. Usually less than the client-side timeout. The default is 5 seconds. The valid range is 1 to 3,600.
Queue Timeout	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.
Client 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.
Geo IP Blocklist	Select a Geo IP block list configuration object. See Using the Geo IP block list.
Geo IP Allowlist	Select an allowlist configuration object. See Using the Geo IP allowlist.
Advanced Settings	
Tune Buffer Size	Adjust the value of the RDP 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.
Response Half Closed Connection	Enable to continue to serve half-closed connections.

Predefined RDP Profile: LB_PROF_RDP

Parameter	Default value
Client Timeout	50
Server Timeout	50
Connect Timeout	5
Queue Timeout	5
Source Address	Disabled
IP Reputation	Disabled
Geo IP Blocklist	None
Geo IP Allowlist	None

Navigate back to Application Profile Types on page 451.

DNS



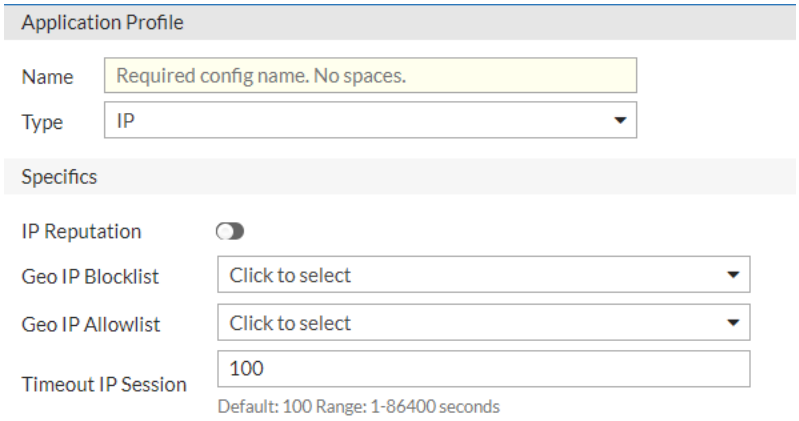
Parameter	Description
Client Address	Enable/disable to use the original client IP address as the source address when connecting to the real server.
DNS Cache Flag	Enable/disable the cache for the DNS virtual server.
DNS Cache Ageout Time	Specify the cache age-out time (in seconds). The default is 3,600. The valid range is 0 to 65,535.
DNS Cache Size	Specify the maximum cache size (in Megabytes). The default is 10. The valid range is 1 to 100.
DNS Cache Entry Size	Specify the maximum cache entry size. The default is 512. The valid range is 256 to 4,096.
DNS Cache Response Type	Select either of the following cache response types: • All Records • Round Robin
DNS Malform Query Action	Select either of the following reactions for the malformed requests: • Drop • Forward
DNS Max Query Length	Specify the maximum query length. The default is 512. The valid range is 256 to 4,096.
DNS Authentication Flag	Enable/disable to authenticate client by redirecting UDP query to TCP.

Predefined DNS Profile: LB_PROF_DNS

Parameter	Default value
Client Address	Off
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
DNS Max Query Length	512
DNS Authentication Flag	Disabled

Navigate back to Application Profile Types on page 451.

IP



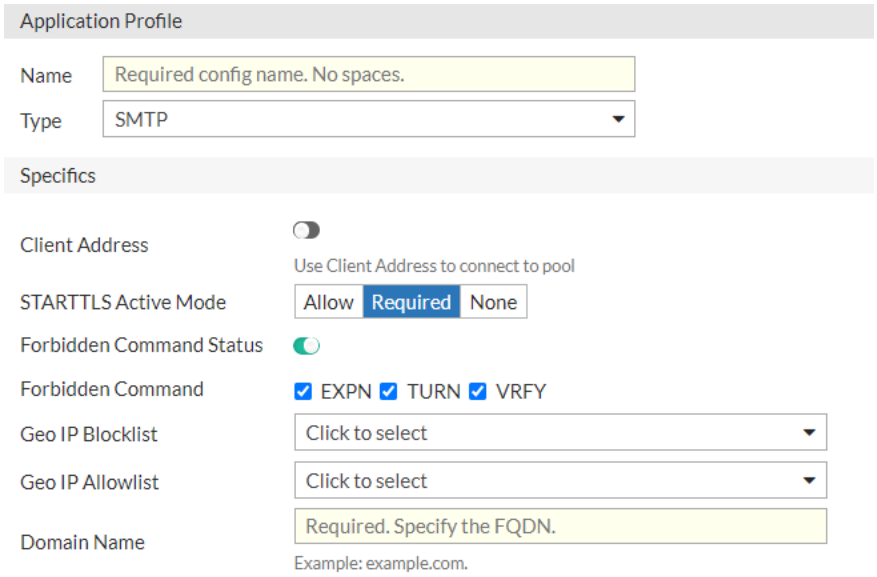
Parameter	Description
IP Reputation	Enable to apply FortiGuard IP reputation service. IP reputation. See Managing IP Reputation policy settings.
Geo IP Blocklist	Select a Geo IP block list configuration object. See Using the Geo IP block list.
Geo IP Allowlist	Select an allowlist configuration object. See Using the Geo IP allowlist.
Timeout IP Session	Client-side session timeout. The default is 100 seconds. The valid range is 1 to 86,400.

Predefined IP Profile: LB_PROF_IP

Parameter	Default value
IP Reputation	Disabled
Geo IP Blocklist	None
Geo IP Allowlist	None
Timeout IP Session	100

Navigate back to Application Profile Types on page 451.

SMTP



Parameter	Description
Client Address	Use the original client IP address as the source address in the connection to the real server. **Note**: When using the NAT Source Pool for SMTP VS, ensure the SMTP application profile is disabled for Client Address. When the SMTP is enabled for Client Address, it will use the original client IP address as the source address when connecting to the real server, which cannot be done when the NAT source pool is used at the same time.
STARTTLS Active Mode	Select one of the following: • Allow—The client can either use or not use the STARTTLS command. • Required—The STARTTLS command must be used to encrypt the connection first. • None—The STARTTLS command is NOT supported.

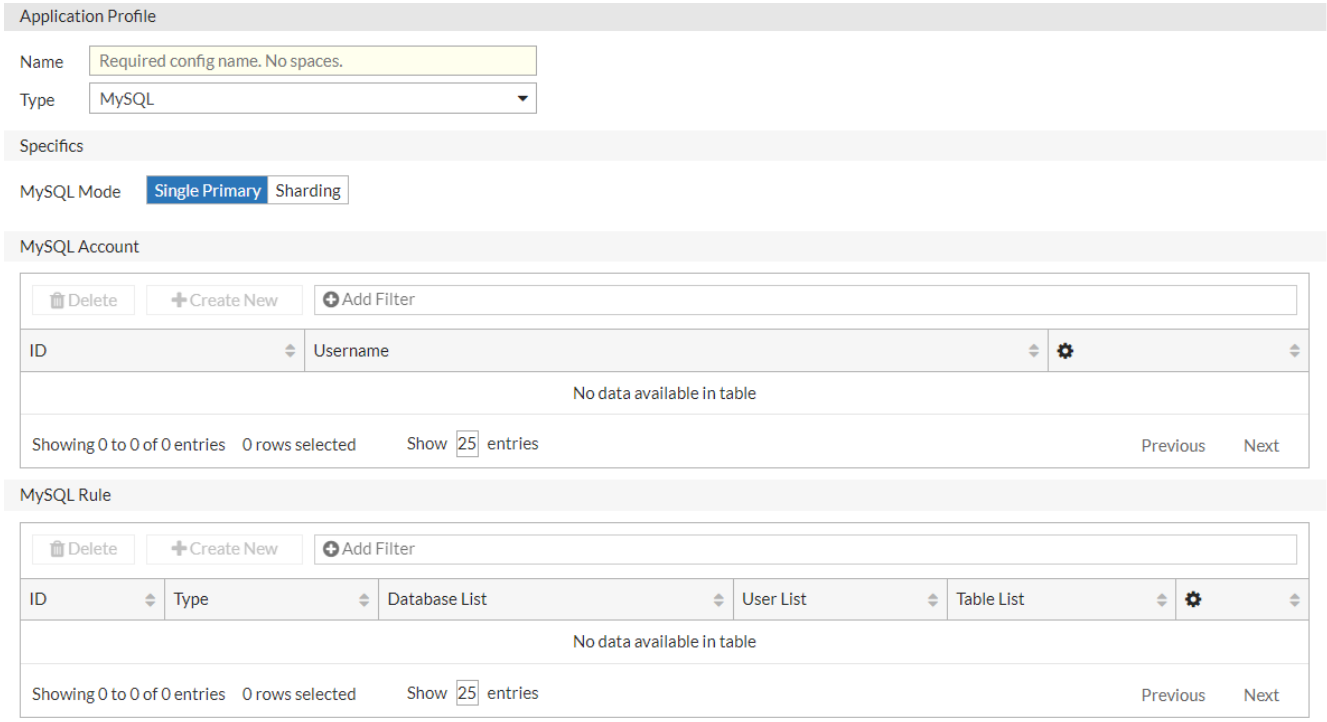
Parameter	Description
Forbidden Command Status	Enable/disable to forbid the command(s) selected in **Forbidden Command**.
Forbidden Command	Select any, all, or none of the commands: • EXPN • TURN • VRFY If selected, the command or commands will be rejected by FortiADC; otherwise, the command or commands will be accepted and forwarded to the back end.
Geo IP Blocklist	Select a Geo IP block list configuration object. See Using the Geo IP block list.
Geo IP Allowlist	Select an allowlist configuration object. See Using the Geo IP allowlist.
Domain Name	Specify the domain name.

Predefined SMTP Profile: LB_PROF_SMTP

Parameter	Default value
Client Address	Disable
STARTTLS Active Mode	Required
Forbidden Command Status	Enable
Forbidden Command	EXPN, TURN, VRFY
Geo IP Blocklist	None
Geo IP Allowlist	None
Domain Name	default.com

Navigate back to Application Profile Types on page 451.

MySQL

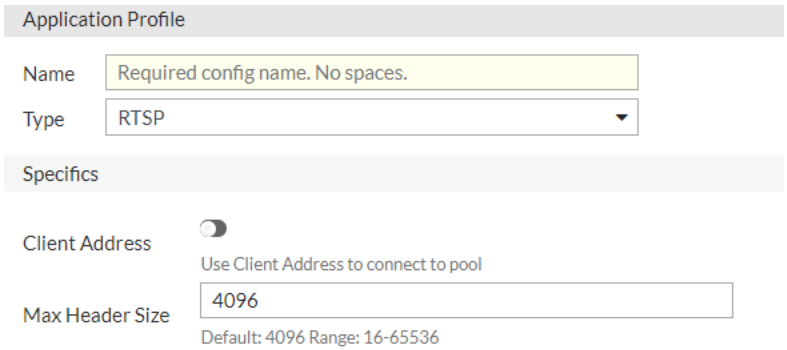


Parameter	Description
MySQL Mode	Select either of the following MySQL modes: • Single Primary — The profile will use the single-primary mode. You will then need to specify and configure the primary server and secondary servers. • Sharding — The profile will use the sharding mode to load-balance MySQL traffic.

For more configuration details, see Configuring MySQL profiles on page 499.

Navigate back to Application Profile Types on page 451.

RTSP



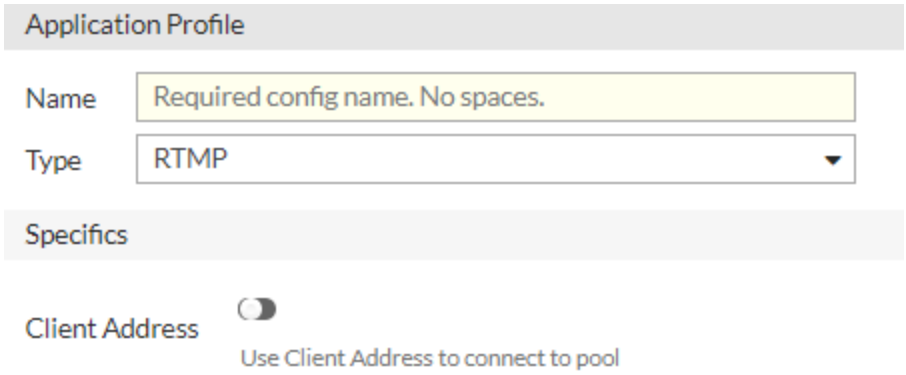
Parameter	Description
Client Address	Disabled by default. When enabled, FortiADC will use the client address to connect to the server pool.
Max Header Size	Specify the maximum size of the RTSP header. Default is 4096. Valid values range from 2048 to 65536.

Predefined RTSP Profile: LB_PROF_RTSP

Parameter	Default value
Client Address	Disabled
Max Header Size	4096

Navigate back to Application Profile Types on page 451.

RTMP



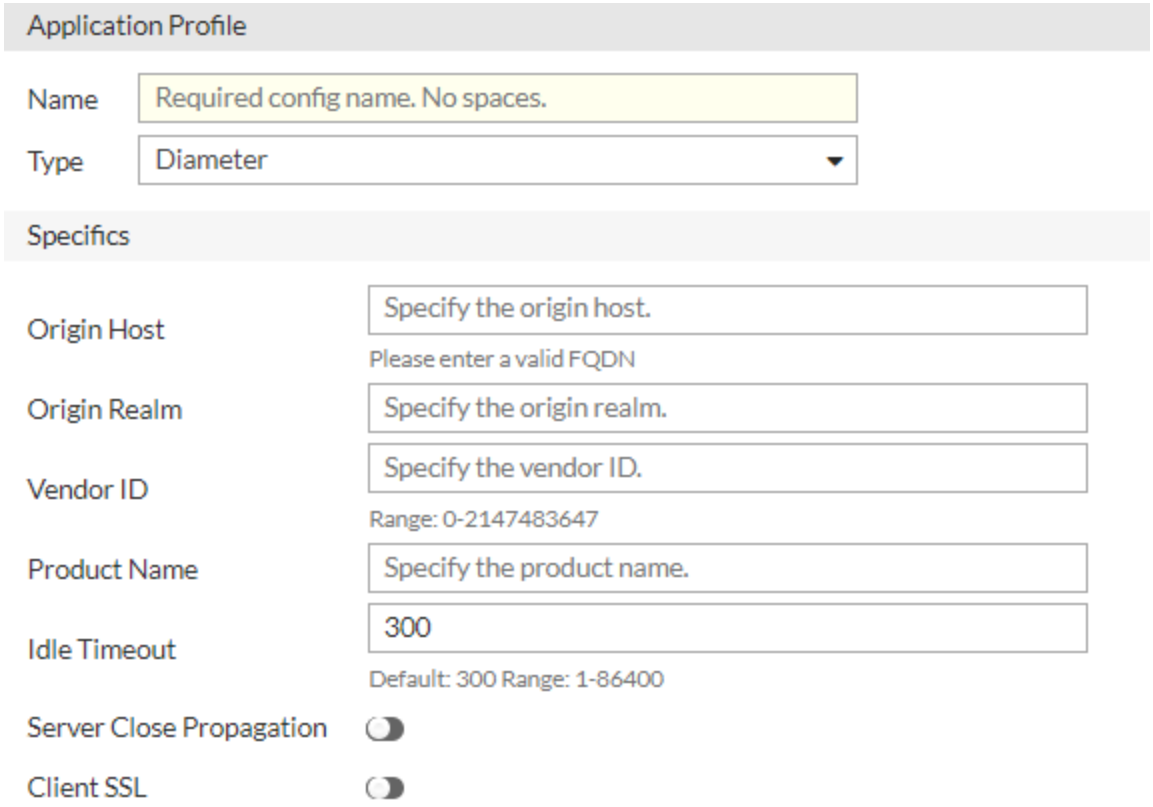
Parameter	Description
Client Address	Disabled by default. When enabled, FortiADC will use the client address to connect to the server pool.

Predefined RTMP Profile: LB_PROF_RTMP

Parameter	Default value
Client Address	Disabled

Navigate back to Application Profile Types on page 451.

Diameter



Parameter	Description
Origin Host	Specify the Origin Host. Once defined, FortiADC will change the Origin-Host AVP of the Diameter packet.
Origin Realm	Specify the Origin Realm. Once defined, FortiADC will change the Origin-Realm AVP of the Diameter packet.
Vendor ID	Specify the Vendor ID. Default is 0. Range is 0-2147483647. Once defined, FortiADC will change the Vendor-ID AVP of the Diameter packet.
Product Name	Specify the Product Name. Once 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. This refers to the built-in session ID persistence timeout.
Server Close Propagation	Disabled by default, this means that the connection on the client side stays open when the server closes the connection on its side.
Client SSL	Enable or disable Client SSL. This is disabled by default.

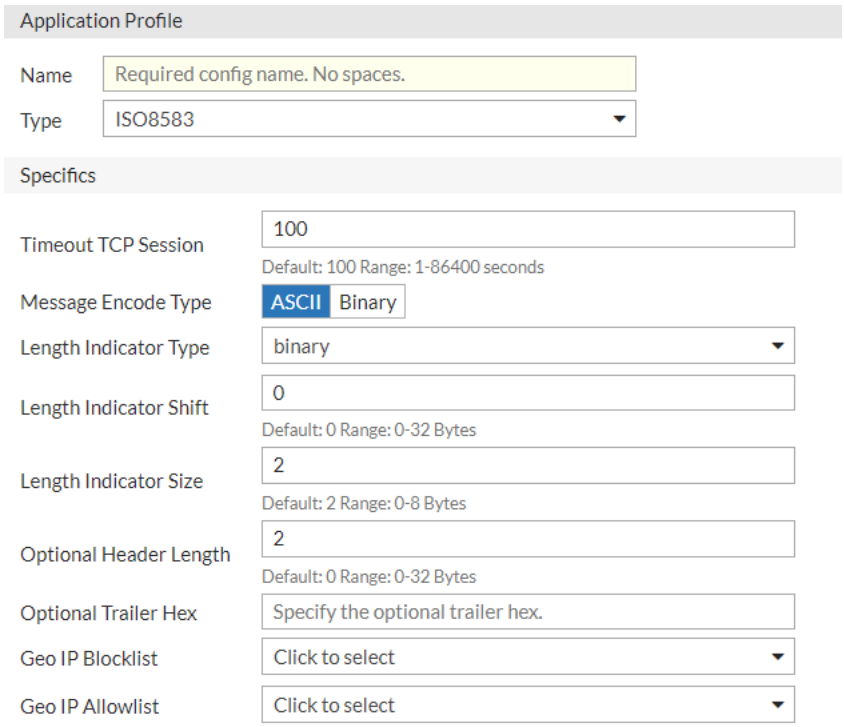
Predefined Diameter Profile: LB_PROF_DIAMETER

Parameter	Default value
Origin Host	Blank
Origin Realm	Blank
Vendor ID	0
Product Name	Blank
Idle Timeout	300
Server Close Propagation	Disable
Client SSL	Disable

Note: When the LB_PROF_DIAMETER profile is applied, FortiADC will not change Diameter packets except the host IP address AVP, which means that FortiADC functions as a relay agent.

Navigate back to Application Profile Types on page 451.

ISO8583



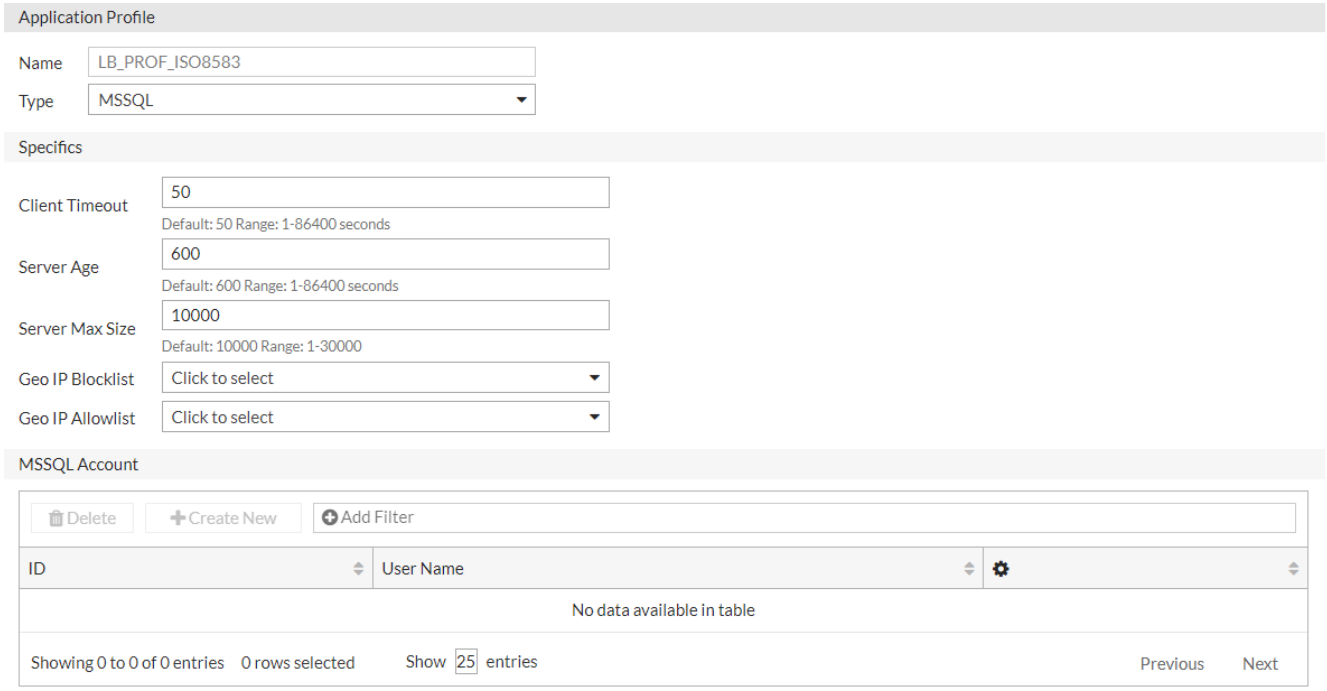
Parameter	Description
Timeout TCP Session	Client-side timeout for connections where the client has not sent a FIN signal, but the connection has been idle. The default is 100 seconds. The valid range is 1 to 86,400 seconds.
Message Encode Type	Specify the encode type for protocol message: • ASCII • Binary The default option is ASCII.
Length Indicator Type	Specify the encode type of length indicator: • binary • BCD • decimal-str • hex-str The default option is binary.
Length Indicator Shift	Specify bytes to shift from the beginning of payload to read length value. Default is 0, range 0-32.
Length Indicator Size	Specify total bytes reading to calculate length. Default is 2, range 0-8.
Optional Header Length	Specify length of optional header before MTI, including the length-indicator. Default is 0, range 0-32.
Optional Trailer Hex	Specify hex string of optional trailer, maximum length 16, i.e. 8 bytes in binary
Geo IP Blocklist	Select a Geo IP block list configuration object. See Using the Geo IP block list.
Geo IP Allowlist	Select an allowlist configuration object. See Using the Geo IP allowlist.

Predefined ISO8583 Profile: LB_PROF_ISO8583

Parameter	Default value
Timeout TCP Session	100
Message Encode Type	ASCII
Length Indicator Type	binary
Length Indicator Shift	0
Length Indicator Size	2
Optional Header Length	2
Optional Trailer Hex	None
Geo IP Blocklist	None
Geo IP Allowlist	None

Navigate back to Application Profile Types on page 451.

MSSQL

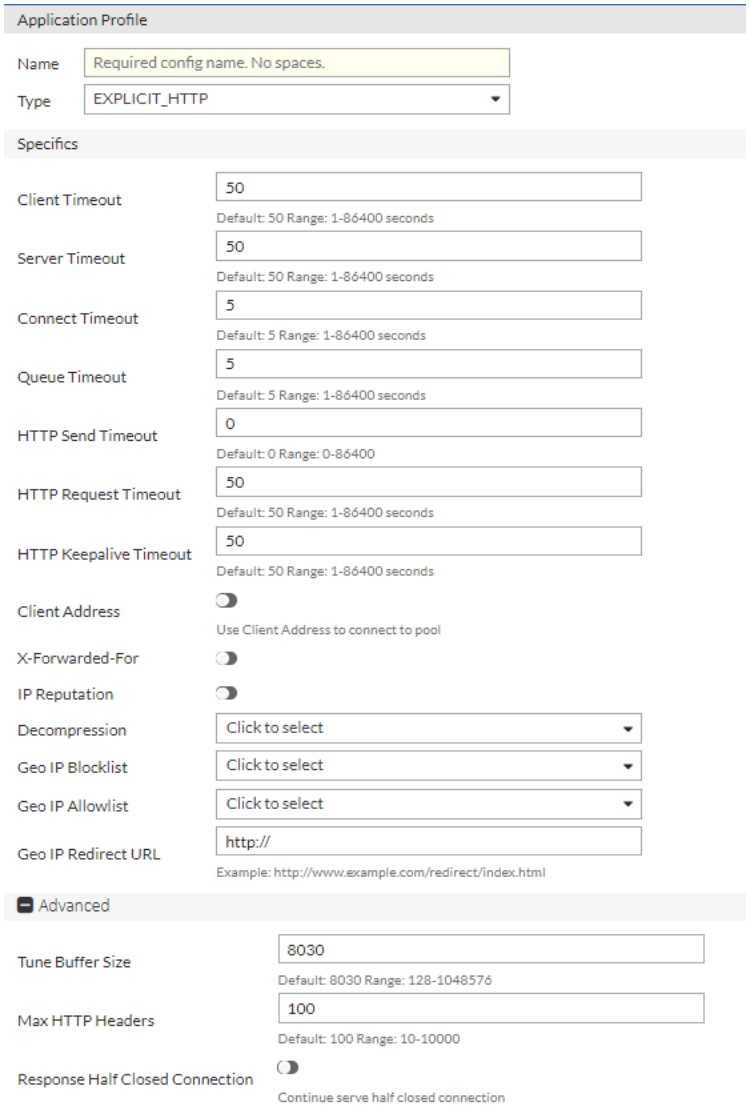


Parameter	Description
Client Timeout	This timeout is counted as the amount of time when the client did not send a complete request HTTP header to the FortiADC after the client connected to the FortiADC. If this timeout expires, FortiADC will send a 408 message to client and close the connection to the client. The default is 50 seconds. The valid range is 1 to 86,400 seconds.
Server Age	Specify the maximum inactivity time for MS SQL server on the server side.
Server Max Size	Specify the maximum connections that can connect to the MS SQL server on the server side.
Geo IP Blocklist	Select a Geo IP block list configuration object. See Using the Geo IP block list.
Geo IP Allowlist	Select an allowlist configuration object. See Using the Geo IP allowlist.

For more configuration details, see Configuring MSSQL profiles on page 506.

Navigate back to Application Profile Types on page 451.

EXPLICIT_HTTP



Parameter	Description
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.

Parameter	Description
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.
Client Address	Use the original client IP address as the source address when connecting to the real server.
X-Forwarded-For	Enable this option to append the client IP address found in IP layer packets to the HTTP header, for example, `X-forwarded-for: 192.168.161.100`. The default header name is `X-forwarded-for`. If you prefer a different name, use **X-Forwarded-For Header** to define a custom name.
X-Forwarded-For Header	Specify a custom name for the HTTP header which carries the client IP address. Do not include the 'X-' prefix. Examples: Forwarded-For, Real-IP, or True-IP.
IP Reputation	Enable to apply the FortiGuard IP reputation service. See Managing IP Reputation policy settings.
IP Reputation Redirect URL	Type a URL including the FQDN/IP and path, if any, to which a client will be redirected if the request violates the IP reputation policy.
Decompression	Select a compression configuration object. See Configuring compression rules.
Geo IP Blocklist	Select a Geo IP block list configuration object. See Using the Geo IP block list.
Geo IP Allowlist	Select an allowlist configuration object. See Using the Geo IP allowlist.
Geo IP Redirect URL	For HTTP, if you have configured a Geo IP redirect action, specify a redirect URL.
Advanced Settings	

Parameter	Description
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. • 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.
Response Half Closed Connection	Enable to continue to serve half-closed connections.

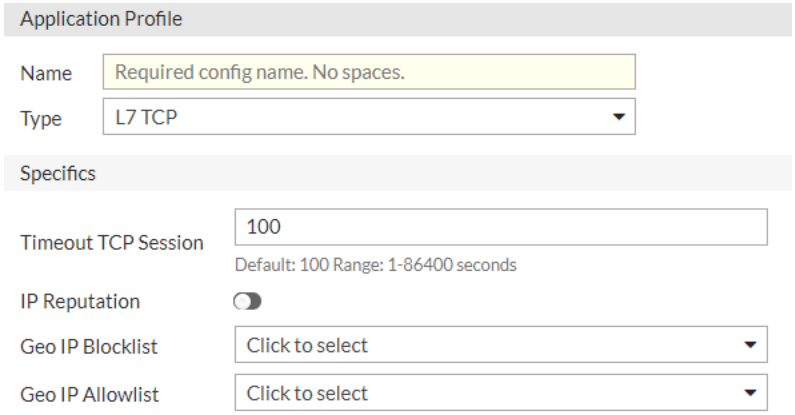
Predefined EXPLICIT_HTTP Profile: LB_PROF_EXPLICIT_HTTP

Parameter	Default value
Client Timeout	50
Server Timeout	50
Connect Timeout	50
Queue Timeout	50
HTTP Send Timeout	0
HTTP Request Timeout	50
HTTP Keepalive Timeout	50
Client Address	Disabled
X-Forwarded-For	Disabled
X-Forwarded-For Header	None
IP Reputation	Disabled
Decompression	None
Geo IP Blocklist	None
Geo IP Allowlist	None
Geo IP Redirect URL	http://
Tune Buffer Size	8030

Parameter	Default value
Max HTTP Headers	100
Response Half Closed Connection	Disabled

Navigate back to Application Profile Types on page 451.

L7 TCP



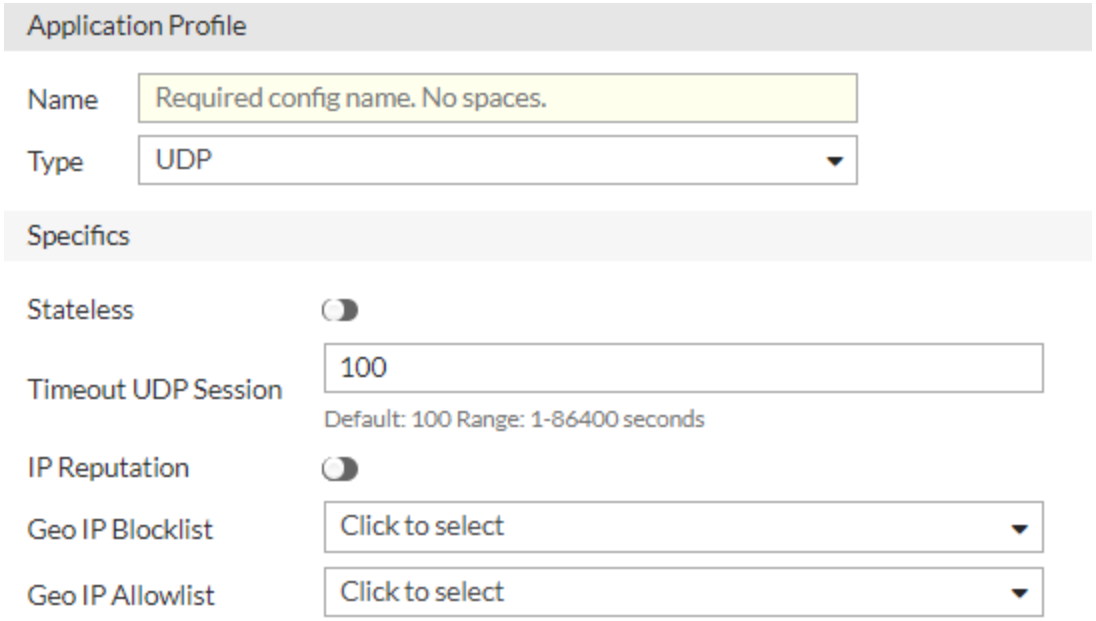
Parameter	Description
Timeout TCP Session	Client-side timeout for connections where the client has not sent a FIN signal, but the connection has been idle. The default is 100 seconds. The valid range is 1 to 86,400.
IP Reputation	Enable to apply the FortiGuard IP reputation service. See Managing IP Reputation policy settings.
Geo IP Blocklist	Select a Geo IP block list configuration object. See Using the Geo IP block list.
Geo IP Allowlist	Select an allowlist configuration object. See Using the Geo IP allowlist.

Predefined L7 TCP Profile: LB_PROF_L7_TCP

Parameter	Default value
Timeout TCP Session	100
IP Reputation	Disabled
Geo IP Blocklist	None
Geo IP Allowlist	None

Navigate back to Application Profile Types on page 451.

L7 UDP



Parameter	Description
Timeout UDP Session	Client-side session timeout. The default is 100 seconds. The valid range is 1 to 86,400 seconds.
IP Reputation	Enable to apply the FortiGuard IP reputation service. See Managing IP Reputation policy settings.
Geo IP Blocklist	Select a Geo IP block list configuration object. See Using the Geo IP block list.
Geo IP Allowlist	Select an allowlist configuration object. See Using the Geo IP allowlist.

Predefined L7 UDP Profile: LB_PROF_L7_UDP

Parameter	Default value
Timeout UDP Session	100
IP Reputation	Disabled
Geo IP Blocklist	None
Geo IP Allowlist	None

Navigate back to Application Profile Types on page 451.

See also:

- Configuring MySQL profiles on page 499
- Configuring MSSQL profiles on page 506

- WebSocket load-balancing on page 509
- TCP multiplexing on page 510

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.

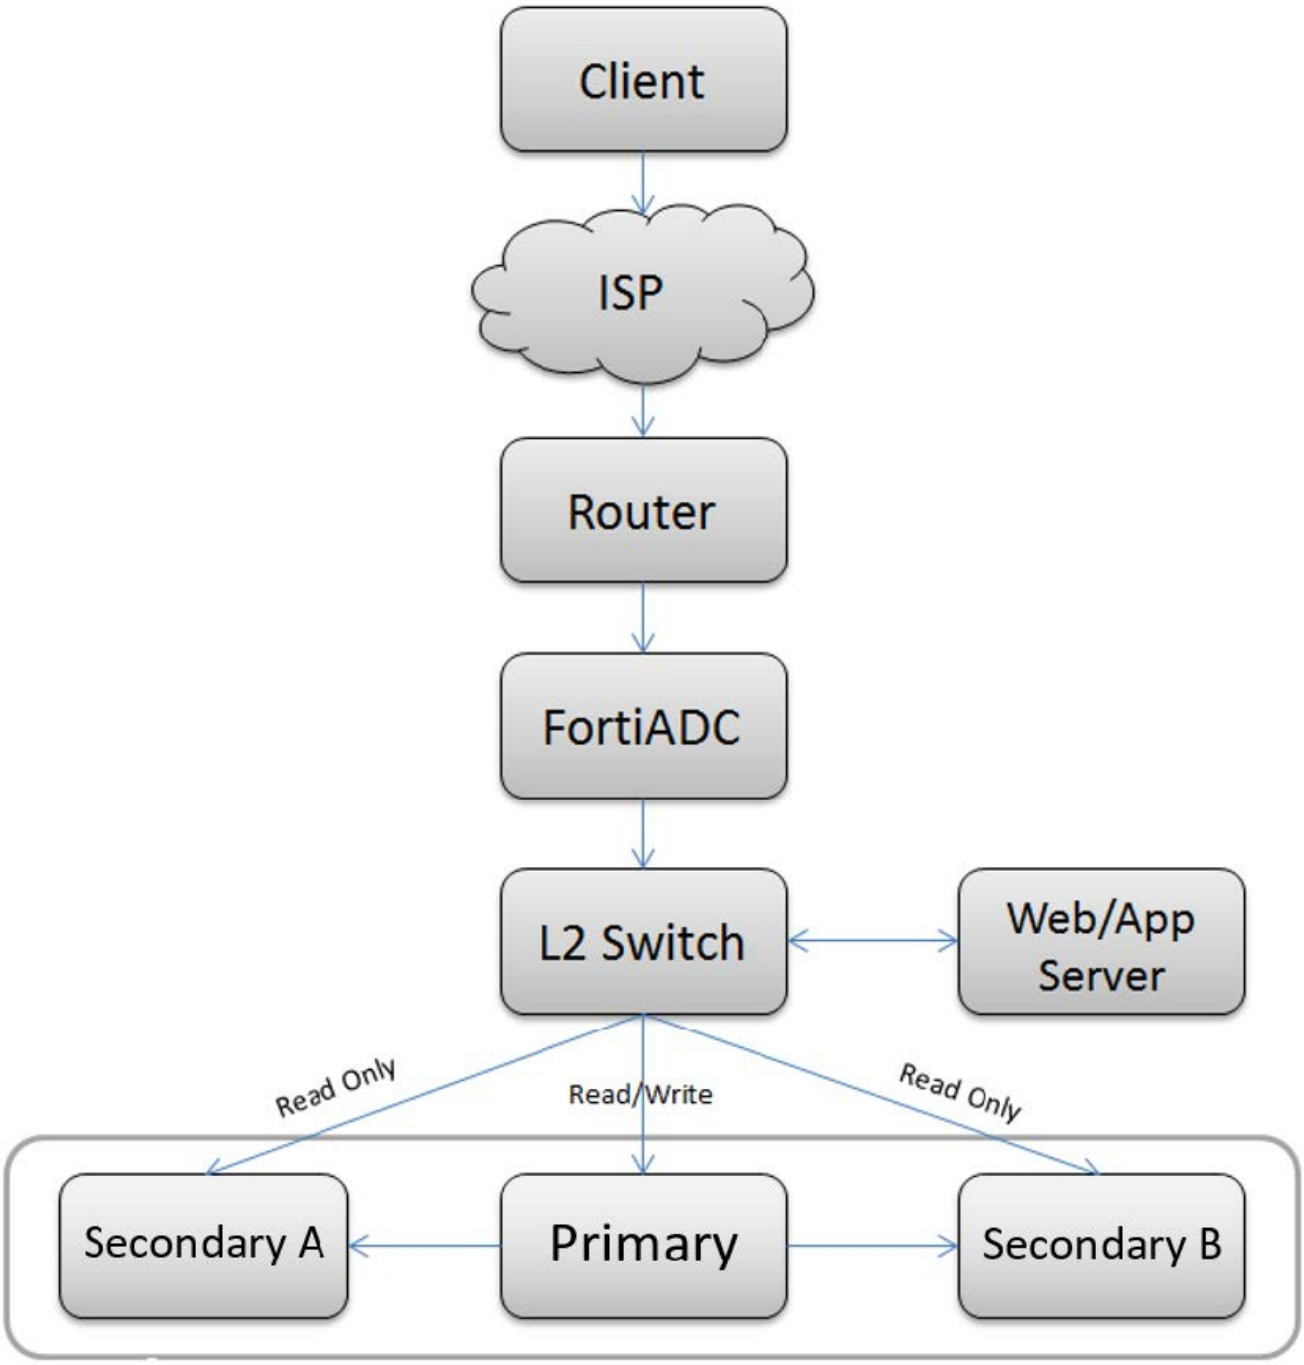
FortiADC supports two MySQL database load-balancing modes: single primary and data sharding.

Single-primary mode

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

Single-primary mode on page 499 illustrates the network topology of database server load-balancing in single-primary mode.

Single-primary mode



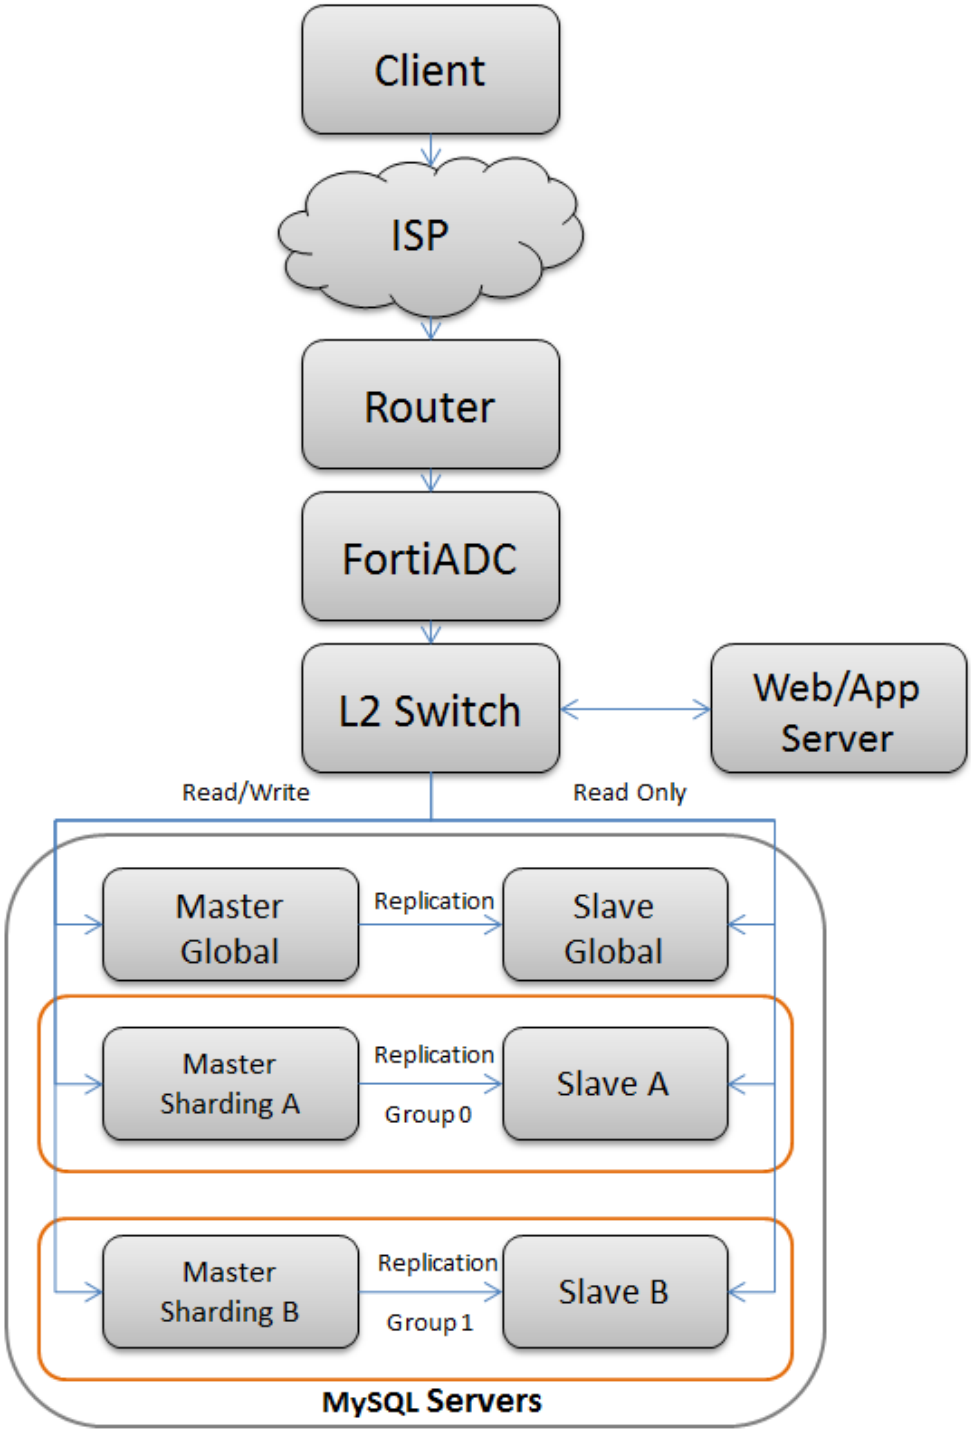
By default, FortiADC passes all write requests to the primary server and all read requests (such as select) to the secondary servers. So once you have created a MySQL server load-balancing profile, FortiADC will automatically apply this default mode when load-balancing MySQL traffic on the network. However, 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 504.

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 501 illustrates MySQL server load-balancing in data-sharding mode.

Sharding mode



In sharding mode, FortiADC stores global data on the primary 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) 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-primary 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 CLI Reference.

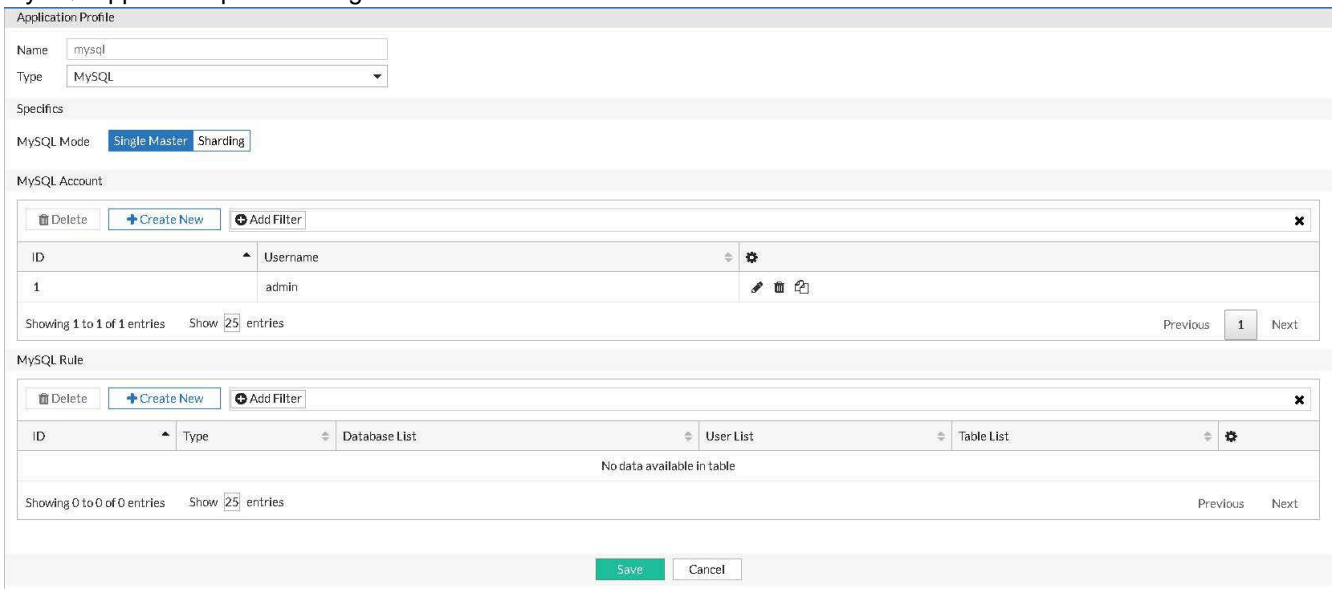
Before you begin:

- 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 primary or Sharding**. Refer to MySQL profile configuration guidelines on page 505.
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 503.

MySQL application profile configuration



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 primary database server or the secondary 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 505.
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 505.
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 primary
next
edit 2
set pool_member_service_port 3306
set pool_member_cookie rs2
set real-server primary2
set mysql-group-id 1
next
edit 3
set pool_member_service_port 3306
set pool_member_cookie rs3
set real-server secondary
set mysql-read-only enable
next
edit 4
set pool_member_service_port 3306
set pool_member_cookie rs4
set real-server secondary2
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 primary—If selected, FortiADC will configure the MySQL profile in single-primary mode. See Single-primary 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.

Parameter	Description
Password	The password for the MySQL user name you've entered above.
MySQL Rule	
Type	Select either of the following: • Primary—If selected, FortiADC will send all data specified in the MySQL rule to the primary MySQL database server. • Secondary—If selected, FortiADC will send all data specified in the MySQL rule to the secondary MySQL database server.
Database List	A list of up to eight MySQL database names separated by space
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.

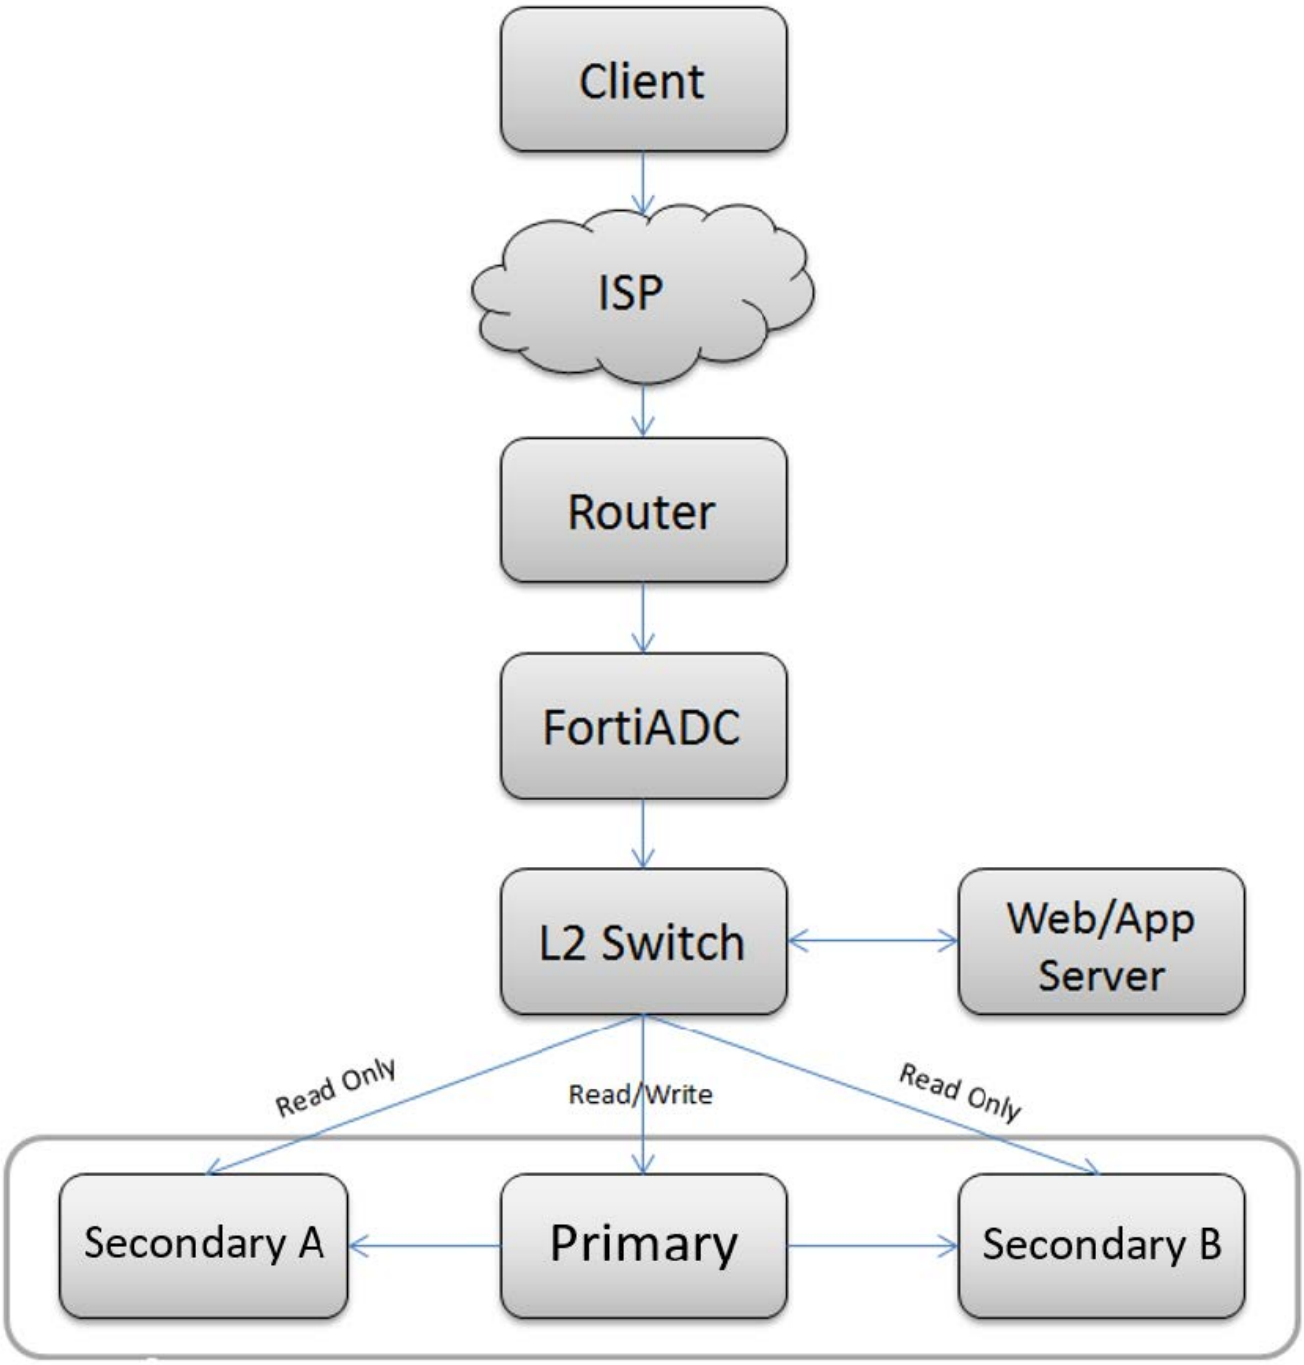
Configuring MSSQL profiles

FortiADC (6.0.0 and later) supports MSSQL server load-balancing.

MSSQL 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 MSSQL application profiles that you can use out of the box. You must configure your own MSSQL load-balancing application profiles to take advantage of this feature.

Single-primary mode

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



By default, FortiADC passes all write requests to the primary server and all read requests (such as select) to the secondary servers. Once you have created a MSSQL server load-balancing profile, FortiADC will automatically apply this default mode when load-balancing MSSQL traffic on the network.

Creating a MSSQL profile

Creating a MSSQL profile involves the following steps:

1. Create a MSSQL configuration object.
2. Specify the existing user name and password of the MSSQL database to be used by the MSSQL profile configuration object.

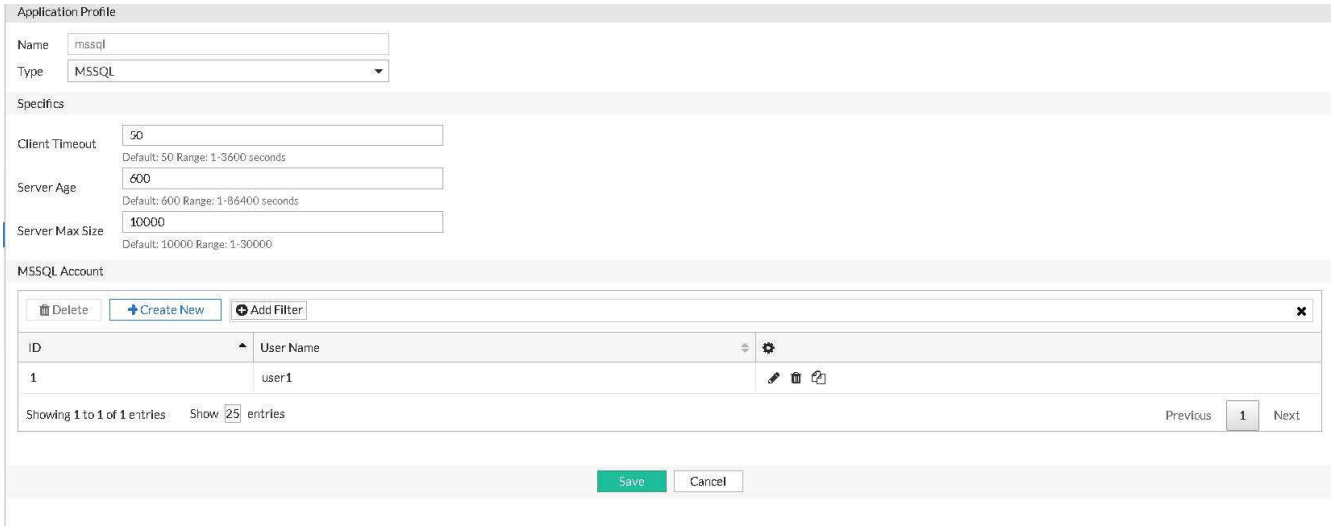
Note: You can create MSSQL profiles from either the GUI or the CLI. The following paragraphs discuss how to configure a MSSQL profile using the GUI. For instructions on how to create MSSQL profiles from the CLI, refer to the CLI Reference.

Before you begin:

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

Creating a MSSQL 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 MSSQL from the drop-down menu.
6. Click **Save**. Your newly created MSSQL profile configuration object is automatically appended to the bottom of the **Server Load Balancing > Application Resources > Application Profile** page.
7. Click the newly created MSSQL profile to open it to see the MSSQL application profile configuration.



Specifying the MSSQL user account

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

Note that you are asked to provide the user name and password of an existing MSSQL account, so do not try to create a new user account here.

To specify a MSSQL user account:

1. In the **MSSQL User Password** pane, click **Create New**. The Edit MSSQL User Password dialog opens.
2. Enter the user name and password of the MSSQL database account,

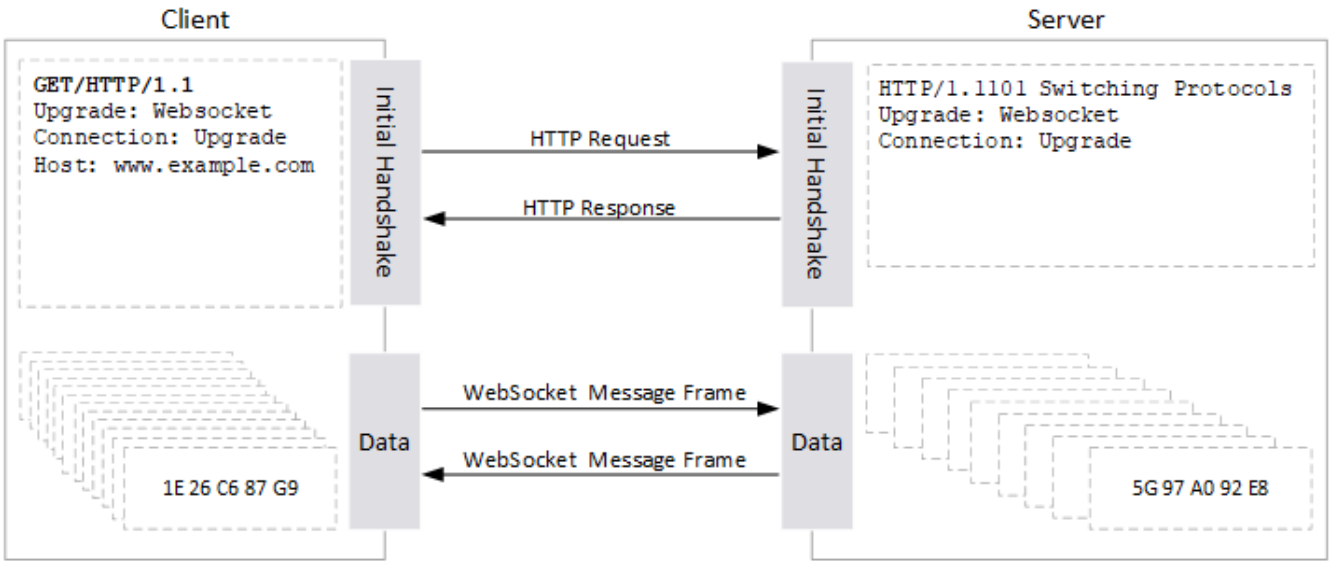
3. Click **Save**.

Parameter	Description
Application Profile	
Name	A unique name for the MSSQL profile you are creating.
Type	MSSQL
MSSQL Account	
User Name	The user name of the MSSQL database.
Password	The password for the MSSQL user name you've entered above.
Specifics	
Client Timeout	Client connection timeout
Server Age	Server connection timeout
Server Max Size	The maximum size of server connection

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 509.

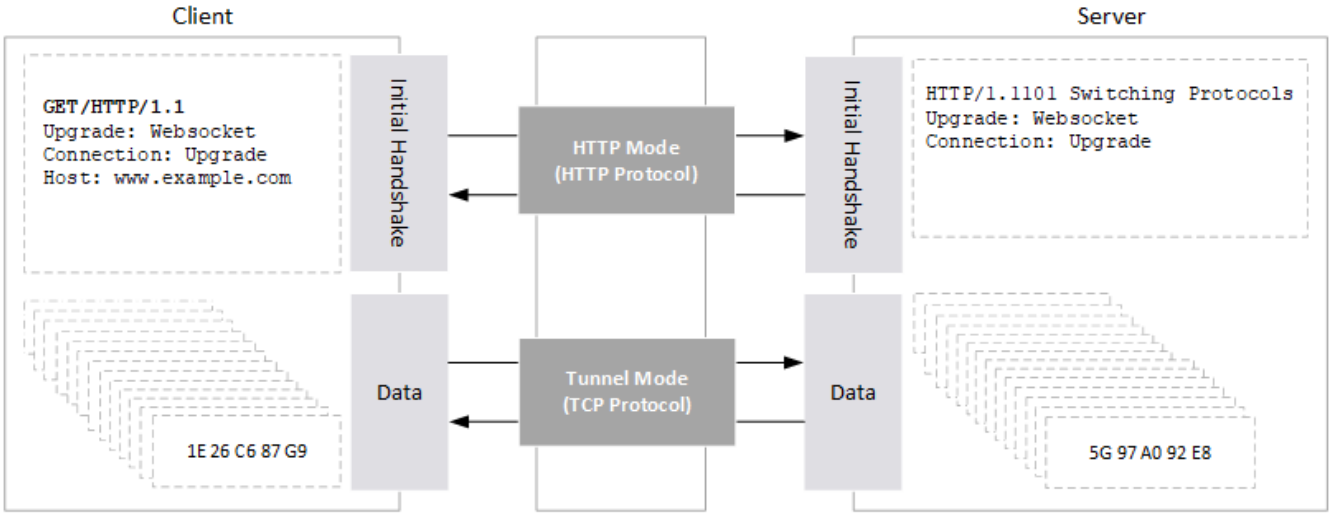
WebSocket load-balancing



You can configure FortiADC in such as 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 510.

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.

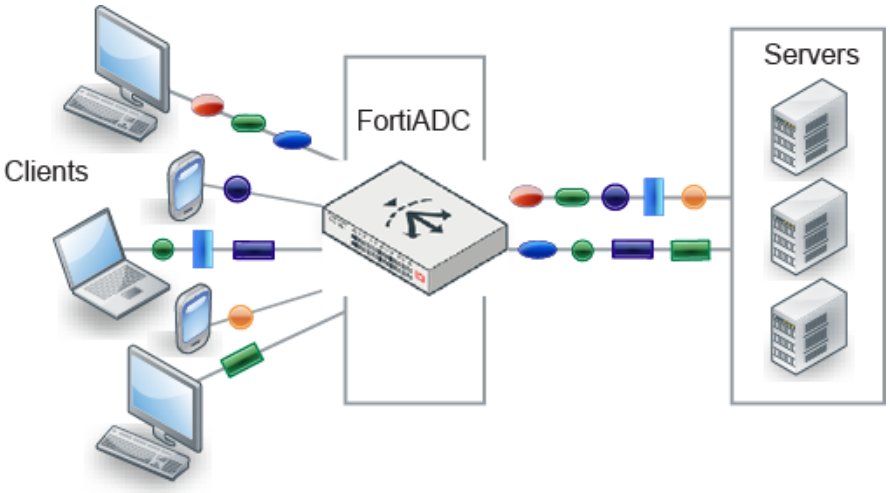
If you only want WebSocket load-balancing, use a Layer-4 or Layer-7 virtual server with a TCP profile.

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

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 510.

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.

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:

- 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:

1. Go to **Server Load Balance > Application Resources**.
2. Click the **Client SSL Profile** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in Client SSL profile configuration guidelines on page 512.
5. 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-CAMELLIA256-SHA384 • *ECDHE-ECDSA-AES256-SHA • ECDHE-ECDSA-AES128-GCM-SHA256 • *ECDHE-ECDSA-AES128-SHA256 • ECDHE-ECDSA-CAMELLIA128-SHA256 • *ECDHE-ECDSA-AES128-SHA

Type	Profile Configuration Guidelines
	• ECDHE-ECDSA-DES-CBC3-SHA • ECDHE-ECDSA-RC4-SHA • ECDHE-RSA-AES256-GCM-SHA384 • ECDHE-RSA-AES256-SHA384 • ECDHE-RSA-CAMELLIA256-SHA384 • *ECDHE-RSA-AES256-SHA • DHE-RSA-AES256-GCM-SHA384 • *DHE-RSA-AES256-SHA256 • DHE-RSA-CAMELLIA256-SHA256 • *DHE-RSA-AES256-SHA • DHE-RSA-CAMELLIA256-SHA • AES256-GCM-SHA384 • *AES256-SHA256 • *AES256-SHA • ECDHE-RSA-AES128-GCM-SHA256 • *ECDHE-RSA-AES128-SHA256 • ECDHE-RSA-CAMELLIA128-SHA256 • DHE-RSA-AES128-GCM-SHA256 • *DHE-RSA-AES128-SHA256 • DHE-RSA-CAMELLIA128-SHA256 • *DHE-RSA-AES128-SHA • AES128-GCM-SHA256 • *AES128-SHA256 • *AES128-SHA • ECDHE-RSA-RC4-SHA • RC4-SHA • RC4-MD5 • EDH-RSA-DES-CBC3-SHA • DES-CBC3-SHA • eNULL *These ciphers are fully supported by hardware SSL (in 400F, 420F, 1200F, 2200F, 4200F and 5000F). Note: We recommend retaining the default list. If necessary, you can deselect the SSL ciphers that you do not want to support.
TLSv1.3 Cipher Suite List	TLSv1.3 ciphers are listed as following: • TLS_AES_256_GCM_SHA384 • TLS_AES_128_GCM_SHA256 • TLS_CHACHA20_POLY1305_SHA256 • TLS_AES_128_CCM_SHA256 • TLS_AES_128_CCM_8_SHA256 Note: This option only available if the TLSv1.3 is checked.
Allowed SSL Versions	You have the following options: • SSLv3

Type	Profile Configuration Guidelines
	• 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. Please make sure that the SSL versions are continuous. IF not, an error message should be returned. • RFC 7919 Comply cannot support SSLv3 and TLS 1.3. If RFC 7919 Comply is enabled and **SSLv3** or **TLSv1.3** is selected in **Allowed SSL Versions**, an error message will display.
Client Certificate Verify	Select the client certificate verify configuration object. **Note**: For VS configurations that reference a ZTNA Profile, ensure the corresponding EMS CA certificate is selected for the corresponding Client SSL profile.
Client Certificate Verify Mode	The **Client Certificate Verify Mode** option is available if the **Client Certificate Verify** is selected. Select one of the following: • Required (default) • Optional
SSL Session Cache Flag	Allows to the same SSL client attempts to reconnect to this SSL server and requests a resumption of a previous SSL session. Note: This feature doesn't support TLSv1.3
Use TLS Tickets	Allows resuming TLS sessions by storing key material encrypted on the clients. Note: This feature doesn't support TLSv1.3
Client Certificate Forward	The **Client Certificate Forward** option is available if the **Client Certificate Verify** is selected. Disabled by default. When enabled, you must specify the client certificate forward header.
Client Certificate Forward Header	The **Client Certificate Forward Header** option is available if the **Client Certificate Verify** is selected and **Client Certificate Forward** is enabled. 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. **Note**: RFC 7919 Comply is not supported for Forward Proxy. If RFC 7919 Comply is enabled and Forward Proxy is enabled, the RFC 7919 Comply feature will not apply to Forward Proxy functionality.

Type	Profile Configuration Guidelines
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: • The feature is disabled by default. • When enabled, you must configure the options below.
Renegotiation Interval	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. **Note**: The SSL DH Parameter Size option is not available when **RFC 7919 Comply** is enabled.
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 whenFortiADC 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 thatFortiADC 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

Type	Profile Configuration Guidelines
	clients that do not support secure renegotiation.
RFC 7919 Comply	Enable/disable parameters to comply with RFC 7919. **Note**: • RFC 7919 Comply is not supported for Forward Proxy. If RFC 7919 Comply is enabled and Forward Proxy is enabled, the RFC 7919 Comply feature will not apply to Forward Proxy functionality. • RFC 7919 Comply cannot support SSLv3 and TLS 1.3. If RFC 7919 Comply is enabled and **SSLv3** or **TLSv1.3** is selected in **Allowed SSL Versions**, an error message will display. • When RFC 7919 Comply is enabled the **SSL DH Parameter Size** option becomes unavailable.
Supported Groups	The **Supported Groups** option is available if **RFC 7919 Comply** is enabled. Specify the supported group objects from the following: • secp256r1 • secp384r1 • secp521r1 • x25519 • x448 • ffdhe2048 • ffdhe3072 • ffdhe4096 • ffdhe6144 • ffdhe8192 At least one item from the FFDHE group must be selected. **Note**: The RFC 7919 Comply feature requires certain cipher selections to correspond with the Supported Group selection. • If a FFDHE group is selected (for example, ffdhe2048), then at least one cipher must be DHE-RSA (for example, DHE-RSA-AES256-SHA256). • If the Supported Group includes groups other than FFDHE (such as a SECP group, secp256r1), then at least one cipher must be ECDHE (for example, ECDHE-ECDSA-AES256-GCM-SHA384). • If a ECDHE cipher is selected (for example, ECDHE-ECDSA-AES256-GCM-SHA384), then the Supported Group must include at least one group that is not FFDHE (such as a SECP group, secp256r1).
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.
Forward Proxy Certificate Caching	The **Forward Proxy Certificate Caching** option is available if **Forward Proxy** is **enabled**.

Type	Profile Configuration Guidelines
	Select a Forward Proxy Certificate Caching rule.
Forward Proxy Local Signing CA	The **Forward Proxy Local Signing CA** option is available if **Forward Proxy** is **enabled**. Select a Forward Proxy Local Signing CA.
Forward Proxy Intermediate CA Group	The **Forward Proxy Intermediate CA Group** option is available if **Forward Proxy** is **enabled**. Select a Forward Proxy Intermediate CA Group.
Forward Proxy Resign Cert by SNI	The **Forward Proxy Resign Cert by SNI** option is available if **Forward Proxy** is **enabled**. Enabled by default, the **Forward Proxy Resign Cert by SNI** option allows SSL Forward Proxy to return re-signed local certificates with modified CN (or SAN) matching the ClientHello Server Server Name Indication (SNI). This allows the proxy to act automatically when the SNI is detected in the ClientHello message to return a re-signed local certificate to the client. The Common Name (CN) and/or Subject Alternative Name (SAN) of this certificate will be adeptly modified to align with the SNI, ensuring a seamless and trustworthy SSL handshake process. **Note**: This function is supported HTTP/S virtual servers and real server pools with HTTP/S topology.
Backend SSL SNI Forward	The **Backend SSL SNI Forward** option is available if **Forward Proxy** is **enabled**. 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	The **Backend Customized SSL Ciphers Flag** option is available if **Forward Proxy** is **enabled**. Enabled by default. In this case, you must specify the backend customized SS ciphers. See below.
Backend Customized SSL Ciphers	The **Backend Customized SSL Ciphers** option is available if **Forward Proxy** and **Backend Customized SSL Ciphers Flag** are **enabled**. Specify the customized SSL ciphers to be supported at the back end.
Backend SSL Cipher Suite List	The **Backend SSL Cipher Suite List** option is available if **Forward Proxy** is **enabled** and **Backend Customized SSL Ciphers Flag** is **disabled**. Select the cipher from the list to be supported at the back end.
Backend TLSv1.3 Cipher Suite List	The **Backend TLSv1.3 Cipher Suite List** option is available if **Forward Proxy** is **enabled** and **Backend Allowed SSL Versions** selected **TLSv1.3**. TLSv1.3 ciphers are listed as following: TLS_AES_256_GCM_SHA384 TLS_AES_128_GCM_SHA256 TLS_CHACHA20_POLY1305_SHA256 TLS_AES_128_CCM_SHA256 TLS_AES_128_CCM_8_SHA256

Type	Profile Configuration Guidelines
	Note: This option only available if the backendTLSv1.3 is checked.
Backend Allowed SSL Versions	The **Backend Allowed SSL Versions** option is available if **Forward Proxy** is **enabled**. 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.
Backend SSL OCSP Stapling Support	The **Backend SSL OCSP Stapling Support** option is available if **Forward Proxy** is **enabled**. Disabled by default. Enable it to let FortiADC support OCSP stapling at the backend.

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.

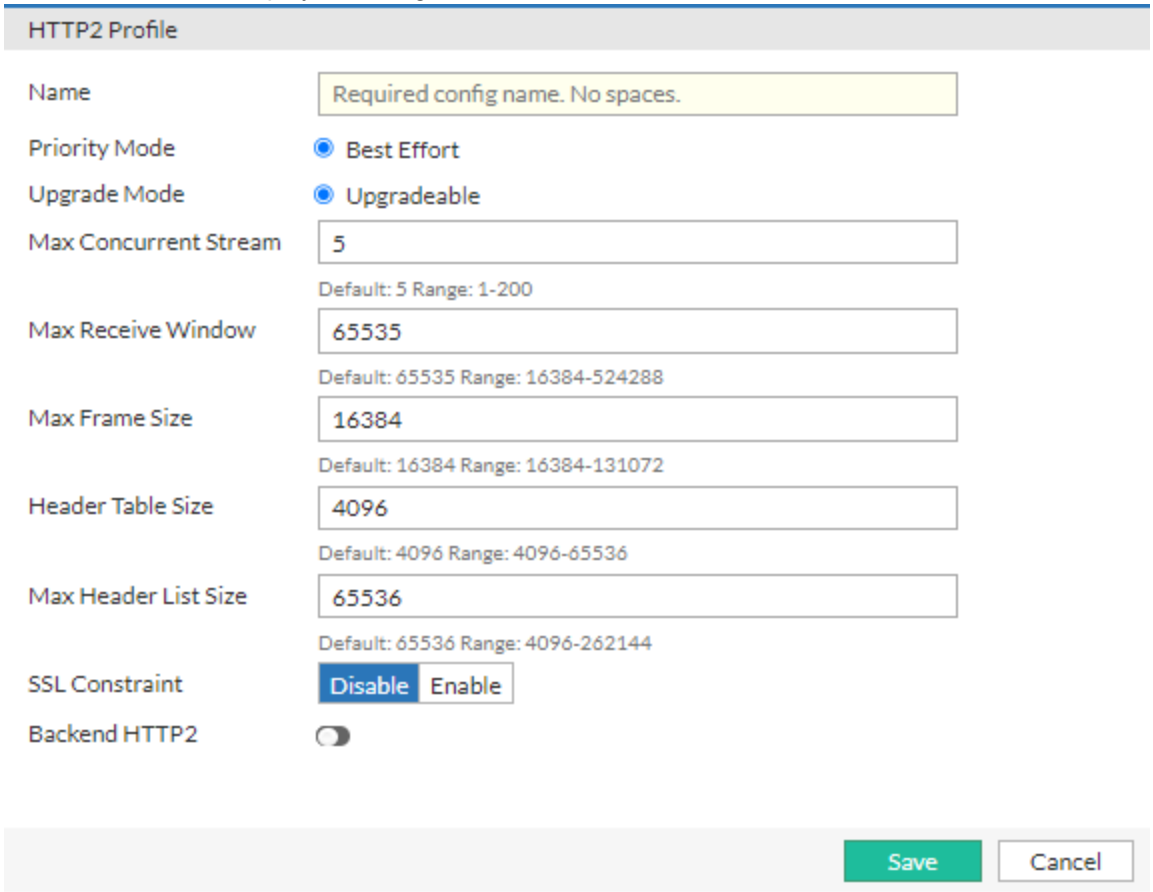
Alternatively, predefined profiles are available to be referenced in HTTP/HTTPS application profiles. All values in the predefined profile is view-only and cannot be modified.

Profile	Description
LB_HTTP2_PROFILE_ DEFAULT	Priority Mode — Best Effort Upgrade Mode — Upgradeable Max Concurrent Stream — 5 Max Receive Window — 65535 Max Frame Size — 16384 Header Table Size — 4096 Max Header List Size — 65536 SSL Constraint — Enable Backend HTTP2 — Disable
LB_HTTP2_PROFILE_ END2END_H2	Priority Mode — Best Effort Upgrade Mode — Upgradeable Max Concurrent Stream — 5 Max Receive Window — 65535 Max Frame Size — 16384 Header Table Size — 4096 Max Header List Size — 65536 SSL Constraint — Enable Backend HTTP2 — Enable Backend Max Receive Window — 65535 Backend Concurrent Stream — 5

Profile	Description
	Backend Proto Mode HTTPS — ALPN Backend Proto Mode HTTP — Force H1 Backend Multiplex Mode — Multi Connection

To configure HTTP2 profiles:

1. Go to **Server Load Balance > Application Resources**.
2. Click the **HTTP2 Profile** tab.
3. Click **Create New** to display the configuration editor.



4. Configure the following settings:

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, and the valid range is 1-200.

Type	Profile Configuration Guidelines
Max Receive Window	Specify the maximum number of bytes that can be received without sending an acknowledgment response. The default value is 65535 bytes, and the valid range is 16384-524288.
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 value is 16384 bytes, and the valid range is 16384-131072. **Note**: When **Backend HTTP2** is **enabled**, the **Max Frame Size** is not supported, as this cannot be set independently for the frontend and backend. Instead, the **HTTP2 Profile Max Frame Size** will override the **Tune Buffer Size** in the **Application Profile**.
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 value is 4096, and the valid range is 4096-65536.
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 value is 65536, and the valid range is 4096-262144.
SSL Constraint	Enable or disable SSL constraint. If enabled, the following conditions must be met: • The TLS implementation supports Server Name Indication. • 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.
Backend HTTP2	Enable/disable support for the backend HTTP/2 functionality. When enabled, the related virtual server will switch to httproxy3 for support. This is disabled by default. **Note**: The backend HTTP/2 implementation is built on HTTP/3, which introduces specific limitations. For details, see Configuring HTTP3 profiles on page 521.
Backend Max Receive Window	The **Backend Max Receive Window** option is available if **Backend HTTP2** is **enabled**. Specify the init-windows-size configuration for the backend HTTP/2 connection. The default value is 65535, and the valid range is 16384-524288.
Backend Concurrent Stream	The **Backend Concurrent Stream** option is available if **Backend HTTP2** is **enabled**.

Type	Profile Configuration Guidelines
	Specify the maximum limit for concurrent streams that the backend server can handle to ensure optimal performance and prevent overloading. The default value is 5, and the valid range is 1-200.
Backend Proto Mode HTTPS	The **Backend Proto Mode HTTPS** option is available if **Backend HTTP2** is **enabled**. Select the HTTPS server backend HTTP/2 protocol mode. • ALPN — Use Application-Layer Protocol Negotiation (ALPN). • Force H1 — Enforce HTTP/1. • Force H2 — Enforce HTTP/2. The default is ALPN.
Backend Proto Mode HTTP	The **Backend Proto Mode HTTP** option is available if **Backend HTTP2** is **enabled**. Select the HTTP server backend HTTP/2 protocol mode. • Force H1 — Enforce HTTP/1. • Force H2 — Enforce HTTP/2. The default is Force H1.
Backend Multiplex Mode	The **Backend Multiplex Mode** option is available if **Backend HTTP2** is **enabled**. Select the backend multiplexing mode. • Multi Connection — Multiple streams from the frontend are mapped to multiple backend connections. • Single Connection — All requests from multiple frontend connections are sent through a single backend connection. The default is Multi Connection.

5. Click **Save**.

Once the HTTP2 Profile configuration is saved, it can be referenced in an HTTP/HTTPS Application Profile configuration.

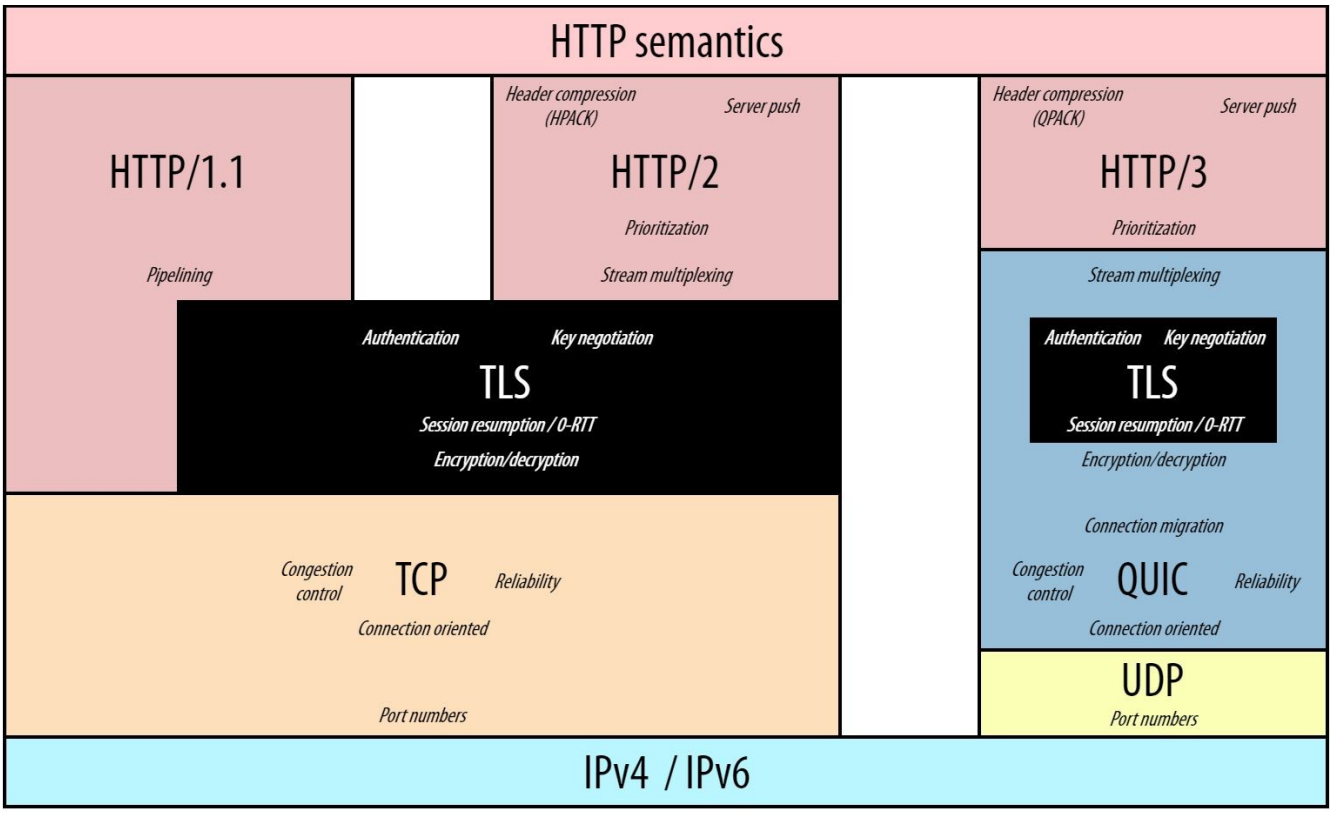
Configuring HTTP3 profiles

HTTP/3 Protocol Overview

HTTP/3 is the latest version of the HTTP protocol and unlike previous versions which relied on TCP to handle streams in the HTTP layer, HTTP/3 uses QUIC (Quick UDP Internet Connections), a multiplexed transport protocol built on UDP. The HTTP/3 protocol has a lower latency as a result of using QUIC, allowing it to have a quicker handshake for establishing a secure session compared to HTTP/2 which achieves this using TCP and TLS.



In version 7.4.0, FortiADC is introducing HTTP/3 support as an experimental feature with limited HTTP/3 functionality, so it is not recommended to be used in production environments. For details, see HTTP/3 supported functionality and limitations on page 525.



Advantages of HTTP/3 using QUIC

As a result of using QUIC, HTTP/3 is designed to improve the speed, reliability, and security of data transfer over the internet.

Faster connection setup (lower latency)

The HTTP/3 protocol has a lower latency as a result of using QUIC, allowing it to have a quicker handshake for establishing a secure session compared to HTTP/2 which achieves this using TCP and TLS.

Customizable congestion control

QUIC allows implementers to customize the congestion control algorithm used within the protocol.

Enhanced security

QUIC encrypts almost all of its packet header fields by default default (via TLS).

Connection migration

QUIC identifies connections based on the ConnectionID (CID) instead of Connection table, which allows migration between different network interfaces or IP addresses without disconnecting.

Less head-of-line (HoL) blocking

In TCP, if a packet is lost or delayed, it can cause subsequent packets to be blocked. QUIC mitigates this issue by using multiple independent byte streams so that the loss or delay of a single packet does not impact the delivery of other packets.

FortiADC HTTP/3 support

You can now configure an HTTP3 Profile that can then be referenced in HTTPS application profiles. Once referenced, the HTTPS profile becomes a HTTP/3 load balance profile and the virtual server that references the profile becomes a HTTP/3 VS. This HTTP/3 VS can only operate under L7-HTTPS VS.

HTTP/3 VS listens to the TCP port and the corresponding UDP port at the same time.



FortiADC does not support server side HTTP/3, instead support is provided for client HTTP/3 to the ADC and then converted to HTTP/1 (conversion to HTTP/2 is not supported).

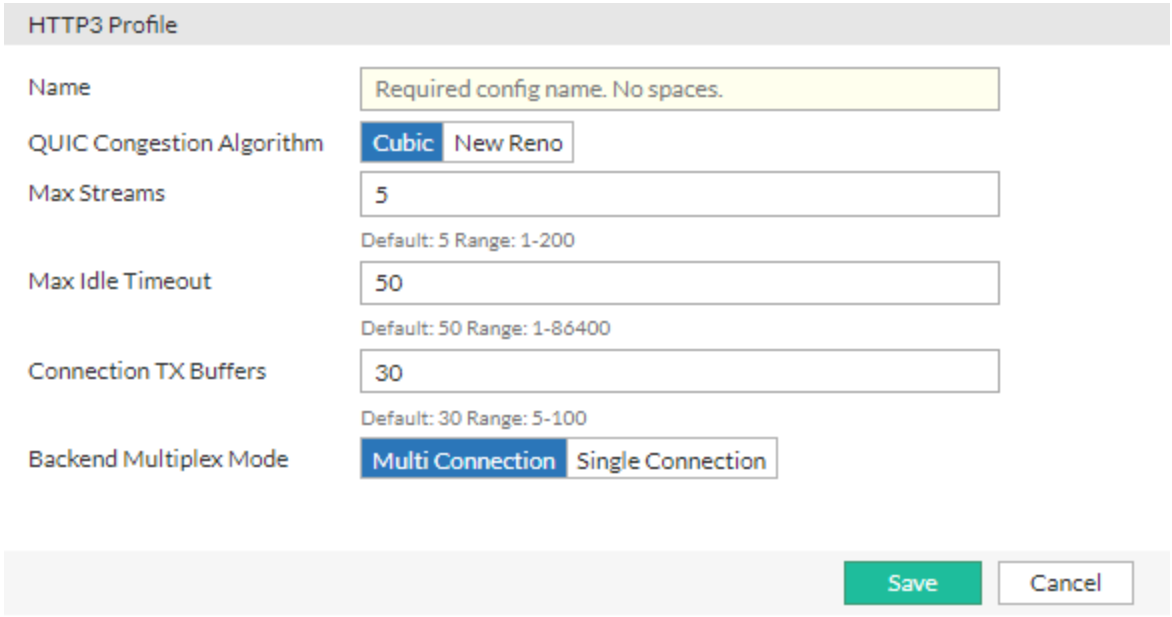
A predefined profile is available to be referenced in HTTPS application profiles. All values in the predefined profile is view-only and cannot be modified.

Profile	Description
LB_HTTP3_PROFILE_ DEFAULT	QUIC Congestion Algorithm — Cubic Max Streams — 5 Max Idle Timeout — 50 Connection TX Buffers — 30 Backend Multiplex Mode — Multi Connection

To configure an HTTP3 Profile:

1. Go to **Server Load Balance > Application Resources**.
2. Click the **HTTP3 Profile** tab.

3. Click **Create New** to display the configuration editor.



4. Configure the following settings:

Setting	Description
Name	Specify a unique name for the HTTP3 profile. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No spaces. After you initially save the configuration, you cannot edit the name.
QUIC Congestion Algorithm	FortiADC supports Cubic and New Reno loss-based congestion control for QUIC, where the congestion control responds to packet loss events. Select the QUIC congestion algorithm to use: • Cubic • New Reno Cubic is the default congestion control algorithm.
Max Streams	Specify the maximum allowable number of HTTP/3 QUIC streams. The default value is 5, and the range is 1-200.
Max Idle Timeout	Specify the HTTP/3 QUIC connection idle timeout in seconds. When no data is transmitted over the HTTP/3 connection after the specified time has elapsed, the HTTP/3 connection will timeout. The HTTP/3 connection is tracked using a unique connection-ID instead of a UDP session. The default value is 50 seconds, and the range is 1-86400 seconds.
Connection TX Buffers	Specify the number of buffers to send on the HTTP/3 QUIC connection. This parameter significantly affects the performance of the HTTP/3 response direction. The higher the number of buffers are sent, the higher the performance will be. However, the memory usage increases. The default value is 30, and the range is 5-100.
Backend Multiplex Mode	Select the backend multiplexing mode.

Setting	Description
	• Multi Connection — Multiple streams from the frontend are mapped to multiple backend connections. • Single Connection — All requests from multiple frontend connections are sent through a single backend connection. The default is Multi Connection.

5. Click **Save**.

Once the HTTP3 Profile configuration is saved, it can be referenced in an HTTPS Application Profile configuration.

HTTP/3 supported functionality and limitations

HTTP/3 support is currently an experimental feature with limited HTTP/3 functionality, so it is not recommended to be used in production environments.

Key limitations:

- HTTP/3 only operates under L7-HTTPS VS.
- HTTP/3 VS does not support dynamic configuration.
- HTTP/3 VS does not support session and persistence table display.
- HTTP/3 VS does not support HTTP detailed information statistics.
- HTTP/3 is only supported on VS, and the backend (RS) only supports HTTP/1.1.

The current iteration of the HTTP/3 feature is supported in limited or conditional capacity. The following lists the configurations that currently support HTTP/3 functionality and in what capacity.

Configuration	Supported HTTP/3 functionality
Application Profile	Profile type must be HTTPS to reference HTTP3 profiles.
Virtual Server	• VS type must be Layer 7 to reference HTTP3 profiles. • Number of ports must be set to one port only, multiple ports is not supported.
LB Method	Supported load balancing methods: • Round Robin • Least Connection • URI Hash • Full URI Hash • Host Hash • Host Domain Hash • Dynamic Load
Persistence	Supported persistence types: • Source Address Hash • Source Address-port Hash • HTTP Header Hash • HTTP Request Hash • Cookie Hash

Configuration	Supported HTTP/3 functionality
	• Persistent Cookie • Insert Cookie • Rewrite Cookie • Embedded Cookie • SSL Session ID
Client SSL Profile	Allowed SSL Versions — SSLv3, TLSv1.0, TLSv1.1, TLSv1.2 and TLSv1.3
Real Server SSL Profile	Allowed SSL Versions — SSLv3, TLSv1.0, TLSv1.1, TLSv1.2 and TLSv1.3

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 526 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 M**ethod** 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 FortiADC 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 528 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_SRC_ADDR	Persistence based on source IP address or subnet.
LB_PERSIS_HASH_SRC_ADDR	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 is based on a hash of a cookie provided by client request.
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 529.
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 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. **Note:** The persistence rule with **Match Across Virtual Servers** enabled works only with L4 virtual servers or the L7 virtual server whose Profile is LB_PROF_RADIUS.
NAT Source Pool Persistence	OFF (disabled) by default. Click the button to enable it. If enabled, the IP selected from the IP pool will persist for the session. **Note**: NAT Source Pool Persistence only works with L4 virtual servers.
Passive Cookie	
Session Keyword Type	Persistence is based on the cookie which generated from the server, including six options: auto/PHPSESSID/JESSIONID/CFID+CFTOKEN/ASP.NET_SessionId/custom. When type is auto, PHPSESSID/JESSIONID/CFID+CFTOKEN/ASP.NET_SessionId can be all checked. When type is custom, user could define the cookie's keyword at will.
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.
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:

Settings	Guidelines
	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 intial 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.
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 client request.
Keyword	Specifies the cookie name. If the specified cookie keyword is a valid cookie name from which the cookie value can be extracted, it will be used to calculate the hash. Without the keyword, the hash will be calculated using the whole cookie. Otherwise, the default round robin method will be used.
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.

Settings	Guidelines
Keyword	Backend server cookie name.
Timeout	Timeout for an inactive persistence session table entry. The default is 300 seconds. The valid range is 0-86,400. If the timeout is set to 0 seconds, the persistence cookie will function as a session cookie, expiring at the end of the session. Typically, when the timeout value is set to 1 or more seconds, a timestamp is applied to the cookie to track the timeout. When the timeout is set to 0, the cookie will not have a timestamp and instead will expire along with the session.
Httponly	Enable/disable to add the "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). If the cookie attribute does not match or is not applicable, the user agent will ignore the cookie and will not enforce persistence.
Secure	Enable/disable 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)). If the cookie attribute does not match or is not applicable, the user agent will ignore the cookie and will not enforce persistence.
Samesite	Add a SameSite attribute to prevent the browser from sending cookies along with cross-site requests, to mitigate the risk of cross-origin information leakage. It provides Strict, Lax, and None values for this attribute: • Nothing — Do not add Samesite attribute to cookies. The default value is Nothing. • None — set the value as none if a cookie is required to be sent by cross origin. **Note**: If **Secure** is enabled, then **Samesite** should be set to **None**. • Lax — any request from the third parties will not carry such cookies except for GET requests that navigate to the destination URL. • Strict — any request from the third parties will not carry such cookies If the cookie attribute does not match or is not applicable, the user agent will ignore the cookie and will not enforce persistence.
Custom Attribute	Enable to specify custom attributes. When Custom Attribute is enabled, the following options become unavailable: Domain, Httponly, Secure, and Samesite.
Custom Attribute Value	The **Custom Attribute Value** option appears if **Custom Attribute** is enabled. Specify the full cookie attributes, including any of the standard attributes and any of the custom attributes. If the cookie attribute does not match or is not applicable, the user agent will ignore the cookie and will not enforce persistence.

Settings	Guidelines
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 `sessid` and the real server pool member Cookie value is `rs1`, FortiADC sends the cookie `sessid=rs1|U6iFN` to the client, where `U6iFN` is the expiration date as a base64 encoded string.
Keyword	Specifies the backend server cookie name.
Timeout	Timeout for an inactive persistence session table entry. The default is 300 seconds. The valid range is 0-86,400. If the timeout is set to 0 seconds, the insert cookie will function as a session cookie, expiring at the end of the session. Typically, when the timeout value is set to 1 or more seconds, a timestamp is applied to the cookie to track the timeout. When the timeout is set to 0, the cookie will not have a timestamp and instead will expire along with the session.
Domain	Specifies the domain attribute of the cookie. If the cookie attribute does not match or is not applicable, the user agent will ignore the cookie and will not enforce persistence.
Httponly	Enable/disable to add the "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). If the cookie attribute does not match or is not applicable, the user agent will ignore the cookie and will not enforce persistence.
Secure	Enable/disable 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)). If the cookie attribute does not match or is not applicable, the user agent will ignore the cookie and will not enforce persistence.
Samesite	Add a SameSite attribute to prevent the browser from sending cookies along with cross-site requests, to mitigate the risk of cross-origin information leakage. It provides Strict, Lax, and None values for this attribute: • Nothing — Do not add Samesite attribute to cookies. The default value is Nothing. • None — set the value as none if a cookie is required to be sent by cross origin. **Note**: If **Secure** is enabled, then **Samesite** should be set to **None**. • Lax — any request from the third parties will not carry such cookies except for GET requests that navigate to the destination URL. • Strict — any request from the third parties will not carry such cookies

Settings	Guidelines
	If the cookie attribute does not match or is not applicable, the user agent will ignore the cookie and will not enforce persistence.
Custom Attribute	Enable to specify custom attributes. When Custom Attribute is enabled, the following options become unavailable: Domain, Httponly, Secure, and Samesite.
Custom Attribute Value	The **Custom Attribute Value** option appears if **Custom Attribute** is enabled. Specify the full cookie attributes, including any of the standard attributes and any of the custom attributes. If the cookie attribute does not match or is not applicable, the user agent will ignore the cookie and will not enforce persistence.
Rewrite cookie	
Rewrite Cookie	Persistence is based on the cookie provided in the backend server response, but the system rewrites the cookie. The system checks the HTTP response for a `Set-Cookie:` 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 `sessid`. The real server pool member Cookie value is `rs1`. After an initial client request, the response from the server contains `Set-Cookie: sessid=666`, which FortiADC changes to `Set-Cookie: sessid=rs1`. FortiADC uses this rewritten value to forward subsequent requests to the same backend server as the original request.
Keyword	Specifies a `Set-Cookie:` 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 `Set-Cookie:` 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 `sessid`. The real server pool member Cookie value is `rs1`. After an initial client request, the response from the server contains `Set-Cookie: sessid=666`, which the system changes to `Set-Cookie: sessid=rs1~666`. It uses this rewritten value to forward subsequent requests to the same backend server as the original request.
Keyword	Specifies a `Set-Cookie:` value to match.
RADIUS Attribute	
Type	Select RADIUS Attribute.

Settings	Guidelines
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. **Note:** The persistence rule with **Match Across Virtual Servers** enabled works only with L4 virtual servers or the L7 virtual server whose Profile is LB_PROF_RADIUS.
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.
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.
ISO8583 Bitmap	
Timeout	Timeout for an inactive persistence session table entry. The default is 300 seconds. The valid range is 1-86,400.
ISO8583 Bitmap Relation	Relation among the bitmap type, be AND/OR. Default is OR.
Keyvalue Relation	Relation of keyvalue, be AND/OR. Default is AND.
Type	Persistence is based on bitmap. Support 30 bitmap type.

Configuring error pages

When backend real servers are unavailable or another status code for other module (ex: WAF/DoS/Auth), FortiADC can respond to clients attempting HTTP/HTTPS connections with a customized HTML error page. From the **Application Resources > Error Page**, you can create a customized error page by uploading an HTML error page file. You can then edit and preview the HTML file from the GUI. Once the HTML error page has been created, you can select it in virtual server configurations.

To aid you in customizing your HTML error page, FortiADC provides all default error page files that can be downloaded from the predefined profile **LB_ERROR_PAGE_DEFAULT**. You may use any of these default error page files as a template for customization.

The current error page file package only requests index.html to replace 503 error message when there are no servers in the pool. We also extend the support to these files listed below:

File Name	MUST	Guidelines
Index.html	Yes	This page will replace 503 error message page.
200.html	No	This page will replace 200 error message page.
202.html	No	This page will replace 202 error message page.
205.html	No	This page will replace 205 error message page.
400.html	No	This page will replace 400 error message page.
401.html	No	This page will replace 401 error message page.
403.html	No	This page will replace 403 error message page.
404.html	No	This page will replace 404 error message page.
405.html	No	This page will replace 405 error message page.
406.html	No	This page will replace 406 error message page.
408.html	No	This page will replace 408 error message page.
410.html	No	This page will replace 410 error message page.
413.html	No	This page will replace 413 error message page.
500.html	No	This page will replace 500 error message page.
501.html	No	This page will replace 501 error message page.
502.html	No	This page will replace 502 error message page.
504.html	No	This page will replace 504 error message page.
waf_deny.html	No	This page will replace all response to a WAF deny action. The error page will show the Message ID, Signature ID, and Client IP of the attack in the message as recorded in the attack log. The error page file does not include the related response-code.html.
default.html	No	This page will replace all other error page doesn't include in the package (excluding 503).

Alternatively, 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 396.

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 page file must be named `index.html` and contained in a tar, tar.gz, or zip file. The maximum file size is 1 MB.

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. The maximum file size is 1MB.
6. Enter the Virtual Path of the error page. This virtual path will conflict with the custom authentication form base page's virtual path and also with SAML's server URL configuration and Captcha path.
7. Click **Save**.
   The newly created error page will be listed in the **Error Page** tab.

To modify an error page:

1. Go to **Server Load Balance > Application Resources**.
2. Click the **Error Page** tab.
3. Double-click the error page or select the (edit) icon in the row of the error page that you want to modify.
   The Error Page configuration editor displays.
4. From the configuration editor, you can make the following modifications:
   - Upload a new error message file.
   - If the uploaded file is a zip file, edit the file directly through the text editor. The GUI text editor supports HTML, CSS, and JS file types.
5. Optionally, click **Preview** to test and view your HTML error page.
   Note: The preview function only supports HTML files and cannot execute any JavaScript contained in the HTML.
6. Click **Save**.

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 537.
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. HTPP/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.

Settings	Guidelines
	HTPP/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
```

Configuring Captcha

FortiADC allows administrators to validate incoming users with CAPTCHAs to determine whether a client is a regular user or an attacker. FortiADC can configure the WAF/DoS Policy to issue CAPTCHAs only to clients who meet the attack rules.

Select a FortiADC default captcha profile from within the virtual server configuration or upload a customized captcha page if you want to use your own captcha verification page for when an WAF/DoS attack detected.

Before you begin:

- You must have Read-Write permission for Server Load Balance settings.
- Copy the captcha file to a location you can reach from your browser; the captcha file must be named
- index.html it must include a tag called “%%FORTIADC_CAPTCHA_IFRAME%%” and be compressed as tar, tar.gz, or zip file. The maximum file size is 1 MB.

To upload a Captcha page file:

1. Go to **Server Load Balance > Application Resources**.
2. Click the **Captcha** tab.
3. Click **Create New** to display the configuration editor.
4. Enter the name of the captcha. You will use this name to select the captcha profile in virtual server configurations. No spaces.
5. Toggle the **Customized Captcha Page** and then click **Choose File** and browse and select the captcha page tar,

tar.gz, or zip file. The maximum file size is 1MB.

6. Save the configuration.

Captcha Configuration

Parameter	Description
Name	Configuration name. Valid characters are `A-Z, a-z, 0-9, _,` and `-`. No spaces are allowed. Maximum length 63. Note: After you initially save the configuration, you cannot edit the name.
Virtual Path	Virtual path of captcha function. This path is running on VS, so it will conflict with other configurations like errorpage's vpath and custom auth page. String type, not empty, maximum length 63, the default value is "/fortiadc_captcha/"
Max Attempts	Maximum attempts for Captcha verification. Integer type, range 1-100, default 5. The client will be blocked upon exceeding max attempts.
Max Picture Changes	The maximum number of times you can change another picture. Integer type, range 1- 100, default 5. Exceed change times change picture action won't success.
Picture Difficulty	There are two difficulty level here can be selected, hard and easy. hard level picture may fight AI picture recognition, but may cause difficulty in human identification. Default value is hard.
Max Block Period	Once client is blocked, how long it will be blocked. Integer type, range 10-2592000, default 86400. Exceed this time client will be reset to untracked state.
Max Verify Period	The longest verification time from captcha verify action start. Unit second, range 20- 86400, default 1200. Exceed this time the client will be blocked.
Customized Captcha Page	Switch for customize captcha page, default disable. If disable, the custom captcha package file option won't valid.
File	File package for customize captcha page. Click 'Choose File' to upload. The file package must include index.html file, and in the index page, it must include a tag called "%%FORTIADC_CAPTCHA_IFRAME%%", that we will insert the verify page box on it. Note: This option is only available when the 'Customized Captcha Page' is enabled.

Application Optimization

The Application Optimization module contains tools to further optimize application response times and server capacity, such as defining compression and caching rules.

The **Server Load Balance > Application Optimization** sub-menu contains the following configurations:

- **Page Speed** — The Page Speed configuration allows you to set rules that FortiADC follows when rendering web pages using sets of controls. Configuration parameters such as setting the File Cache Inode Limit allows you to limit the number of inodes a host has on a shared server to reduce unnecessary system resource consumption when opening multiple files. For details, see Creating a Page Speed configuration on page 543.

- **Page Speed Profile** — The Page Speed Profile is referenced in the Page Speed configuration. The Page Speed Profile allows you to specify the options for optimizing the delivery of web content such as HTML, CSS, and image. For details, see Creating Page Speed profiles on page 545.
- **Compression** — Applicable to HTTP/HTTPS load-balancing, you can configure Compression rules to offload HTTP/HTTPS response compression to FortiADC to save resources on your backend servers. This then allow FortiADC to decompress the compressed HTTP/HTTPS client requests for WAF inspection before passing them to your backend servers. FortiADC supports HTTP/HTTPS response compression in GZIP and DEFLATE formats. For details, see Configuring compression rules on page 546.
- **Caching** — Applicable to HTTP/HTTPS load-balancing, you can configure Caching rules to enable the system RAM cache to store HTTP content and serve subsequent HTTP requests for that content without forwarding the requests to the backend servers which reduces the load on the backend servers. Once Caching is enabled, the system evaluates HTTP responses to determine whether or not to cache the content. For details, see Using caching features on page 548.

Creating a Page Speed configuration

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

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

To create a Page Speed configuration object:

1. Click **Server Load Balance > Application Optimization**.
   The configuration page displays the **Page Speed** tab.
2. Click **Create New** to display the configuration editor.
3. Make the entries or selections as described in PageSpeed configuration on page 543.
4. Click Save when done.

PageSpeed configuration

Parameter	Description
PageSpeed	
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 *inode* 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.

Parameter	Description
	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 Page Speed profiles on page 545
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, `(http(s)://)*example.com/*/htmls/*.html` 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. 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 `"*.jsp*"` to match `http://example.com/index.jsp?test=xyz`.
Resource Control	
Origin Domain Patten	Specify the original domain pattern in regular expression in alphanumeric characters. For example, `(http(s)://)*.example.com` 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.

Parameter	Description
	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, `http://www.example.com` 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, `http://www.example.com` 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 Page Speed profiles

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

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

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

- HTML
- CSS
- Image

To create a Page Speed profile:

1. Go to **Server Load Balance > Application Optimization**.
2. Click the **Page Speed Profile** tab.
3. Make the entries or selections as described in Application optimization parameters on page 545.
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.

Parameter	Description
	Note: This ensures that the CSS styes 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 `<img>` tag or in the inline `"style=attibute".` 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

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.

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

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 547.
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. HTPP/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. HTPP/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.

Settings	Guidelines
URI Rule	Click **Add** and specify the URI to create the rule. Note: You must use a regular expression, e.g., `https://example.com/tmp/test.txt`.
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 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
```

Using caching features

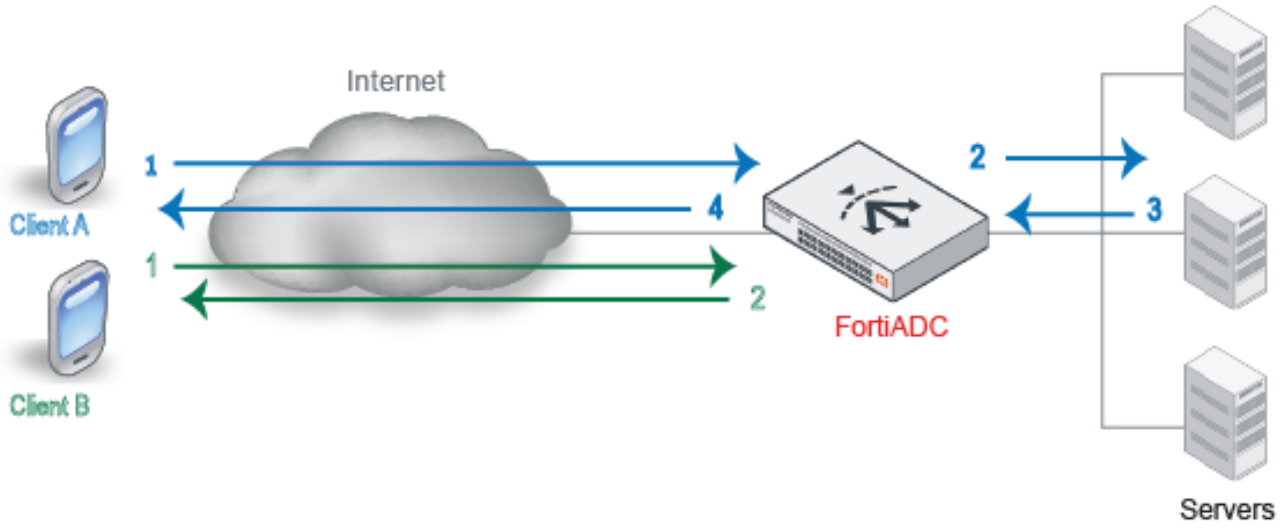
The system RAM cache can store HTTP content and serve subsequent HTTP requests for that content without forwarding the requests to the backend servers, thereby reducing the load on the backend servers.

You can configure basic static caching or dynamic caching rules.

Static caching

Static caching feature on page 548 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 551.
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 precence 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.

Real Server Pool

The Real Server Pool module contains configurations to define the destination real servers to distribute the incoming network traffic from the virtual server.

The **Server Load Balance > Real Server Pool** sub-menu contains the following configurations:

- **Real Server Pool** — The Real Server Pool configuration contains groups of real servers that host the applications that you load-balance. The incoming traffic received from the virtual server is load-balanced across this pool of real servers. For details, see Using real server pools on page 552.

- **Real Server** — Real servers are physical servers that are used to form real server pools. The Real Server configuration object specifies the status and, type, and IP address or FQDN of the physical server. The Real Servers configuration object can then be used as a member to form real server pools. For details, see Configuring real servers on page 559.
- **Server SSL** — 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. For details, see Configuring real server SSL profiles on page 562.

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:

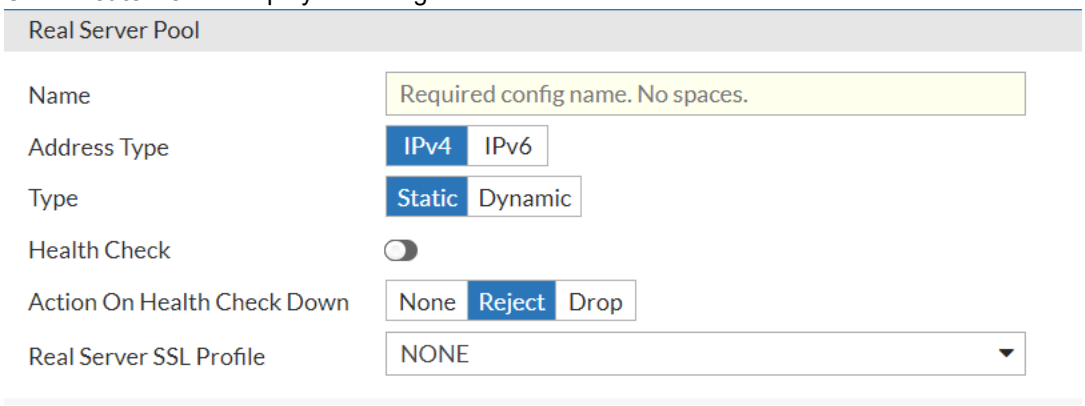
- You must have a good understanding and knowledge of the backend server boot behavior, for example, how many seconds it takes to “warm up” after a restart before it can process traffic.
- You must know the IP address and port of the applications.
- If you want to select user-defined health checks, you 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 real server pool:

1. Go to **Server Load Balance > Real Server Pool**.
   The configuration page displays the **Real Server Pool** tab.

2. Click **Create New** to display the configuration editor.



3. Configure the following settings for the **Real Server Pool**:

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.
Address Type	• IPv4 • IPv6
Type	Select the type of IP address the real servers should use: • Static • Dynamic
SDN Connector	The **SDN Connector** option is available if **Type** is **Dynamic**. Select an existing SDN connector configuration. If a cloud platform SDN connector is selected, you have the option to use the cloud service to create a dynamic real server pool that can be automatically populated with the cloud instances. However, once this dynamic real server pool is deleted or the SDN connector becomes invalid, all real servers that were automatically populated by the dynamic pool will also be deleted automatically.
Service	The **Service** option is available if **Type** is **Dynamic**. Specify the service for the SDN connector. **New in FortiADC 7.6.0:** For the **AWS** SDN connector, you now have the option to select services with the autoscaling group tag (**AutoScaleGroup=xxx**) to dynamically populate the real server pool with the instances tagged as part of the autoscaling group. When a scale-in or scale-out event occurs on the AWS side, the SDN connector will update the real server and pool based on the scale-in/scale-out result.
Service Port	The **Service Port** option is available if **Type** is **Dynamic** and **SDN Connector** is a an **AWS SDN connector**. Specify the service port. The default is 80, and the range is 0-65535.

Settings	Guidelines
IP Address Type	The **IP Address Type** option is available if **Type** is **Dynamic** and when FortiADC is deployed on public cloud platforms. Select whether the SDN connector should get the private address or public address of the instances. FortiADC will get the IP addresses of the real servers from the FortiADC instance deployed on the public cloud platform, and dynamically populate the objects in the server pool configuration. The resulting real server pool configurations will have the server type **Dynamic Auto**.
Health Check	Enable health checking for the pool. You can override this for individual servers in the pool.
Health Check Relationship	The **Health Check Relationship** option is available if **Health Check** is **enabled**. • AND—All of the selected health checks must pass for the server to the considered available. • OR—One of the selected health checks must pass for the server to be considered available.
Health Check List	The **Health Check List** option is available if **Health Check** is **enabled**. Select one or more user-defined or predefined Health Check configuration objects. Shift-click to select multiple objects at the same time.
Action On Health Check Down	Select the action to take when Health Check fails. • None — New connections are scheduled to other available pool members. For existing active connections, they will expire after the session times out. In which case FortiADC will send TCP RST or ICMP unreachable message to clients. If a pool member passes the health check before the existing active connections expire, those connections will remain and will not be reset. • Reject — This is the default setting. FortiADC removes any current connections that associate with the pool member that failed the health check. For existing active connections, FortiADC will send TCP RST or ICMP unreachable message to clients. If there are other available pool members, FortiADC resets existing active connections and sends new arriving connections to the available pool member. • Drop — FortiADC silently removes the connections that associate with the pool member that failed the health check. All new connections are scheduled to other healthy pool members.
Direct Route Mode	Enable Direct Route Mode to send health check packets directly to the specified backend server IP address. This is disabled by default.
Direct Route IPv4/ Direct Route IPv6	The **Direct Route IPv4** or **Direct Route IPv6** option is available if **Direct Route Mode** is enabled. Specify the IPv4 or IPv6 address of the backend server to directly send health check packets.
Real Server SSL Profile	Select a Real Server SSL Profile to apply.

4. Click **Save** to save the real server pool parent configuration.
   The **Member** section becomes available to add real server pool member configurations.

5. Under the **Member** section, click **Create** to display the configuration editor.
6. Configure the following settings for the real server pool **Member**:

Settings	Guidelines
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: • 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.
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. 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.

Settings	Guidelines
	• 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
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.
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.
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	The **Health Check** option is available if **Health Check Inherit** is **disabled**. Enable this option to specify a health check configuration object for this server.
Health Check Relationship	• AND—All of the selected health checks must pass for the server to the 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	The **RS Profile** option is available if **RS Profile Inherit** is **disabled**. If RS Profile Inherit 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.
MySQL Read Only	Disabled by default. Select the button to enable it.
MySQL Group ID	Specify the MySQL group ID.
MSSQL Read Only	Disabled by default. Select the button to enable it.
Proxy Protocol	This is a protocol of application layer, which is located upper layer of HTTP and SSL, and it contains a head to description the real IP address of client. There two major version of this protocol v1 and v2. Support none, v1, v2. None will disable this function and it’s the default value, v1 and v2 is different version of this protocol, the v1 version is human readable. You need co-deployment with FortiWeb, and because X-Forword-For option isn’t valid for them they demand use proxy protocol to deliver the real client’s IP address to them.

Settings	Guidelines
	Only support : HTTP/HTTPS/TCPS/RDP, Either L7 and L2 VS of these type can support it.
Modify Host	Enable to allow FortiADC to modify the HTTP header according to the "host" field of the real server. This is disabled by default.
Host	The **Host** option is available if **Modify Host** is enabled. Specify the host as a string. This field cannot be left empty. The input validation regex is `([0-9A-Za-z+/?%$#&\=.:_-]|\[|])+$`, maximum of 255 characters. The default value is `host`.

7. Click **Save**.
   The newly created real server pool member is listed under the **Member** section.
   Optionally, you can configure the column settings to view and sort the real server pool members, such as adding the **FQDN** column that is not included in the default view.
   Note: Real Server Pool members that are generated from the auto populate function cannot be cloned or edited.



When cloning a Real Server Pool that has auto-populated pool members, the cloned configuration will initially appear without the auto-populated pool members. However, the auto-populated pool members will appear again once the cloning operation is complete.

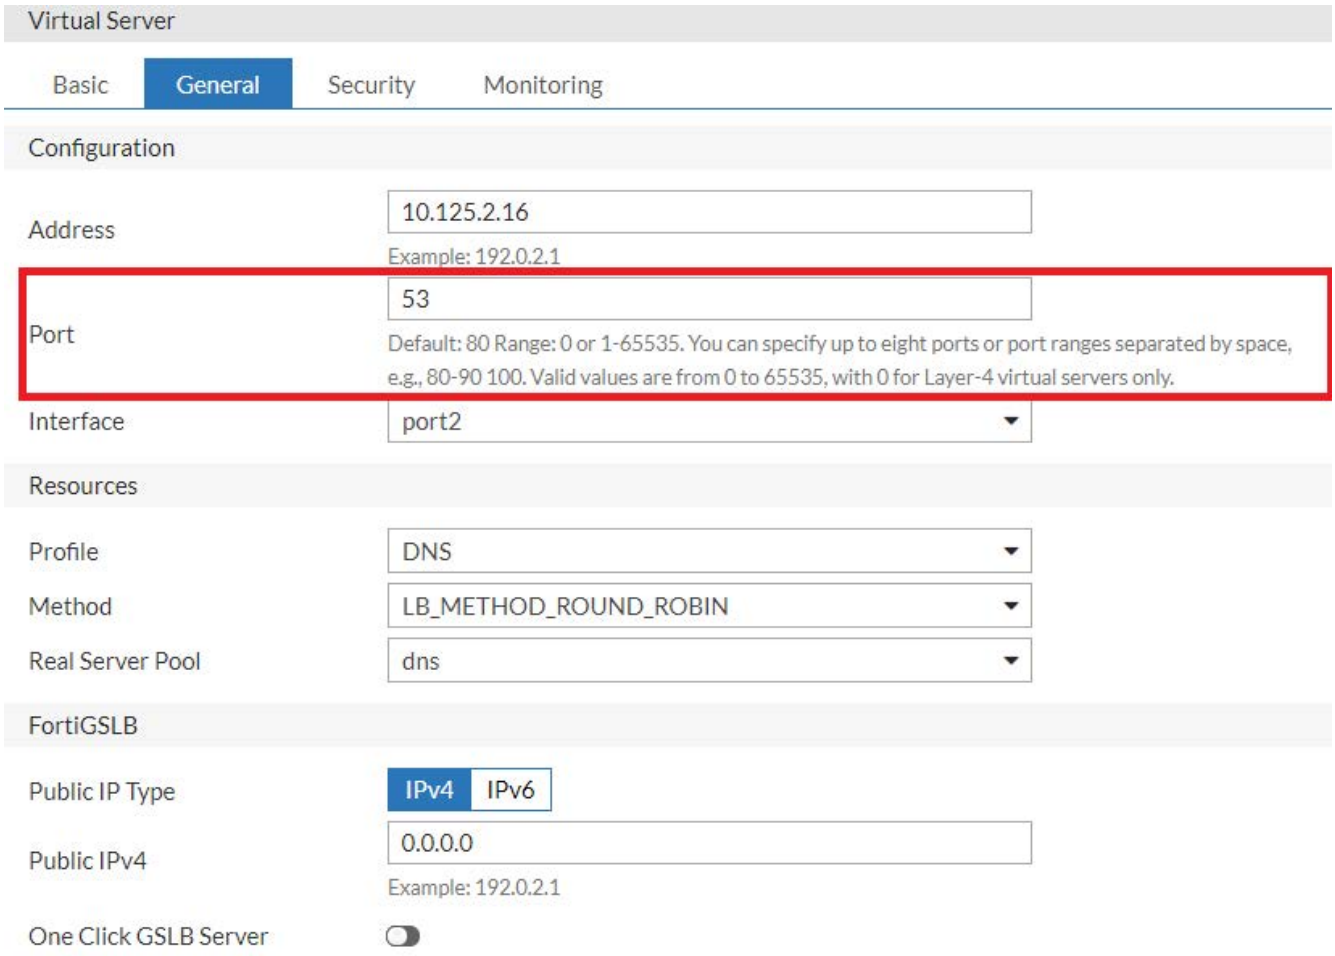
Example: Using port ranges and the port 0 configuration

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.

Key FortiADC configuration elements



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 Using real server pools on page 552.

To configure a real server configuration object:

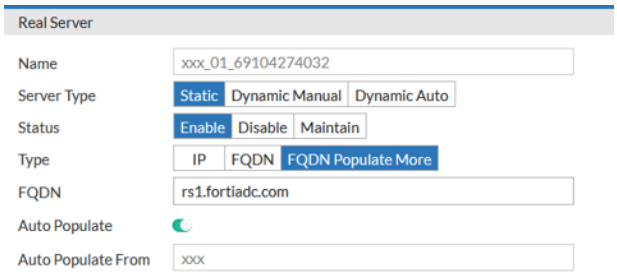
1. Go to **Server Load Balance > Real Server Pool**.
2. Click the **Real Server** tab.

3. Click **Create New** to display the configuration editor.
4. Configure the following settings:

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.
Server Type	Select the real server type: • Static • Dynamic Manual • Dynamic Auto* ***Note**: The **Dynamic Auto** server type is not a selectable option. The Dynamic Auto server type is used to mark real servers that have been dynamically populated using the IP addresses from the FortiADC instances deployed on public cloud platforms. As these real servers are automatically populated by the dynamic pool, they are read-only and cannot be modified from the real server configuration editor. Dynamic Auto real servers will be deleted automatically once their dynamic real server pool is deleted or the SDN connector becomes invalid.
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.
Type	The **Type** option is available if the **Server Type** is **Static**. Select the static server type to use for DNS resolution: • IP — Use the specified IPv4 Address or IPv6 Address. • FQDN — Use the specified FQDN of the Real Server for DNS resolution. The address will be resolved by the DNS query. Only one IPv4 or IPv6 address will be used to form the real server. • FQDN Populate More — Use the specified FQDN of the Real Server for DNS resolution. The address will be resolved by the DNS query. Multiple IPv4 or IPv6 addresses will be used to form the real server. A maximum of 16 IPv4 or IPv6 addresses is supported for one domain. If the DNS query gets more than 1 IP or IPv6 address for the Real Server FQDN, then the system will automatically generate real servers based on the specified FQDN. Only select the FQDN Populate More option when creating a new Real Server configuration. If an existing real server configuration set to IP or FQDN type is changed to FQDN Populate More, the auto-population function will fail to populate pool members.
Address	The **Address** option is available if the **Server Type** is **Static** and the **Type** is **IP**. For IPv4 real server, enter the real server's IP address in IPv4 address format.

Settings	Guidelines
Address6	The **Address6** option is available if the **Server Type** is **Static** and the **Type** is **IP**. For IPv6 real server, enter the real server's IP address in IPv6 address format.
FQDN	The **FQDN** option is available if the **Server Type** is **Static** and the **Type** is **FQDN** or **FQDN Populate More**. Specify a fully qualified domain name, such as "www.example.com". This is required.
Auto Populate	The **Auto Populate** field appears if the **Server Type** is **Static** and the **Type** is **FQDN Populate More**. Auto Populate is a read-only status that indicates whether or not a real server was auto-generated through FQDN Populate More. • Disabled — By default, Auto Populate is disabled. This indicates that the current real server configuration is not auto-generated. • Enabled — This indicates that the current real server configuration was automatically generated by the system based on the "parent" real server's FQDN. If the DNS query gets more than 1 IP or IPv6 address for the Real Server FQDN of the parent" real server, then the system will automatically populate real servers based on the specified FQDN, with the "child" real server name based on the "parent" real server name as specified in the **Auto Populate From** field (read-only). The auto-generated real servers will be read-only and cannot be cloned or edited.
SDN Connector	The **SDN Connector** is available if the **Server Type** is **Dynamic Manual**. Specify a previously configured SDN Connector. For details, see External connectors.
Instance	The **Instance** is available if the **Server Type** is **Dynamic Manual**. Specify the instance of your SDN Connector.

5. Click **Save**.
The new real server configuration appears on the Real Server page.
For real servers with the type **FQDN Populate More**, real servers will automatically generate if there are more than 1 IP/IPv6 addresses associated with the specified FQDN. The auto-generated real servers will be read-only and cannot be cloned or edited.
Auto-generated real servers will be named after the "parent" real server and a combination of sequential numbers and random numbers, in the following format:
Real Server Name_\d\d_random number
For example:
If the parent Real Server Name is "xxx", then the name of the auto-generated real server can be "xxx_01_69104274032".



The **Auto Populate** and **Auto Populate From** fields are read-only information to indicate the origin of the auto-generated real server.

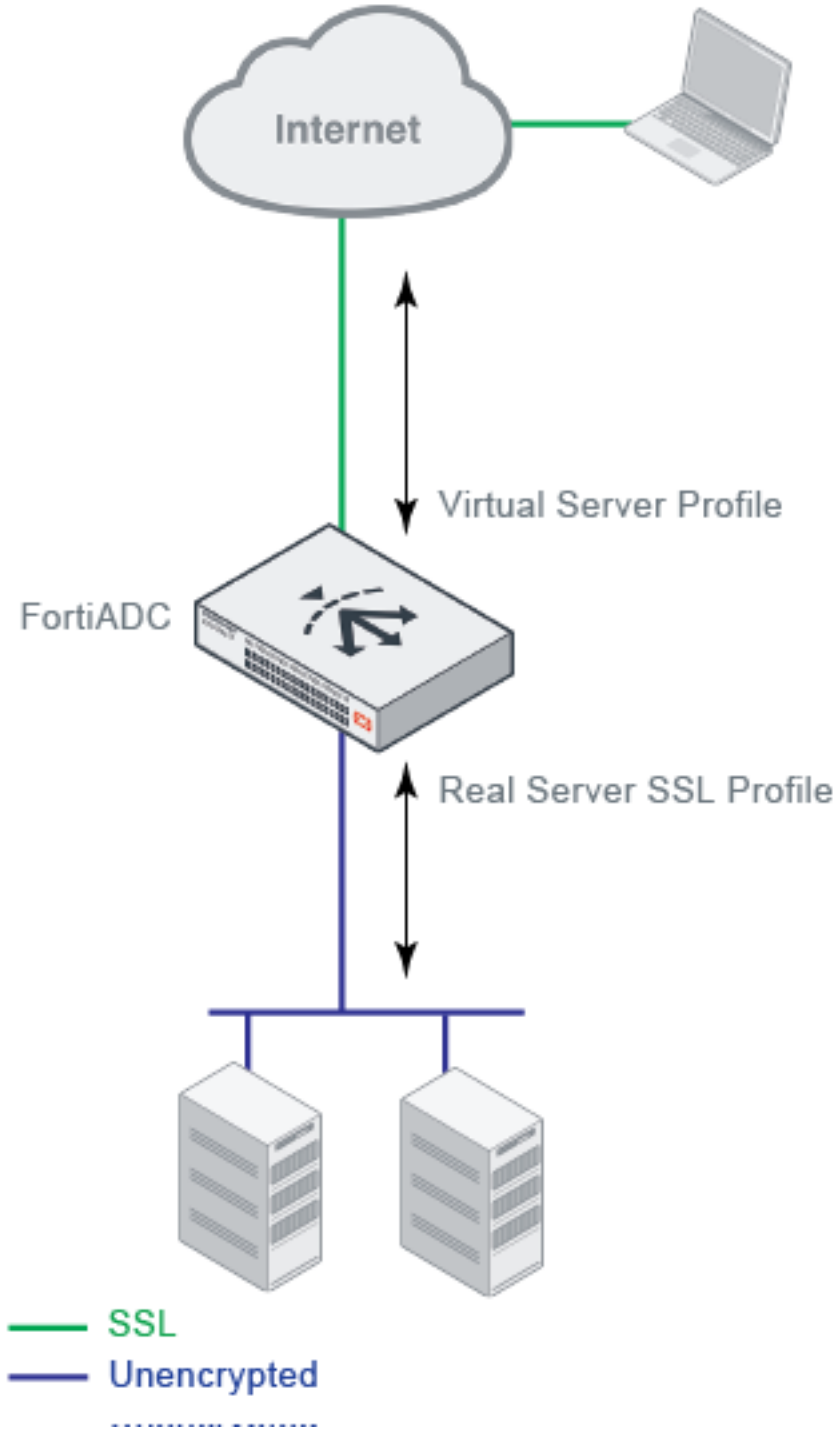
The FQDN Populate More real server type is only applicable to newly created real server configurations. Do not edit an existing real server configuration set to IP or FQDN type to FQDN Populate More as the auto-population will not function and pool members will not auto-populate.

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 562 illustrates the basic idea of client-side and server-side profiles.

SSL profiles



Predefined real server profiles on page 564 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: custom
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_SSLV3	• Allow version SSLv3 • Cipher suite list: 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-MD5 ECDHE-RSA-DES-CBC3-SHA EDH-RSA-DES-CBC3-SHA DES-CBC3-SHA EDH-RSA-DES-CBC-SHA DES-CBC-SHA
LB_RS_SSL_PROF_MEDIUM	• Allow version: TLSv1.0, TLSv1.1, and TLSv1.2 • Cipher suite list: 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 RC4-SHA EDH-RSA-DES-CBC3-SHA DES-CBC3-SHA

Profile	Defaults
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 565.
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/manprimary/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-CAMELLIA256-SHA384

Settings	Guidelines
	• *ECDHE-ECDSA-AES256-SHA • ECDHE-ECDSA-AES128-GCM-SHA256 • *ECDHE-ECDSA-AES128-SHA256 • ECDHE-ECDSA-CAMELLIA128-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-CAMELLIA256-SHA384 • *ECDHE-RSA-AES256-SHA • DHE-RSA-AES256-GCM-SHA384 • *DHE-RSA-AES256-SHA256 • DHE-RSA-CAMELLIA256-SHA256 • *DHE-RSA-AES256-SHA • DHE-RSA-CAMELLIA256-SHA • AES256-GCM-SHA384 • *AES256-SHA256 • *AES256-SHA • ECDHE-RSA-AES128-GCM-SHA256 • *ECDHE-RSA-AES128-SHA256 • ECDHE-RSA-CAMELLIA128-SHA256 • *ECDHE-RSA-AES128-SHA • DHE-RSA-AES128-GCM-SHA256 • *DHE-RSA-AES128-SHA256 • DHE-RSA-CAMELLIA128-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 *These ciphers are fully supported by hardware SSL (in 400F, 420F, 1200F, 2200F, 4200F and 5000F). Note: We recommend retaining the default list. If necessary, you can deselect the SSL ciphers that you do not want to support.

Settings	Guidelines
TLSv1.3 Cipher Suite List	TLSv1.3 ciphers are listed as following: • TLS_AES_256_GCM_SHA384 • TLS_AES_128_GCM_SHA256 • TLS_CHACHA20_POLY1305_SHA256 • TLS_AES_128_CCM_SHA256 • TLS_AES_128_CCM_8_SHA256 Note: This option only available if the TLSv1.3 is checked.
Allowed SSL Versions	You have the following options: • SSLv3 • TLSv1.0 • TLSv1.1 • TLSv1.2 • TLSv1.3 **Note**: • Please make sure that the SSL version is continuous. If not, an error message should be returned. • RFC 7919 Comply cannot support SSLv3 and TLS 1.3. If RFC 7919 Comply is enabled and **SSLv3** or **TLSv1.3** is selected in **Allowed SSL Versions**, an error message will display.
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.
Local Certificate	Select a local certificate object. This should be the backend servers' certificate, NOT the appliance's GUI web server certificate. See Manage Certificates.
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: • This option is enabled by default. • When disabled, you must select an option for Renegotiation-Deny-Action.
Renegotiation Period	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.

Settings	Guidelines
	• 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.
RFC 7919 Comply	Enable/disable parameters to comply with RFC 7919. **Note**: RFC 7919 Comply cannot support SSLv3 and TLS 1.3. If RFC 7919 Comply is enabled and **SSLv3** or **TLSv1.3** is selected in **Allowed SSL Versions**, an error message will display.
Supported Groups	The **Supported Groups** option is available if **RFC 7919 Comply** is enabled. Specify the supported group objects from the following: • secp256r1 • secp384r1 • secp521r1 • x25519 • x448 • ffdhe2048 • ffdhe3072 • ffdhe4096 • ffdhe6144 • ffdhe8192 At least one item from the FFDHE group must be selected. **Note**: The RFC 7919 Comply feature requires certain cipher selections to correspond with the Supported Group selection. • If a FFDHE group is selected (for example, ffdhe2048), then at least one cipher must be DHE-RSA (for example, DHE-RSA-AES256-SHA256).

Settings	Guidelines
	• If the Supported Group includes groups other than FFDHE (such as a SECP group, secp256r1), then at least one cipher must be ECDHE (for example, ECDHE-ECDSA-AES256-GCM-SHA384). • If a ECDHE cipher is selected (for example, ECDHE-ECDSA-AES256-GCM-SHA384), then the Supported Group must include at least one group that is not FFDHE (such as a SECP group, secp256r1).

Scripting

The Scripting module contains tools that allow you to apply Lua Scripting for HTTP/HTTPS virtual servers and Layer 7 TCP and UDP virtual servers. These Lua Scripts allow you to perform actions that are not supported by the current built-in feature set.

The **Server Load Balance > Scripting** sub-menu contains the following configurations:

- **HTTP** — From the HTTP page, you can import HTTP scripts directly in the GUI. Applicable to Layer 4 and 7 HTTP/HTTPS virtual servers, the HTTP scripts are event-triggered, allowing you to manipulate HTTP requests and responses, redirection, and dynamically change backend routing. This functionality can be combined with other HTTP related functions such as WAF, SSL, and Authentication. To get you started, FortiADC provides system predefined HTTP scripts that can be cloned for customization. For details, see Using HTTP scripting on page 569.
- **Stream** — From the Stream page, you can import Stream scripts directly in the GUI. Applicable to Layer 7 TCP and UDP virtual servers, the Stream scripts are event-triggered, allowing you to manipulate the TCP/UDP request and response. For details, see Using Stream scripting on page 588.

Using HTTP scripting

Enable scripting for Layer 2 and Layer 7 HTTP/HTTPS virtual servers to perform actions that are not supported by the current built-in feature set. You can import HTTP scripts from the GUI, **Server Load Balance > Scripting > HTTP** tab. To get you started, FortiADC provides system predefined HTTP scripts that can be cloned for customization. The HTTP scripts are event-triggered, allowing you to manipulate HTTP requests and responses, redirection, and dynamically change backend routing. This functionality can be combined with other HTTP related functions such as WAF, SSL, and Authentication.

FortiADC HTTP scripts are based on Lua 5.3.



FortiADC does not support all Lua functions. For the full list of supported functions, see the Script Reference Guide.

Please note that Scripting is an *as is* feature and does not come with any functionality or performance guarantees.

This section includes the following:

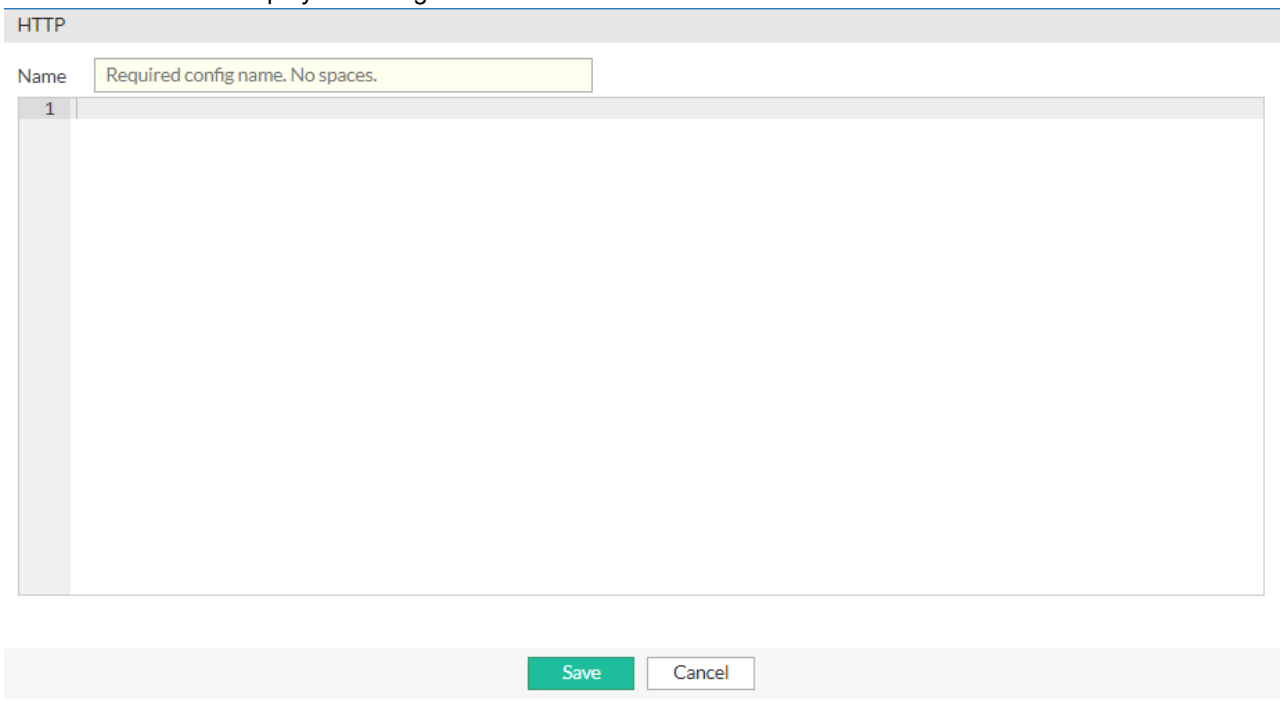
- Create a script object on page 570
- Import a script on page 570
- Export a script on page 571

- Delete a script on page 571
- Predefined HTTP scripts on page 572
- Multi-script support on page 580

Create a script object

To create a script configuration object:

1. Go to **Server Load Balance > Scripting**.
   The configuration page displays the **HTTP** tab.
2. Click **Create New** to display the configuration editor.



3. Enter a unique name for the HTTP script configuration. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No spaces.
   After you initially save the configuration, you cannot edit the name.
4. In the text box, type or paste your HTTP script.
   If you want to include this script as part of a multi-script configurations that allows you to execute multiple scripts in a certain order, ensure to set its priority. For more information, see Multi-script support on page 580.
5. Click **Save**.
   Once the HTTP script configuration is saved, you can specify it in the virtual server.

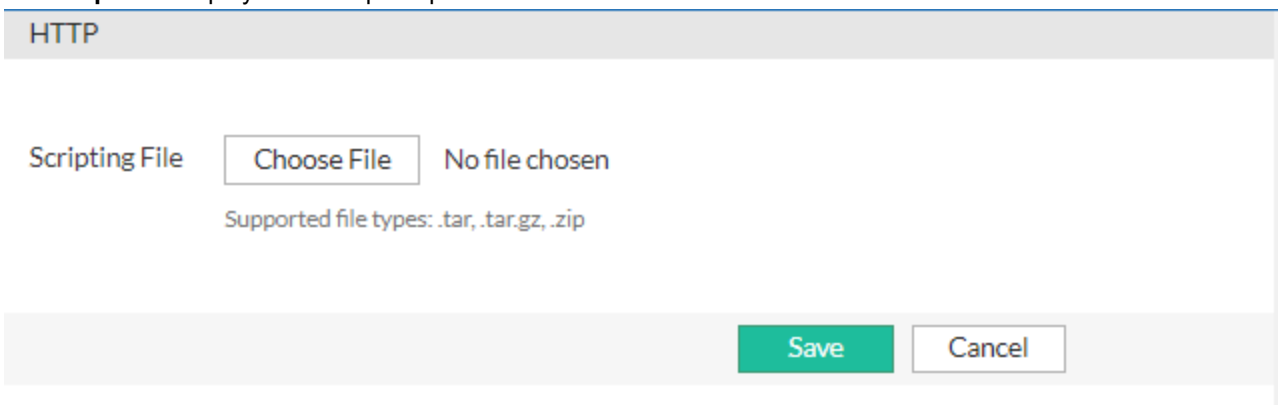
Import a script

HTTP scripts can be imported as TAR, TAR GZ, or ZIP files.

To import a script:

1. Go to **Server Load Balance > Scripting**.
   The configuration page displays the **HTTP** tab.

2. Click **Import** to display the file import options.



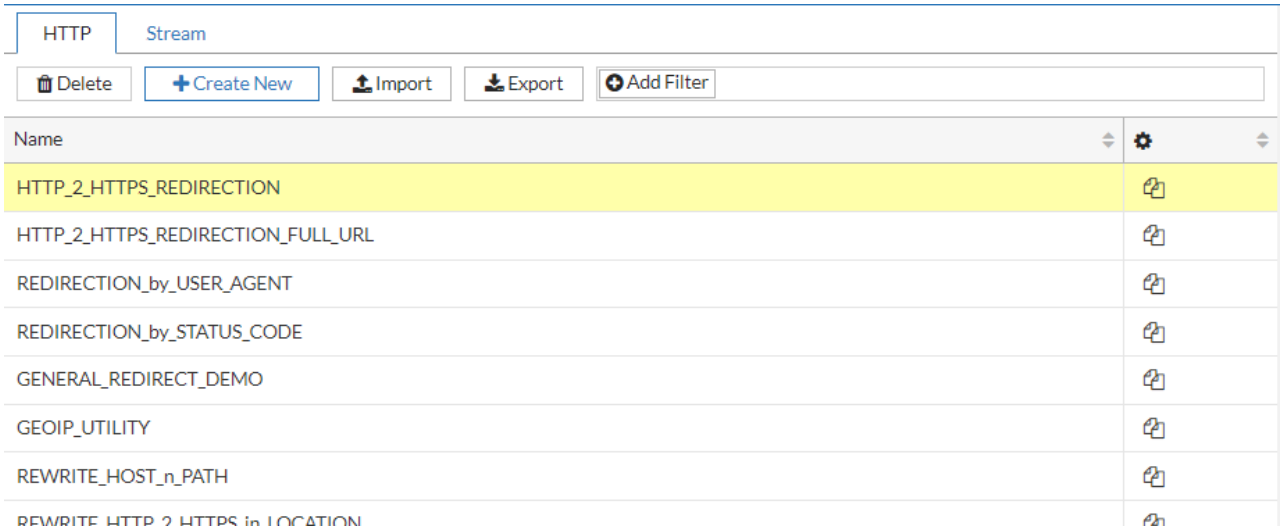
3. Click **Choose File** and browse for the script file. Supported file types are .tar, .tar.gz, and .zip.
4. Click **Save**.
   Once the file is successfully imported, it will be listed in the **Scripting > HTTP** page.

Export a script

HTTP script configurations can be exported as TAR files.

To export a script:

1. Go to **Server Load Balance > Scripting**.
   The configuration page displays the **HTTP** tab.
2. From the **HTTP** page, select an HTTP script configuration.
   In the example below, the HTTP_2_HTTPS_REDIRECTION script is selected.



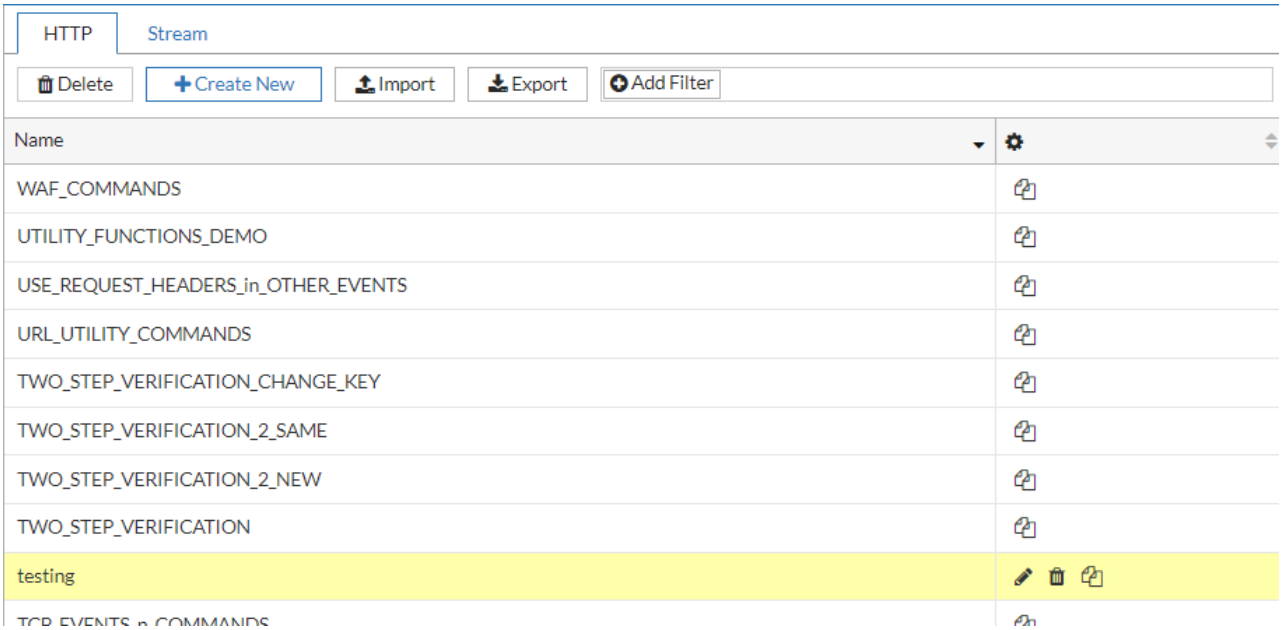
3. Click **Export** initiate the file download.
   The selected script configuration will be exported as a .tar file.

Delete a script

User-defined HTTP script configurations can be deleted.

To delete a script:

1. Go to **Server Load Balance > Scripting**.
   The configuration page displays the **HTTP** tab.
2. From the **HTTP** page, select a user-defined HTTP script configuration. System predefined scripts cannot be deleted.
   In the example below, the **testing** script configuration is selected.



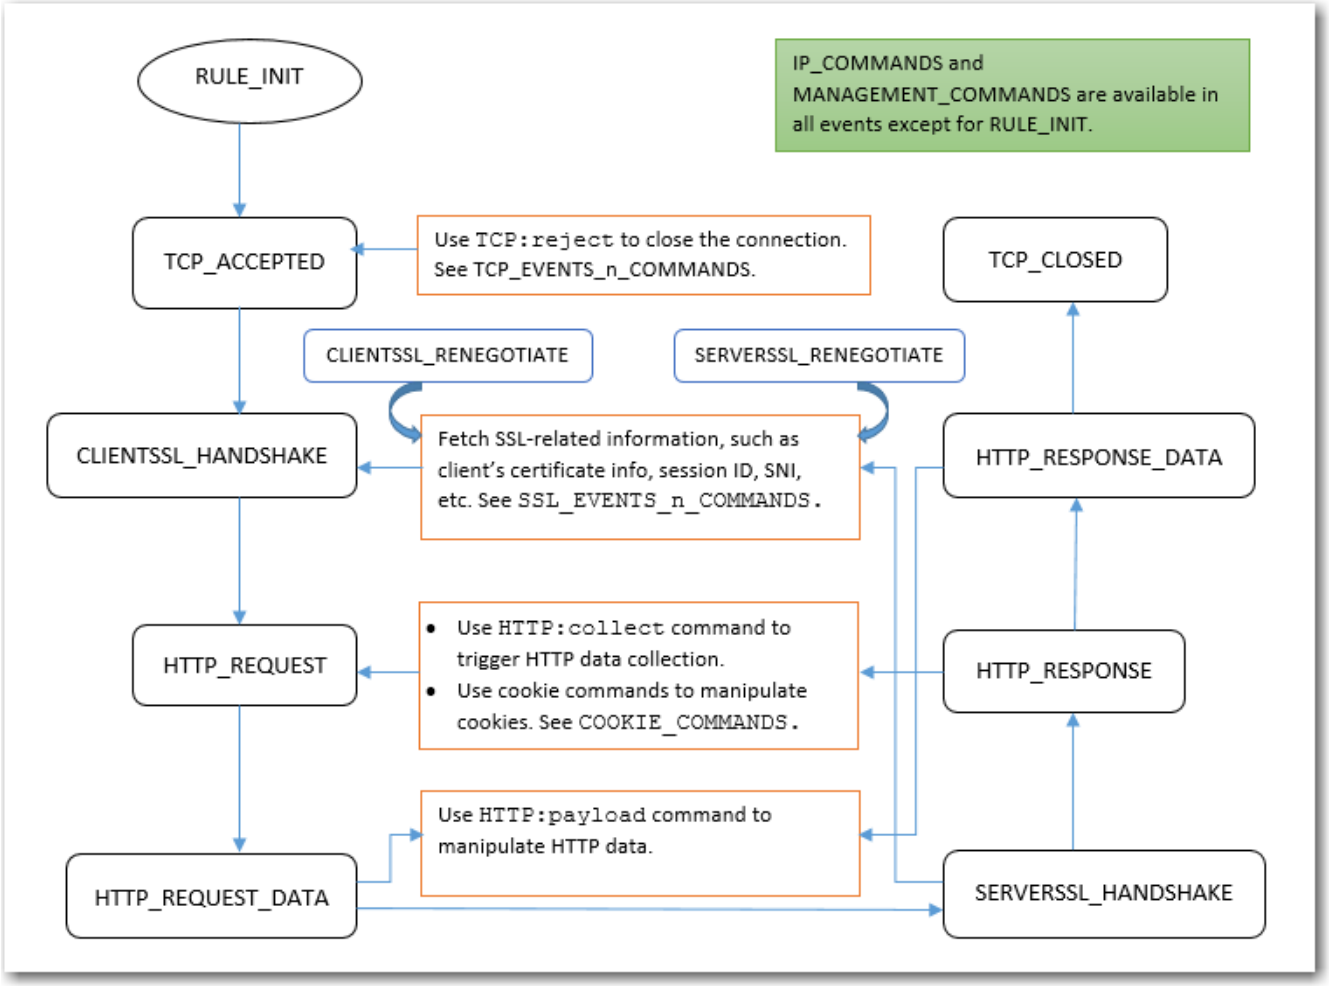
3. Click **Delete** from the top navigation, or click (delete icon) of the configuration.
   Multiple script configurations can be deleted using the **Delete** button on the top navigation.

Predefined HTTP scripts

You can view and use predefined HTTP scripts by going to **Server Load Balance > Scripting > HTTP**.

Scripts and predefined commands on page 572 highlights the functions of these scripts and commands and shows how to use them.

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.

The following table lists the FortiADC predefined scripts available for users to apply and customize.

Group	Predefined script	Usage
Authentication	AUTH_COOKIE_BAKE	Allows you to retrieve the baked cookie and edit the cookie content.
	AUTH_EVENTS_n_COMMANDS	Lists the authentication event and commands.
	CUSTOMIZE_AUTH_KEY	Demonstrates how to customize the crypto key for authentication cookie.
	TWO_STEP_VERIFICATION	Demonstrates how to perform 2-Step Verification using FortiToken. One needs have authentication policy configured and selected in a virtual-server.
	TWO_STEP_VERIFICATION_2_NEW	Demonstrates how to perform 2-Step Verification using FortiToken for the second authentication group.
	TWO_STEP_VERIFICATION_2_SAME	Demonstrates how to perform 2-Step Verification for the second authentication group using the same token group.
	TWO_STEP_VERIFICATION_ CHANGE_KEY	Demonstrates how to change the AES key and its size for stored token group.
Cookie	COOKIE_COMMANDS	Lists the two cookie commands and shows how to use them.
	COOKIE_COMMANDS_USAGE	Demonstrates 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.
Feature	WAITING_ROOM	The sample Waiting Room script demonstrates how you can place visitors in a virtual queue instead of denying them service directly when the server side reaches its configured capacity limit during high-demand situations. In this virtual Waiting Room, visitors can see their position in line and when their turn arrives, they are redirected to the requested page. Configuration parameters include the waiting room name, total resource limit threshold (default is 1000), and the Resource URL applicable to the waiting room. You can also customize the message displayed to users when they are placed in the waiting room by editing the HTML page section of the script.

Group	Predefined script	Usage
		Required data structures such as atomic counters and shared tables are already built into the script; however, you have the option to apply user-defined atomic counters and shared tables to customize the script.
HTTP	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. **Note**: This script can be used directly without making any changes.
	HTTP_2_HTTPS_REDIRECTION_FULL_URL	Redirects requests to the specified HTTPS URL. **Note**: This script can be used directly without making any 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.
	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**: This script can be used directly without making any changes.
	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). **Note**: 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.

Group	Predefined script	Usage
	REDIRECTION_by_USER_AGENT	Redirects requests based on User Agent (for example, a redirect to the mobile version of a site). **Note**: 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**: This script can be used directly without making any changes.
	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**: This script can be used directly without making any changes.
	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".` **Note**: This script can be used directly without making any changes.
	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**: This script can be used directly without making any changes.
	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.

Group	Predefined script	Usage
		Note: Do **not** use this script "as is". Instead, copy it and customize the content you want to store, use `collect() in HTTP_REQUEST` to trigger `HTTP_DATA_REQUEST`,or use `collect() in HTTP_RESPONSE` to trigger `HTTP_DATA_ RESPONSE`.
IP	IP_COMMANDS	Used to get various types IP Address and port number between client and server side.
Optimization	MULTIPLE_SCRIPT_CONTROL_DEMO_1	Uses `demo_1` and `demo_2` script to show how multiple scripts work. `Demo_1` with priority 12 has a higher priority. 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 can enable or disable other events. Do **not** use this script "as is". Instead, copy it and customize the operation.
	RAM_CACHING_COMMANDS	Lists the RAM caching event and commands.
	RAM_CACHING_DYNAMIC	Demonstrates how to use script to do dynamic RAM caching. **Note**: Dynamic caching is identified by a configured ID. Ensure the RAM caching configuration is selected in the HTTP or HTTPS profile.
	RAM_CACHING_GROUPING	Demonstrates how to create multiple variations based on client IP address. The sort of grouping applies to both regular caching and dynamic caching. **Note**: Ensure the RAM caching configuration is selected in HTTP or HTTPS profile.
Routing	CONTENT_ROUTING_by_URI	Routes to a pool member based on URI string matches. **Note**: 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.

Group	Predefined script	Usage
		Note: You should not use this script as is. Instead, copy it and customize the X-Fowarded-For header values and pool member names.
	PERSIST_COMMANDS	Demonstrates how to use persistence commands and event. The PERSISTENCE event is triggered when FortiADC receives the HTTP REQ and is ready to dispatch to the real server. You can set the entry in PERSISTENCE, then look up it in POST_PERSIST. FortiADC will dispatch to the dedicated server according to your entry set in PERSISTENCE if this session has not been assigned to the real server before.
SSL	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.
	SSL_EVENTS_n_COMMANDS	Demonstrates how to fetch the SSL certificate information and some of the SSL connection parameters between server and client side.
TCP	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.
	SOCKOPT_COMMANDS	Demonstrates how to the TCP:sockopt with usage examples.
	TCP_EVENTS_n_COMMANDS	Demonstrates how to reject a TCP connection from a client in TCP_ACCEPTED event.

Group	Predefined script	Usage
Utility	AES_DIGEST_SIGN_2F_COMMANDS	Demonstrates how to use AES to encryption/decryption data and some tools to generate the digest.
	ATOMIC_COUNTER_COMMANDS	Allows you to create and configure shared atomic counters that are accessible by multiple httproxy processes within one VS. The stored data is located in shared memories. In the Waiting Room setup, the atomic counters track variables at running time, including the current resource count, the current position in line, and the current total number of users in the waiting queue.
	CLASS_SEARCH_n_MATCH	Demonstrates how to use the `class_match` and `class_search` utility function.
	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.
	GEOIP_UTILITY	Used to fetch the GEO information country and possible province name of an IP address.
	MANAGEMENT_COMMANDS	Allow you to disable/enable rest of the events from executing.
	SHARED_TABLE_COMMANDS	Allows you to create and configure shared hash tables that are accessible by multiple httproxy processes within one VS. Both the table and stored data are located in shared memories. In the Waiting Room setup, the shared table is used to track current active resource occupiers such as active sessions.
	URL_UTILITY_COMMANDS	Demonstrates 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
WAF	WAF_COMMANDS	Demonstrates how to use WAF related functions and events.

Multi-script support

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 allow the system to execute the individual scripts in the order they are presented in the multi-script file.

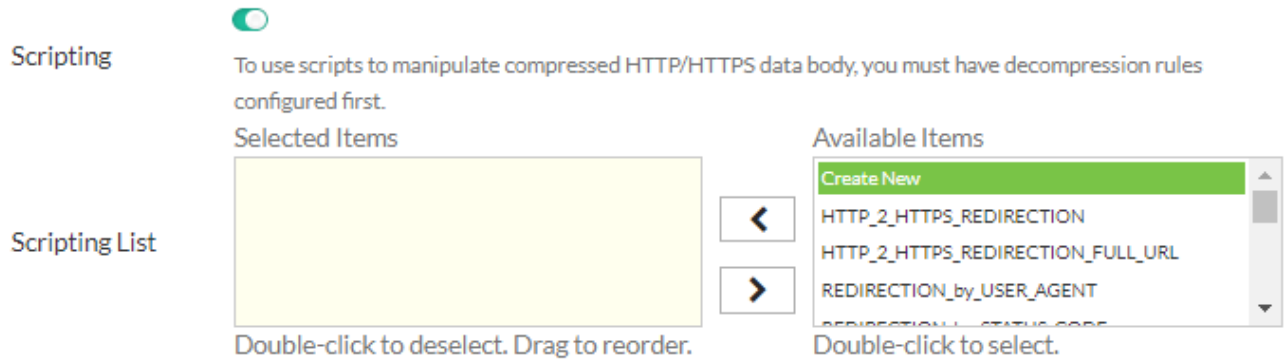
Currently, up to 16 individual scripts can be added to create a large multi-script file.

If desired, you can disable the processing of remaining scripts in the multi-script, or you can even complete disable the processing of certain events (for example, you can disable the processing of 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 call prevails.

In practice, instead of creating a single large and complex script containing all necessary logic, it's often more advantageous to decompose it into smaller functional components represented by individual scripts. This approach offers several benefits. Firstly, executing multiple scripts concurrently is more efficient than running them sequentially. Additionally, breaking down a massive script into smaller units enhances flexibility, particularly when applying scripts to various virtual servers. Some servers may require only specific scripts, while others may utilize all available ones. With smaller, modular scripts, you have the flexibility to select and combine only the necessary components to construct a comprehensive multi-script file, each with its designated priority, and apply them collectively to a virtual server.

Apply multiple scripts on page 580 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 information 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 581 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.

Waiting Room for virtual queuing

The Waiting Room feature allows you to place visitors in a virtual queue instead of denying them service directly when the server side reaches its configured capacity limit during high-demand situations. In this virtual Waiting Room, visitors can see their position in line and when their turn arrives, they are redirected to the requested page. The queuing mechanism uses the FIFO (First In, First Out) servicing order to allow visitors who first enter the waiting room to be the first to leave the queue.

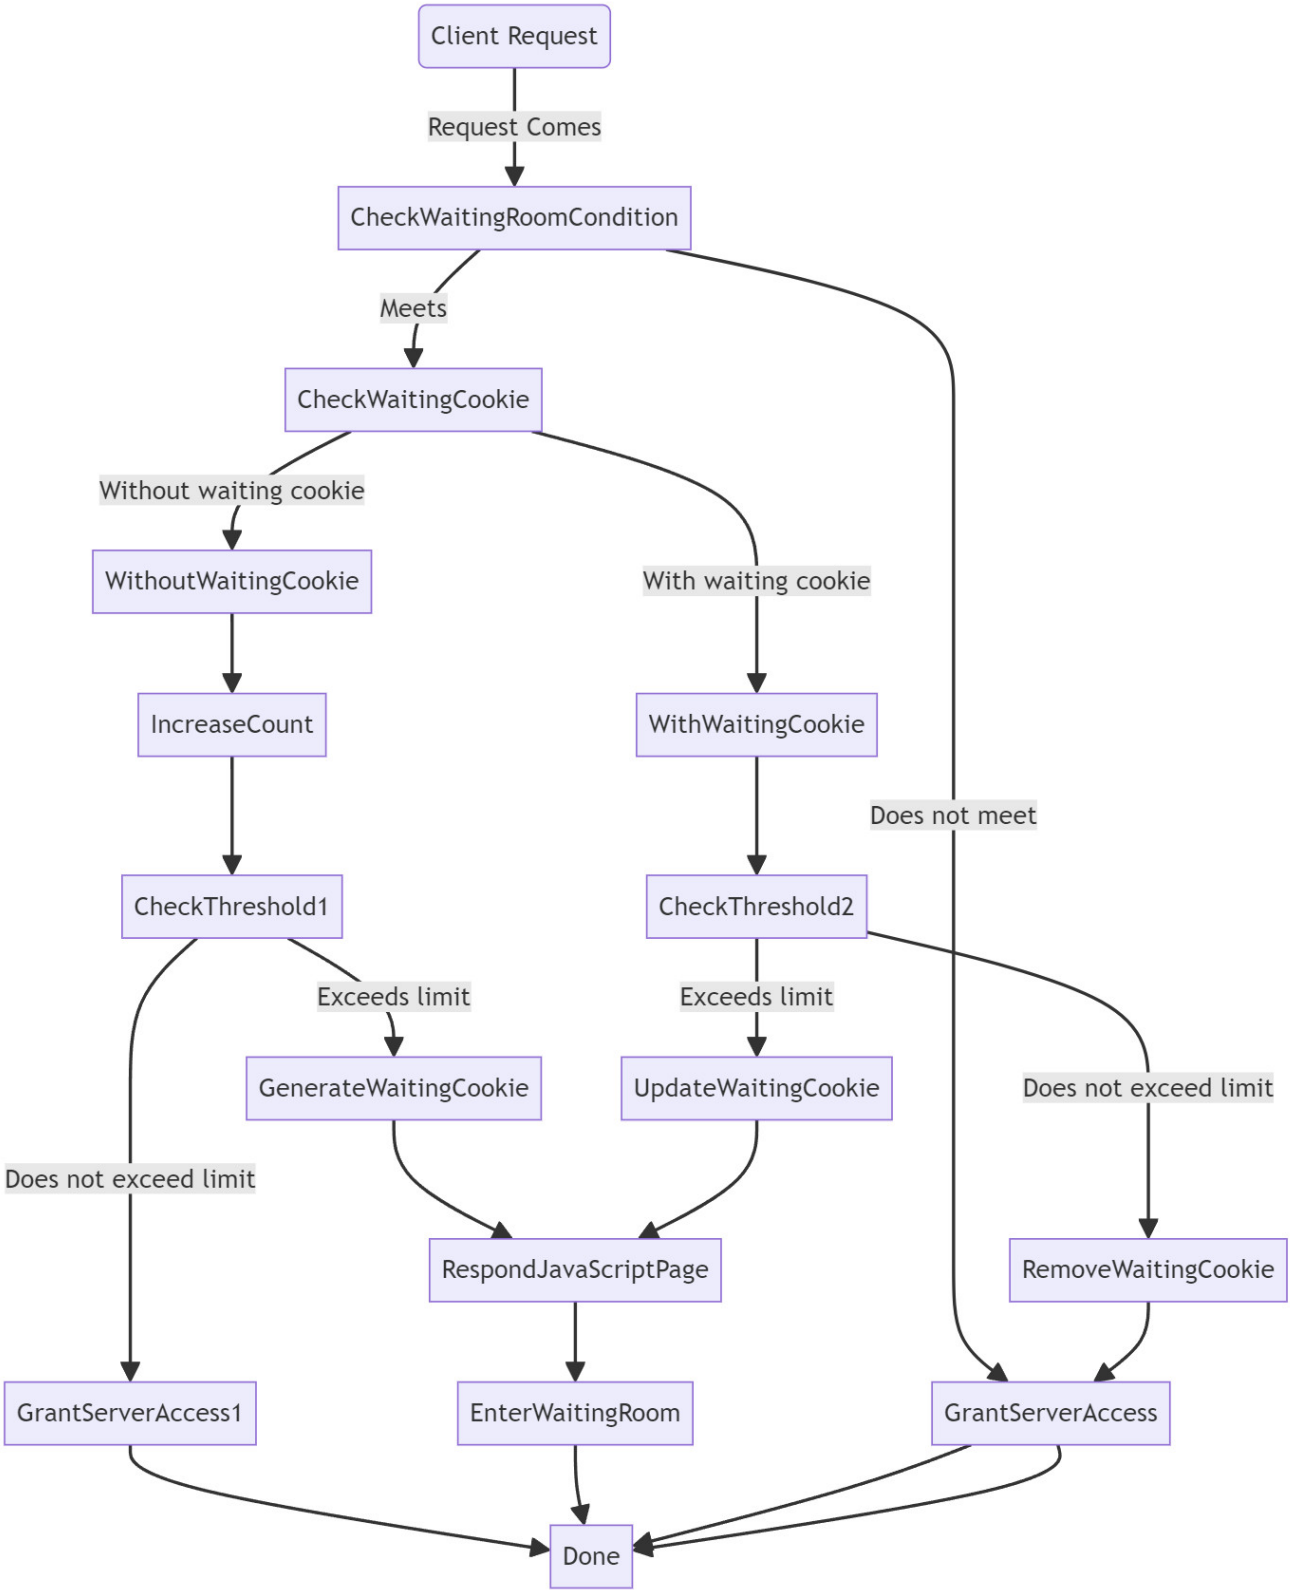
The Waiting Room feature is supported in Layer 7 HTTP/HTTPS virtual servers through Lua scripting.

Overview



The Waiting Room page displays the visitor's queue position, and a JavaScript refreshes the page automatically on a regular interval as specified in the Waiting Room script (such as every 10 seconds). Upon each refresh cycle, a waiting room cookie is returned, at which point FortiADC will examine the cookie to determine whether the visitor will stay in the queue. As such, the Waiting Room page must remain open.

The Waiting Room flow chart below illustrates how FortiADC decides whether visitors remain in the queue, advance their queue position, or exit from the waiting room.



When a client request is made, FortiADC considers the following factors:

- Waiting room condition — if the request matches the waiting room conditions specified in the script, such as the resource URI.
- Resource threshold limit — if the admission of the visitor will exceed the resource threshold limit.
- Waiting cookie — If a waiting cookie is present.

Based on these factors, FortiADC determines the path to take visitors, of which there are five possible paths that lead to either the waiting room page or the intended destination. First, the request must meet the conditions to activate the waiting room, such as matching the resource URI. If the conditions are not met, then visitors will be directed to their intended destination instead of the waiting room. If the waiting room conditions are met, then visitors will be set onto the next stage to either queue in the waiting room or be directed to the requested page based on the resource threshold limit. In this second stage, FortiADC will check for the presence of a waiting cookie — visitors without a waiting cookie will be assigned one to enter the waiting room queue if their admission will exceed the resource threshold limit; however, if the limit is not exceeded then the visitor will be directed to the requested page instead. When FortiADC detects the presence of a waiting cookie, this indicates that the visitor is already in the waiting room queue. FortiADC will then evaluate the threshold limit and either advance the visitor's queue position or direct them to the requested page.

Predefined HTTP scripts for implementing the Waiting Room:

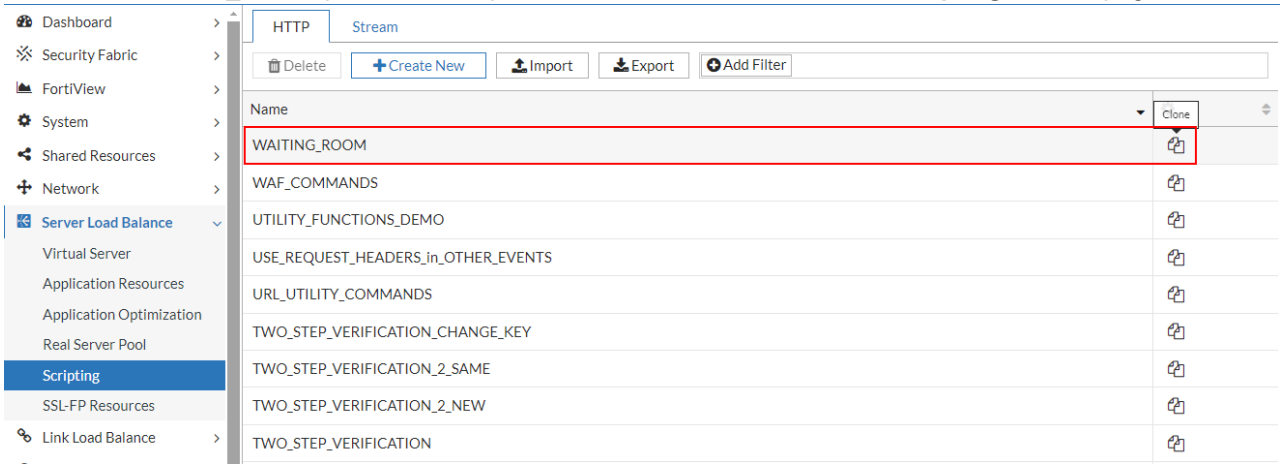
Predefined script	Description
WAITING_ROOM	The sample Waiting Room script demonstrates how you can place visitors in a virtual queue instead of denying them service directly when the server side reaches its configured capacity limit during high-demand situations. In this virtual Waiting Room, visitors can see their position in line and when their turn arrives, they are redirected to the requested page. Configuration parameters include the waiting room name, total resource limit threshold (default is 1000), and the Resource URL applicable to the waiting room. You can also customize the message displayed to users when they are placed in the waiting room by editing the HTML page section of the script. Required data structures such as atomic counters and shared tables are already built into the script; however, you have the option to apply user-defined atomic counters and shared tables to customize the script.
ATOMIC_COUNTER_ COMMANDS	Includes commands that allow you to create and configure shared atomic counters that are accessible by multiple httproxy processes within one VS. The stored data is located in shared memories. In the Waiting Room setup, the atomic counters track variables at running time, including the current resource count, the current position in line, and the current total number of users in the waiting queue.
SHARED_TABLE_COMMANDS	Includes commands that allow you to create and configure shared hash tables that are accessible by multiple httproxy processes within one VS. Both the table and stored data are located in shared memories. In the Waiting Room setup, the shared table is used to track current active resource occupiers such as active sessions.

Setting up the Waiting Room

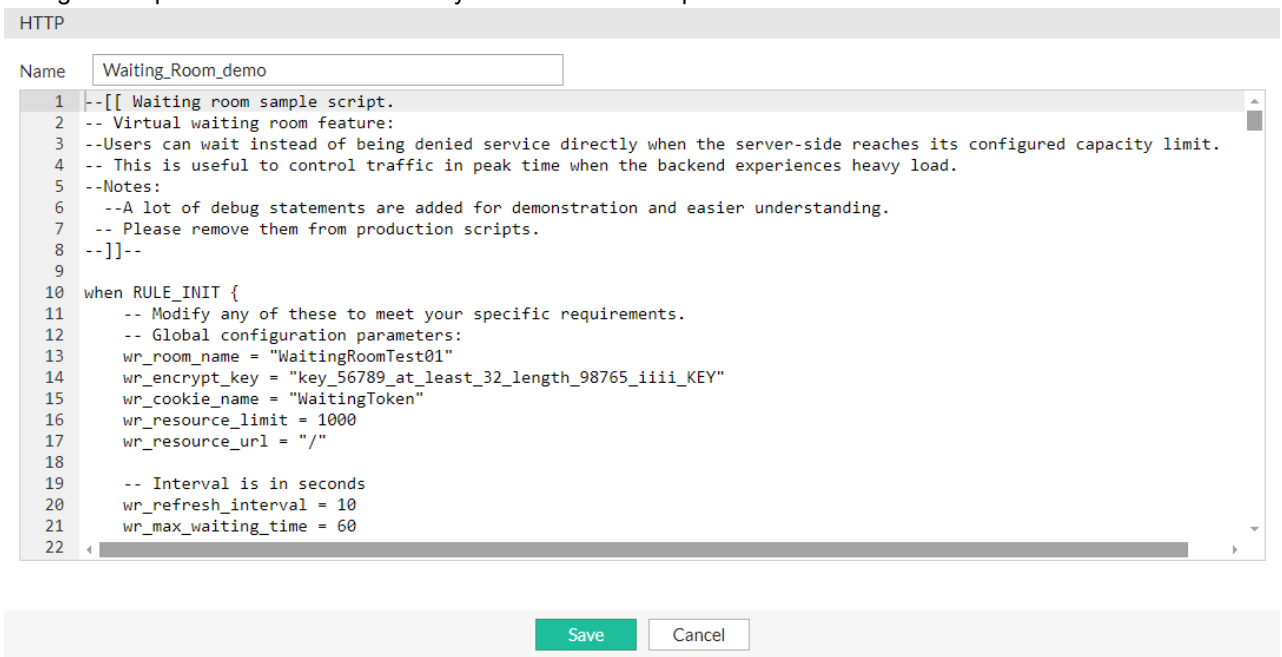
The WAITING_ROOM predefined HTTP script is designed to be used as a template for users to clone and modify as needed. Built-in instructions indicate which configuration parameters can be modified in the WAITING_ROOM script. All required data structures are already included in the template, such as atomic counters and shared tables; however, you have the option to create user-defined objects to use in your custom waiting room. The following steps describes the basic workflow to setting up a virtual waiting room to direct visitor overflow.

To set up a virtual waiting room:

1. Clone the **WAITING_ROOM** predefined script from the **Server Load Balance > Scripting > HTTP** page.



2. Name the cloned waiting room script and customize the parameters as needed, such as those under "Global configuration parameters" section. Save your customized script.

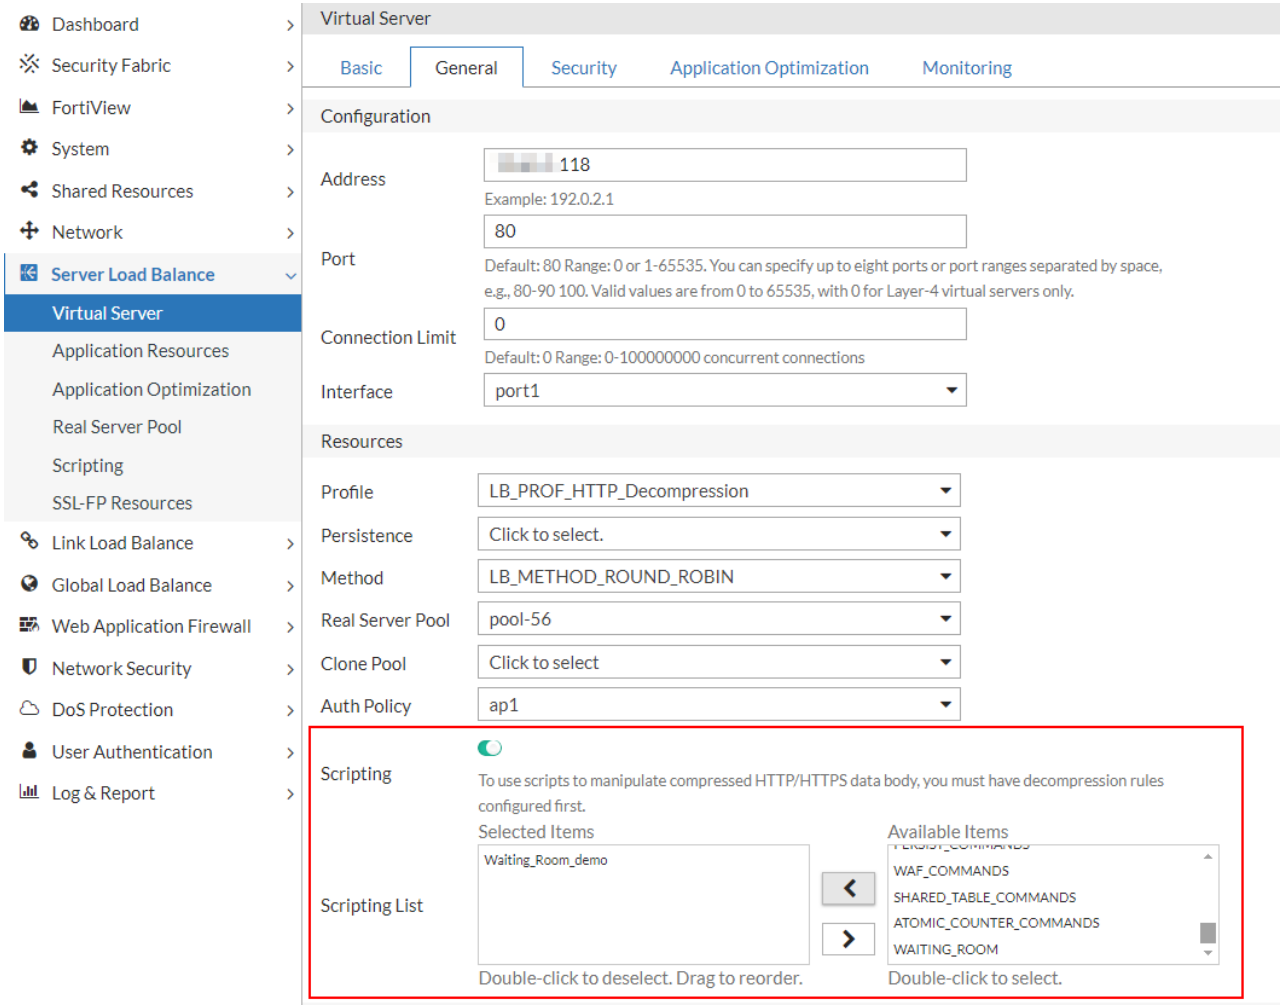


The resource limit, "wr_resource_limit" is often modified for testing purposes. A smaller threshold can make testing easier. For example, if the "wr_resource_limit = 2", then the threshold is three client sessions to trigger the waiting room. Otherwise, if the resource limit is 1000 instead, then it would require 1001 sessions to trigger the waiting room.

a. Optionally, you can customize and edit the HTML response page within the "waiting room HTML page" section of the script (line 39-71).

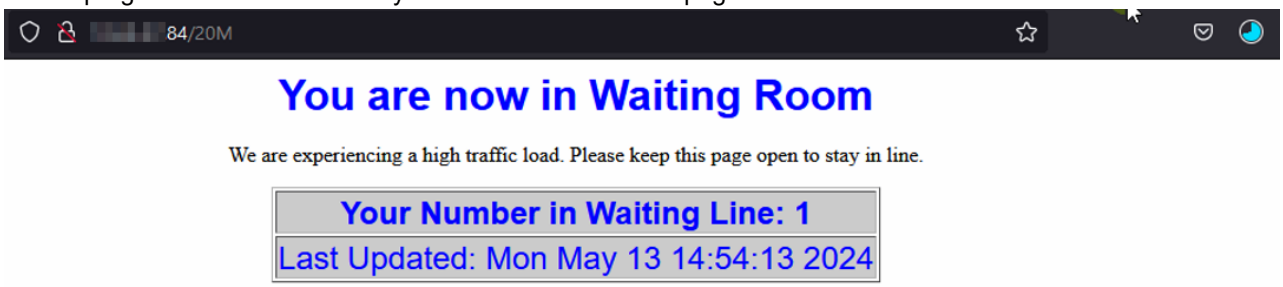
```
35  -- We reset the line num when all the current waiting clients get serviced,
36  -- Similar to repeat waiting numbers only till next day in real world.
37  wr_waiting_total_cnt_name = "WaitingTotalCounterDemo1"
38
39  -- This is the waiting room HTML page, the script will fill a waiting number and a last updated timestamp.
40  -- The embedded javascript is to refresh automatically. Other contents are customizable.
41  wr_payload_template =
42      [[
43      <html>
44      <head>
45      <meta http-equiv="Content-Language" content="en-us">
46      <meta http-equiv="Content-Type" content="text/html; charset=us-ascii">
47      <title>Waiting Room Test Page</title>
48          <style>
49              table {
50                  margin-left: auto;
51                  margin-right: auto;
52              }
53          </style>
54      </head>
55      <script type="text/javascript">
56
```

3. Apply your waiting room script in the virtual server. Ensure this virtual server is a Layer 7 HTTP or HTTPS VS and is enabled.



4. Use a browser on the client side to test the waiting room setup by making requests to the configured VS to ensure the waiting room behaves as intended.
You should be directed to your virtual waiting room and be placed in the queue if you have exceeded the resource

limit. After some time has passed, when some of the HTTP connections has been closed, your queue position should progress and then eventually be directed to the server page.



Using Stream scripting

Stream Scripting in FortiADC allows you to perform actions that are not supported by the current built-in feature set for Layer 7 TCP and UDP virtual servers. Stream scripts enable you to use predefined script commands and variables to manipulate the TCP/UDP request and response. You can import Stream scripts from the FortiADC GUI. To get started, FortiADC provides system predefined scripts that can be cloned for customization.

FortiADC HTTP scripts are based on Lua 5.3.



FortiADC does not support all Lua functions. For the full list of supported functions, see the Script Reference Guide.

Please note that Scripting is an *as is* feature and does not come with any functionality or performance guarantees.

This section includes the following:

- Configuring Stream Scripting on page 588
- Predefined Stream scripts on page 590
- Multi-script support for Stream Scripting on page 591

Configuring Stream Scripting

From the FortiADC GUI, you can type or paste the script content into the configuration page. Alternatively, you can clone a system predefined script to customize. For details, see . From the HTTP Script page, you also have the option to import, export, and delete scripts.

Before you begin:

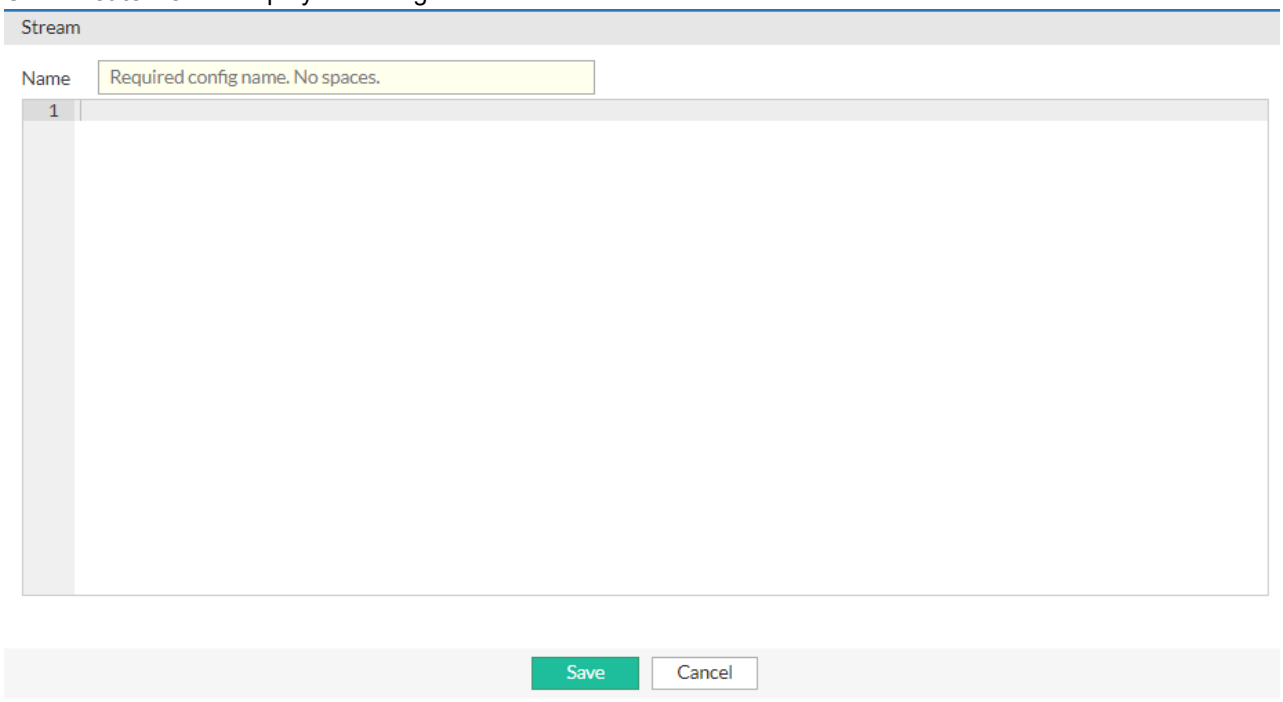
- You must have Read-Write permission for Server Load Balance settings.

After you have created a script configuration object, you can specify it in the virtual server configuration.

To create a Stream script configuration object:

1. Go to **Server Load Balance > Scripting**.
2. Click the **Stream** tab.

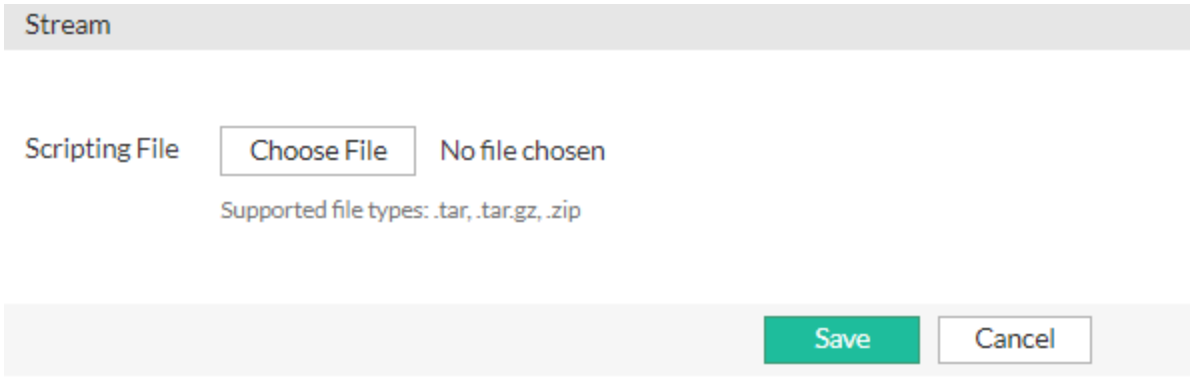
3. Click **Create New** to display the configuration editor.



4. Enter a unique name for the Stream script configuration. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No spaces. After you initially save the configuration, you cannot edit the name.
5. In the text box, type or paste your Stream script.
   If you want to include this script as part of a multi-script configurations that allows you to execute multiple scripts in a certain order, ensure to set its priority. For more information, see Multi-script support for Stream Scripting on page 591.
6. Click **Save**.
   Once the Stream script configuration is saved, you can specify it in the virtual server.

To import a Stream script:

1. Go to **Server Load Balance > Scripting**.
2. Click the **Stream** tab.
3. Click **Import** to display the file import options.

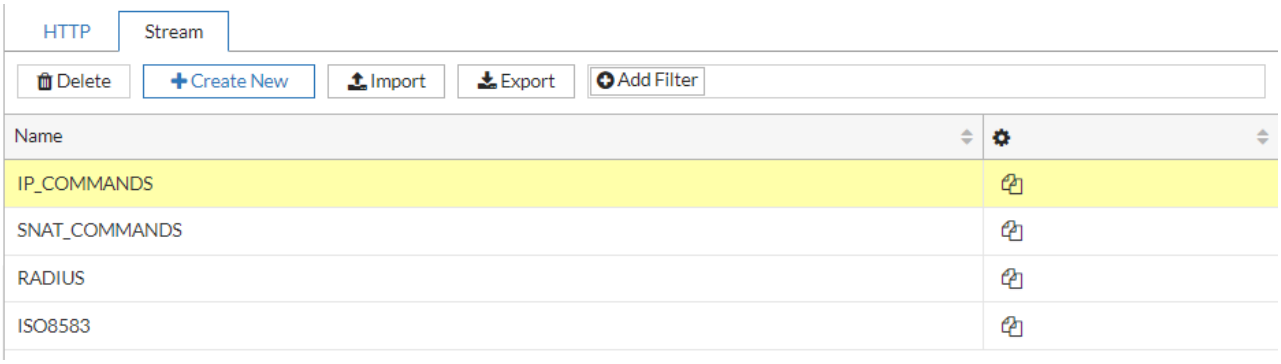


4. Click **Choose File** and browse for the script file. Supported file types are .tar, .tar.gz, and .zip.

5. Click **Save**.
   Once the file is successfully imported, it will be listed in the **Scripting > Stream** page.

To export a Stream script:

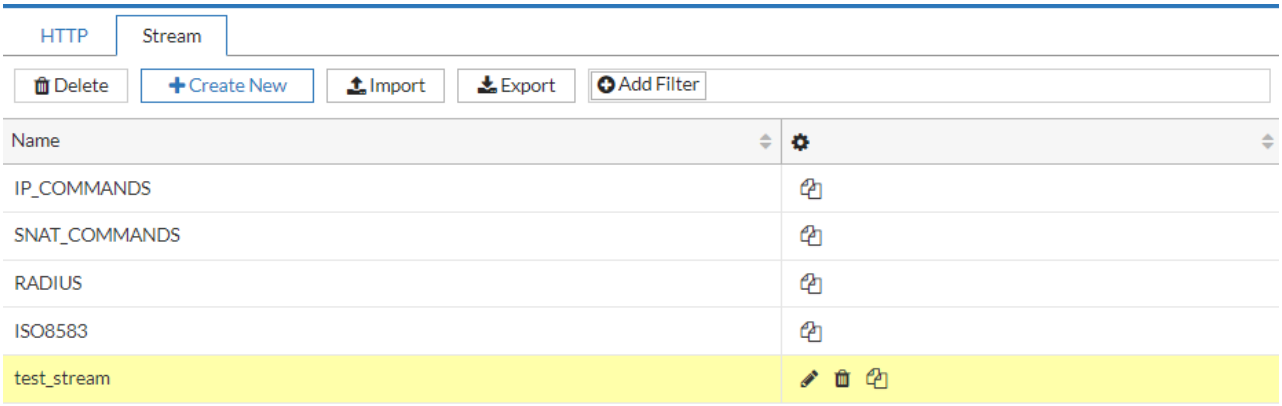
1. Go to **Server Load Balance > Scripting**.
2. Click the **Stream** tab.
3. From the **Stream** page, select a Stream script configuration.
   In the example below, the IP_COMMANDS script is selected.



4. Click **Export** initiate the file download.
   The selected script configuration will be exported as a .tar file.

To delete a Stream script:

1. Go to **Server Load Balance > Scripting**.
2. Click the **Stream** tab.
3. From the **Stream** page, select a user-defined Stream script configuration. System predefined scripts cannot be deleted.
   In the example below, the **test_stream** script configuration is selected.



4. Click **Delete** from the top navigation, or click (delete icon) of the configuration.
   Multiple script configurations can be deleted using the **Delete** button on the top navigation.

Predefined Stream scripts

FortiADC provides system predefined scripts for Stream Scripting.

The following table lists the FortiADC predefined scripts available for users to apply and customize.

Predefined script	Usage
IP_COMMANDS	Used to get various types of IP addresses and port numbers between the client and server side.
SNAT_COMMANDS	Allows you to overwrite client source address to a specific IP for certain clients.
ISO8583	Implements ISO8583 load balancing based on message content.
RADIUS	Implements RADIUS load balancing based on message content.

Multi-script support for Stream Scripting

Linking multiple scripts to the same virtual server

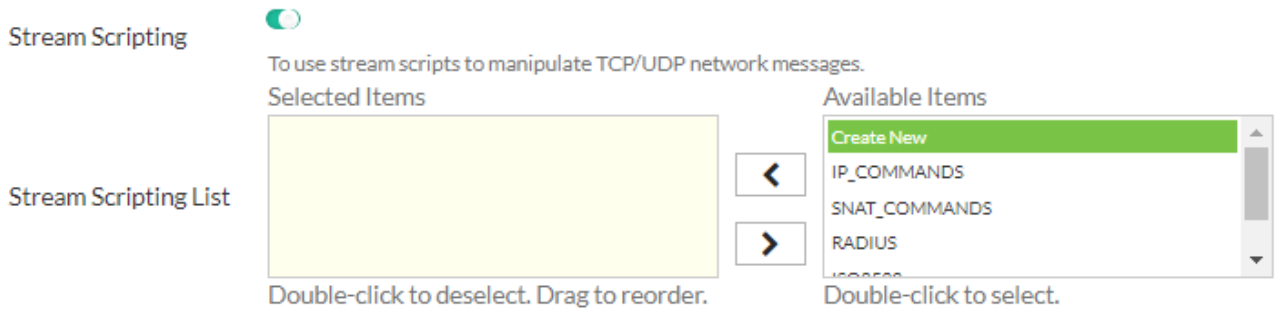
FortiADC supports the use of a single script file containing multiple scripts and applies them to a single virtual server in one execution. Different scripts can contain the same event. You can specify the priority for each event in each script file to control the sequence in which multiple scripts are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

Currently, up to 16 individual scripts can be added to create a large multi-script file.

In practice, instead of creating a single large and complex script containing all necessary logic, it's often more advantageous to decompose it into smaller functional components represented by individual scripts. This approach offers several benefits. Firstly, executing multiple scripts concurrently is more efficient than running them sequentially. Additionally, breaking down a massive script into smaller units enhances flexibility, particularly when applying scripts to various virtual servers. Some servers may require only specific scripts, while others may utilize all available ones. With smaller, modular scripts, you have the flexibility to select and combine only the necessary components to construct a comprehensive multi-script file, each with its designated priority, and apply them collectively to a virtual server.

Apply multiple scripts on page 591 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 STREAM_CLIENT_INIT priority 100 {
    local header = TCP:peek(20)
}
```

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

Script 1:

```
when STREAM_REQUEST_DATA priority 500 {
    local header = TCP:receive(20)
}
```

Script 2:

```
when STREAM_REQUEST_DATA priority 500 {
    local header = TCP:receive(20)
    local len = string.byte(string.sub(header, 3, 3))

    if (len < 2) then
        TCP:reject()
    end
}
```

Script 3:

```
when STREAM_REQUEST_DATA priority 400 {
    TCP:set_src_key(get_key())
}
```

Script 4:

```
when STREAM_RESPONSE_DATA priority 600 {
    TCP:set_dst_key(get_key())
}
```

Individual script files are loaded separately into the Lua stack. A numeric value (starting from 1) is appended to each event (e.g., for STREAM_REQUEST_DATA event, there are functions `STREAM_REQUEST_DATA1`, `STREAM_REQUEST_DATA2`, and so on so forth).

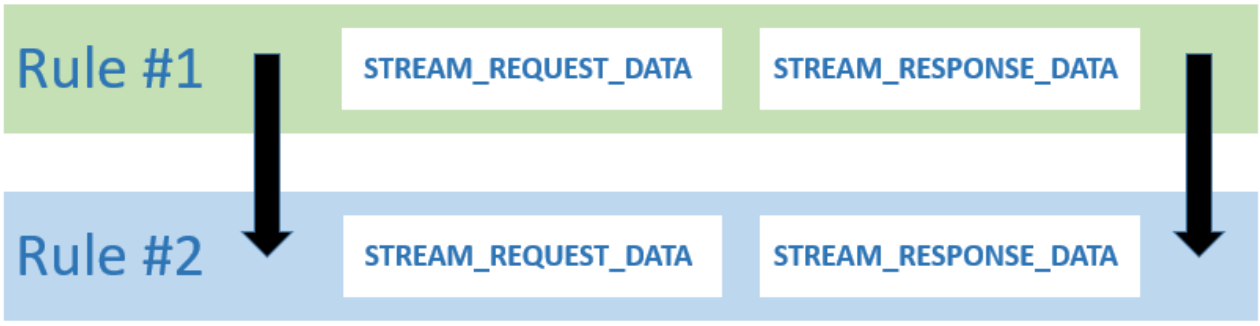
To support multiple scripts, FortiADC:

- 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.

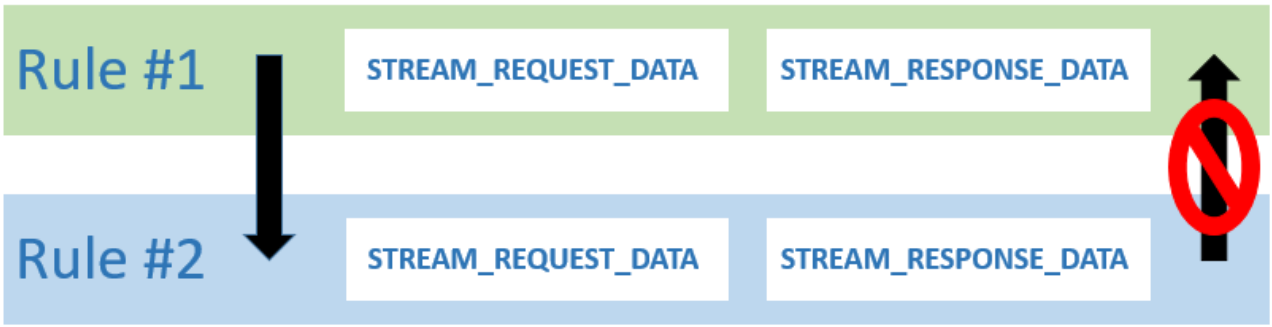
Compiling principles

- 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 7 virtual servers on TCP or UDP protocols 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 configuration can contain up to 256 individual scripts, each being no more than 32 kilobytes.

SSL-FP Resources

The SSL-FP Resources module contains configurations to support SSL Forward Proxy functionality, including SSL decryption exceptions using L2 Exception Lists and Web Filtering, and Certificate Caching that enhance performance by storing certificates to avoid repeated retrievals.

The **Server Load Balance > SSL-FP Resources** sub-menu contains the following configurations and pages:

- **L2 Exception List** — The L2 Exception List configures exceptions for SSL decryption, allowing exclusion of specific websites based on FortiGuard categories or custom host name/destinations IP/source IP, ensuring compliance with jurisdictional restrictions. For details, see Configuring an L2 exception list on page 594.
- **Web Filter Profile** — The Web Filter Profile allows you to group FortiGuard categories for inclusion in the L2 Exception List. For details, see Creating a Web Filter Profile configuration on page 595.
- **Web Category** — The Web Category page allows you to view the Web Categories imported from FortiGuard used to create the Web Filter Profile. For details, see Viewing Web Filter categories on page 596.
- **Certificate Caching** — Certificate Caching enables the system to store presented certificates for future use which can enhance performance by avoiding repeated retrievals from the remote server and locally resigned certificate by SNI. For details, see Configuring certificate caching on page 596.

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 L2 Exception List configuration editor.
4. Configure the following L2 Exception List settings:

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.

Settings	Guidelines
Web Filter Profile	Select a Web Filter Profile configuration.

5. Click **Save**.
   The Member section becomes available to configure.
6. In the **Member** section, click **Create New** to display the Member configuration editor.
7. Configure the following Member settings:

Settings	Guidelines
Type	How you want to define the exception: • Host • IP/Netmask
Host Pattern	The **Host Pattern** option is available if the **Type** is **Host**. Specify a wildcard pattern, such as `*.example.com`.
IP/Netmask	The **IP/Netmask** option is available if the **Type** is **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.
IP/Netmask Role	The **IP/Netmask** option is available if the **Type** is **IP/Netmask**. Specify the role of the IP/Netmask: • Destination — The IP/Netmask is set as the Destination IP, and the L2 SSL Forward Proxy VS will be bypassed based on this Destination IP. • Source — The IP/Netmask is set as the Source IP, and the L2 SSL Forward Proxy VS will be bypassed based on this Source IP.
Comments	Optionally, enter a comment to describe this L2 Exception List Member.

8. Click **Save**.
   The L2 Exception List Member configuration is saved and added to the Member section.
9. Click **Save** to commit the changes made for the Member configuration.

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 596.
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.

Viewing Web Filter categories

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.

To 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 597.
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.

Link Load Balance

The **Link Load Balance** menu contain features and configurations that allow you to link load balance.

This section is organized into the following sub-menu topics:

- Link Policy on page 603
- Link Group on page 604
- Virtual Tunnel on page 611

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.

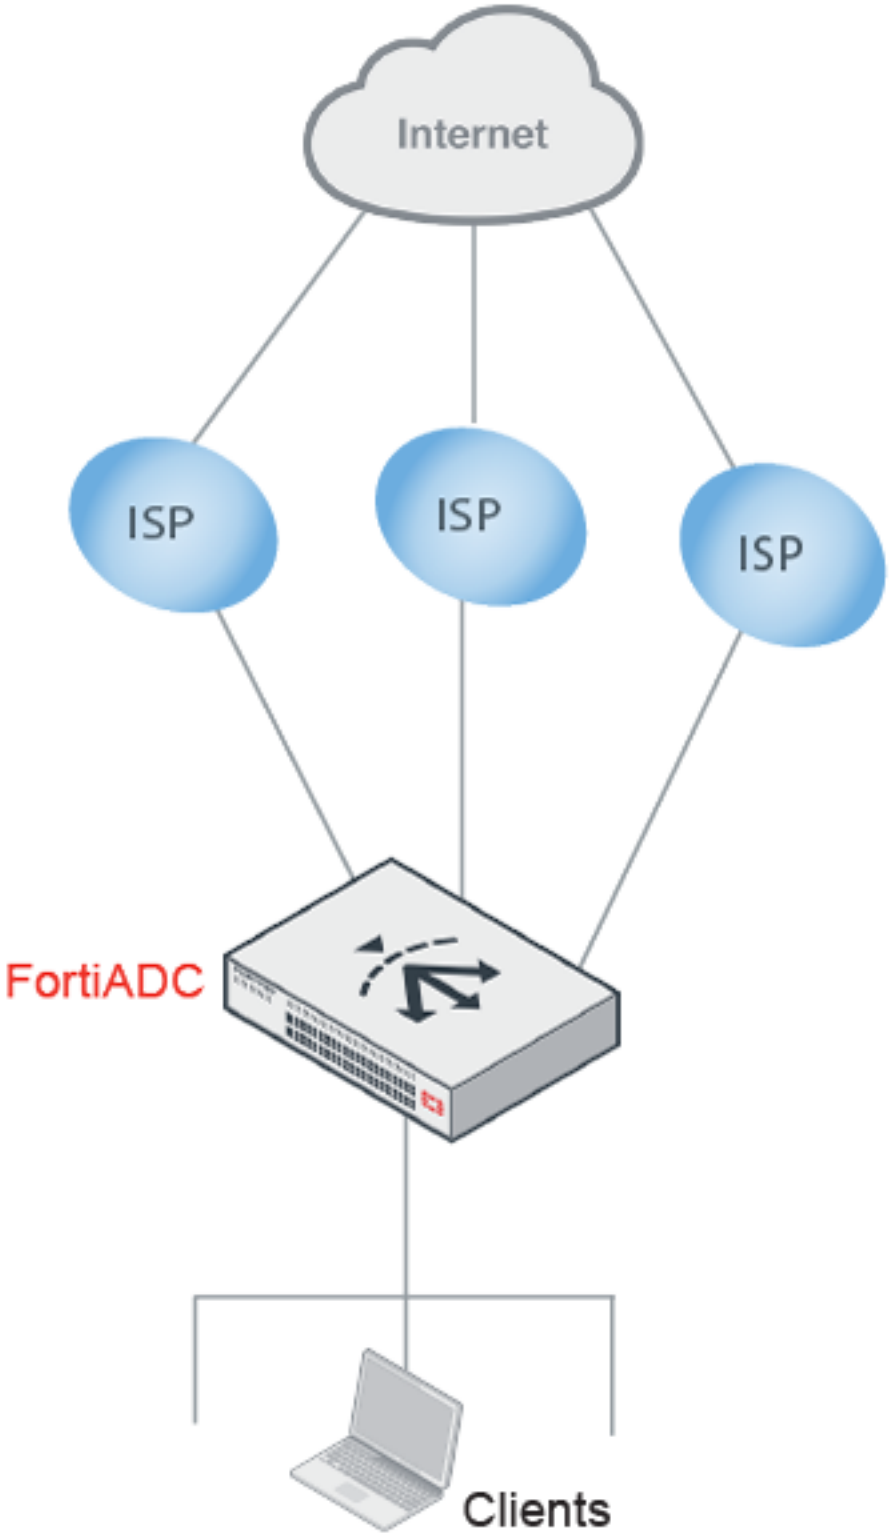
The LLB feature supports load balancing among link groups or among virtual tunnel groups.

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 598 shows an example topology when FortiADC is deployed to support link groups.

LLB link groups

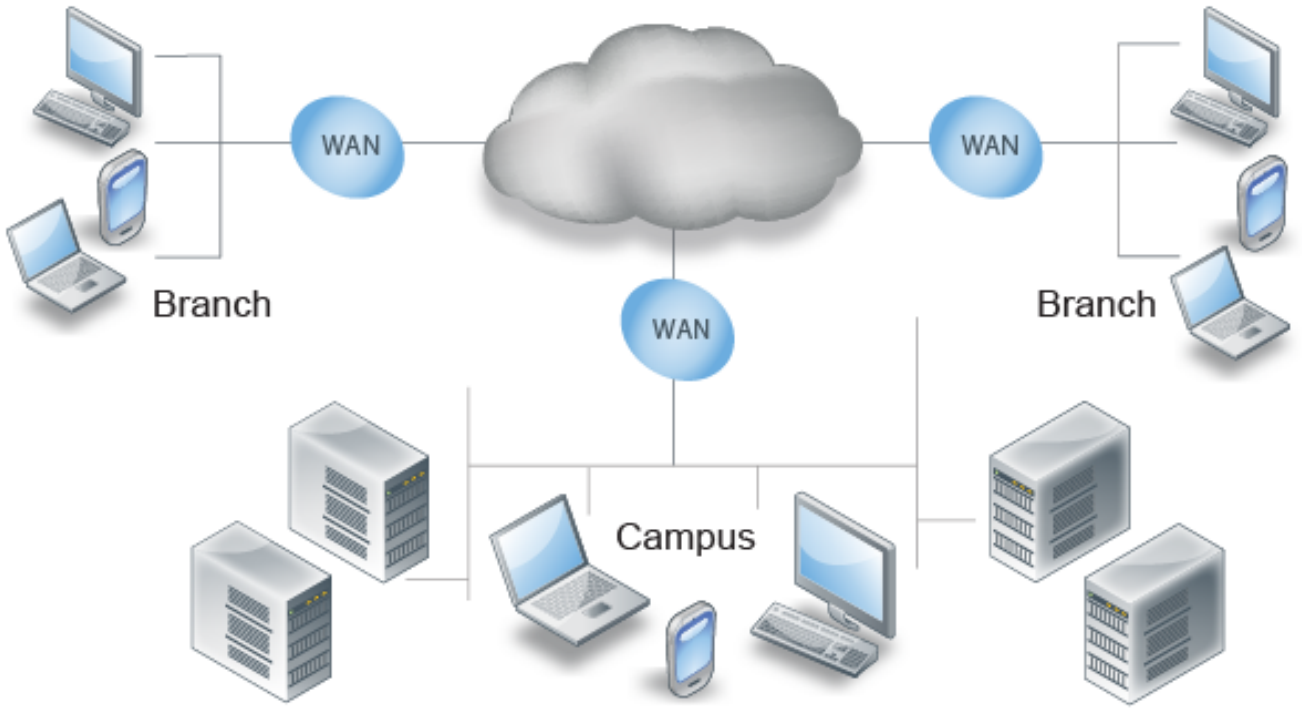

Internet
ISP
ISP
ISP
FortiADC
Clients

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.

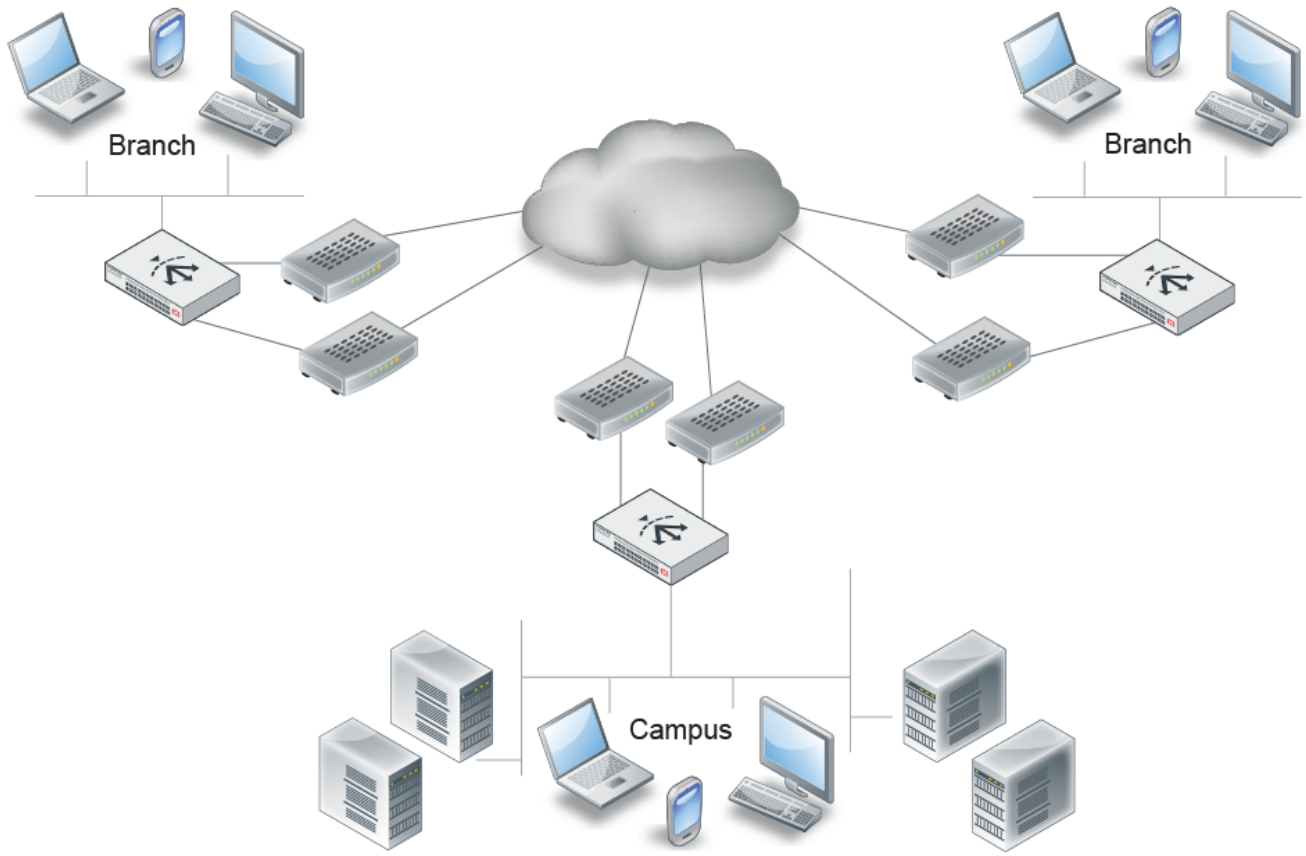
WAN connectivity over single leased lines on page 600 shows an example of a deployment that does not use LLB. It uses dedicated leased lines for its WAN links, which are reliable, but expensive.

WAN connectivity over single leased lines



LLB virtual tunnels on page 600 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:

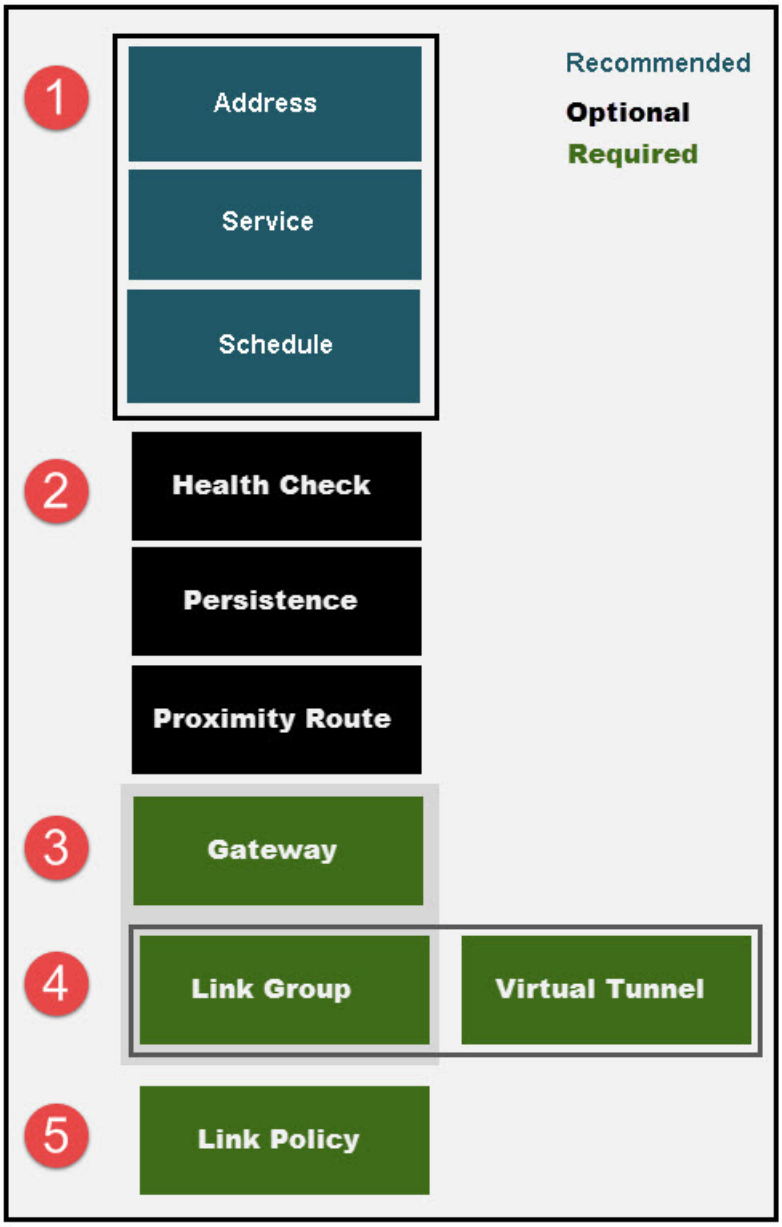
1. LLB link policy
2. Policy route
3. Static/Dynamic route
4. LLB default link group

Link load balancing configuration overview

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

LLB configuration summary on page 601 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.

Link Policy

From the **Link Load Balance > Link Policy** sub-menu, you can configure link policies to match traffic to rules that use a link group or virtual tunnel.

The link 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 604.
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 Schedule Group.
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 Virtual Tunnel.
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.

Link Group

The **Link Load Balance > Link Group** sub-menu allows you to perform the following tasks relating to grouping links for link load balancing:

- Configuring a link group on page 605
- Configuring gateway links on page 607

- Configuring LLB persistence rules on page 608
- Configuring proximity route settings on page 610

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 605.
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).

Settings	Guidelines
	• 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 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 607.
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 the 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.
Outbound Spillover Threshold	Maximum outbound bandwidth rate for a link in a spillover load balancing pool.

Settings	Guidelines
	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 LLB 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 608 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. Source address based persistence consumes a significant amount of memory, as calculated using the following formula per VS: *(max persistence entry size) x (size per entry of the table) x (content routing or pool count)* For example: If the max persistence entry size is 262144 (default), the per entry of the table is 44 bytes, and there is no content routing. 262144 x 44 = 11534336 bytes (which is ≈11MB)
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 609.
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 following the persistence rule.
Destination IPv4 Netmask Bits	Number of bits in a subnet mask to specify a network segment that should following 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 following 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 following 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 610.
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.

Virtual Tunnel

From the **Link Load Balance > Virtual Tunnel** sub-menu, you can configure virtual tunnels that allow reliable, site-to-site connectivity using Generic Routing Encapsulation (GRE) to tunnel traffic between pairs of FortiADC appliances.

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:

1. Go to **Link Load Balance > Virtual Tunnel**.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration and add members as described in Virtual tunnel configuration on page 611.
4. 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.

Global Load Balance

The **Global Load Balance** menu contain features and configurations that allow you to global load balance.

This section is organized into the following sub-menu topics:

- GLB Wizard on page 617
- FQDN on page 619
- Zone Tools on page 626
- Global Object on page 653

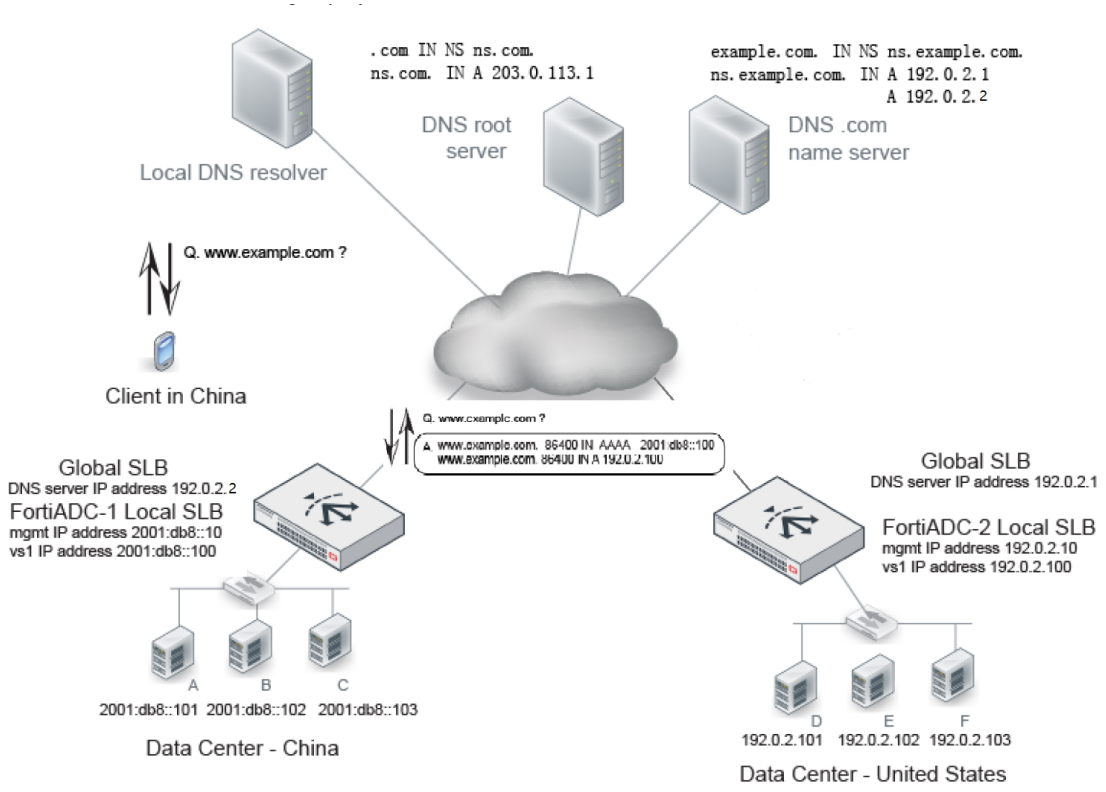
Global load balancing basics

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 613 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-1 and FortiADC-2 are also the GSLB. They host the DNS servers that are 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 them. 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 (in the example above, both FortiADC-1 and FortiADC-2 are GSLBs). 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 615 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


DNS Server and Policy Settings
1
Address Group
2
Remote DNS Server
DSSET List
3
Zone
4
DNS64
Response Rate Limit
5
Global DNS Policy
6
General Settings
Trust Anchor Key
Global Load Balancing
Servers and Locations Settings
1
Data Center
Servers
Virtual Server Pool
Host
Dynamic Proximity
Persistence
2
Zone
3
Global DNS Policy
Required
Required with predefined (any or none)
Optional with defaults
Optional
Zone is generated

Basic steps (DNS server)

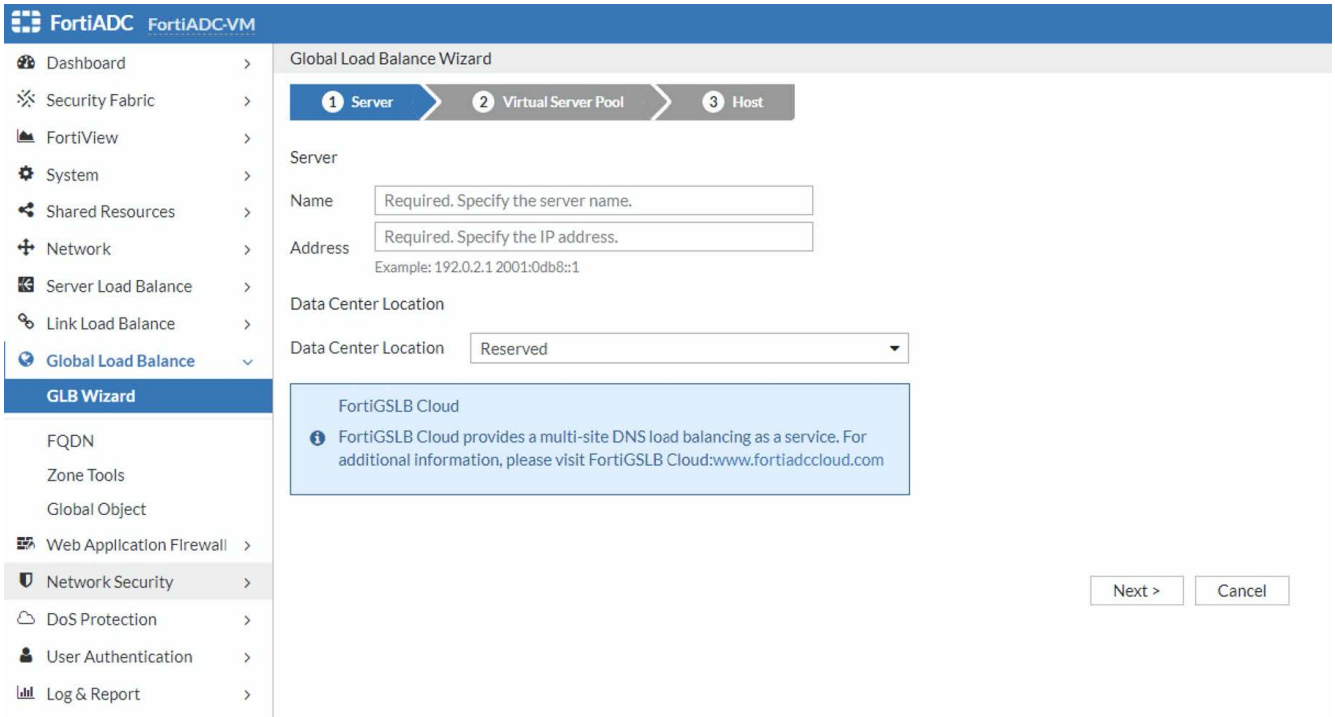
1. Configure address groups to use in your DNS policy matching rules. The system includes the predefined address groups **any** and **none**.
2. Configure remote DNS servers (forwarders) and the DSSET list that you might reference in the zone configuration.
3. Complete the zone configuration. The global 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.

GLB Wizard

From the **Global Load Balance > GLB Wizard** sub-menu, you can configure a GLB object in the GUI using a step-by-step tool to easily guide you through the GLB configuration process.



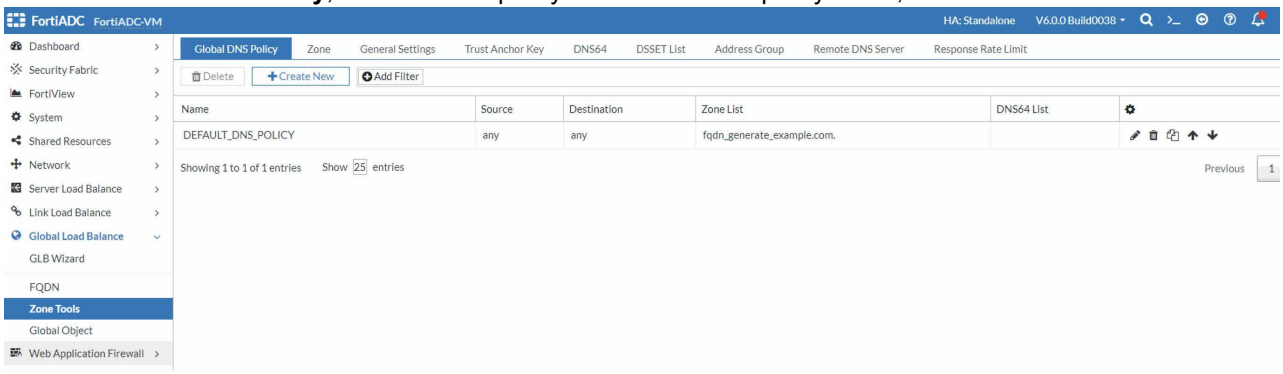
To configure a GLB wizard:

1. Go to **Global Load Balance > GLB Wizard**.
   The configuration page displays the Global Load Balance Wizard configuration editor.
2. In **Server**, configure the Name, Address, and Data Center Location.
3. 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.
4. 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.
5. 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.

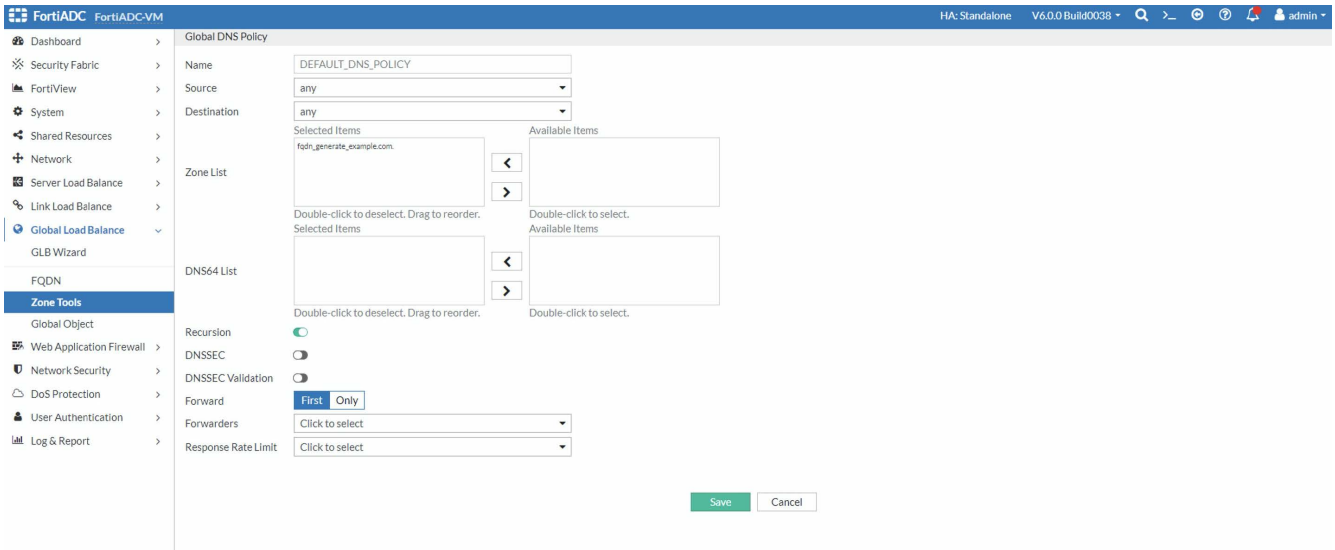
In **Zone Tools**, three things will happen:

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**.



FQDN

The **Global Load Balance > FQDN** sub-menu allows you to perform the following tasks relating to the FQDN (fully qualified domain name):

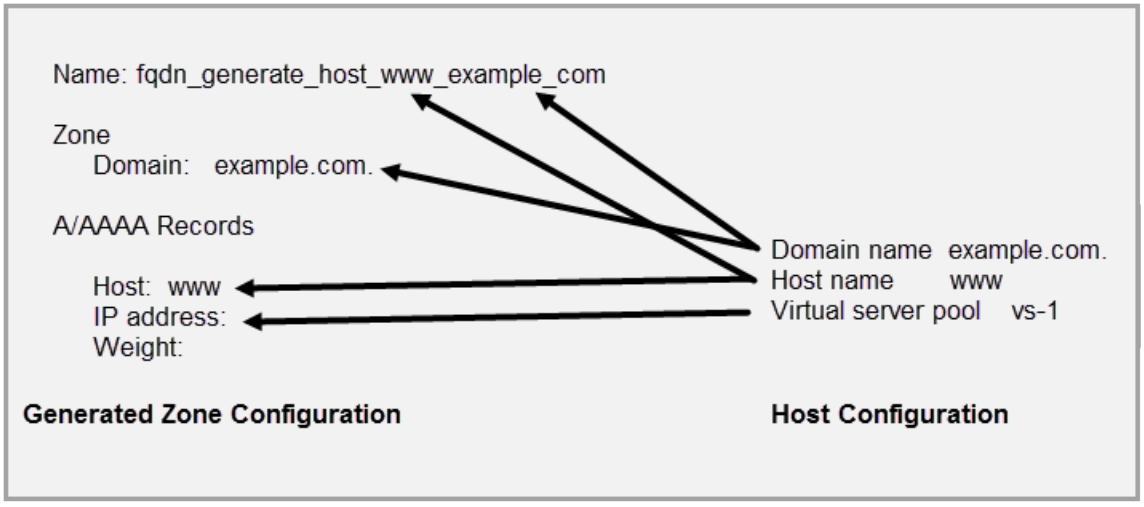
- Configuring hosts on page 619
- Configuring virtual server pools on page 621
- Configuring location lists on page 623
- Configuring GLB settings on page 623

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 619 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:

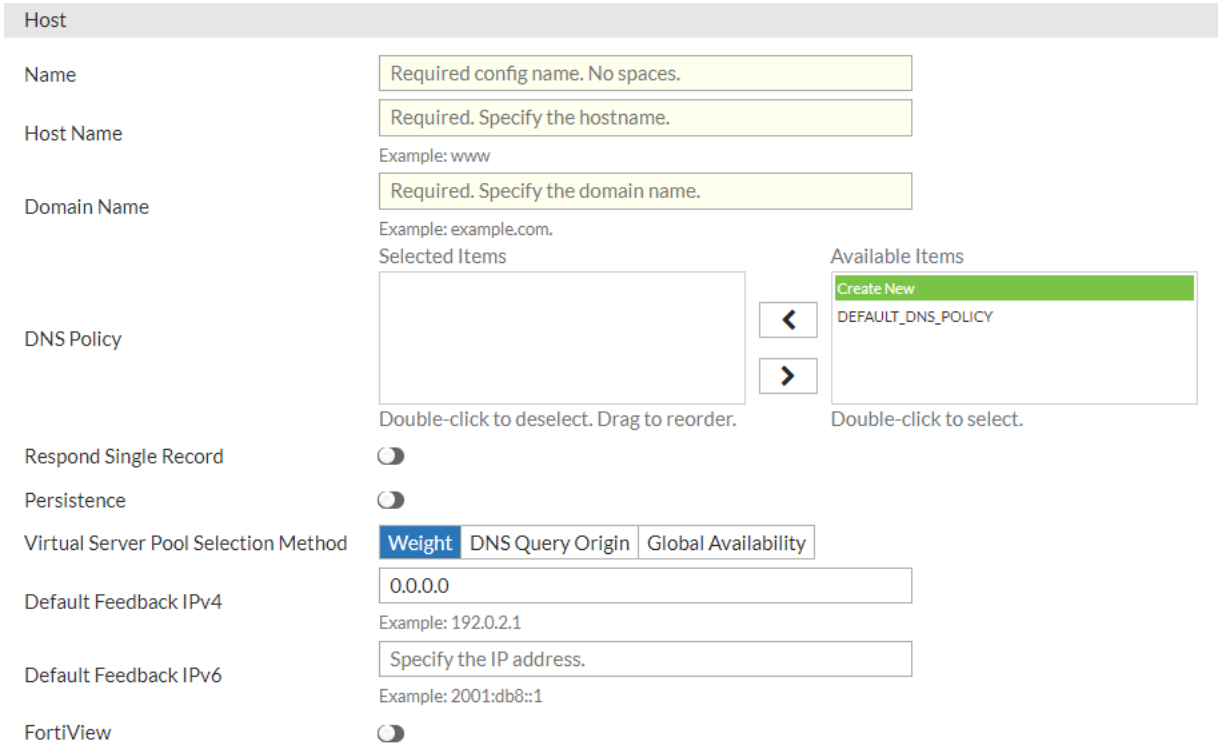
- You must have created the global virtual server pools you want to use.
- You must have Read-Write permission for Global Load Balance settings.

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**.
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 620.
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 `www`. **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: `example.com`.
DNS Policy	Select one or more DNS policies 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

Settings	Guidelines
	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 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**.
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 622.
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`, `_`, and `-`. 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-Connnections—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. • Bytes-Per-Second—Virtual servers with the lowest traffic will respond. • Server-Performance—Virtual servers with better server-performancec 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**.
2. Click the **Location List** tab.
3. Complete the configuration as described in Location List settings on page 623.
4. Click **Save**.

Location List settings

Settings	Guidelines
Name	Specify a unique name for the location list.
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 "tp1"
      set member ZZ US CN65
   next
end
```

Where ZZ stands for Reserved, US for United States, and CN65 for China, Xingjiang

Configuring GLB settings

Configure the global load balancing settings of the FQDN, including the following:

- Listening interface — The listen port is used for communication for GLB and SLB server.
- Authentication — Configure the authentication between the GLB and the server.
- Dynamic proximity — This is used to order DNS lookup results based on the shortest application response time (RTT) for ICMP or TCP probes sent by the local SLB to the DNS resolver that sent the DNS request. The system caches the RTT results for the period specified by the timeout. When there are subsequent requests from clients that have a source IP address within the same network (as specified by the netmask affinity), the RTT is taken from the results table instead of a new, real-time probe. This reduces response time.
- Persistence — Configure the source address affinity and a timeout for GSLB persistence. You enable persistence per host in the GSLB host configuration.
  If the DNS query is for a host that has persistence enabled, the DNS server replies with a response that has the

virtual server IP addresses listed in the order determined by the GSLB proximity algorithms, and the client source IP address (for example 192.168.1.100) is recorded in the persistence table. If source address affinity is set to 24 bits, subsequent queries for the host from the 192.168.1.0/24 network 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).
Persistence is required for applications that include transactions across multiple hosts, so the persistence table is also used for queries for other hosts with the same domain. For example, a transaction on a banking application might include connections to login.bank.com and transfer.bank.com. To support persistence in these cases, the GSLB persistence lookup accounts for domain as well. The first query for login.bank.com creates a mapping for the source address network 192.168.1.0/24 and the domain bank.com. When the DNS server receives subsequent requests, it consults the persistence table for a source network match, then a domain match and a hostname match. In this example, as long as you have created host configurations for both login.bank.com and transfer.bank.com, and persistence is enabled for each, the persistence table can be used to ensure the DNS answers to queries from the same network list the resource records in the same order.

Before you begin:

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

To configure the GLB settings of the FQDN:

1. Go to **Global Load Balance > FQDN**.
2. Click the **GLB Setting** tab.
3. Configure the following GLB settings:

Setting	Description
Auth Type	Select the authentication type: • None — No password. • TCP MD5SIG — With password, but cannot 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	The **Password** option is available if **Auth Type** is **TCP MD5SIG** or **Auth Verify**. Enter the password to authenticate the key. This password is used for authentication between the GLB and the server. The same password must be set on both, otherwise the two will not be able to synchronize.
CA Verify	Enable/disable the certificate verification when synchronizing the SLB information to the GSLB server.
Trusted CA Group	The **Trusted CA Group** option is available if **CA Verify** is **enabled**. Select a trusted CA group to verify the peer certificate.

Setting	Description
Trusted Intermediate CA Group	The **Trusted Intermediate CA Group** option is available if **CA Verify** is **enabled**. Select a trusted intermediate CA group to verify the peer certificate.
IPv4 Accessed Status	Enable/disable listening for DNS requests on the interface IPv4 address.
IPv6 Accessed Status	Enable/disable listening for DNS requests on the interface IPv6 address.
Listen on All Interfaces	Enable/disable IPv4/IPv6 network access status on all interfaces.
Listen on Interface List	The **Listen on Interface List** option is available if **Listen on All Interfaces** is **disabled**. Specify specific interfaces for IPv4/IPv6 network access.
Listen on Port	Specify the port to listen on. Default: 5858 Range: 1-65535.

4. Configure the following Proximity settings:

Setting	Description
Protocol	Select the proximity detection protocol: • ICMP • ICMP and TCP
Retry Number	Specify the number of retries if the probe fails. The default is 3. The valid range is 1-10 times.
Retry Interval	Specify the interval between retries if the probe fails. The default is 3. The valid range is 1-3600 seconds.
IPv4 Prefix Length	Specify the number of IPv4 netmask bits that define network affinity for the RTT table. The default is 24. For example, if the GLB records an RTT for a client with source IP address 192.168.1.100, the record is stored and applies to all requests from the 192.168.1.0/24 network.
IPv6 Prefix Length	Specify the number of IPv6 netmask bits that define network affinity for the RTT table. The default is 64.
Aging Timeout	Specify the for how long RTT results are cached. This setting specifies the length of time in seconds for which the RTT cache entry is valid. The default is 86400. The valid range is 60-2,592,000 seconds.

5. Configure the following Persistence settings:

Setting	Description
IPv4 Mask Length	Specify the number of IPv4 netmask bits that define network affinity for the persistence table. The default is 24.
IPv6 Mask Length	Specify the number of IPv6 netmask bits that define network affinity for the persistence table. The default is 64.
Aging Period	This setting specifies the length of time in seconds for which the entry is

Setting	Description
	maintained in the persistence table. The default is 86400. The valid range is 60-2,592,000 seconds.

6. Click **Save**.

Zone Tools

The **Global Load Balance > Zone Tools** sub-menu allows you to perform the following tasks relating to the zone tools for global load balancing:

- Configuring a Global DNS policy on page 626
- Configuring DNS zones on page 627
- Configuring general settings on page 634
- Configuring the trust anchor key on page 648
- Configuring DNS64 on page 649
- Configuring the DSSET list on page 650
- Configuring an address group on page 651
- Configuring remote DNS servers on page 652
- Configuring the response rate limit on page 652

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.
- You must have configured address objects, remote servers, DNS zones, and optional configuration objects you want to specify in your policy.
- You must have 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 627.
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.
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. **Note**: DNS forwarding functionality requires Recursion to be enabled.
DNSSEC Validation	Enables/disables DNSSEC validation.
Forward	DNS forwarding functionality requires **Recursion** to be **enabled**. Select from the following DNS forwarding options: • First—The DNS server queries the forwarders list before performing 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 primary or a forwarder.
- DNSSEC—Whether to use DNSSEC and the DNSSEC algorithm/key size.
- 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 primary for one zone and a forwarder for another zone.

Before you begin:

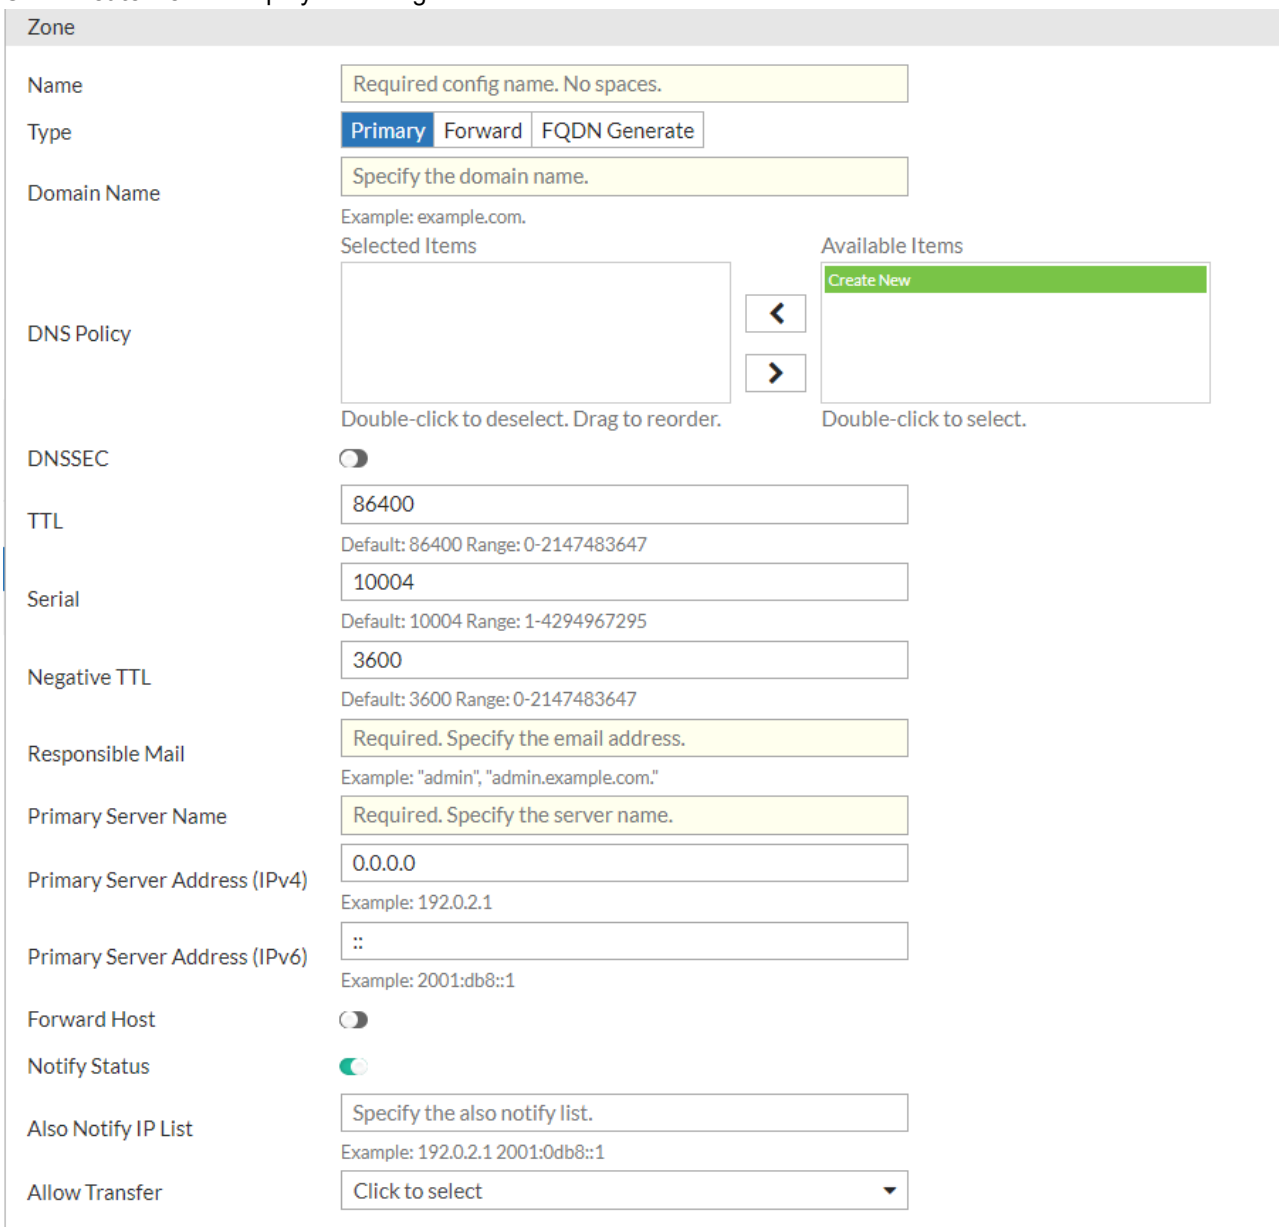
- You must have a good understanding of DNS and knowledge of the DNS deployment in your network.
- You must have authority to create authoritative DNS zone records for your network.
- You must have Read-Write permission for Global Load Balance settings.

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 629.

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	• Primary—The configuration contains the “primary” copy of data for the zone and is the

Settings	Guidelines
	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: `example.com.`
DNS policy	Select one or more DNS policies you want the zone to use.
Primary Options	
DNSSEC	Enable/disable DNSSEC. The Domain Name System Security Extensions (DNSSEC) is a feature of the Domain Name System (DNS) that authenticates responses to domain name lookups. Only when a DNS policy has been set, *and* DNSSEC is enabled, will the **Back Up DSSET Key**, **Regenerate DNSSEC Key**, and **Restore DNSSEC Key** appear. Back Up DSSET Key includes the following types of keys: • **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. Restore DNSSEC Key should be a tar type file.
DNSSEC Algorithm	The **DNSSEC Algorithm** option is available if **DNSSEC** is enabled. Select the cryptographic algorithm to use for authenticating DNSSEC. • RSA-SHA1 • NSEC3-RSA-SHA1 • RSA-SHA256 • RSA-SHA512 • ECDSAP256-SHA256 • ECDSAP384-SHA384
DNSSEC Key Size	The **DNSSEC Key Size** option is available if **DNSSEC** is enabled. Select the key size (number of bits) for the encryption algorithm. • 1024 bits • 2048 bits • 4096 bits **Note**: Prior to FortiADC 7.4.0, the DNSSEC key size only supported 512 bits, so configurations carried over from previous versions can continue using the 512-bit key. However, we recommend updating to the new 1024/2048/4096 bit keys as the 512-bit key is less secure and is no longer supported in the latest BIND 9 version.
DSSET List	The **DSSET List** option is available if **DNSSEC** is enabled. Select a DSSET configuration object. See Configuring the DSSET list.

Settings	Guidelines
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.
Serial	Set the serial number of the zone. Default 10004. Range 1-4294967295.
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 `hostprimary.example.com.`. **Note**: Format is `mailbox-name.domain.com.` (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 hostprimary@example.com.
Primary Server Name	Sets the server name in the SOA record. If the name does not end with a trailing dot, the DNS system will automatically append the root domain, potentially resulting in an incorrect entry. For example, `admin.` is valid, but `admin` will be treated as `admin.<root domain>`.
Primary Server Address	The IP address of the primary server.
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."
Forward Options	
Forward Host	Enable Forward Host to allow DNS queries to be forwarded to remote servers at the zone level. This is disabled by default. This only requires the forwarded DNS query to match the zone and no other information is required to match such as the hostname. It is recommended to verify the DNS forwarding success by sending the DNS response and checking the packet data. For details, see Verify DNS forwarding to remote server on page 633.
Forward	The **Forward** option is available if **Forward Host** is **enabled**. • First—The DNS server queries the forwarder before doing its own DNS lookup. This is the default option. • 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	The **Forwarders** option is available if **Forward Host** is **enabled**. Select a remote server configuration object.

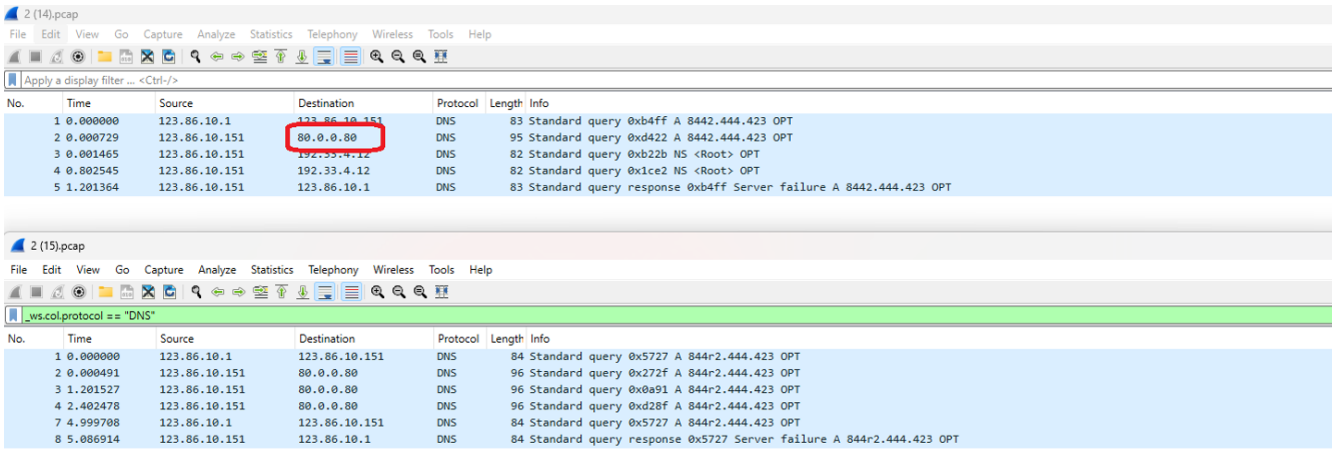
Settings	Guidelines
FQDN Record	
FQDN Record table	Displays a summary of all DNS RR for the zone, including generated and manually configured RR.
A/AAAA Record	
Hostname	The hostname part of the FQDN, such as `www`. **Note**: You can specify the @ symbol to denote the zone root. The value substituted for @ is the preceding $ORIGIN directive.
Type	• IPv4 • IPv6
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, `www.example.com` is an alias for example.com.
Target	The true or canonical domain name. For instance, `example.com`.
NS Record	
Domain Name	The domain for which the name server has authoritative answers, such as `example.com`. Note: FortiADC supports third-party domain names.
Hostname	The hostname part of the FQDN, such as `ns`.
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 `mail`.
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.

Settings	Guidelines
	An example SPF record has the following form: `v=spf1 +mx a:colo.example.com/28 -all` If you complete the entry from 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., www.
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 10.168.192.in-addr.arpa. or 1
FQDN	A fully qualified domain name, such as "www.example.com".
CAA Record	
Hostname	The hostname of CAA record
Value	Specify the value
Flag	Range 0-255. Default is 0.
Tag	Issue/Issuewild/Iodef

Verify DNS forwarding to remote server

To verify if the DNS queries are being successfully forwarded to the remote server at the zone level, you can send a test query and check its packet data. Using tools such as "nslookup" or "dig", you can send a DNS query, and then collect and view its packet data using PCAP tools such as "wireshark".

The example below, we will analyze the packet data to verify whether the DNS query was sent to the forwarder.



From this packet data, you can extract the following information:

1. Client 123.86.10.1 sent the DNS query (8442.444.423 A record) to 123.86.10.151.
2. Request for GLB zone domain matches record; no match for hostname.
3. The DNS query was sent to the forwarder (80.0.0.80).

In this example, it can be confirmed that the DNS query was successfully sent to the forwarder after the DNS server confirms the GLB zone domain match without requiring the hostname to match as well.

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.

From general settings, you can also enable DNS over HTTP/HTTPS (DoH) and DNS over TLS (DoT) to encrypt the DNS query.

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.
- If enabling DNS over HTTPS/TLS, you must have prepared a dedicated DNS server domain and a certificate pair for your DNS over HTTPS/TLS service. For details, see Configuring DNS over HTTPS and DNS over TLS on page 636.

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 635.
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. **Note**: DNS forwarding functionality requires **Recursion** to be **enabled**.
DNSSEC Validation	Enables/disables DNSSEC validation.
Minimal Responses	Enables/disables Minimal Responses to hide the **Authority Section** and **Additional Section** of DNS queries. When the DNS query only shows the minimal response, it can significantly increase the performance of the FortiADC DNS service by increasing the QPS.
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.
Interface List	The **Interface List** option is available if **Listen on All Interface** is disabled. If not listening to all interfaces, select one or more ports to listen on.
DNS over HTTPS	Enables/disables DNS over HTTPS to encrypt DNS queries using the HTTPS protocol.
DNS over HTTPS Port	The **DNS over HTTPS Port** option is available if **DNS over HTTPS** is enabled. Specify the port to listen on DNS over HTTPS. Default: 443 Range: 1-65535.
DNS over HTTPS Interface List	The **DNS over HTTPS Interface List** option is available if **DNS over HTTPS** is enabled. Select the interface(s) to listen on for DNS over HTTPS.
DNS over HTTP	Enables/disables DNS over HTTP to encrypt DNS queries using the HTTP protocol.
DNS over HTTP Port	The **DNS over HTTP Port** option is available if **DNS over HTTP** is enabled. Specify the port to listen on DNS over HTTP. Default: 80 Range: 1-65535.
DNS over HTTP Interface List	The **DNS over HTTP Interface List** option is available if **DNS over HTTP** is enabled. Select the interface(s) to listen on for DNS over HTTP.
DNS over TLS	Enables/disables DNS over TLS to encrypt DNS queries using the TLS protocol.
DNS over TLS Port	The **DNS over TLS Port** option is available if **DNS over TLS** is enabled. Specify the port to listen on DNS over TLS. Default: 853 Range: 1-65535.
DNS over TLS Interface List	The **DNS over TLS Interface List** option is available if **DNS over TLS** is enabled. Select the interface(s) to listen on for DNS queries for DNS over TLS.
Certificate	The **Certificate** option is available if **DNS over HTTPS** or **DNS over TLS** is enabled.

Settings	Guidelines
	Select the certificate object to apply for DNS over HTTPS or DNS over TLS. This certificate must refer to the DNS server domain or IP address. For details, see Configuring DNS over HTTPS and DNS over TLS on page 636.
Forward	DNS forwarding functionality requires **Recursion** to be **enabled**. Select from the following DNS forwarding options: • 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 DNS over HTTPS and DNS over TLS

DNS over HTTPS (DoH) and DNS over TLS (DoT) are protocols used to encrypt communications with DNS resolvers. DoH encrypts the DNS traffic by passing DNS queries through an HTTPS encrypted session. Whereas DoT adds TLS encryption on top of the UDP that is used for DNS queries.

The primary difference between the DoH and DoT standards is what port they use. DoT only uses port 853, whereas DoH uses port 443 (which is the port that all other HTTPS traffic uses as well).

To configure DNS over HTTPS or DNS over TLS on FortiADC, follow the basic steps below:

Step 1: Prepare a full domain name or an IP address for your DoH/DoT service on page 636

Step 2: Prepare a certificate pair for the DoH/DoT service on page 637

Step 3: Enable DoH/DoT service on FortiADC on page 637

Step 4: Enable DoH/DoT service on your browser or local application on page 638

Step 1: Prepare a full domain name or an IP address for your DoH/DoT service

Before you can configure DoH/DoT on FortiADC, you must first prepare a full domain name or an IP address for the DoH/DoT service. This will then be used for your custom DoH/DoT server URL and to sign the certificate.

You can prepare the full domain using either of the following methods:

- If your organization manages its own public domain, you can add a new record to the domain.
  a. Login to you DNS service provider and go to your DNS Domain management page.
  b. Add a new record to the existing public domain as `dns.yourdomain.com`. The IP address is the DNS-over-HTTPS or DNS-over-TLS service public IP.
- If you want to test your own domain, you can add the full domain name resolution to the local hosts file.
  For example: The Ubuntu local hosts file would be located at `/etc/hosts`, and in Windows it would be at `c:\Windows\System32\Drivers\etc\hosts`.

Step 2: Prepare a certificate pair for the DoH/DoT service

Use the full domain name or IP address for the DoH/DoT service previously prepared to create a certificate pair.

You can prepare the certificate pair for your DoH/DoT service using either of the following methods:

- Apply for a public certificate from a public CA with your full domain name or IP address.
- Generate a self-signed certificate. **Note**: A self-signed certificate cannot be generated through FortiADC.
  For example: Generating a self-signed certificate in Ubuntu with OpenSSL
  a. Prepare the prerequisites:

```
mkdir demoCA
mkdir demoCA/newcerts
echo 01 > demoCA/serial
touch demoCA/index.txt
```

  b. Add the following lines to the file `/usr/lib/ssl/openssl.cnf` under the `[v3_req]` section.

```
subjectAltName = @alt_names
[alt_names]
DNS.1=dns.yourdomain.com
IP.1=yourdomain IP
```

  c. Generate the root CA key.

```
openssl genrsa -out rootca.key 2048
```

  d. Generate the root CA cert.

```
openssl req -new -x509 -days 3650 -key rootca.key -out rootca.crt -subj
"/C=Country/ST=State/L=Location/O=Company/OU=Department/CN=yourdomain.com/emailAddre
ss=admin@yourdomain.com"
```

  e. Generate the DNS server private key.

```
openssl genrsa -out dns-doh.key 2048
```

  f. Generate the DNS server cert (enter `y` if prompted).

```
openssl req -new -key dns-doh.key -out dns-doh.csr -subj
"/C=Country/ST=State/L=Location/O=Company/OU=Department/CN=yourdomain.com/emailAddre
ss=admin@yourdomain.com"

openssl ca -in dns-doh.csr -out dns-doh.crt -cert rootca.crt -keyfile rootca.key -
days 365 -extensions v3_req
```

  g. Verify the certificate.

```
openssl verify -CAfile rootca.crt dns-doh.crt
```

Step 3: Enable DoH/DoT service on FortiADC

After preparing the certificate pair for the DoH/DoT service, you can import the certificate and then enable the DoH/DoT function in FortiADC.

1. Go to **System > Manage Certificates**.
2. Click the **Local Certificate** tab.
3. Click **Import** to upload the prepared certificate in FortiADC.
4. Go to **Global Load Balance > Zone Tools**.

5. Click the **General Settings** tab.
6. Configure the following relevant settings to enable DNS over HTTPS service and save the configuration:

Setting	Guidelines
Global DNS Configuration	Enable Global DNS Configuration.
DNS over HTTPS	Enable DNS over HTTPS service.
DNS over HTTPS port	Default port is 443. Change the HTTPS service port number if it is not the default HTTPS service port or if there is an IP/port conflict.
DNS over HTTPS Interface List	Select the interfaces that allow the DNS over HTTPS service.
Certificate	Select the matching certificate.

   a. Test your DNS over HTTPS configuration.
   For example, you can use a DNS lookup tool such as Dig from a remote system.
   `dig @yourdomain.com example.com +https`
   If the DNS over HTTPS is successfully configured, you should get the IP address of example.com: 93.184.216.34.
   From Dig's output, you should also see the following:
   `;; SERVER: 18.217.127.135#443(dns.yourdomain.com) (HTTPS)`
   This confirms that the query/response operation was performed successfully over HTTPS (TCP port 443) rather than the traditional UDP port 53.
7. Configure the following relevant settings to enable DNS over TLS service and save the configuration:

Setting	Guidelines
Global DNS Configuration	Enable Global DNS Configuration.
DNS over TLS	Enable DNS over TLS service.
DNS over TLS port	Default port is 853.
DNS over TLS Interface List	Select the interfaces that allow the DNS over TLS service.
Certificate	Select the matching certificate.

   a. Test your DNS over TLS configuration.
   For example, you can use a DNS lookup tool such as Dig from a remote system.
   `dig @yourdomain.com example.com +tls`
   If the DNS over HTTPS is successfully configured, you should get the IP address of example.com: 93.184.216.34.
   From Dig's output, you should also see the following:
   `;; SERVER: 18.217.127.135#853(dns.yourdomain.com) (TLS)`
   This confirms that the query/response operation was performed successfully over TLS (TCP port 853) rather than the traditional UDP port 53.

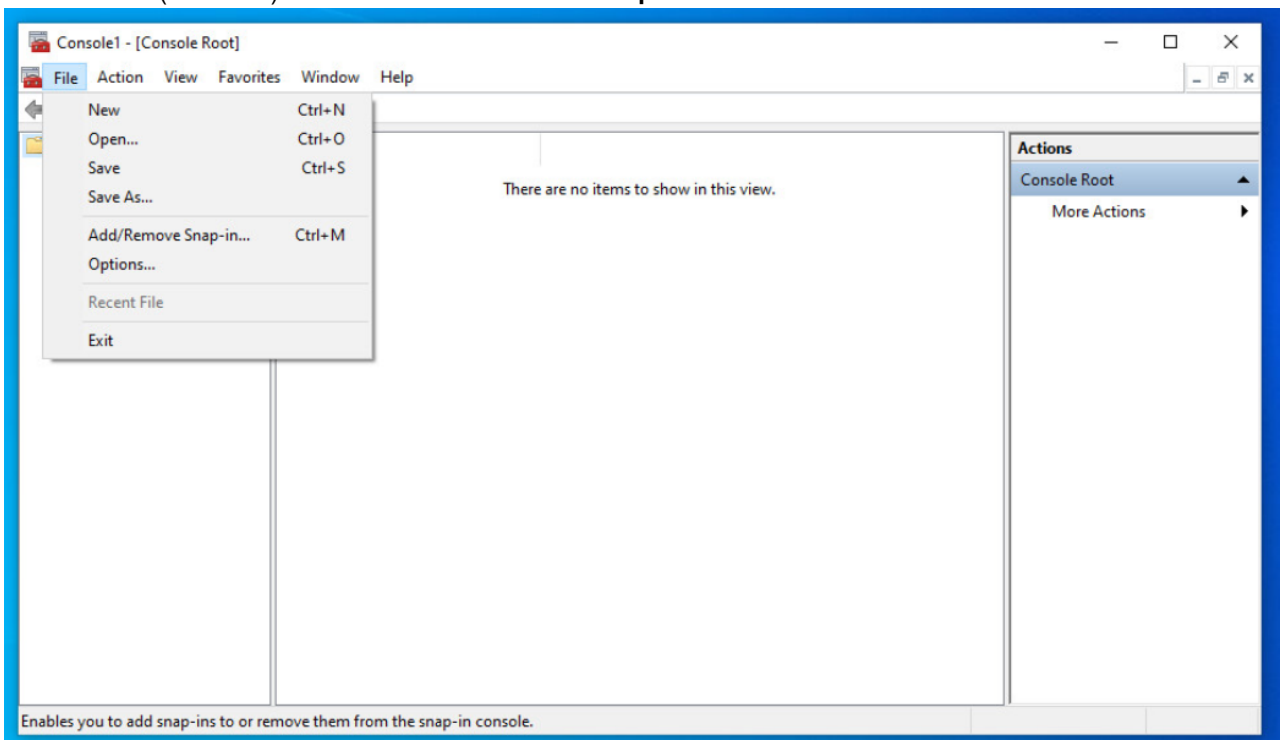
Step 4: Enable DoH/DoT service on your browser or local application

After you have configured DoH/DoT on FortiADC, enable DoH/DoT service on your browser or local application. However, if you have used a self-signed certificate pair, you need to first import that certificate into the local system before enabling DoH/DoT on the browser or local application.

Importing the self-signed certificate pair to the local system for DoH/DoT

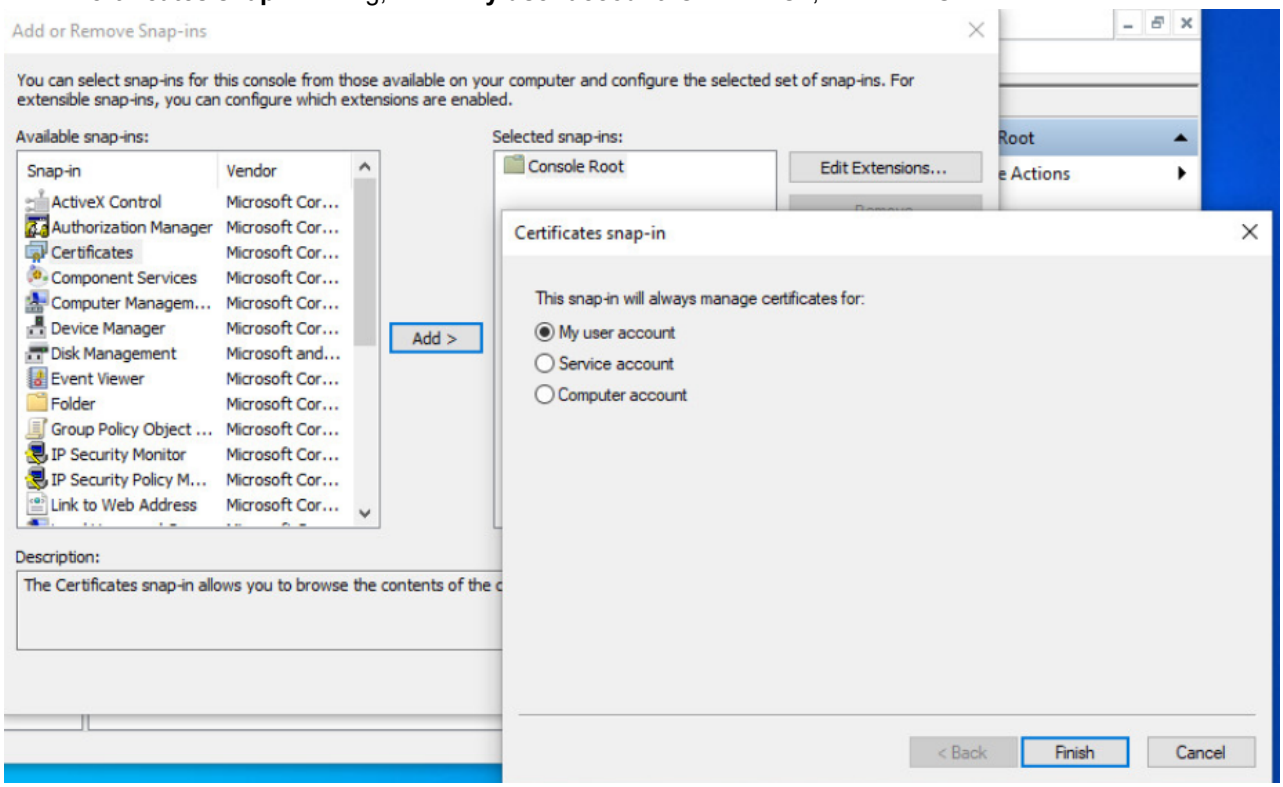
Example 1: Importing the self-signed certificate to Windows

1. Launch MMC (mmc.exe). Go to **File > Add/Remove Snap-ins**.

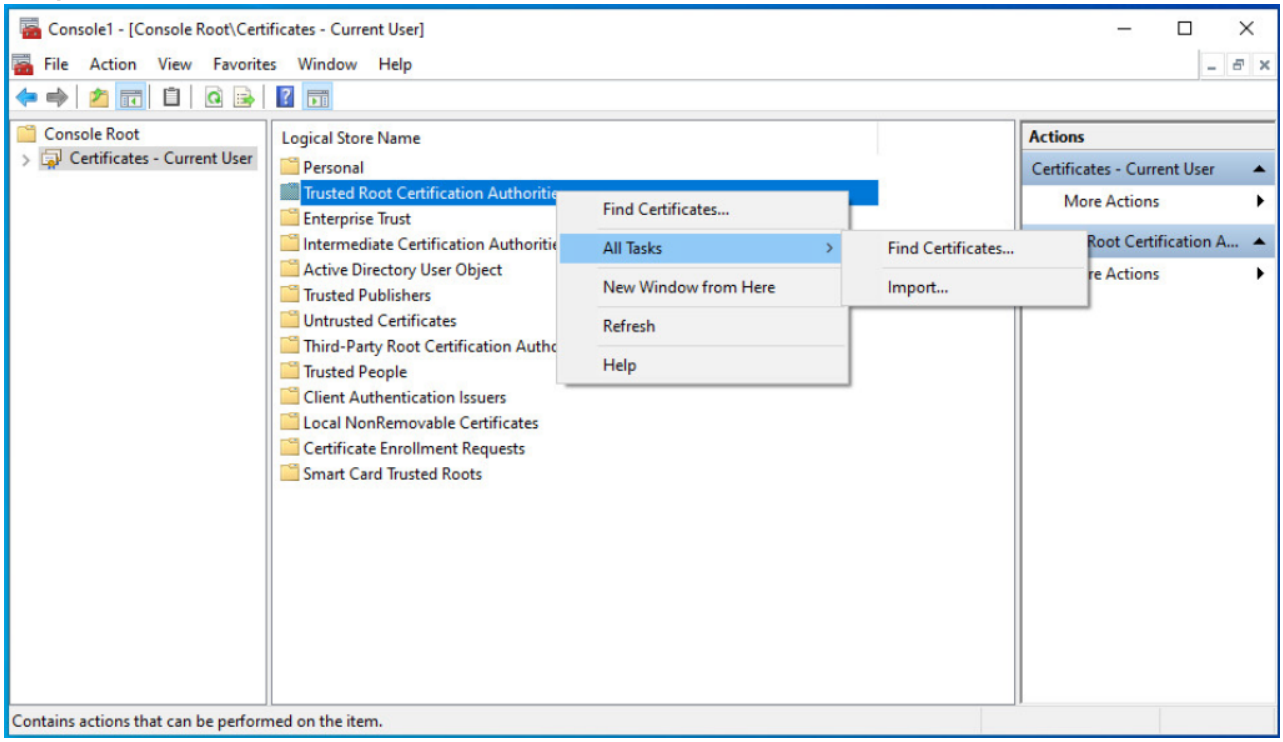


2. Select **Certificates**, then click **Add**.
   The Certificates snap-in dialog displays.

3. In the **Certificates snap-in** dialog, select **My user account**. Click **Finish**, then click **OK**.

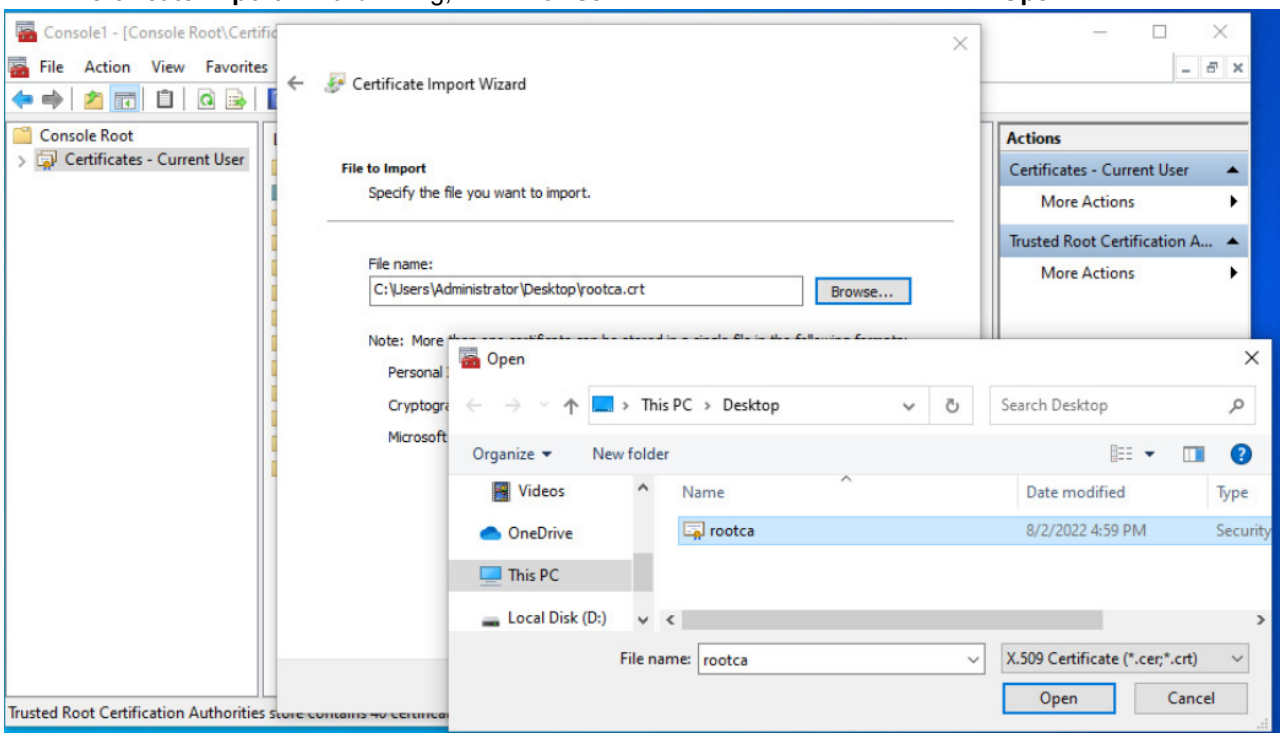


4. Select **Certificates - Current User**, right-click on **Trusted Root Certification Authorities**. Then select **All Tasks > Import**.



The Certificate Import Wizard dialog displays.

5. In the **Certificate Import Wizard** dialog, click **Browse** to select the rootca.crt file and click **Open**.



6. Select **Place all certificates in the following store** and set **Trusted Root Certification Authorities** as the **Certificate store**. Click **Next** and **Next** again, then click **Finish**.
7. Click **Yes** and **Finish** for the prompt windows.

Example 2: Importing the self-signed certificate to Ubuntu

1. Install or update ca-certificates.

```
$ sudo apt-get install -y ca-certificates
```

2. Copy your certificate in PEM format (the format that has `----BEGIN CERTIFICATE----` in it) into `/usr/local/share/ca-certificates` and name it with a .crt file extension.

```
$ sudo cp rootca.crt /usr/local/share/ca-certificates
```

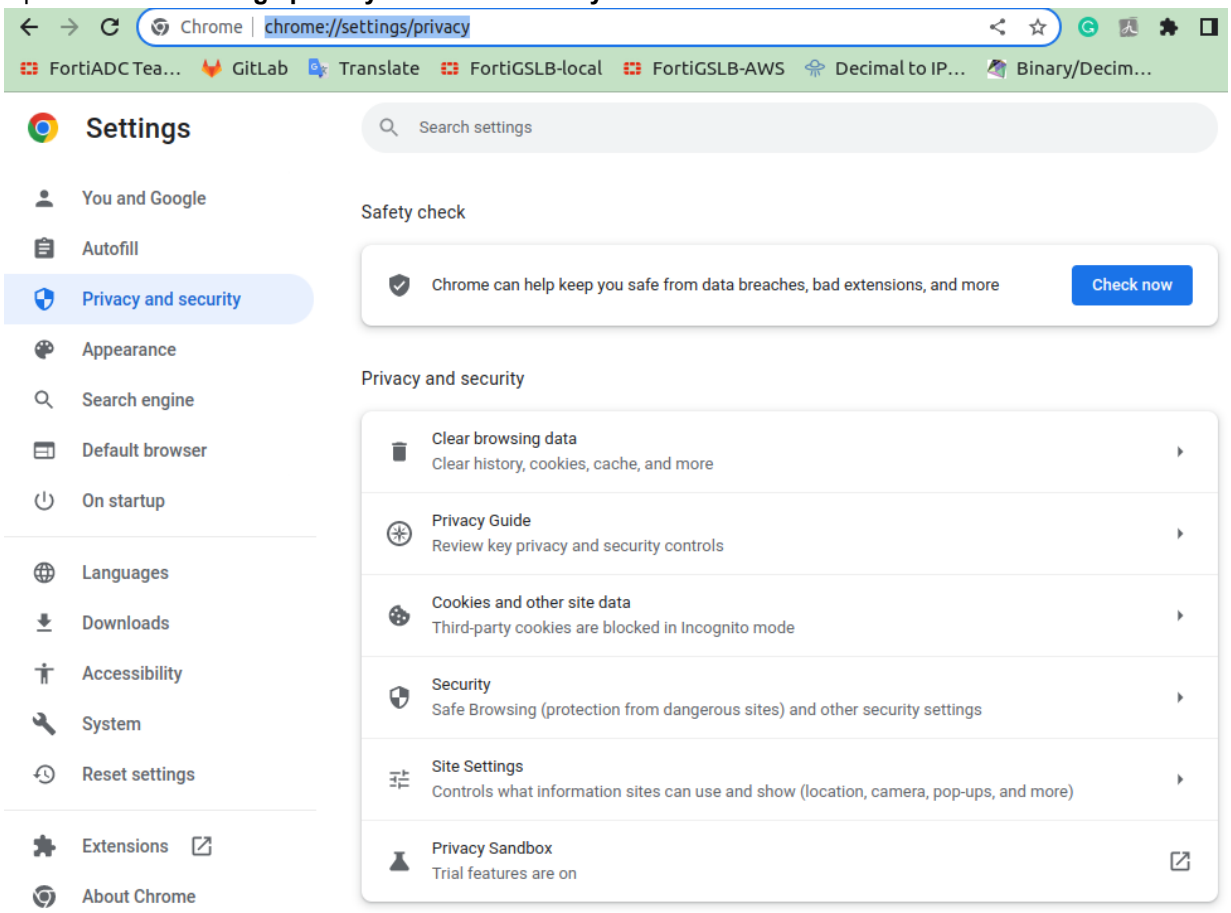
3. Update your ca chain list.

```
$ sudo update-ca-certificates
```

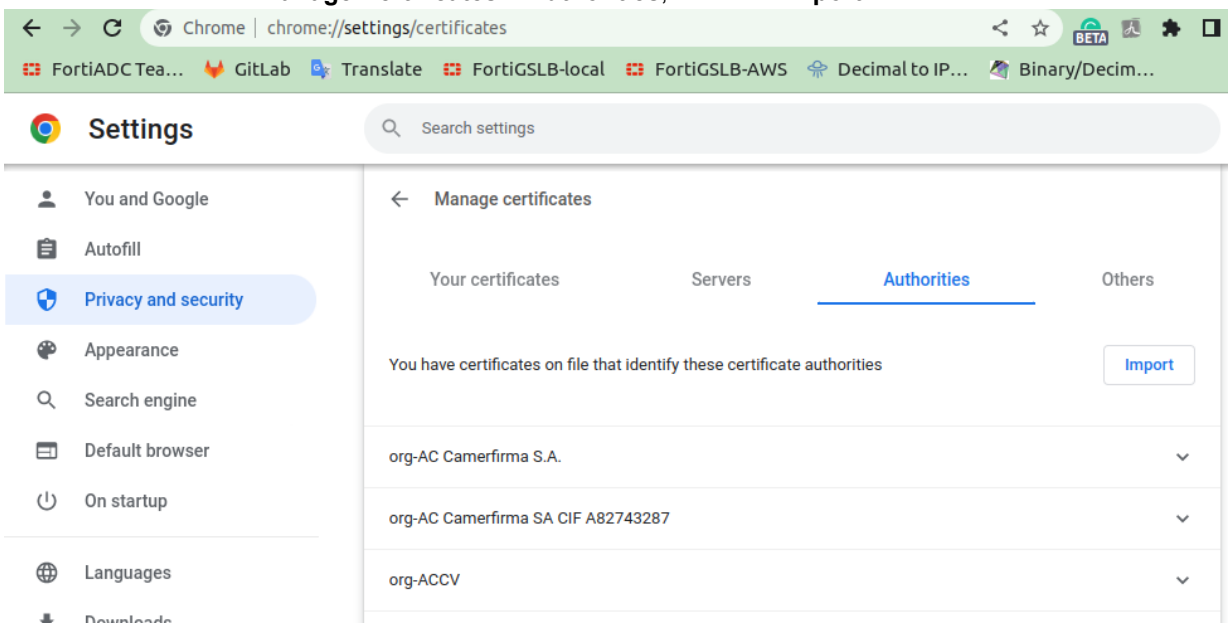
Enabling DNS over HTTPS on the browser

Example 1: Enable DoH in Chrome (version 105.0.5195.102)

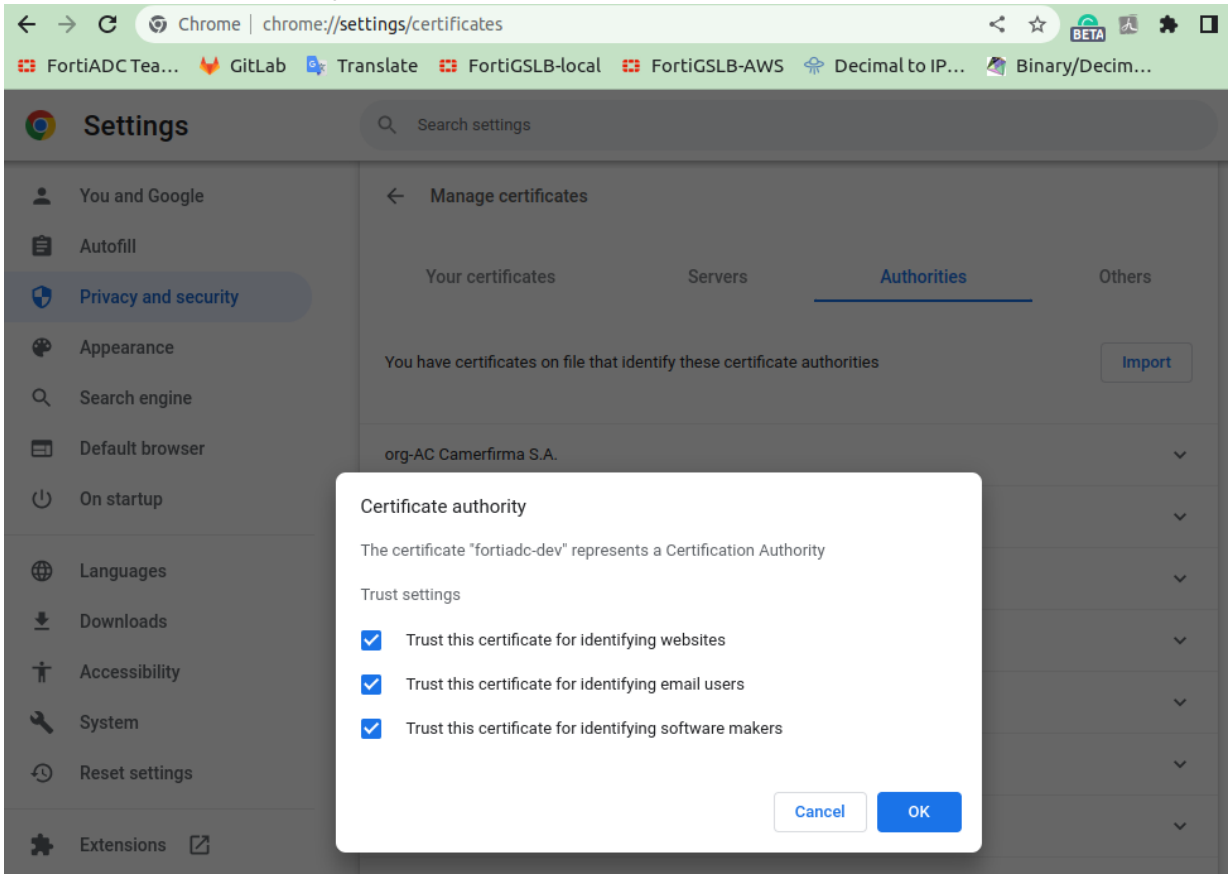
1. Import self-signed certificate to Chrome. You may skip this step if you use a public certificate.
   a. Open **chrome://settings/privacy** and select **Security**.



b. Scroll down and select **Manage Certificates > Authorities**, then click **Import**.

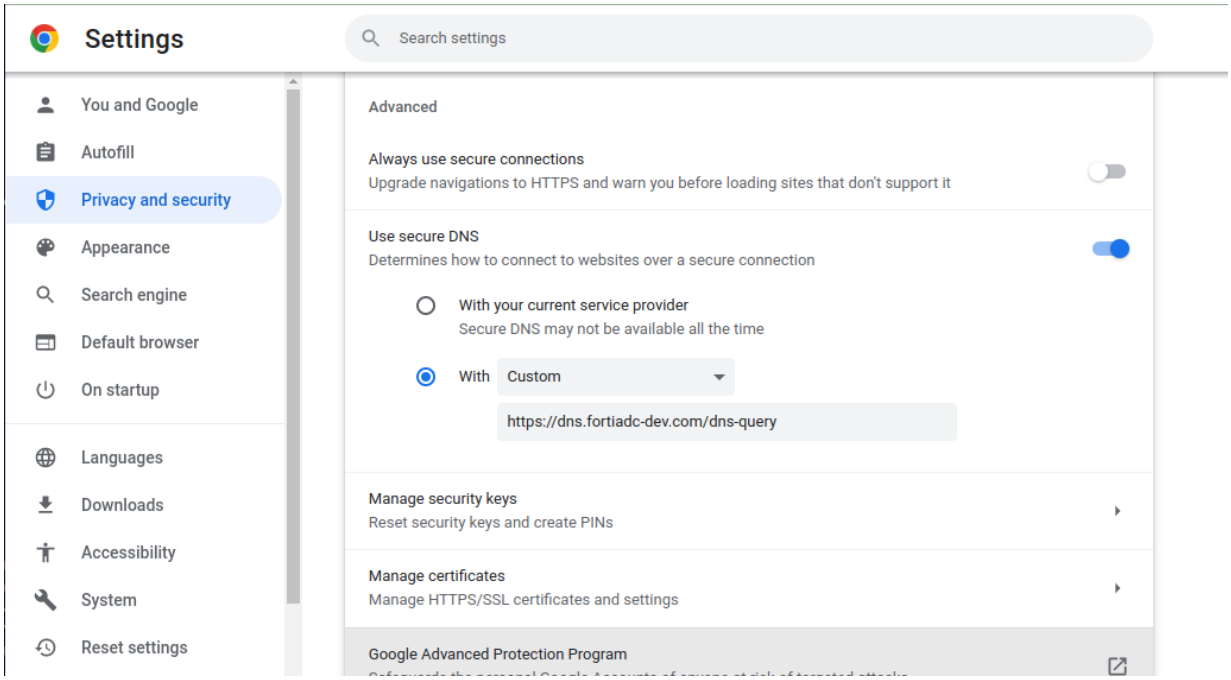


c. Select the file rootca.crt and open. Check all the boxes and click **OK**.



2. Enable custom DNS over HTTPS settings.
   a. Go to **Settings > Privacy and security > Security**.
   b. Enable **Use secure DNS**, select **With**, then select **Custom** from the drop-down list.

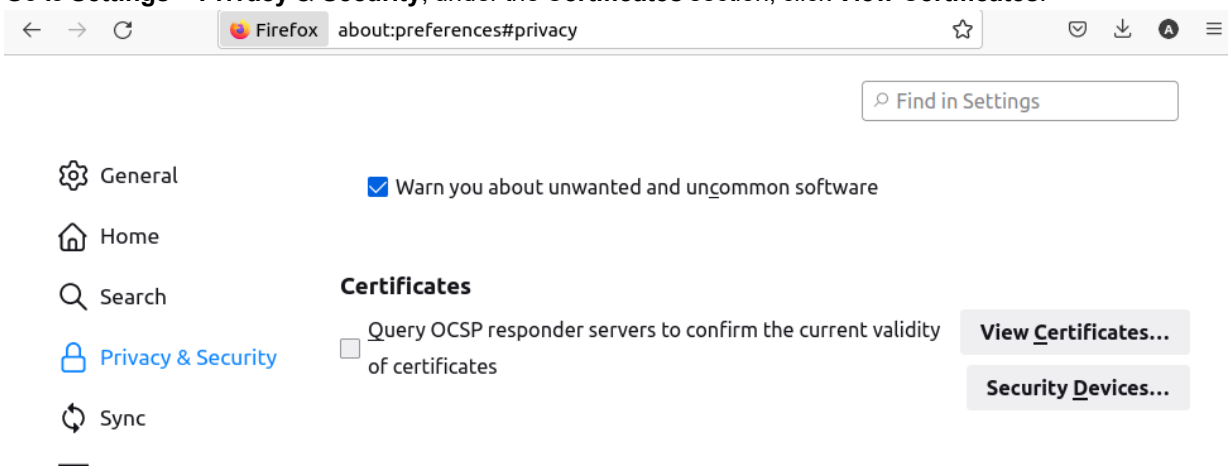
c. Input the URL that matches your FortiADC DNS server domain and your certificate as `https://yourdomain/dns-query` or `https://yourIP:port/dns-query`.



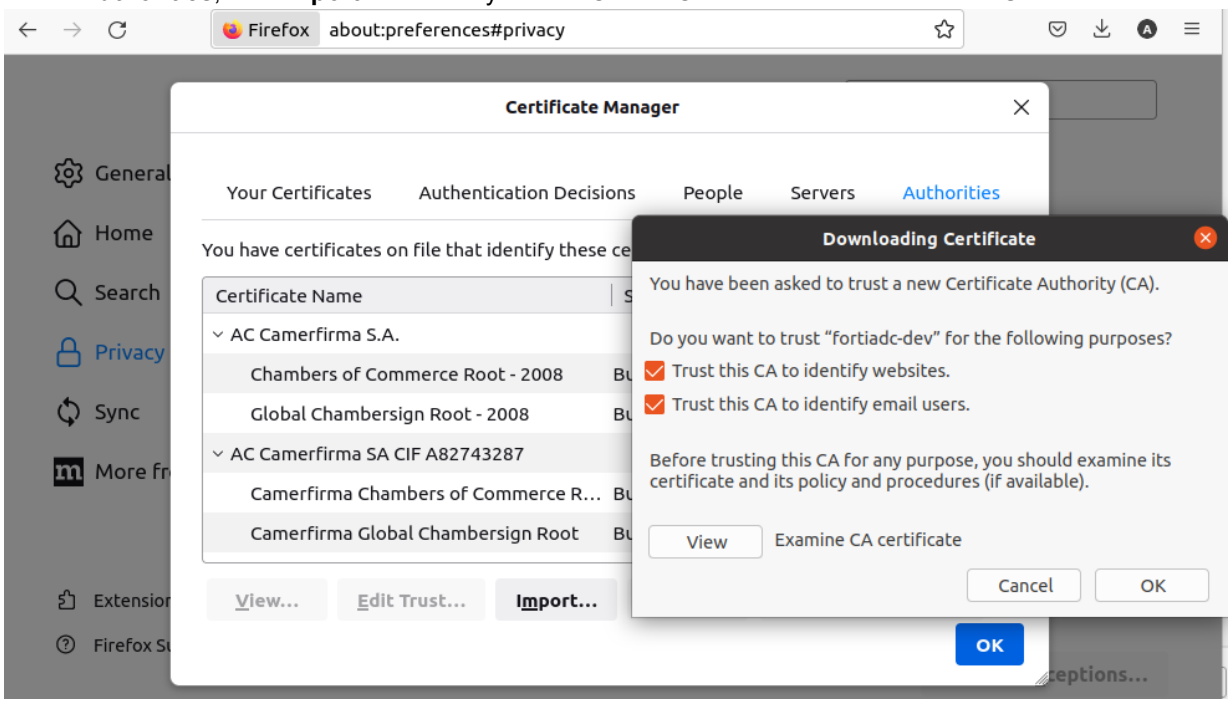
3. Debugging.
   Try the following URL and see if there is a file downloading.
   `https://yourdomain/dns-query?dns=q80BAAABAAAAAAAAA3d3dwdleGFtcGxlA2NvbQAAAQAB`
   If downloaded, then all configurations are correct.
   Otherwise, the most common error is `ERR_CERT_COMMON_NAME_INVALID`, which means your server cert CN is invalid or does not match your FortiADC server.

Example 2: Enable DoH in Firefox (version 104.0.2)

1. Import self-signed certificate to Firefox. You may skip this step if you use a public certificate.
   a. Go to **Settings > Privacy & Security**, under the **Certificates** section, click **View Certificates**.

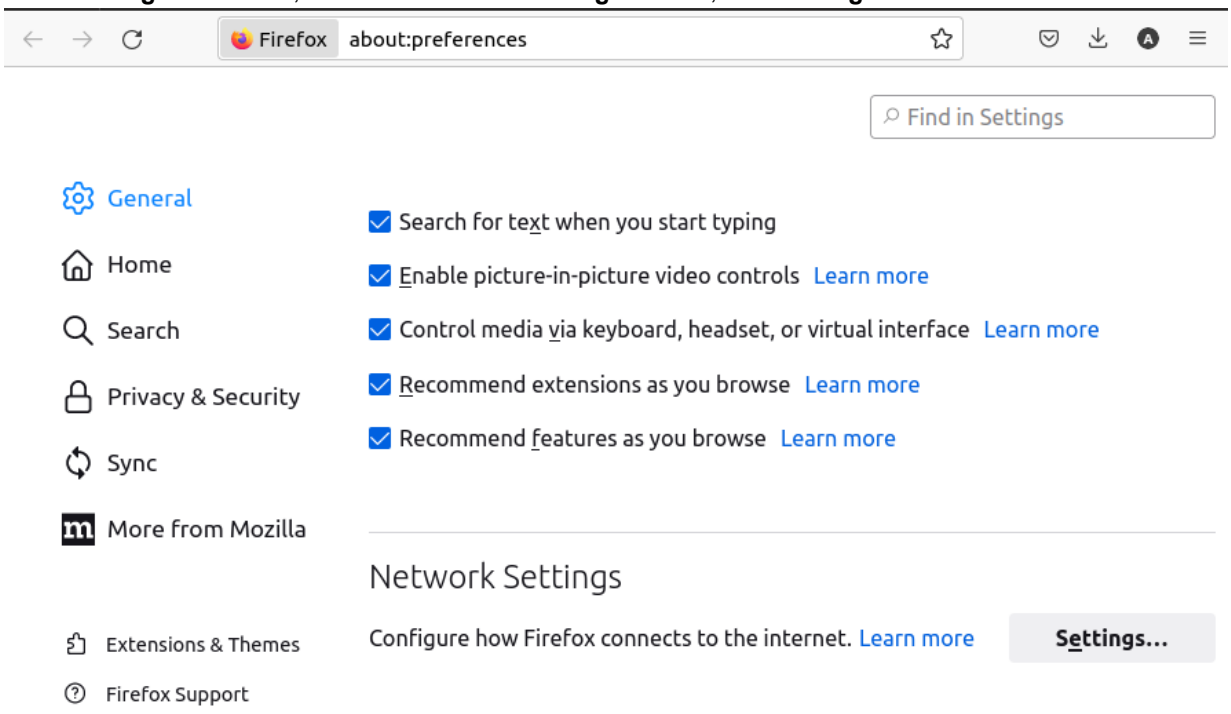


b. Select **Authorities**, click **Import** and select your root CA file. Check all the boxes and click **OK** to save.



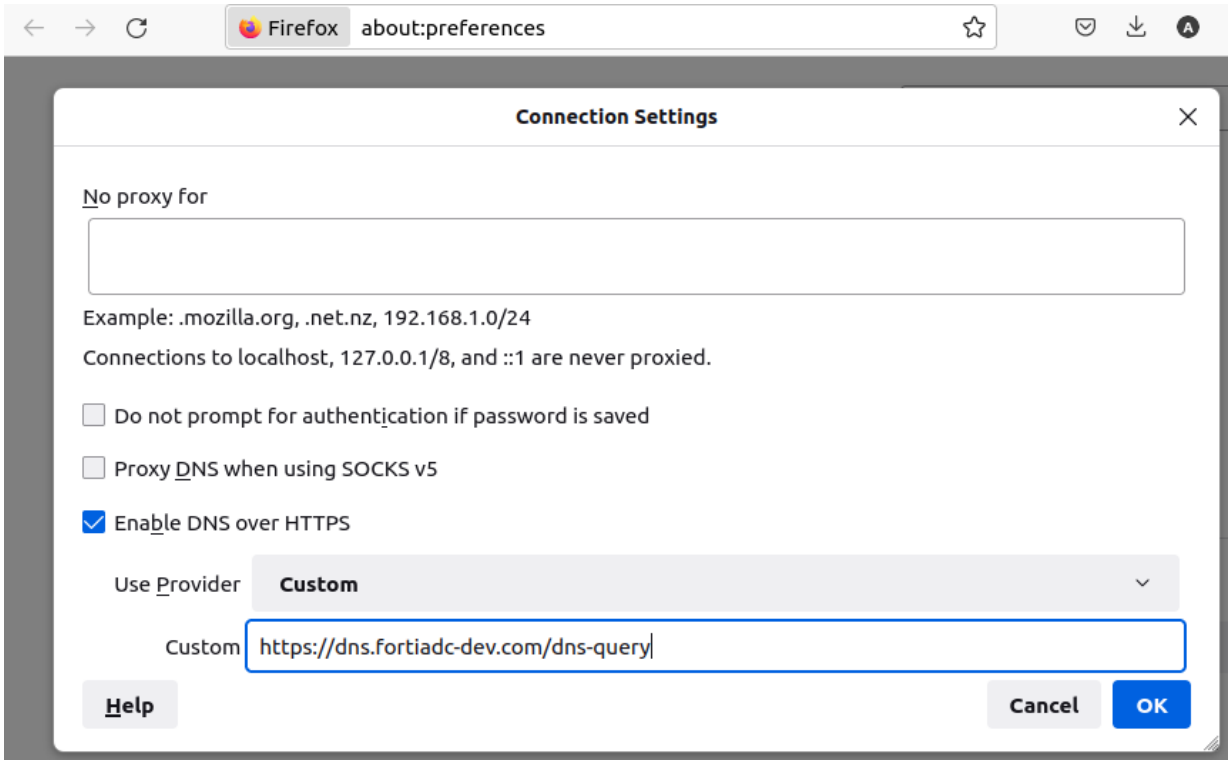
2. Enable custom DNS over HTTPS settings.

a. Go to **Settings > General**, under the **Network Settings** section, click **Settings**.



b. Select **Enable DNS over HTTPS**. In the **Use Provider** field, select **Custom** from the drop-down list.

c. Input the URL that matches your FortiADC DNS server domain and your certificate as `https://yourdomain/dns-query` or `https://yourIP:port/dns-query`. Click **OK** to save.



3. Debugging.
   Try the following URL and see if there is a file downloading.
   `https://yourdomain/dns-query?dns=q80BAAABAAAAAAAAA3d3dwdleGFtcGxlA2NvbQAAAQAB`
   If downloaded, then all configurations are correct.
   Otherwise, the most common error is `ERR_CERT_COMMON_NAME_INVALID`, which means your server cert CN is invalid or does not match your FortiADC server.

Enabling DNS over TLS on your local application

Example 1: Enable DoT in Ubuntu with systemd

1. Check the systemd version.
   systemd start to support strict DNS over TLS mode from version 243. Use the following command to check your version and update it if the version is too old.

   ```
   $ systemd --v
   systemd 249 (249.11-0ubuntu3.6)
   ```

2. Set up the systemd configuration.
   Modify `/etc/systemd/resolved.conf` so that it is similar to what is shown below. Be sure to enable DNS over TLS and to configure the IP addresses of the DNS servers you want to use.

   ```
   $ cat /etc/systemd/resolved.conf
   DNS=10.106.210.81
   FallbackDNS=8.8.8.8
   #Domains=
   #DNSSEC=no
   DNSOverTLS=yes
   ```

```
#LLMNR=yes
#MulticastDNS=yes
#Cache=yes
#DNSStubListener=yes
#ReadEtcHosts=yes
```

3. Restart services.
   To make the settings configured in the previous steps take effect, restart `systemd-resolved`.

```
$ sudo systemctl restart systemd-resolved
```

4. Check that everything is running correctly.

```
$ resolvectl status
Global
           Protocols: -LLMNR -mDNS +DNSOverTLS -DNSSEC
    resolv.conf mode: foreign
  Current DNS Server: 10.106.210.81
         DNS Servers: 10.106.210.81
Fallback DNS Servers: 8.8.8.8
```

5. Verify the configuration.
   Use the following command to perform a DNS query.

```
$ sudo resolvectl flush-caches
$ resolvectl query google.com
google.com: 142.250.72.206                    -- link: ens160
            2607:f8b0:4005:801::200e          -- link: ens160

-- Information acquired via protocol DNS in 23.8ms.
-- Data is authenticated: no; Data was acquired via local or encrypted transport: yes
-- Data from: network
```

   Use `tcpdump` to capture the traffic in another terminal. You will find that the traffic goes to DNS server 10.106.210.81 port 853 instead of the regular DNS service port 53.

Example 2: Enable DoT in Ubuntu with Unbound

1. Install Unbound.

```
# sudo apt install -y unbound
# sudo systemctl enable unbound
```

2. Edit the configuration file.

```
# cat /etc/unbound/unbound.conf.d/pihole.conf
server:
    port: 53
    tls-upstream: yes
    tls-cert-bundle: "/etc/ssl/certs/ca-certificates.crt"

forward-zone:
    name: "."
    forward-addr: 10.106.210.81@853

# unbound-checkconf
unbound-checkconf: no errors in /etc/unbound/unbound.conf
```

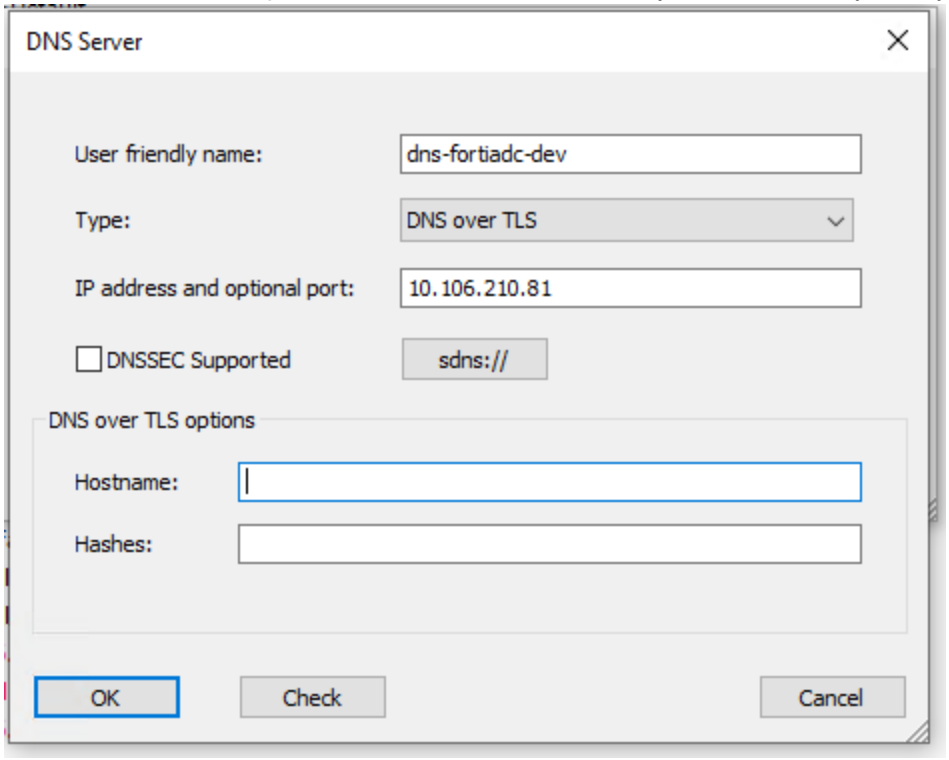
3. Restart services.

```
# sudo systemctl restart unbound
```

4. Verify the configuration.
Perform a simple Dig test and use `tcpdump` to capture the traffic in another terminal. You will find that all the DNS traffic goes to Server 10.106.210.81 port 853 instead the regular DNS service port 53.

Example 3: Enable DoT in Windows with YogaDNS

1. Go to https://dns.sb/guide/dot/windows/
2. Set the **IP address and optional port** to match your FortiADC configuration.
3. Under **DNS over TLS** options, set the **Hostname** to match your certificate or you may leave it blank. Click **OK**.



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 649.
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:
	`\"<domainname>\" <num1> <num2> <num3> \"<content>\"`
	The following is an example:
	`\".\" 256 3 5 \"AwEAAbDrWmiIReotvZ6FObgKygZwUxSUJW9z5pjiQMLH0JBGXooHrR16 pdKhI9mNkM8bLUMtwYfgeUOYXIvfagee8rk=\"`

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. For details, 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 650.

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 651.

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: `dns.example.com. IN DS 13447 5 1 A5AD9EFB6840F58CF817F3CC7C24A7ED2DD5559C`

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 651

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.
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.

Settings	Guidelines
	• 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:

1. Go to **Global Load Balance > Zone Tools**.
2. Click the **Remote DNS Server** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration and add members as described in Remote DNS server configuration on page 652.

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
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 653.

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.

Global Object

The **Global Load Balance > Global Object** sub-menu allows you to perform the following tasks relating to the global object for global load balancing:

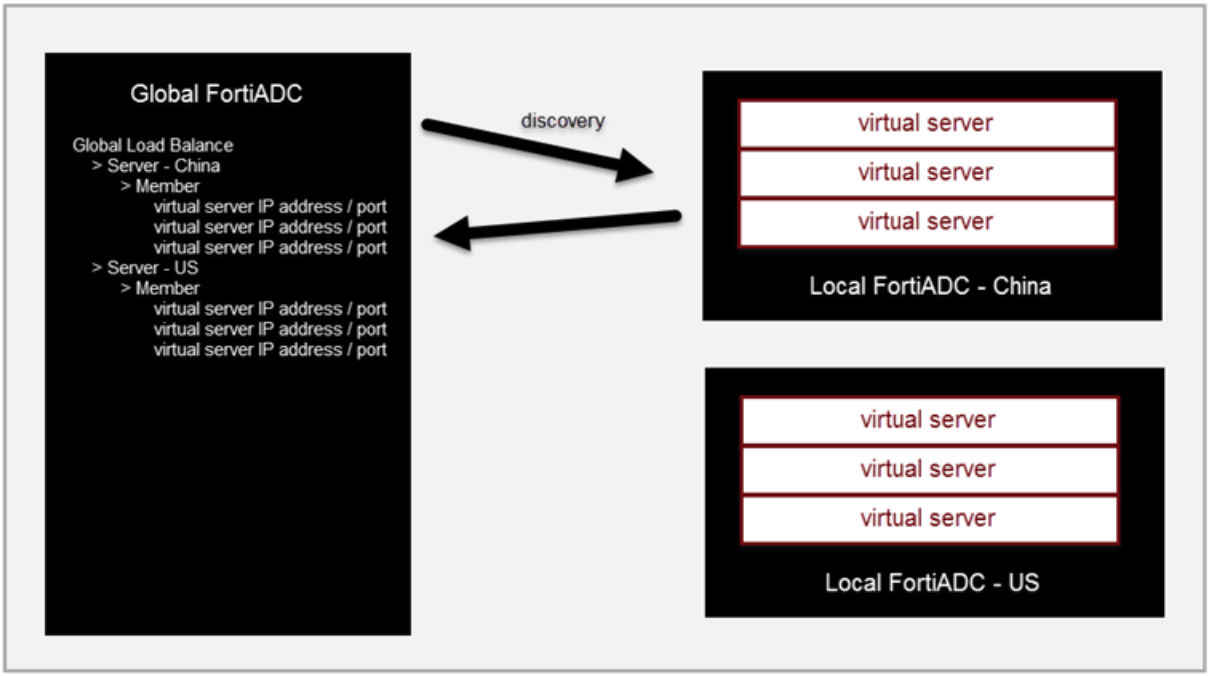
- Configuring servers on page 653
- Configuring a global load balance link on page 658
- Configuring data centers on page 658

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 653 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 created an SDN connector configuration.
  Note: Currently, the SDN Connector option only supports AWS Connectors
- You must have Read-Write permission for Global Load Balance settings.

After you have created a server configuration object, you can specify it in 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. Configure the following GLB Server settings:

Setting	Description
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. **Note**: After you initially save the configuration, you cannot edit the name.
Type	Select the remote server to use for global server load balancing: • FortiADC SLB — use a FortiADC instance.

Setting	Description
	• Generic Host — use a third party ADC or server. • SDN Connector — use an existing external connector that is connected to the FortiADC Security Fabric. **Note**: Currently, the SDN Connector option only supports AWS Connectors.
Auth Type	The **Auth Type** option is available if **Type** is **FortiADC SLB**. Select the authentication type: • None — No password. • TCP MD5SIG — With password, but cannot 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	The **Password** option is available if **Type** is **FortiADC SLB** and **Auth Type** is **TCP MD5SIG**. Enter the password to authenticate the key. The password you enter here must match the password configured on the FortiADC appliance in a global sever load-balancing configuration.
User Defined Certificate	The **User Defined Certificate** option is available if **Type** is **FortiADC SLB**. Enable to use a self-defined certificate for authentication.
Certificate	The **Certificate** option is available if **Type** is **FortiADC SLB** and **User Defined Certificate** is **enabled**. Select the local certificate object to use for the GSLB server.
Address Type	The **Address Type** option is available if **Type** is **FortiADC SLB**. IPv4 or IPv6.
IP Address/IPv6 Address	The **IP Address** or **IPv6 Address** option is available if **Type** is **FortiADC SLB**. Specify the IPv4 or IPv6 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 responses.
Port	The **Port** option is available if **Type** is **FortiADC SLB**. Specify the port. Default: 5858 Range: 1-65535.
SDN Connector	The **SDN Connector** option is available if **Type** is **SDN Connector**. Select the SDN Connector to synchronize to the GSLB server. For public SDN type servers, GSLB can update the public IP dynamically. **Note**: Currently, only AWS connectors are supported.
Use SDN Private IP	The **Use SDN Private IP** option is available if **Type** is **SDN Connector**.

Setting	Description
	Enable to use the SDN Private IP address.
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 automatic synchronization with the remote server. When enabled, Global load balancing will synchronize automatically with the server member. If auto-sync is enabled for SDN Connector type servers, all instances from the SDN connector will be added as server members. **Note**: When disabling auto-sync, the server member will be cleared and re-synced.
Health Check Control	The **Health Check Control** option is available if **Type** is **Generic Host** or **SDN Connector**. 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**: Health checking is built-in, and you can optionally configure a gateway health check.
Health Check Relationship	The **Health Check Relationship** option is available if **Type** is **Generic Host** and **Health Check Control** is **enabled**. • 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	The **Health Check List** option is available if **Type** is **Generic Host** and **Health Check Control** is **enabled**. Select one or more health check configuration objects.

5. Click **Save**.
   After the GLB Server configuration is saved, the Member section becomes available to configure.
6. Under the Member section, configure the Member list configuration according to the GLB server type.

Setting	Description
FortiADC SLB	
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.
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).

Setting	Description
	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/IPv6 Address	Virtual server IPv4 or IPv6 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.
Generic Host	
Health Check Inherit	Enable to inherit the health check settings from the parent configuration. The Health Check Inherit option is enabled by default. Disable to specify health check settings in this member configuration.
Health Check Control	The **Health Check Control** option is available if **Health Check Inherit** is **disabled**. Enable health checking for the virtual server.
Health Check Relationship	The **Health Check Relationship** option is available if **Health Check Inherit** is **disabled** and **Health Check Control** is **enabled**. • 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	The **Health Check List** option is available if **Health Check Inherit** is **disabled** and **Health Check Control** is **enabled** Specify one or more health check configuration objects.
SDN Connector	
SDN Instance	Select an instance from the SDN's instance list.
Health Check Inherit	Enable to inherit the health check settings from the parent configuration. The Health Check Inherit option is enabled by default. Disable to specify health check settings in this member configuration.
Health Check Control	The **Health Check Control** option is available if **Health Check Inherit** is **disabled**. Enable health checking for the virtual server.
Health Check Relationship	The **Health Check Relationship** option is available if **Health Check Inherit** is **disabled** and **Health Check Control** is **enabled**. • 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.

Setting	Description
Health Check List	The **Health Check List** option is available if **Health Check Inherit** is **disabled** and **Health Check Control** is **enabled** Specify one or more health check configuration objects.

Note:
- If Health Check is disabled for SDN connector server members, their health check status will always appear as available since the status cannot be verified through health check.
- If the SDN instance is changed in the SDN connector, the SDN connector server member status will not be affected. For example, if the EC2 instance is terminated in AWS, that instance will still remain a server member in the SDN connector GLB server.

7. Save the Member list configuration and then save the GLB Server configuration to commit the Member list changes.

Configuring a global load balance 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 658.
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. **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 659.
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: Stating 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 "tp1"
      set member ZZ US CN65
   next
end
```

Where ZZ stands for Reserved, US for United States, and CN65 for China, Xingjiang

Web Application Firewall

The **Web Application Firewall** menu contain features and configurations that allow you to use web application firewall policies to scan HTTP requests and responses against known attack signatures and methods and filter matching traffic.

This section is organized into the following sub-menu topics:

- WAF Profile on page 667
- Known Web Attacks on page 680
- Common Attacks Detection on page 688
- Sensitive Data Protection on page 697
- Input Validation on page 718
- Access Protection on page 729
- CORS Protection on page 733
- API Protection on page 739
- Bot Mitigation on page 756
- Web Vulnerability Scanner on page 771

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 661 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.

- Web Attack Signature policy —The signature database includes signatures that can detect known attacks and exploits that can be found in 29 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.
- Cookie Security policy — This policy enables you to create rules that prevent cookie-based attacks and apply them in a protection profile.
- Data Leak Prevention policy — This policy enables you to create rules that prevent information leaks, damages and loss.
- HTTP Header Security policy — This policy enables you to create rules to prevent or mitigate known XSS, clickjacking, and MIME sniffing security vulnerabilities. These response headers define security policies to client browsers so that the browsers avoid exposure to known vulnerabilities when handling requests.

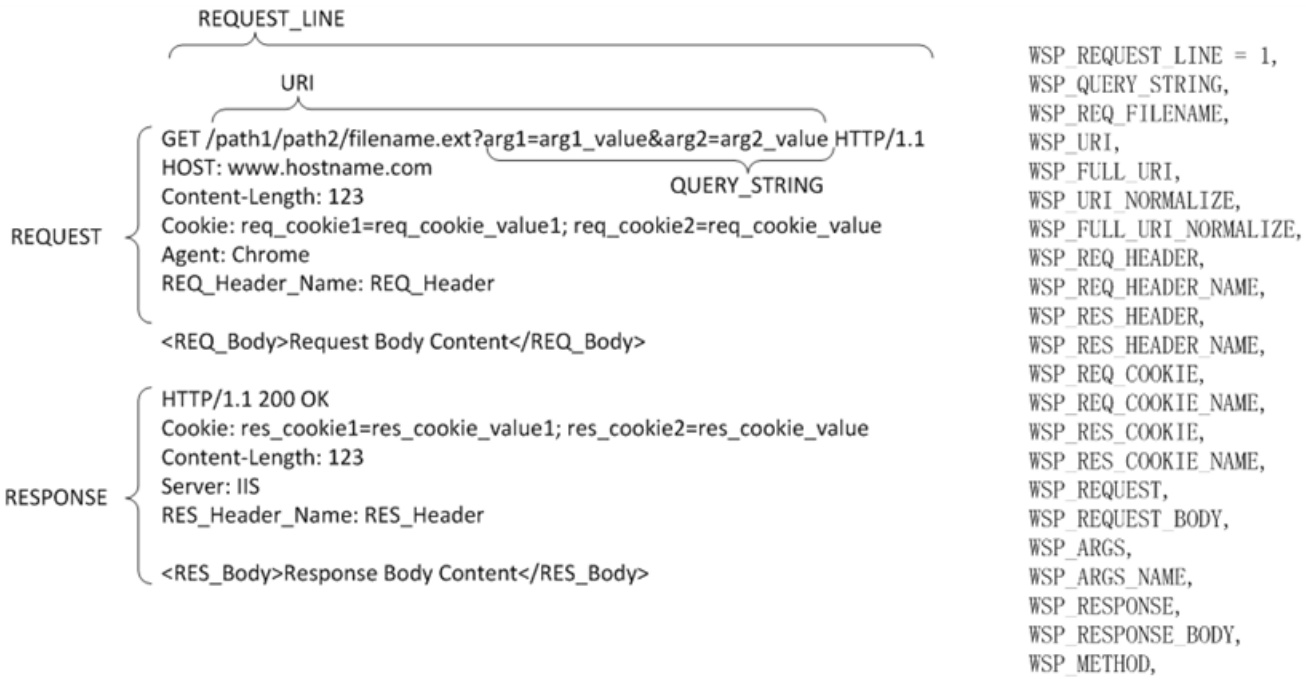
- Input Validation Policy — This policy enables you to create rules to prevent suspicious HTTP requests by verifying the user input from scan points like URL parameter, HTML form, hidden fields, and upload file.
- Brute Force Attack Detection policy — This policy enables you to create rules to prevent too many login tests.
- Credential Stuffing Defense policy — This policy enables you to create rules to identify login attempts using username and password that have been compromised using an always up-to-date feed of stolen credentials.
- JSON Detection policy — This policy enables you to create rules that enforce security checks that examine client HTTP requests for anomalies in JSON data in HTTP POST operations.
- XML Detection policy — This policy enables you to create rules that examine client requests for anomalies in XML code.
- OpenAPI Detection policy — This policy enables you to create rules through defining a standard, language-agnostic interface to RESTful APIs, which allows both humans and computers to discover and understand the capabilities of the service without access to source code, documentation, or through network traffic inspection.
- API Gateway policy — This policy includes an API management tool that sits between a client and a collection of backend services. It acts as a reverse proxy to accept all API calls and return the appropriate result.
- 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 allowlists to exclude known trusted sources (good Bots) from detection.
- Threshold Based Detection — This policy enables you to create rules to detect bad bots, such as web crawlers, content scraping, and attack bots.
- Biometrics Based Detection — This policy enables you to create rules that detect bots using behavioral biometrics such as mouse movement, keyboard, screen touch, and scroll.
- Advanced Protection policy — This policy enables you to create rules that detect web crawlers and content scraping.
- CSRF Protection policy — This policy enables you to create rules that protect backend servers from CSRF attacks.

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



Number	Attack	Solution
1	SQL Injection	FortiADC Supports it in two ways to prevent the SQL injection attack: • Signatures • SQL/XSS Injection Detection (Under WAF > Common Attacks Detection) Enable FortiADC's exclusive SQL/XSS Injection Detection function for XSS attacks prevention .
2	Cross Site Scripting	FortiADC Supports it in two ways to prevent the XSS injection attack. • Signatures • SQL/XSS Injection Detection (Under WAF > Common Attacks Detection) Enable FortiADC's exclusive SQL/XSS Injection Detection function for XSS attacks prevention .

3	Parameter/HTTP Tampering	FortiADC Supports it with "request-body-detection" signature profile.
		• If the signature profile was created by being cloned from "High-Level-Security" profile, the "request-body-detection" is enabled already.
		• If the signature profile was created by "Create New", the "request-body-detection" is disabled.
		You can enable it through CLI, for example:
		`config security waf web-attack-signature`
		`edit "Poc_Test"`
		`set request-body-detection enable`
		`next`
		`end`

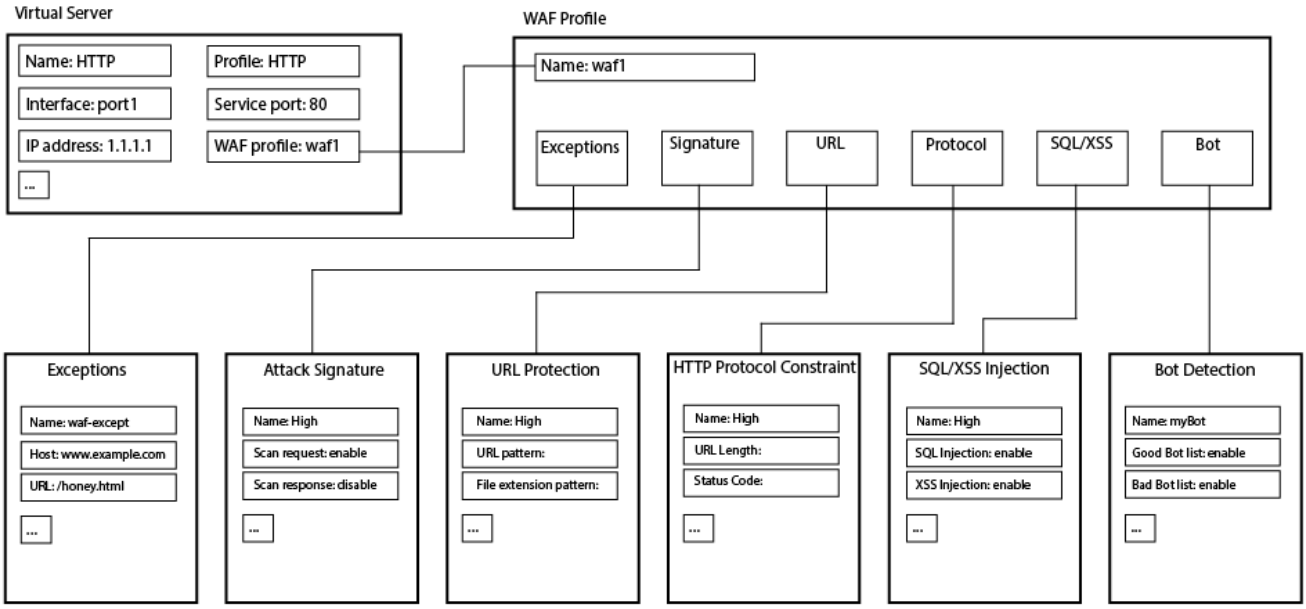
4	Sensitive information	It can be protected by configuring the Sensitive Data Type in Data Leak

Number	Attack	Solution
		Prevention (DLP) policy.
5	Cross Site Request Forging (CSRF)	It can be prevented by configuring ".*" in Parameter Value.
6	Session Hijacking	It can be prevented by enabling Cookie Security and configuring Authentication policy.
7	Blind SQL Injection	FortiADC supports it in two ways to prevent the SQL injection attack. • Signatures • SQL/XSS Injection Detection (Under WAF > Common Attacks Detection) It's recommended to enable the Exclusive SQL/XSS Injection Detection function for SQL attack prevention.
8	Request Smuggling	FortiADC strictly follows RFC 7230, section 3.3.3. If both Content-Length and Transfer-Encoding HTTP Header exist in the request, Content-Length will be removed. This ensures that the HTTP Request Smuggling attack can be blocked by FortiADC without any additional settings.
9	Web Scraping	FortiADC provides three ways to prevent the Web Scraping attacks. • WAF Signatures • Content Detection (Under WAF > Threshold Based Detection) • Content Scraping (Under WAF > Common Attacks Detection > Advance Protection)

Web application firewall configuration overview

WAF configuration overview on page 663 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, Bot Detection policy, and more. 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: WAF profiles, Web Attack Signature policies, HTTP Protocol Constraint policies, SQL/XSS Injection Detection policies, JSON Detection and XML Detection.

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.

Action

You can create an action which FortiADC takes when the conditions are fulfilled for WAF.

Basic Steps

1. Create configuration objects that define the action.
2. Select this action to a WAF rule configuration.

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.

OWASP TOP10

From the **Web Application Firewall > OWASP TOP10** sub-menu, you can configure a WAF profile based on OWASP Top 10 attacks. In the configuration wizard, you can select one or more OWASP Top 10 attacks, then FortiADC will aggregate all the WAF policies that can protect against the selected attacks. After you complete the OWASP TOP10 wizard, it will be listed in the WAF Profile table.



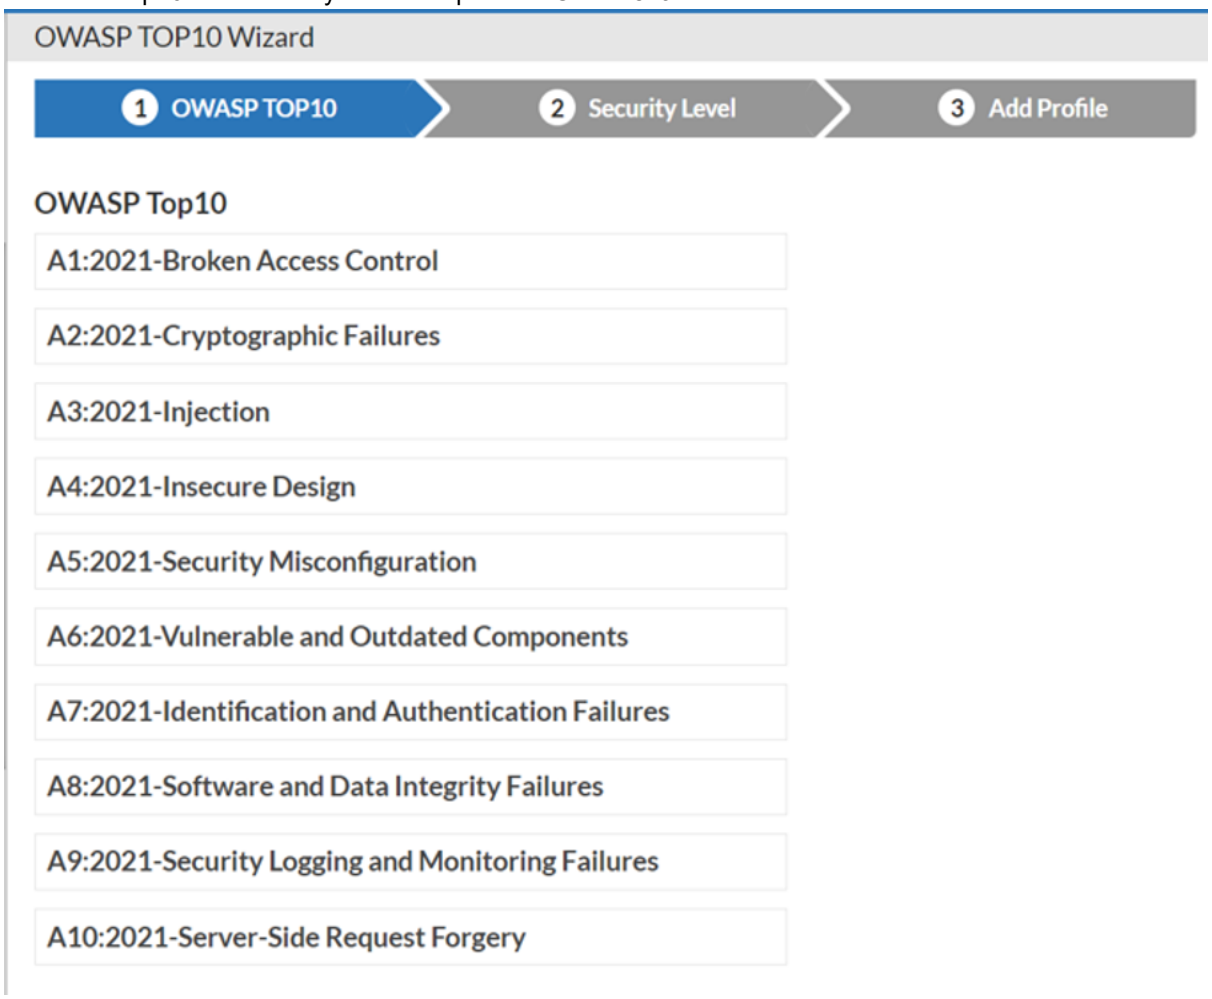
From FortiADC 7.1.0, the OWASP Top 10 list has been updated to the latest 2021 version. The OWASP Top 10 Wizard is automatically updated to the 2021 list, and the OWASP Top 10 2021 log data will be displayed through FortiView.

Log data from OWASP Top 10 2017 can still be accessed through the Security log.

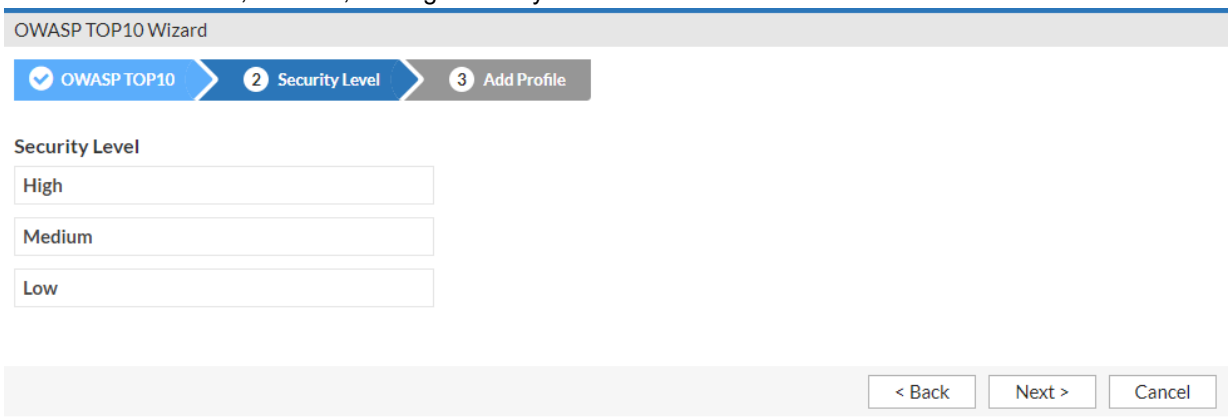
To create a OWASP TOP10 profile:

1. Go to **Web Application Firewall > OWASP TOP10 Wizard**
   To access this part of the web UI, you must have Read-Write permission for Security settings.

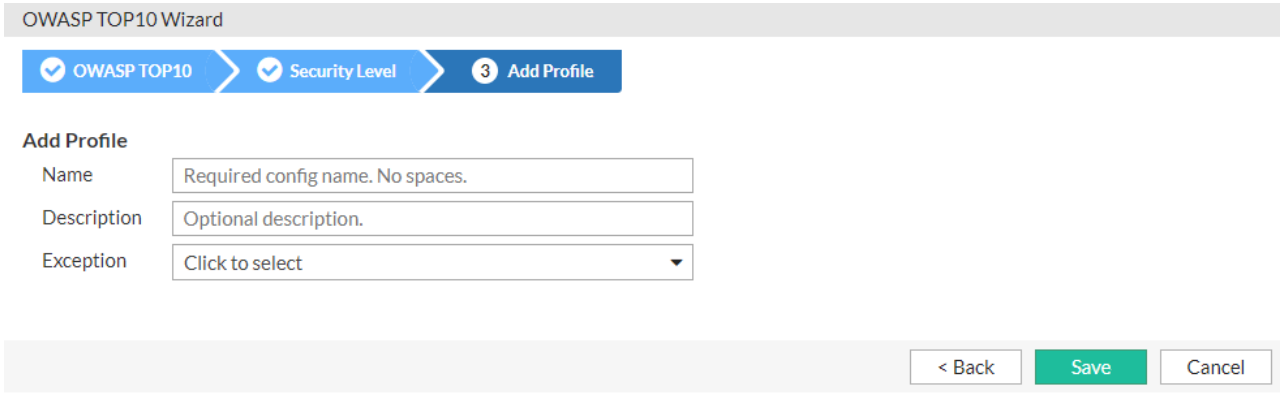
2. Select the top10 attacks that you want to prevent. Click **Next**.



3. Select the Security Level, so that this OWASP Top10 profile will protect against the attacks with the corresponding security level.
   a. High: Only the attacks with high security level will be screened out.
   b. Medium: Attacks with medium and high security levels will both be screened out.
   c. Low: Attacks with low, medium, and high security levels will all be screened out.



4. Enter a name and brief description for the profile. Select an exception configuration object. Exceptions identify specific hosts or URL patterns that are not subject to processing by this rule.



5. Save the configuration.

You can view this profile in **Web Application Firewall > WAF Profile**.

WAF Profile

The **Web Application Firewall > WAF Profile** sub-menu allows you to perform the following tasks relating to the WAF profile for the Web Application Firewall:

- Configuring a WAF Profile on page 667
- Configuring WAF Action objects on page 671
- Configuring WAF Exception objects on page 673

Configuring a WAF Profile

A WAF profile references the WAF policies that are to be enforced.

Predefined WAF profiles on page 667 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

Predefined Profiles	Description
	• 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:

1. Go to **Web Application Firewall > WAF Profile**.
2. Click the **WAF Profile** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in WAF Profile configuration on page 668.
5. 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.
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.
Description	A string to describe the purpose of the configuration, to help you and other administrators more easily identify its use.
Rule Match Record	Enable to allow the Security Log to display the part of the rule that is matched when the security event is logged. This is disabled by default.
Use Original IP	Enable to use the client's original IP address instead of the source IP for attack identification. This is disabled by default. When enabled, the client's original IP address from the "X-Forwarded-For" header is used in place of the standard source IP, impacting specific WAF functions. This setting affects WAF Blocked IP matching, Known Good Bots in Bot Detection, and source IP-based rules in the Exception feature. Log records are updated with the original client IP for detected attacks, ensuring accurate logging. This feature applies to both HTTP and HTTPS traffic and supports IPv4 and IPv6 addressing.

Settings	Guidelines
	Note: If the HTTP request packet does not contain the "X-Forwarded-For" information or contains invalid data, the system will default to using the HTTP source IP address instead. For further details, see Using the Original IP on page 670.
Standard Protection	
Adaptive Learning	Select a predefined or user-defined Adaptive Learning configuration object.
Web Attack Signature	Select a predefined or user-defined Web Attack Signature configuration object.
HTTP Protocol Constraint	Select a predefined or user-defined HTTP Protocol Constraint configuration object.
Sensitive Data Protection	
Cookie Security	Select a user-defined Cookie Security configuration object.
Data Leak Prevention	Select a user-defined Data Leak Prevention configuration object.
HTTP Header Security	Select a user-defined HTTP Header Security configuration object.
Input Protection	
SQL/XSS Injection Detection	Select a predefined or user-defined SQL/XSS Injection Detection configuration object.
Input Validation Policy	Select a user-defined Input Validation Policy configuration object.
CORS Protection	Select a user-defined CORS Protection configuration object.
Access Protection	
Brute Force Attack Detection	Select a user-defined Brute Force Attack Detection configuration object.
URL Protection	Select a user-defined URL Protection configuration object.
Credential Stuffing Defense	Select a user-defined Credential Stuffing Defense configuration object.
API Protection	
JSON Detection	Select a predefined or user-defined JSON Detection configuration object.
XML Detection	Select a predefined or user-defined XML Detection configuration object.
OpenAPI Detection	Select a user-defined OpenAPI Detection configuration object.
API Gateway	Select a user-defined API Gateway configuration object.
Bot Mitigation	
Bot Detection	Select a user-defined Bot Detection configuration object.
Threshold Based Detection	Select a predefined or user-defined Threshold Based Detection configuration object.
Biometrics Based Detection	Select a user-defined Biometrics Based Detection configuration object.

Settings	Guidelines
Fingerprint Based Detection	Select a user-defined Fingerprint Based Detection configuration object.
Advanced Bot Protection	Select a user-defined Advanced Bot Protection configuration object.
Advanced Protection	
Advanced Protection	Select a user-defined Advanced Protection configuration object.
CSRF Protection	Select a user-defined CSRF Protection configuration object.

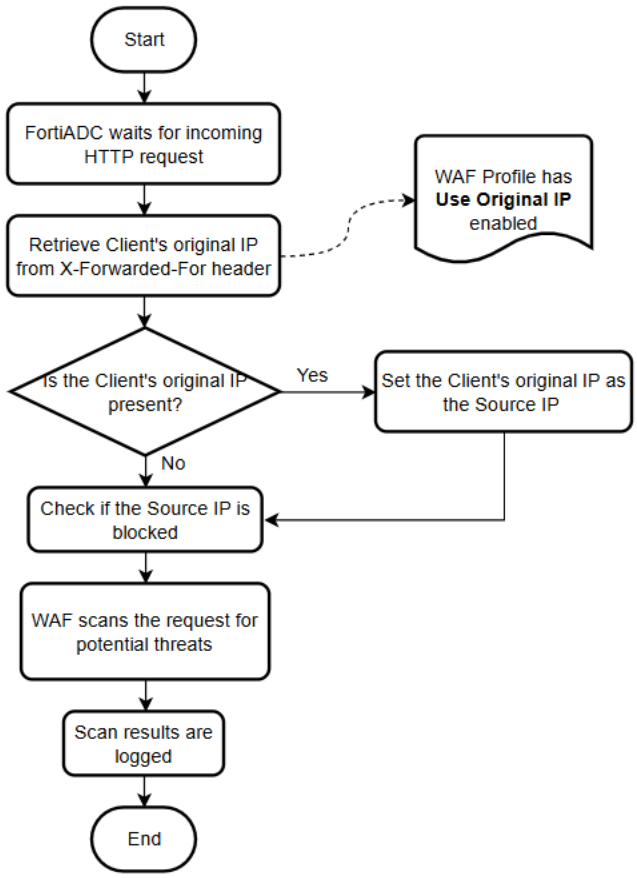
Using the Original IP

When FortiADC is deployed behind a proxy, incoming requests appear to originate from the proxy's IP address. If FortiADC identifies malicious requests, it may block the proxy's IP, leading to service disruptions for all legitimate clients routed through it. To address this issue, the **Use Original IP** option in the WAF Profile enables WAF modules to inspect the client's original IP address instead of the HTTP source IP for attack source identification. This is accomplished by leveraging the X-Forwarded-For HTTP extension header, which contains the original client's IP address appended by the proxy. By using this information, FortiADC can accurately identify clients and apply relevant security policies and actions, allowing for more precise logging of attack incidents based on the original client IP. This enhances the effectiveness of threat detection and response mechanisms.

When **Use Original IP** is enabled, log records are updated with the original client IP upon attack detection. It applies to both HTTP and HTTPS traffic and supports both IPv4 and IPv6 addressing, ensuring comprehensive inspection across protocols. However, if the HTTP request packet does not contain the X-Forwarded-For information or contains invalid data, the system will default to using the HTTP source IP address instead.

Process Flow:

The following process outlines the steps taken by FortiADC when handling an incoming HTTP request with the **Use Original IP** option is enabled in the WAF profile. This workflow ensures accurate client identification and effective threat detection while maintaining network security.



1. FortiADC awaits for an incoming HTTP request.
2. Retrieve the client's original IP:
   If **Use Original IP** is enabled in the WAF profile, retrieve the client's original IP address from the X-Forwarded-For header.
3. Check for the client's original IP in the X-Forwarded-For header:
   a. **Yes**: If the client's original IP is present, set it as the source IP.
   b. **No**: If absent, proceed to the next checkpoint.
4. Verify if the source IP is blocked:
   When **Use Original IP** is enabled, use the client's original IP from the X-Forwarded-For header instead of the standard Source IP for this check.
5. Scan the request for potential threats:
   The WAF scans the request for potential threats. Certain rules, such as the **Source IP** in Exception lists and Known Good Bots in Bot Detection, will use the client's original IP if **Use Original IP** is enabled.
6. Log the scan results:
   The results are logged, and if **Use Original IP** is enabled, the client's original IP from the X-Forwarded-For header is recorded in place of the standard source IP.

Configuring WAF Action objects

Configure what action FortiADC should take when it meets the WAF conditions.

After you have created an action object, you can specify it in individual WAF feature rules.

Before you begin:

- You must have Read-Write permission for Security settings.

In many cases, you can use predefined profiles to get started.

Predefined actions	Description
alert	WAF policies will allow the traffic to pass and log the event.
block	WAF policies will drop the current attack session by HTTP 403 message and block the attacker (according the attacker's IP address) for 1 hour, and log the event.
captcha	WAF policies will allow the traffic to pass if the client successfully fulfills the CAPTCHA request, and log the event.
deny	WAF policies will the drop current attack session by HTTP 403 message, and log the event.
silent-deny	WAF policies will drop the current attack session by HTTP 403 message, without logging the event.

To configure a WAF Action object:

1. Go to **Web Application Firewall > WAF Profile**.
2. Click the **Action** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration of WAF Action objects.
5. Save the 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.
Action Type	Select which action FortiADC takes when the conditions are fulfilled for WAF: • Pass — Allow the request. • Deny — Block the request. • Period Block — Deny all the HTTP requests from a source IP within a period which specified by Period Block. • Redirect — Send a redirect. You must specify the redirect URL. • Captcha — Requires the client to successfully fulfill the CAPTCHA request.
Deny Code	The **Deny Code** option is available if the **Action Type** is **Deny** or **Period Block**. Select the HTTP response code, Default: 403. 200, 202, 204, 205, 400, 403, 404, 405, 406, 408, 410, 500, 501, 502, 503, 504
Period Block	The **Period Block** option is available if the **Action Type** is **Period Block**. Specify a time period when action blocks the client. Default: 60 seconds, Range: 1- 3600 seconds.

Settings	Guidelines
Redirect URL	The **Redirect URL** option is available if the **Action Type** is **Redirect**. Specify the URL that you want to redirect.
Log Status	Enable/Disable log of events
Comment	Enter comment or description of the action for your records.

Configuring WAF Exception objects

WAF exceptions identify specific patterns that are not subject to processing by WAF rules. Use WAF exception rules to reduce false-positives triggered by legitimate HTTP requests that match an attack signature rule. FortiADC supports URL, hosts and source IP patterns matching in the WAF exception rules.

You can create and configure WAF Exception objects using either of the following methods:

- From the **Web Application Firewall > WAF Profile > Exceptions** tab, you can create/configure exception objects to then apply to specific WAF profiles and individual WAF feature rules. For detailed steps, see Configuring WAF exception rules from the WAF Profile > Exceptions tab on page 673.
- From the WAF security log (**Log & Report > Security Log** or **FortiView > Security Logs**), you can create/configure exception objects to directly apply to the specific WAF log. For detailed steps, see Configuring WAF exception rules from the WAF log on page 675.

Before you begin:

- You must have Read-Write permission for System settings.



For optimal functionality, we recommend keeping the number of WAF exception rules configured to a minimum. If a large number of WAF exception rules are configured, none may work effectively due to limitations of the shared memory (maximum total is 256.0 MBs in the VM platform).

Configuring WAF exception rules from the WAF Profile > Exceptions tab

You can create or configure exception objects to then apply to specific WAF profiles and individual WAF feature rules.

To configure a WAF exception rule from the WAF Profile > Exceptions tab:

1. Go to **Web Application Firewall > WAF Profile**.
2. Click the **Exceptions** tab.
3. Click **Create New** to display the configuration editor.
4. Configure the following:

Parameter	Description
Name	Enter a unique name for the WAF Exception. Maximum length is 130 characters. **Note**: Once saved, the name of an Exception cannot be changed.

5. Click **Save**.

6. Under **Exception Rule**, click **Create New** to display the configuration editor.
7. Select the exception pattern **Element Type** from the drop-down menu.
   - URL
   - Source IP
   - Source IPv6
   - HTTP Method
   - HTTP Header
   - Cookie
   - Parameter
8. Configure the following parameters based on the exception pattern **Element Type**.

Parameter	Description
URL	
Exception Host Status	Enable/disable the setting exceptions by host pattern.
Host Pattern	The **Host Pattern** option appears if **Exception Host Status** is enabled. Specify the matching string. Regular expressions are supported. Maximum length is 128 characters. For example, you can specify `www.example.com`, `*.example.com`, or `www.example.*` to match a literal host pattern or a wildcard host pattern.
URL Pattern	Specify the matching string. Must begin with a URL path separator (/). Regular expressions are supported. Maximum length is 128 characters. For example, you can specify path names and files with expressions like `\/admin`, `.*\/data\/1.html`, or `\/data.*`.
Source IP	
IPv4/Netmask	Specify the IPv4 address and netmask. For example: 192.0.2.5/24
Source IPv6	
IPv6/Netmask	Specify the IPv6 address and netmask. For example: 2001:0db8:85a3::8a2e:0370:7334/64
HTTP Method	
HTTP Method	Select the HTTP method(s): • GET • POST • HEAD • TRACE • CONNECT • DELETE • PUT • PATCH • OPTIONS • OTHERS

Parameter	Description
HTTP Header	
Name Pattern	Specify the matching string. Regular expressions are supported. Maximum length is 1024 characters. For example: `. Content*`
Check Value of Specified Element	Enable/disable value checking for the specified element.
Value Pattern	The **Value Pattern** option appears if **Check Value of Specified Element** is enabled. Specify the matching string. Regular expressions are supported. Maximum length is 1024 characters. For example: `. Content*`
Cookie	
Name Pattern	Specify the matching string. Regular expressions are supported. Maximum length is 1024 characters. For example: `. Content*`
Check Value of Specified Element	Enable/disable value checking for the specified element.
Value Pattern	The **Value Pattern** option appears if **Check Value of Specified Element** is enabled. Specify the matching string. Regular expressions are supported. Maximum length is 1024 characters. For example: `. Content*`
Parameter	
Name Pattern	Specify the matching string. Regular expressions are supported. Maximum length is 1024 characters. For example: `. Content*`
Check Value of Specified Element	Enable/disable value checking for the specified element.
Value Pattern	The **Value Pattern** option appears if **Check Value of Specified Element** is enabled. Specify the matching string. Regular expressions are supported. Maximum length is 1024 characters. For example: `. Content*`

9. Click **Save**.

Configuring WAF exception rules from the WAF log

You can create or configure exception objects to directly apply to the specific WAF log. You have the option to create exception rules in basic mode or advanced mode. Using basic mode, you can apply the URL or Source IP/ Source IPv6

from the WAF log directly to the exception rule. The advanced mode enables you to create and configure exception rules for all supported element patterns.

To configure a WAF exception rule from the WAF log using basic mode:

1. Go to **Log & Report > Security Log** or **FortiView > Security Logs**.
2. Navigate to the WAF security log and expand the details for which you want to add an exception rule.
3. In the log details, click **Add Exception** to display the configuration editor.
4. Configure the following parameters.

Parameter	Description
Profile Name	Select the WAF Profile to apply the WAF exception rule.
Advanced Mode	To configure the WAF exception rule using basic mode, ensure **Advanced Mode** is disabled. By default, **Advanced Mode** is disabled, which enables basic mode.
Element Type	Select either of the following: • URL — To apply the HTTP URL from this WAF log for this exception rule. • Source IP — To apply the Source from this WAF log for this exception rule. The **Source IP** option appears if the Source address is IPv4. • Source IPv6 — To apply the Source IPv6 from this WAF log for this exception rule. The **Source IPv6** option appears if the Source address is IPv6.

5. Click **Save**.

To configure a WAF exception rule from the WAF log using advanced mode:

1. Go to **Log & Report > Security Log** or **FortiView > Security Logs**.
2. Navigate to the WAF security log and expand the details for which you want to add an exception rule.
3. In the log details, click **Add Exception** to display the configuration editor.
4. Configure the following:

Parameter	Description
Profile Name	Select the WAF Profile to apply the WAF exception rule. **Note**: The profile name parameter will display according to the WAF subcategory of the log. For example, if the WAF subcategory of the log is "cookie security" then the parameter will display as "Cookie Security Profile Name".
Advanced Mode	Enable Advanced Mode to configure the WAF exception rule using advanced mode. Once **Advanced Mode** is enabled, the **Element Type** drop-down menu will include all supported exception pattern.

5. Select the exception pattern **Element Type** from the drop-down menu.
   - URL
   - Source IP

- Source IPv6
- HTTP Method
- HTTP Header
- Cookie
- Parameter

6. Configure the following parameters based on the exception pattern **Element Type**.

Parameter	Description
URL	
Exception Host Status	Enable/disable the setting exceptions by host pattern.
Host Pattern	The **Host Pattern** option appears if **Exception Host Status** is enabled. Specify the matching string. Regular expressions are supported. Maximum length is 128 characters. For example, you can specify `www.example.com`, `*.example.com`, or `www.example.*` to match a literal host pattern or a wildcard host pattern.
URL Pattern	Specify the matching string. Must begin with a URL path separator (/). Regular expressions are supported. Maximum length is 128 characters. For example, you can specify path names and files with expressions like `\/admin`, `.*\/data\/1.html`, or `\/data.*`.
Source IP	
IPv4/Netmask	Specify the IPv4 address and netmask. For example: 192.0.2.5/24
Source IPv6	
IPv6/Netmask	Specify the IPv6 address and netmask. For example: 2001:0db8:85a3::8a2e:0370:7334/64
HTTP Method	
HTTP Method	Select the HTTP method(s): • GET • POST • HEAD • TRACE • CONNECT • DELETE • PUT • PATCH • OPTIONS • OTHERS
HTTP Header	
Name Pattern	Specify the matching string. Regular expressions are supported. Maximum length is 1024 characters. For example: `. Content*`

Parameter	Description
Check Value of Specified Element	Enable/disable value checking for the specified element.
Value Pattern	The **Value Pattern** option appears if **Check Value of Specified Element** is enabled. Specify the matching string. Regular expressions are supported. Maximum length is 1024 characters. For example: `. Content*`
Cookie	
Name Pattern	Specify the matching string. Regular expressions are supported. Maximum length is 1024 characters. For example: `. Content*`
Check Value of Specified Element	Enable/disable value checking for the specified element.
Value Pattern	The **Value Pattern** option appears if **Check Value of Specified Element** is enabled. Specify the matching string. Regular expressions are supported. Maximum length is 1024 characters. For example: `. Content*`
Parameter	
Name Pattern	Specify the matching string. Regular expressions are supported. Maximum length is 1024 characters. For example: `. Content*`
Check Value of Specified Element	Enable/disable value checking for the specified element.
Value Pattern	The **Value Pattern** option appears if **Check Value of Specified Element** is enabled. Specify the matching string. Regular expressions are supported. Maximum length is 1024 characters. For example: `. Content*`

7. Click **Save**.

Limitations: Escaped Characters

All **Name Pattern** and **Value Pattern** fields support regular expression. However, some characters must be escaped to be a valid regular expression or be functional as an exception rule. If your expression contains characters that require escaping, an error message may be triggered to reject the invalid expression. However, it is also possible that an error may not be triggered by unescaped characters if it is considered syntactically correct — in which case these expressions would not function as exception rules since they will not match any user traffic.

This section lists the most commonly used special characters that need to be escaped to make an input valid or functional as an exception rule.

Brackets: `[]`

Brackets (`[]`) require other characters between the brackets to be a valid exception rule regular expression.

For example:

Invalid: `[]`

Valid: `[123]` — Valid

`[]` is an invalid exception rule regular expression because the input only contain brackets with no other characters in between. Whereas `[123]` is valid because there are number characters between the brackets.

Parentheses: `()`

Parentheses (`()`) require a backslash (`\`) before each parenthesis to be a valid exception rule regular expression — `()` → `\(\)`

For example:

Invalid: `http://x.x.x.x/login?link=mocha:alert('attack%20success')`

Valid: `http://x.x.x.x/login?link=mocha:alert\('attack success'\)`

Focusing on the parameter value, `mocha:alert('attack%20success')` is invalid because there is no backslash before each parenthesis. Whereas `mocha:alert\('attack success'\)` is valid with the backslash inserted before each parenthesis.

Asterisk: `*`

Asterisks (`*`) require a backslash (`\`) before each asterisk to be a valid and functional exception rule — `*` → `\*`

For example:

curl -vv -X POST --cookie "`Cookie123=abcd`"

"`http://x.x.x.x/index.php?n123=v123&p_name1=p_value1`"

Where the cookie name is "`cookie`" and the cookie value is "`a*`"

Invalid: `a*`

Valid: `a\*`"

In this case, both `a*` and `a\*` are both correct in syntax. However, `a*` would not be functional as an exception rule because it would not match any user traffic.

Space: `%20`

Spaces (`%20`) in URLs must be replaced with spaces to be a valid exception rule regular expression.

For example:

Invalid: `http://x.x.x.x/login?link=mocha:alert('attack%20success')`

Valid: `http://x.x.x.x/login?link=mocha:alert(\'attack success'\)`

Focusing on the parameter value, the invalid expression becomes valid when the `%20` is replaced with the space: `mocha:alert('attack%20success')` → `mocha:alert(\'attack success'\)`

Single Quotes: `''`

When the Name Pattern or Value Pattern fields contain single quotes, it will be automatically escaped.

For example:

In the GUI or CLI, you may enter the parameter value pattern as: `alert\('attack%20success'\)`

In the CLI, the value pattern will appear as: `alert\\(\'attack success\'\\)`

Known Web Attacks

The **Web Application Firewall > Known Web Attacks** sub-menu allows you to configure Web Attack Signature policies. The signature database includes signatures that can detect known attacks and exploits that can be found in 29 scanpoints. In your policy configuration, you choose classes of scanpoints to process: HTTP Headers, HTTP Request Body, and HTTP Response Body.

- Configuring a Web Attack Signature policy on page 680
- Using the Signature Creation Wizard on page 685

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 683 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 681 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. Scan HTTP Response Body—Disabled.	High Severity Action—Deny. Medium Severity Action—Alert. Low Severity Action—Alert.
Alert-Only	Scan HTTP header—Enabled. Scan HTTP Request Body—Disabled. Scan HTTP Response Body—Disabled.	High Severity Action—Alert. Medium Severity Action—Alert. 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 681.
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.

Settings	Guidelines
	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. 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. • 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

Settings	Guidelines
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 • CVE Number (if one exists) • Reference (if one exists) • Found In
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 683 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)	Generic XSS Attack (42)
SQL Injection (2)	Generic SQL Injection (43)
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) Xpath Injection (49) XML External Entities (57) Insecure Deserialization (59)

Category (ID)	Subcategory (ID)
	HTTP Header Injection (60) Buffer Overflow (62) Denial Of Service (64)
Trojans (4)	Trojans (44)
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) Generic Malicious Leakage (47) Path Travel (58)
Known Exploits (6)	Oracle 9i (27) Coppermine Photo Gallery (28) Netscape Enterprise Server (29) Cisco IOS HTTP Service (30) Microsoft SQL Server (31) HP OpenView Network Node Manager (32) Best Sofrware 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) CHETCPASSWD (40) SAP (41) Generic Exploit (48) Lighttpd Server (53) Caucho Resin Server (54) JRun Web Server (55)

Category (ID)	Subcategory (ID)
	IBM Lotus Domino (56) WordPress (61) Struts 2 (63) Joomla! (65)
Credit Card Detection (7)	Credit Card Detection (45)
Bad Robot (8)	Bad Robot (46)
Cross Site Scripting (Extended) (9)	Cross Site Scripting (Extended) (50)
SQL Injection (Extended) (10)	SQL Injection (Extended) (51)
Generic Attacks (Extended) (11)	Generic Attacks (Extended) (52)

Using the Signature Creation Wizard

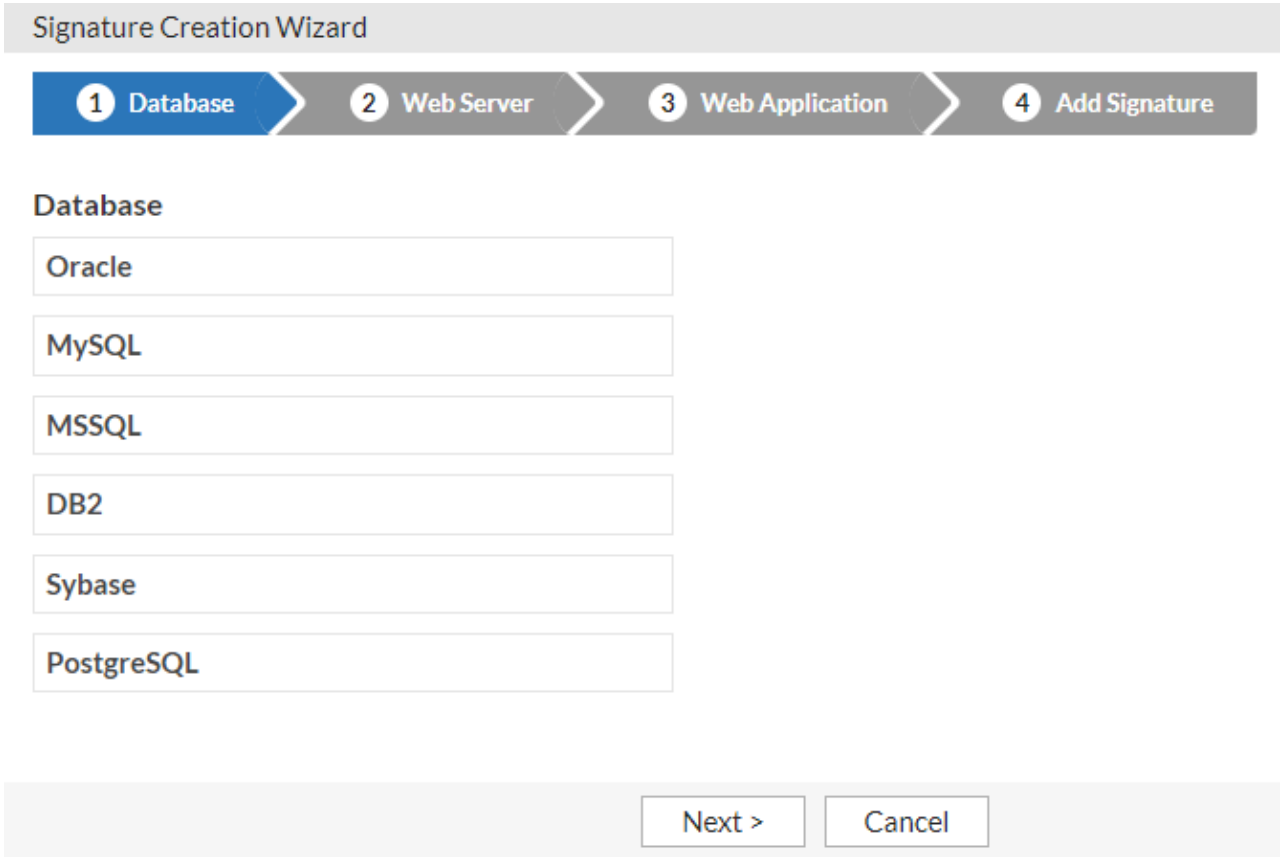
The Signature Creation Wizard is a step-by-step tool available in the GUI that helps you configure a Web Attack Signature policy based on the database, web server, or web application you want to protect. Using your selections, the Signature Creation Wizard automatically compiles a list of the most applicable WAF Signatures for protecting your database, web server, or web application. You can then apply those WAF Signatures in a Web Attack Signature policy.

Alternatively, you can custom configure a Web Attack Signature policy without using the Signature Creation Wizard. For more information, see Configuring a Web Attack Signature policy on page 680.

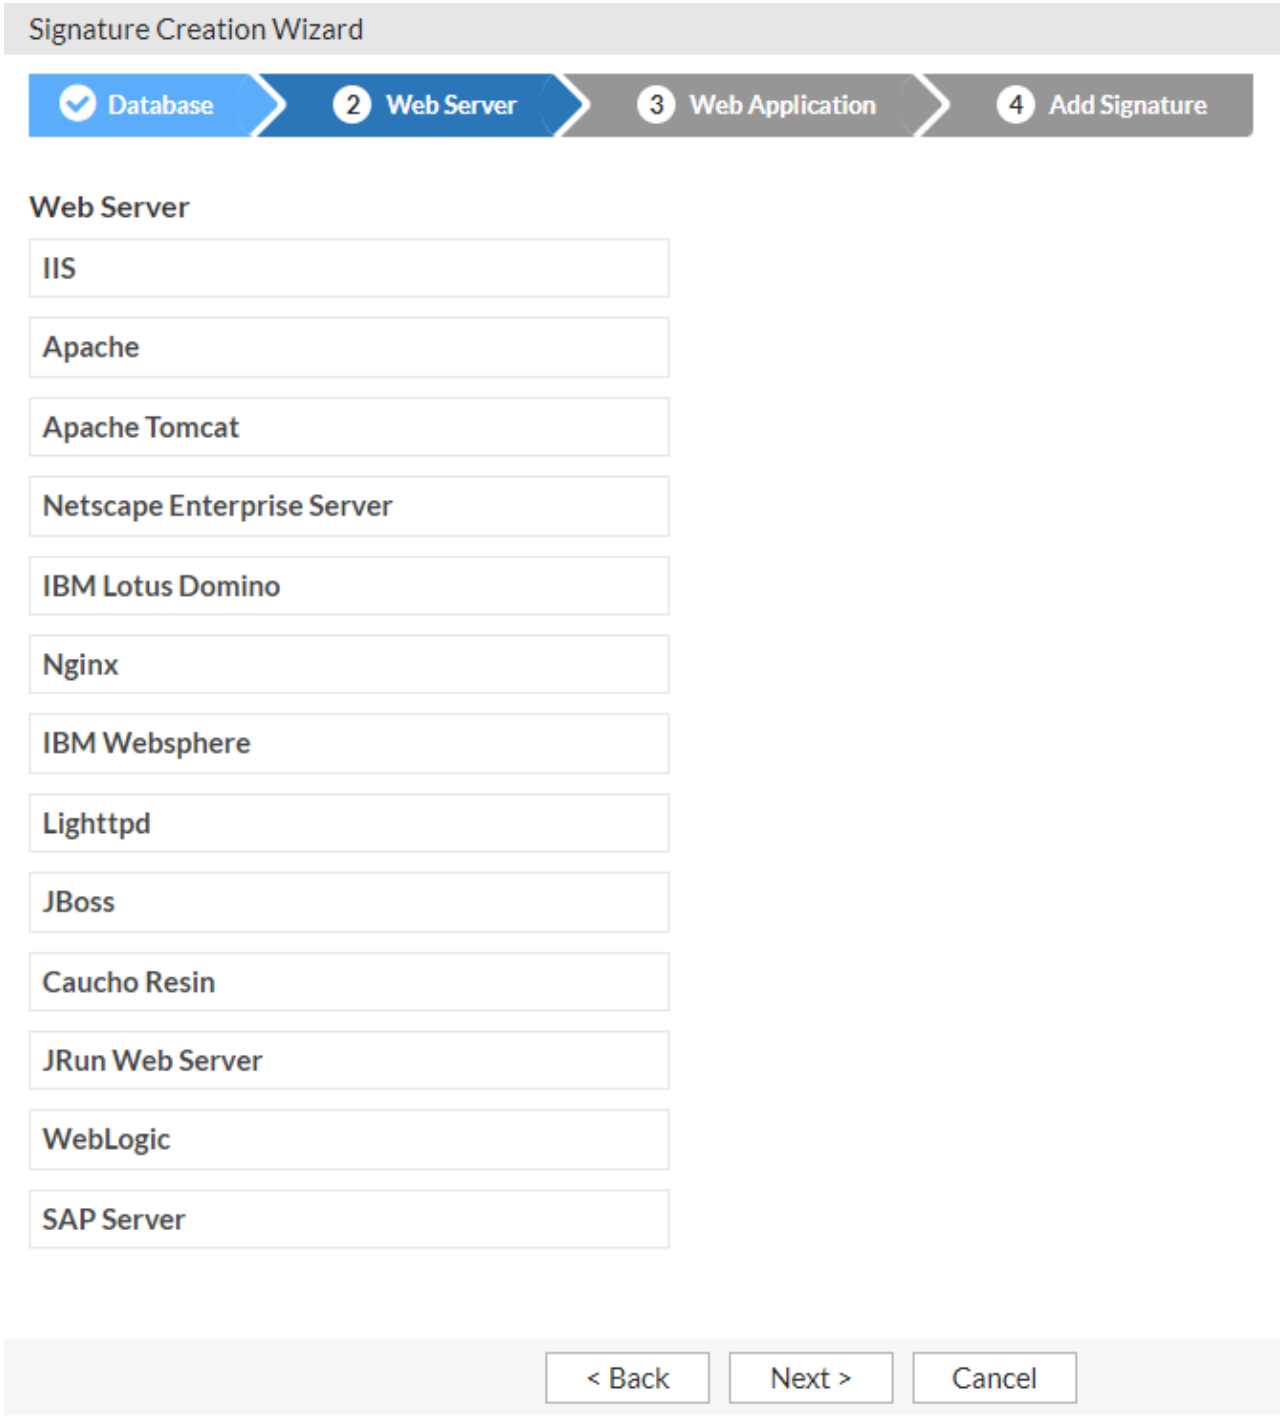
To configure a Web Attack Signature using the Signature Creation Wizard:

1. Go to **Web Application Firewall > Known Web Attacks**.
2. Click the **Web Attack Signature** tab.
3. Click **Signature Wizard** to display the configuration editor.

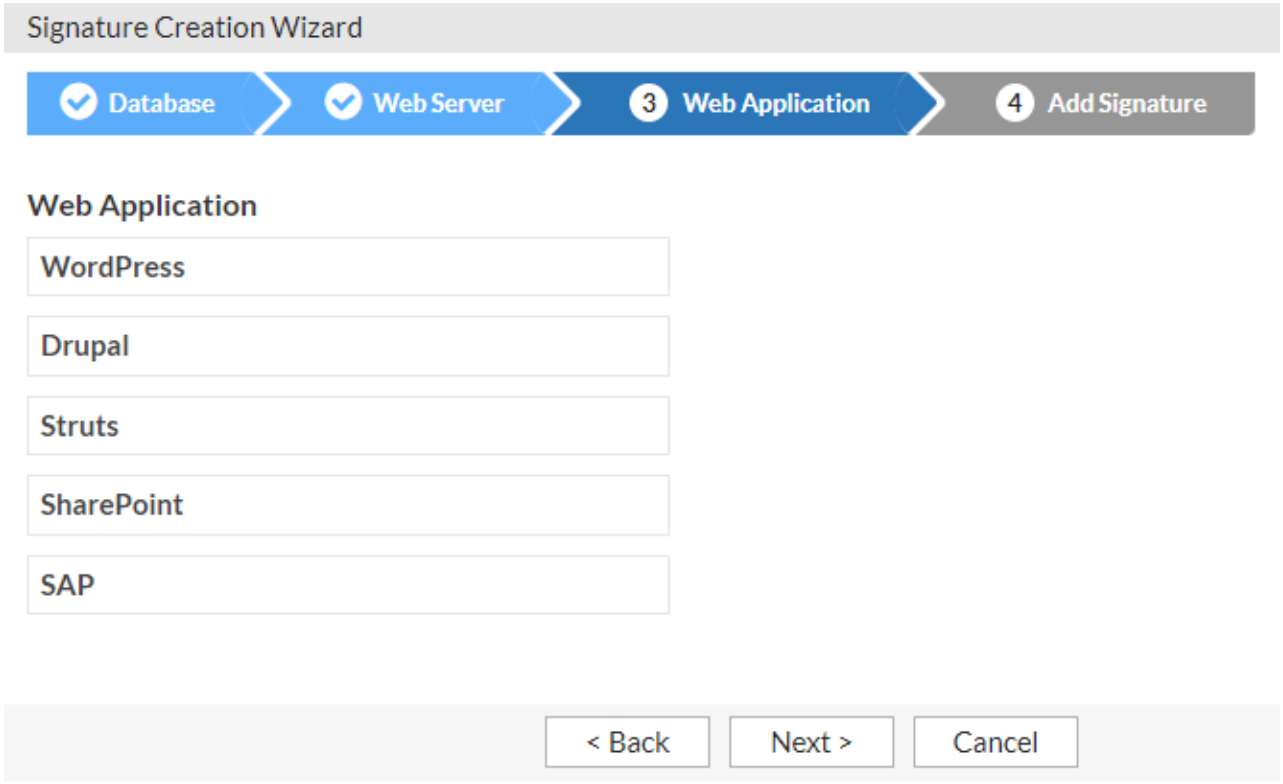
4. From the **Database** tab, select the database(s) you want to protect. Click **Next** to navigate to the **Web Server** tab.



5. From the **Web Server** tab, select the web server(s) you want to protect. Click **Next** to navigate to the **Web Application** tab, or click **Back** to navigate to the **Database** tab.



6. From the **Web Application** tab, select the web application(s) you want to protect. Click **Next** to navigate to the **Add Signature** tab, or click **Back** to navigate to the **Web Server** tab.



7. From the **Add Signature** tab, enter a name for your Web Attack Signature policy and click **Save**.
   The Category and Sub Category sections display the applicable WAF Signatures that you can apply to your policy.
8. Click **Apply** to save your Web Attack Signature policy.
9. Optionally, you can click **Save and Edit Signature Details** to apply WAF Exceptions where applicable.

Common Attacks Detection

The **Web Application Firewall > Common Attacks Detection** sub-menu allows you to perform the following tasks relating to the Common Attacks Detection for the Web Application Firewall:

- Configuring an Advanced Protection policy on page 688
- Configuring an HTTP Protocol Constraint policy on page 690
- Configuring an SQL/XSS Injection Detection policy on page 693
- Configuring a CSRF protection policy on page 696

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 689.



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.
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 profile that you want to apply. See Configuring WAF Action objects on page 671. The default value is Alert.

Settings	Guidelines
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 profile that you want to apply. See Configuring WAF Action objects on page 671. The default value 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.

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 691 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 691.
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

Settings	Guidelines
	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.
Maximum Request Body Length	Maximum length of the HTTP body. The default is 67108864. The valid range is 1-67108864.
Constraint Method Override	Enable/Disable to scan request method and try to match it in request method rule in following override headers: • X-HTTP-Method • X-Method-Override • X-HTTP-Method-Override
Request Method Rule	
Method	Select one or more methods to match in the HTTP request line: • CONNECT • DELETE • GET • HEAD • OPTIONS • POST • PUT • TRACE

Settings	Guidelines
	• 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	Select the action profile that you want to apply. See Configuring WAF Action objects on page 671. 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 an SQL/XSS Injection Detection policy

SQL/XSS Injection Detection Policies are designed to identify and prevent SQL injection and cross-site scripting (XSS) attacks. These types of attacks occur when user-supplied data is improperly processed by an interpreter. In an SQL injection attack, malicious HTTP requests are crafted to manipulate SQL queries, allowing unauthorized access to a web application's database. XSS attacks, on the other hand, exploit vulnerabilities that enable the execution of malicious scripts within a user's web browser.

Unlike traditional signature-based detection methods, the Web Application Firewall (WAF) SQL and XSS injection detector module employs lexical analysis to identify these threats. This approach is both complementary and faster, offering enhanced detection capabilities.

The policy configuration allows for enabling or disabling scanpoints, defining actions when traffic matches specific signatures, and setting the severity of events.

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 694 describes the predefined policies.

Predefined SQL injection and XSS detection policies

	SQL Injection			XSS		
Predefined Rules	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 694.
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.

Settings	Guidelines
	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	Select the action profile that you want to apply. See Configuring WAF Action objects on page 671. The default is alert, but we recommend using 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	Select the action profile that you want to apply. See Configuring WAF Action objects on page 671. The default is alert, but we recommend you deny XSS Injection.
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 a CSRF protection policy

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, 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 an 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. After you initially save the configuration, you cannot edit the name.
Action	Select which action profile that you want to apply. See Configuring WAF Action objects on page 671.

Settings	Guidelines
	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.

Sensitive Data Protection

The **Web Application Firewall > Sensitive Data Protection** sub-menu contains features and configurations for Sensitive Data Protection that allows the Web Application Firewall (WAF) to prevent information leaks, damage and loss, including Data Leak Prevention (DLP), Cookie Security, and HTTP Header Security.

- Configuring a Cookie Security policy on page 697
- Configuring an HTTP Header Security policy on page 700

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**.
   The configuration page displays the **Cookie Security** tab.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in Cookie Security configuration on page 698.



If you want to drop a large number of packets when traffic match the rules, you should set Action to “block” instead of “deny."

4. 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.
Samesite	Add SameSite attribute to prevent the browser from sending cookies along with cross-site requests, to mitigate the risk of cross-origin information leakage. It provides Strict, Lax, and None values for this attribute: • Strict: any request from the third parties will not carry such cookies; • Lax: any request from the third parties will not carry such cookies except for GET requests that navigate to the destination URL. • None: set the value as none if a cookie is required to be sent by cross origin. • Nothing: Do not add Samesite attribute to cookies. The default value is **Nothing**.
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.
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.
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.
Encrypted Cookie Type	**All**—will encrypt all cookies. **List**—will encrypt the cookie that matches with the cookie-list.

Settings	Guidelines
	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.
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 673.
Action	Select the action profile that you want to apply. See Configuring WAF Action objects on page 671. 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 encrpyted_cookie_type is set to "list."

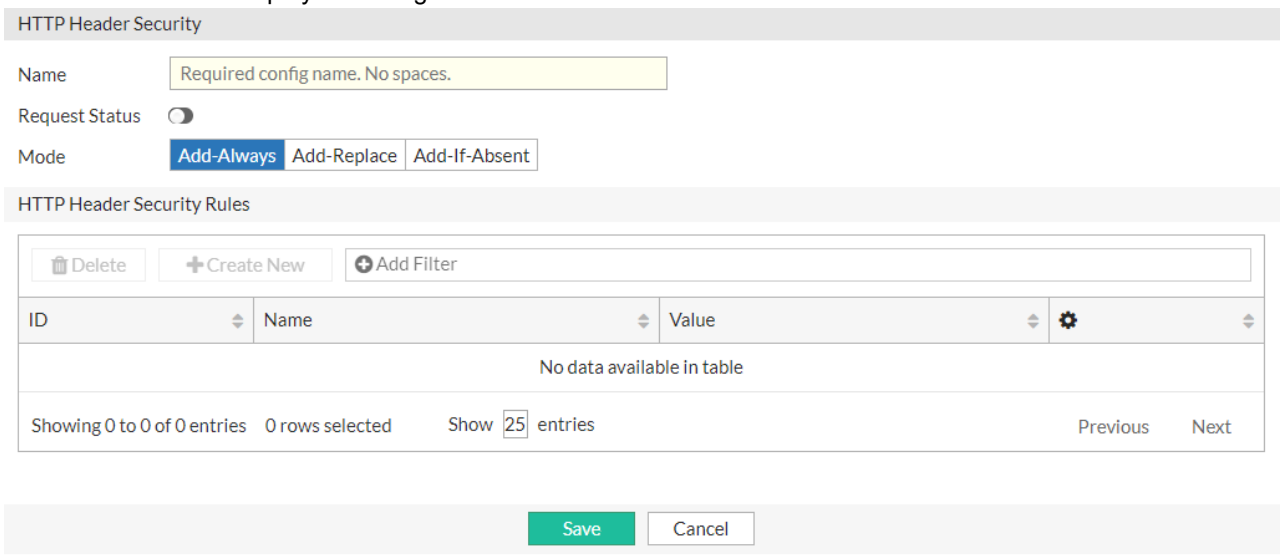
Configuring an HTTP Header Security policy

HTTP response security headers are a set of standard HTTP response headers proposed to prevent or mitigate known XSS, clickjacking, and MIME sniffing security vulnerabilities. These response headers define security policies to client browsers so that the browsers avoid exposure to known vulnerabilities when handling requests.

When the HTTP Security Headers feature is enabled, headers with specified values are inserted into HTTP responses coming from the backend web servers. This is a quick and simple solution to address the security vulnerabilities on user's website without code and configuration changes.

To configure an HTTP Header Security policy:

1. Go to **Web Application Firewall > Sensitive Data Protection**.
2. Click the **HTTP Header Security** tab.
3. Click **Create New** to display the configuration editor.

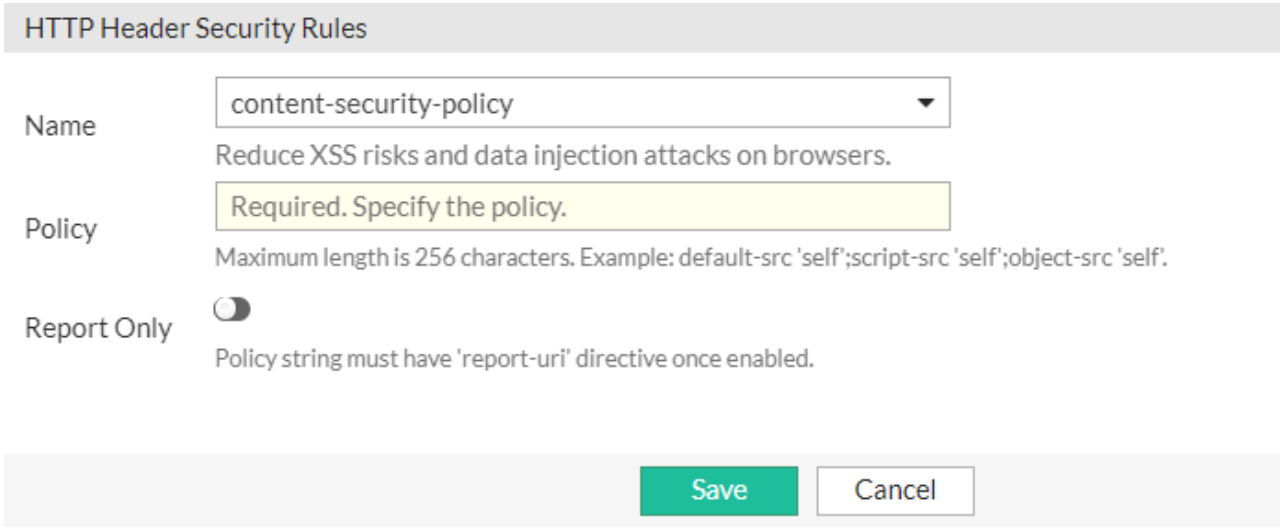


4. Configure the following HTTP Header Security settings:

Setting	Description
Name	Enter a unique HTTP Header Security policy name. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No space is allowed. Note: Once saved, the name of an HTTP Header Security policy cannot be changed.
Request Status	Enable/disable request URL match. This is disabled by default. • Enable — Responses to the request will be processed with the security headers only if the URL of a request matches the specified request URL. • Disable — All responses will be processed with the selected security header(s).
Request URL	The **Request URL** option is available if **Request Status** is **enabled**. Specify the URL used to match requests so that security headers can be applied to responses of the matched requests.

Setting	Description
Mode	Specify header operation mode for the response from the back-end server(s). • Add-Always — always add the specified header(s). • Add-Replace —add the specified header(s) if not exist, replace the value of header(s) which exist already. • Add-If-Absent — only add the specified header(s) if not exist, do nothing if the same header(s) exist The default option is **Add-Always**.

5. Click **Save**.
Once the HTTP Header Security configuration is saved, the HTTP Header Security Rules section becomes available to configure.
6. Under the **HTTP Header Security Rules** section, click **Create New** to display the configuration editor.



7. Configure the following HTTP Header Security Rules settings:

Setting	Description
Name	Select the HTTP security header name from the drop-down menu. • content-security-policy • x-content-type-options • x-frame-options • x-xss-protection • http-strict-transport-security
Policy	The **Policy** option is available if the **Name** is **content-security-policy**. Enter the header value(s) that setting restrictions on resource types and sources. For example, default-src 'self';script-src 'self';object-src 'self'.
Report Only	The **Report Only** option is available if the **Name** is **content-security-policy**.

Setting	Description
	Enabling report-only switches to “Content-Security-Policy-Report-Only” header, which accepts all directives of CSP. However, “report-only” header only monitors the violations. FortiADC will check the existing of “report-uri” directive once “report-only” selected.
Max-age	The time, in seconds, that the browser should remember that a site is only to be accessed using HTTPS. A max-age value of zero (i.e., “max-age=0”) signals the UA cease regarding the host as a Known HSTS Host, including the includeSubDomains directive (if asserted for that HSTS Host).
Include Subdomain	Optional. If enabled, rule will apply to all of the site's subdomains as well.
Preload	Google maintains an HSTS preload service: https://hstspreload.org/. By following the guidelines and successfully submitting your domain, browsers will never connect to your domain using an insecure connection. While the service is hosted by Google, all browsers have stated an intent to use (or actually started using) the preload list. Most major browsers (Chrome, Firefox, Opera, Safari, IE 11 and Edge) also have HSTS preload lists based on the Chrome list. (See the HSTS compatibility matrix.) However, it is not part of the HSTS specification and should not be treated as official.

8. Click **Save**.
   Once the HTTP Header Security Rules configuration is saved, the editor dialog closes.
9. Click **Save** to update the HTTP Header Security configuration.
   Once the HTTP Header Security policy is saved, you can reference it in a WAF Profile.

HTTP Security Header Names

Security Header	Description
Content Security Policy	A Content Security Policy (CSP), is an additional layer of security delivered via an HTTP header. This policy helps prevent attacks such as Cross Site Scripting (XSS) and other code injection attacks by defining content sources which are approved thus allowing the browser to load them. Without a CSP, the browser simply loads all files on a page without considering the source which could be harmful. This puts both the site and it’s visitors at risk of malicious activity. There are multiple directives available to website owners who want to implement a content security policy. A server may also define multiple directives within a CSP security header. For a detailed list of examples and references, visit content-security-policy.com. Additionally, you can use a tool called cspisawesome.com to easily create a CSP specific to your needs. FortiADC also provides a “report-only” flag to switch to “Content-Security-Policy-Report-Only” header, which accepts all directives of CSP, but the difference is that “report-only” header only monitor the violations. FortiADC will check the existing of “report-uri” directive once “report-only” selected.

Security Header	Description
X Content Type options	The X Content Type Options response HTTP header is a marker used by the server to indicate that the MIME types advertised in the Content-Type headers should not be changed and be followed. This helps reduce the danger of drive-by downloads and helps treat the content the proper way. There is only one directive that can be used, which is nosniff. An example of the header looks like: `x-content-type-options: nosniff`

X Frame Options	The X Frame Options header provides clickjacking protection by not allowing iframes to load on your website. It is supported by IE 8+, Chrome 4.1+, Firefox 3.6.9+, Opera 10.5+, Safari 4+.
	There are three directives available for this header: deny, sameorigin and allow-from. But “allow-from” is obsolete and no longer works in modern browers, FortiADC will notsupport it.
	On FortiADC, there are two directive options: deny and sameorigin.
	Once “deny” selected, the header looks like:
	`x-frame-options: DENY`
	Once “sameorigin” selected, the header looks like:
	`x-frame-options: SAMEORIGIN`

X XSS protection	The X XSS Protection header is designed to enable the cross-site scripting (XSS) filter built into modern web browsers. This is usually enabled by default, but using it will enforce it.
	Although these protections are largely unnecessary in modern browsers when sites implement a strong Content-Security-Policy that disables the use of inline JavaScript ('unsafe-inline'), they can still provide protections for users of older web browsers that don't yet support CSP.
	On FortiADC, this function has two modes to be choose: sanitizing-mode and block-mode. Once sanitizing-mode selected (usually default in browsers), an example looks like:
	`x-xss-protection: 1`
	Once block-mode selected, an example looks like:
	`x-xss-protection: 1; mode=block`

HTTP strict transport security	• The HTTP strict-transport-security (HSTS) header is a security enhancement that restricts web browsers to access web servers solely over HTTPS. This ensures the connection cannot be establish through an insecure HTTP connection, would helps to protect websites against protocol downgrade attacks and cookie hijacking. There are three directives for this header: max-age=<expire-time> • includeSubDomains • preload An example looks like: `strict-transport-security: max-age=31536000; includeSubDomains; preload`

Data Loss Prevention

The FortiADC Data Loss Prevention (DLP) module prevents sensitive data from leaving or entering your network by scanning for various patterns while inspecting traffic passing through the FortiADC. Data matching defined sensitive data patterns is blocked, logged, or allowed when it passes through the FortiADC.

The FortiADC DLP module has integrated with the FortiGuard DLP service, allowing FortiADC to download DLP signatures directly from FortiGuard to enrich the FortiADC DLP signature data types. For more information, see FortiGuard DLP service on page 705.

Alternatively, Data Loss Prevention can still function without the FortiGuard DLP service. You can still configure a DLP Policy to defend against data loss using only Sensitive Data Type.

The DLP module is configured based on the following components:

Component	Description
DLP Policy	Define rules for matching a sensor based on file content or an HTTP Payload, and the email protocol being used to attach files. It also allows you to choose the action to allow, log, or block the address.
Sensitive Data Type	Define the type of pattern that DLP is trying to match. For example, this can be a predefined type including credit card or US social security number (SSN), or you can use keyword, regular expression, or a hexadecimal value to match data.
DLP Sensor	Define which dictionaries to check. You can match any or all dictionaries. It can also count the number of dictionary matches to trigger the sensor. **Note**: This DLP component requires the FortiGuard DLP service to be enabled.
DLP Dictionary	Combine multiple data type entries to match all or any. When selecting a data type such as keyword, regex or hex, define the pattern that you are looking for. **Note**: This DLP component requires the FortiGuard DLP service to be enabled.

In the backend, DLP uses Hyperscan to perform a one-parse algorithm for scanning multiple patterns. This allows DLP to scale up without any performance downgrade.

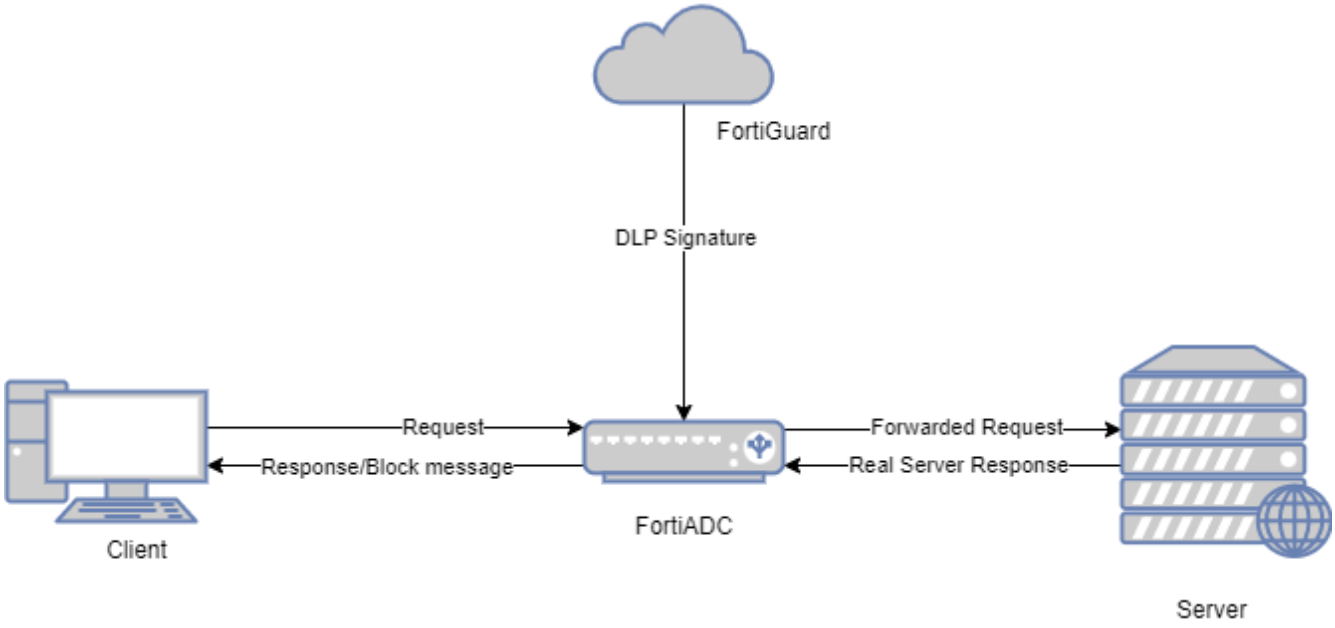
Basic configuration

To deploy Data Loss Prevention, follow the workflow below:

1. Configure the DLP Dictionary to define the collection of data type entries to use in the DLP Sensor. For details, see Configuring a DLP Dictionary object on page 715.
2. Configure the DLP Sensor to define which dictionary to check. For details, see Configuring a DLP Sensor object on page 713.
3. Configure the Sensitive Data Type to define the type of pattern that DLP is trying to match. For details, see Configuring a Sensitive Data Type object on page 710.
4. Configure the DLP Policy to define the rules for matching a sensor or sensitive data type. For details, see Configuring a DLP Policy on page 707.
5. Apply the DLP Policy to a WAF profile. For details, see Configuring a WAF Profile on page 667.

FortiGuard DLP service

The Data Loss Prevention module in FortiADC integrates with the FortiADC DLP service, allowing FortiADC to download DLP signatures directly from FortiGuard to enrich the FortiADC DLP signature data types. It uses a customizable database of more than 500 predefined data patterns and policies to simplify and expedite DLP deployment and integration into existing environments. The FortiGuard DLP service database undergoes continuous maintenance and updates to stay in sync with the latest advancements in network security intelligence.



By enabling the FortiGuard Data Loss Prevention service in FortiADC, you will have access to a database of predefined DLP patterns such as data types, dictionaries, and sensors. If the HTTP payload or files passing through FortiADC contain data that matches the patterns defined in these dictionaries, FortiADC will initiate specified actions to safeguard the data.

Name	Match Type	Comment
fg-EICAR-TEST-FILE	ANY	EICAR Test File for DLP
fg-aus-pass-dict	ALL	Australia Passport Dictionary
fg-can-bank_account-dict	ALL	Canadian Bank Account Dictionary
fg-can-bank_account-pk	ANY	Proximity keywords for Canadian Bank Account Number
fg-can-dl-dict	ANY	Canadian Driver's License Dictionary
fg-can-health_service-dict	ALL	Canadian Health Service Dictionary
fg-can-health_service-pk	ANY	Proximity keywords for Canadian Health Service Number
fg-can-natl_id-pk	ANY	Proximity keywords for Canadian SIN Card Number
fg-can-natl_id-sin-dict	ALL	Canadian SIN Card Number Dictionary
fg-can-pass-dict	ALL	Canadian Passport Dictionary
fg-can-phin-dict	ALL	Canadian Personal Health Identification Number Dictionary
fg-can-phin-pk	ANY	Proximity keywords for Canadian Personal Health Identification Number
fg-fra-pass-dict	ALL	France Passport Dictionary
fg-glb-cc-dict	ANY	Global Credit Card Dictionary
fg-glb-cc-pk	ANY	Proximity keywords for Credit Card Numbers
fg-glb-dl-pk	ANY	Proximity keywords for Driver's Licenses
fg-glb-pass-pk	ANY	Proximity keywords for Passport Number
fg-glb-swift-pk	ANY	Proximity keywords for SWIFT Codes
fg-jpn-pass-dict	ALL	Japan Passport Dictionary
fg-uk-pass-dict	ALL	UK Passport Dictionary
fg-usa-natl_id-pk	ANY	Proximity keywords for USA SSN Card Number
fg-usa-natl_id-ssn-dict	ALL	USA SSN Card Number Dictionary
fg-usa-pass-dict	ANY	USA Passport Dictionary

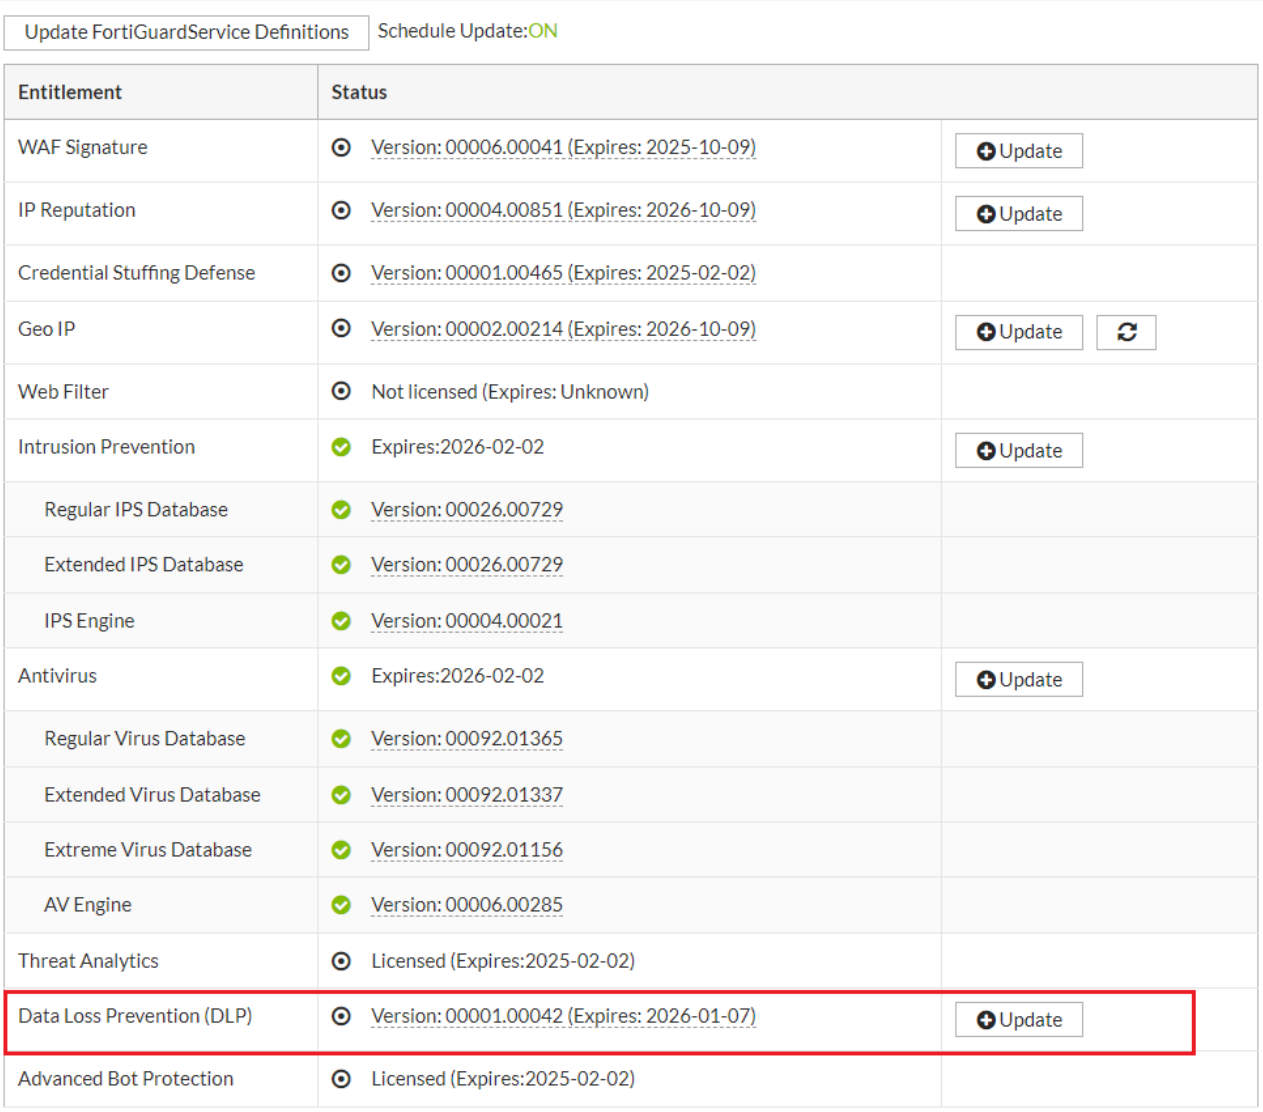
FortiGuard DLP service license

You can contact the Fortinet Sales Team to purchase a separate FortiGuard DLP service license, or a bundled license that includes the FortiGuard DLP service.

Updating the FortiGuard DLP database

1. Register your FortiGuard DLP service license at the Fortinet Customer Service & Support website: https://support.fortinet.com.
2. Log in to FortiADC. Go to **System > FortiGuard** and check the status of the FortiGuard DLP service license.

3. Typically, the system will automatically update the DLP database from FortiGuard. If it is not up-to-date, click **Update** to initiate the DLP database update.



Configuring a DLP Policy

The Data Loss Prevention (DLP) feature allows the Web Application Firewall (WAF) to prevent information leaks, damage and loss. DLP 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.

You can create a DLP Policy to match a sensor based on file content or an HTTP Payload, and the email protocol being used to attach files. It also allows you to choose the action to allow, log, or block the IP address.

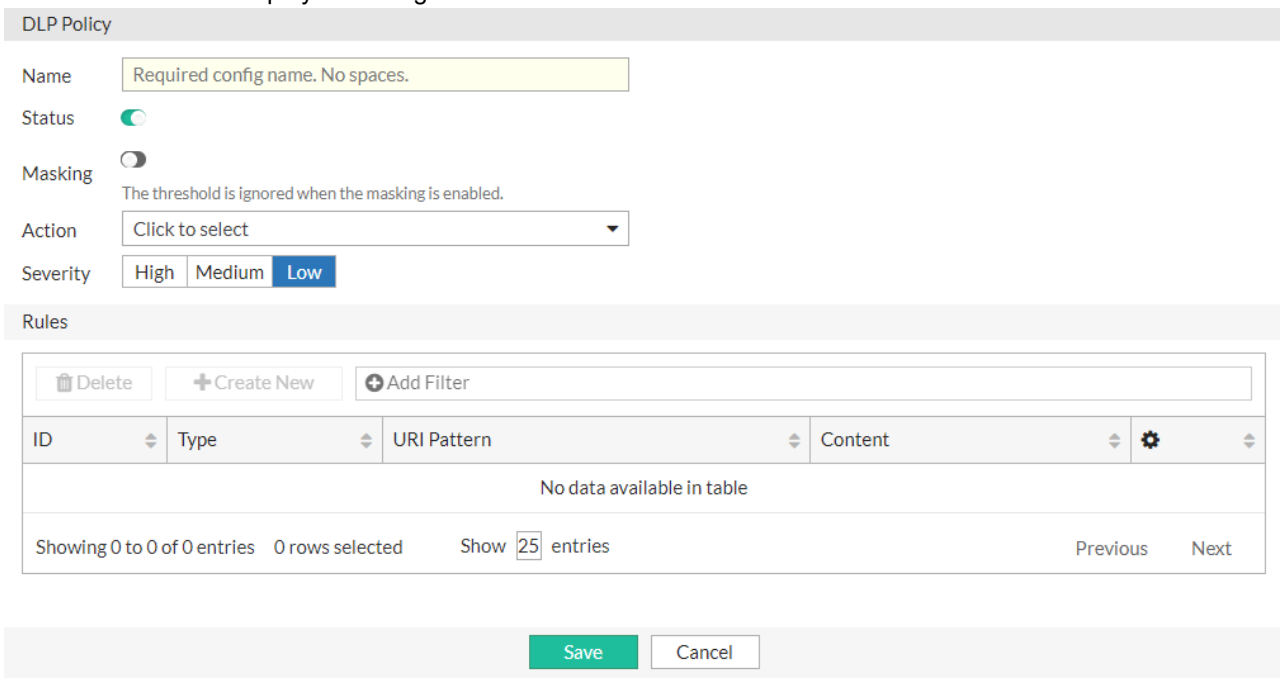
Before you begin:

- Configure a virtual server with a WAF Profile.
- Configure a DLP Dictionary object. For details, see Configuring a DLP Dictionary object on page 715.

- Configure a DLP Sensor object. For details, see Configuring a Sensitive Data Type object on page 710.
- Configure a Sensitive Data Type object. For details, see Configuring a Sensitive Data Type object on page 710.

To configure a DLP Policy:

1. Go to **Web Application Firewall > Data Loss Prevention**.
   The configuration page displays the **DLP Policy** tab.
2. Click **Create New** to display the configuration editor.

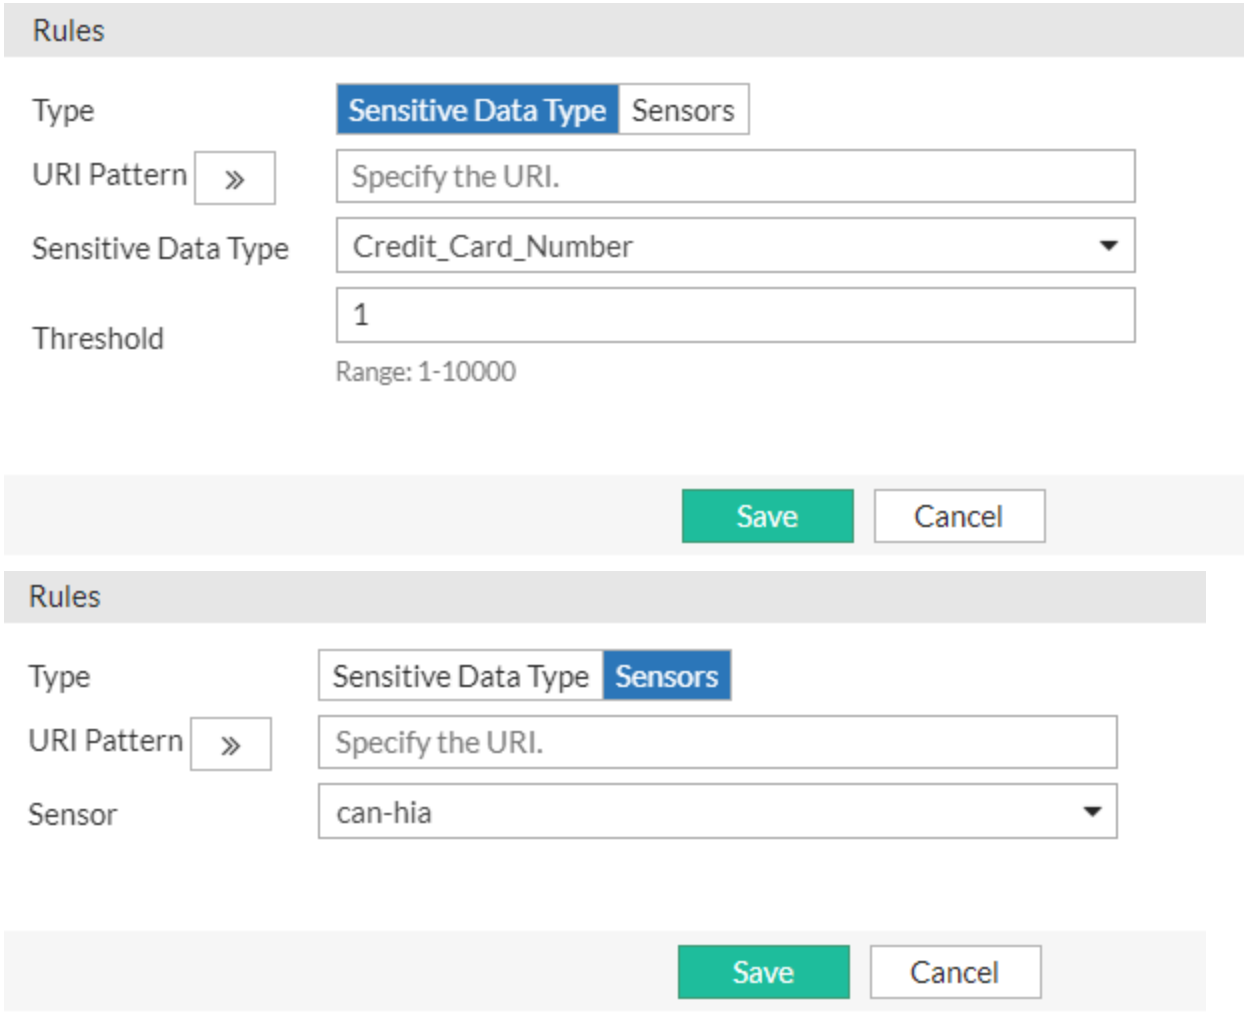


3. Configure the following DLP Policy settings:

Setting	Description
Name	Specify a name for the DLP Policy. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No spaces. The configuration name cannot be edited once it has been saved.
Status	Enable or disable this profile. Default is disable.
Masking	The **Masking** option is available if **Status** is **enabled**. 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	The **Action** option is available if **Status** is **enabled**. Select the action profile that you want to apply. See Configuring WAF Action objects on page 671. The default is Alert.
Severity	The **Severity** option is available if **Status** is **enabled**.

Setting	Description
	Set the severity level in the WAF logs for potential attacks detected by the DLP Policy. • High • Medium • Low The default option is **Low**.

4. Click **Save**.
   After the Data Leak Prevention configuration is saved, the Rule section becomes available to configure.
5. Under the **Rule** section, click **Create New** to display the Data Loss Prevention Rule configuration editor.



6. Configure the following Data Loss Prevention Rule settings:

Setting	Description
Type	Select the DLP data type to match: • Sensitive Data Type • Sensors

Setting	Description
	The default option is **Sensitive Data Type**.
URI Pattern	Specify the URI Pattern in the Data Loss Prevention rules. Scanning and receiving an empty value means this rule is not working.
Sensitive Data Type	The **Sensitive Data Type** option is available if the **Type** is **Sensitive Data Type**. Select the Sensitive Data Type you want to apply. See Configuring a Sensitive Data Type object on page 710.
Threshold	The **Threshold** option is available if the **Type** is **Sensors**. Set a threshold for the Data Loss Prevention rule. The rule will not take effect until the target data exceeds the threshold's specified value. Range 1-10000. Default is 1. This will not work if Masking is enabled.
Sensor	The **Sensor** option is available if the **Type** is **Sensors**. Select the DLP Sensor you want to apply. See Configuring a DLP Sensor object on page 713.

7. Click **Save**.
   Once the Data Loss Prevention Rule configuration is saved, the editor dialog closes.
   You can add more rule items, with a maximum of 256 rules, but detection will stop after matching as many as 8 rules.
8. Click **Save** to update the DLP Policy configuration.
   Once the DLP Policy is saved, you can reference it in a WAF Profile.

Configuring a Sensitive Data Type object

A Sensitive Data Type object is referenced as part of the Data Loss Prevention (DLP) policy to prevent information, damage and loss by specifying strings as sensitive data.

Predefined Sensitive Data Type objects

You can use the following predefined Sensitive Data Type objects or clone to use as a template.

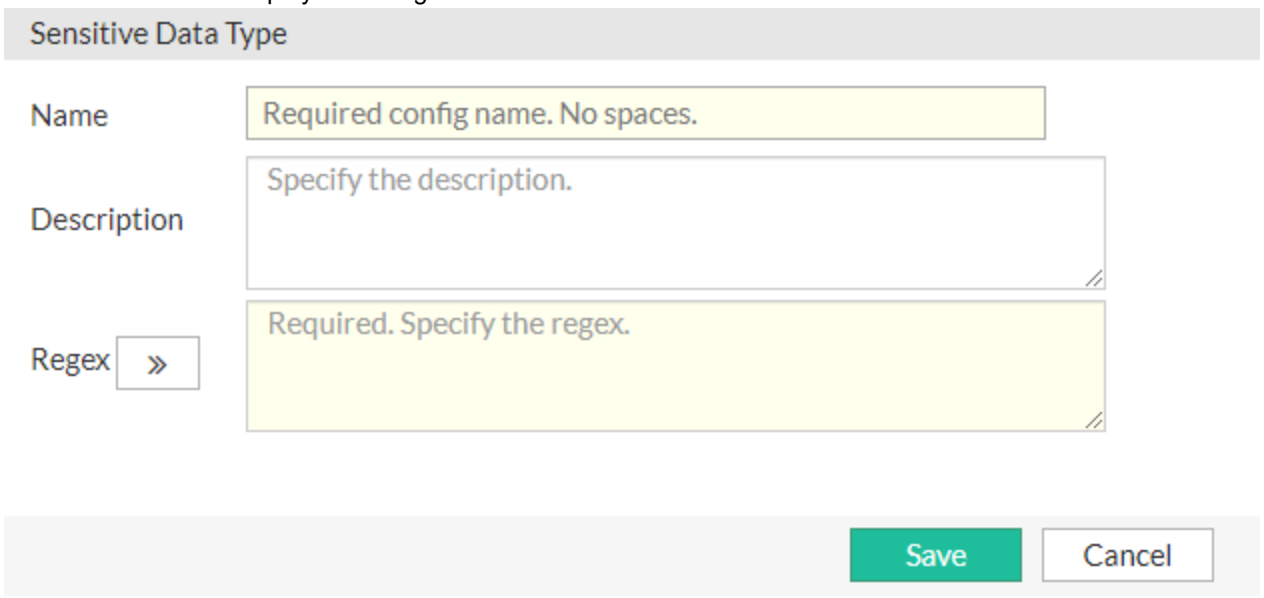
Predefined Sensitive Data Type objects	Description
Credit_Card_Number	For credit card numbers from MC, Visa, Amex, Diners/CarteBlanche, Discover/Novus, Enroute, and JCB. Matches 341-1111-1111-1111 | 5431-1111-1111-1111 | 30569309025904 Non-Matches 30-5693-0902-5904 | 5631-1111-1111-1111 | 31169309025904.
US_Social_Security_Number	This regex validates U.S. social security numbers, within the range of numbers that have been currently allocated. Matches 078-05-1120 | 078 05 1120 Non-Matches 987-65-4320 | 000-00-0000 | (555) 555-5555.
Email	This regex validates email address. Matches example@fortinet.com Non-Matches @fortinet.com.

Predefined Sensitive Data Type objects	Description
URL	This regex validates URL. Matches http://www.fortinet.com | https://127.0.0.1/path/example.php?name=test1 | ftp://user:pass@example.com:123 Non-Matches /fortinet.com
Numbers	This regex validates numbers. Matches 65535 Non-Matches a123.
Strings	This regex validates a string. Matches abc Non-Matches abc123.
Date/Time	This regex validates email address. Matches 29/02/1972 | 5-9-98 | 10-11-2002 | February 29, 2004 | 12:15 | 10:26:59 | 22:01:15 Non-Matches 32/12/2019.
IP Address	This regex validates IPv4 or IPv6 address. Matches 127.0.0.1 | FEDC:BA98:7654:3210:FEDC:BA98:7654:3210 | ::FFFF:129.144.52.38 Non-Matches 256.0.0.1 | FEDC:BA98:7654:3210 | ::
GUID	This regex validates a globally unique identifier. Matches 2064d355-c0b9-41d8-9ef7-9d8b26524751 | 2064D355-C0B9-41D8-9EF7-9D8B26524751 Non-Matches 2064D355.
US Phone	This regex validates a US phone number WITH area code. It is written to all users to enter whatever delimiters they want or no delimiters at all. Matches 111-222-3333 | 111.222.3333 | (111) 222-3333| 1112223333 Non-Matches + 41 111-222-3333 .
US ZIP Code	This regex validates US zip codes. Matches all zip codes of exactly 5 digits except 00000. Optionally, matches zip5+zip4 where zip5 is exactly 5 digits, zip4 is exactly 4 digits, and zip5 and zip4 are, optionally, separated by a single space or hyphen. Captures zip5 and zip4 to named groups to facilitate program manipulation. Matches 12345 | 123456789 | 12345-6789 Non-Matches 123456.
US State Name and Abbrev.	This regex validates 50 US States's Name and Abbrev, case insensitive. Matches California | NewYork | North Carolina | AL.
US Street Address	This regex validates a US Street Address. Matches 123 Lincoln Avenu | 123 West Main St | 12345 Via De La Rosa Non-Matches Lincoln Avenu.
UK Vehicle Registration	This regex validates a UK vehicle registration system currently in use (as defined by the DVLA and put into effect from September 2001, and therefore does not allow registrations prior to this date). Matches AB51DVL | AB 51 DVL Non-Matches AB-51-DVL.
UK Bank Sort Code	This regex validates the format of a UK bank sort code. Matches 20-40-36 | 50-25-48 | 45-85-66 Non-Matches 204036.
Post Office Box	This regex validates a Post Office Box. Matches P. O. Box | p.o. box | PO Box | po box Non-Matches office box.
Chinese ID card	This regex validates a Chinese ID card number. Matches 2064d355-c0b9-41d8-9ef7-9d8b26524751 | 2064D355-C0B9-41D8-9EF7-9D8B26524751 Non-Matches 2064D355.
Chinese phone	This regex validates a Chinese telphone number. Matches 86 13512341234 | +86

Predefined Sensitive Data Type objects	Description
	15812341234 | 86 13612341234 Non-Matches 14012341234.
Australian Phone	This regex validates a Australian telephone number, most Australian telephone numbers including 13, 1300, 1800, 1900, std and international +61- format numbers. It allows optional spaces, dashes and brackets in most cases. Matches 1300 123 123 | 1300123123 | +61212341234 | (02) 1234-1234 | 02 1234 1234 Non-Matches 1400123123.
Canadian Postal Code	Canadian Postal Code format is (A1A 1X1) or (a1a 1x1). Its made up of two parts. Forward Sortation Area (FSA) and Local Delivery Unit (LDU). Read more on wikipedia. The letters D, F, I, O, Q, or U are not used on postal Code. Matches M1R 4B0 | L0R 1B1 | L0R1B9 Non-Matches MDR 4B0.

To configure a Sensitive Data Type object:

1. Go to **Web Application Firewall > Data Loss Prevention**.
2. Click the **Sensitive Data Type** tab.
3. Click **Create New** to display the configuration editor.



4. Configure the following Sensitive Data Type settings:

Setting	Description
Name	Specify a name for the Sensitive Data Type object. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No spaces. The configuration name cannot be edited once it has been saved.
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

Setting	Description
	predefined regex strings named Credit_Card_Number and US_Social_Security_Number.

5. Click **Save**.
After the Sensitive Data Type configuration is saved, you can reference this object in the Data Loss Prevention rule.

Configuring a DLP Sensor object

A DLP Sensor defines which dictionaries to check. You can match any dictionary or all dictionaries. It can also count the number of dictionary matches to trigger the sensor.

Before you begin:

- You must have a valid FortiGuard DLP service license and have enabled the service on FortiADC. For details, see FortiGuard DLP service on page 705.
- Configure a DLP Dictionary object. For details, see Configuring a DLP Dictionary object on page 715.

Predefined DLP Sensor objects

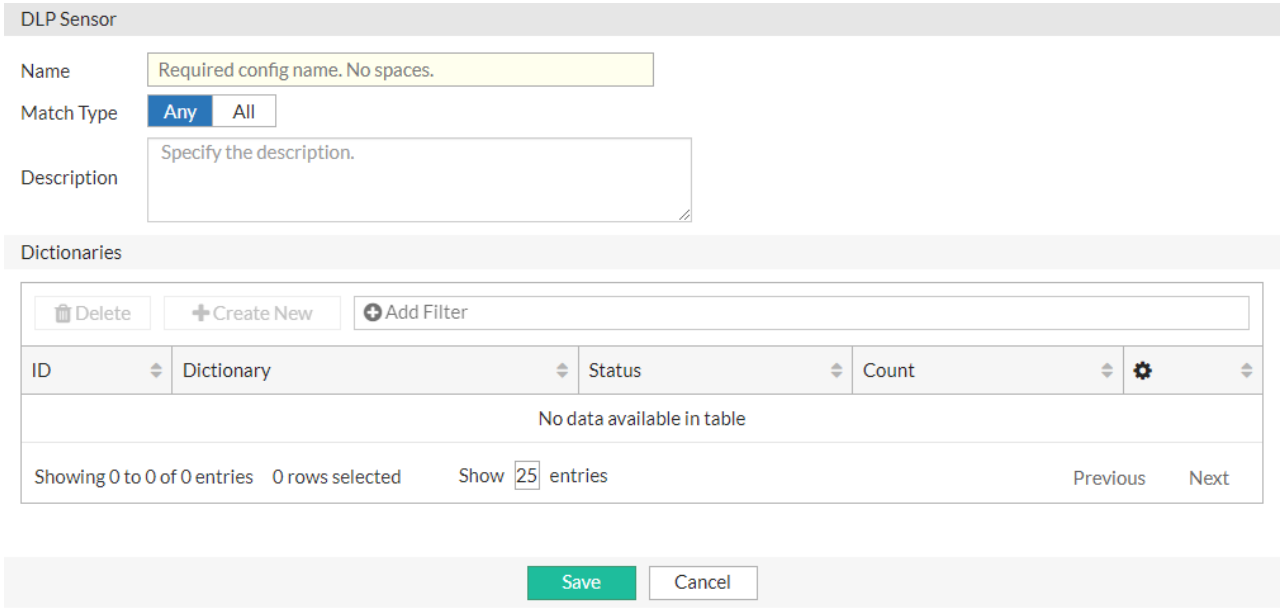
You can use the following predefined DLP Sensor objects in Data Loss Prevention rules.

Predefined DLP Sensor object	Match Type	Description	Dictionaries
can-hia	Any	Canadian Health Information Act (HIA) Sensor	• can-pass-dict • can-natl_id-sin-dict • can-phin-dict • can-health_service-dict
can-pii	Any	Canadian Personal Identifiable Information (PII) Sensor	• can-dl-dict • can-natl_id-sin-dict • can-pass-dict • can-health_service-dict • can-bank_account-dict • can-phin-dict
source_code	Any	Source Code Sensor	• source_code-python • source_code-c • source_code-java

To configure a DLP Sensor:

1. Go to **Web Application Firewall > Data Loss Prevention**.
2. Click the **DLP Sensor** tab.

3. Click **Create New** to display the configuration editor.

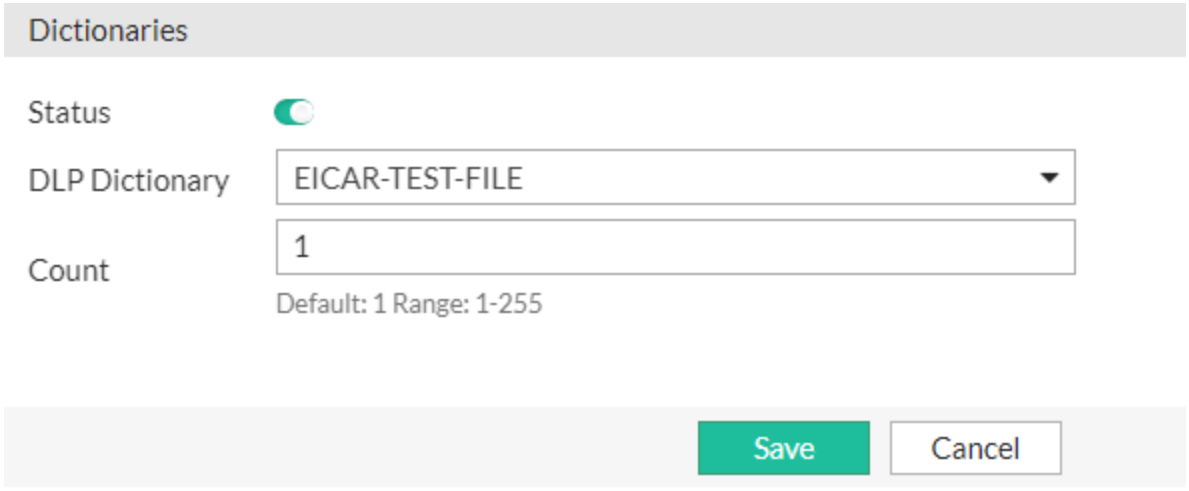


4. Configure the following DLP Sensor settings:

Setting	Description
Name	Specify a name for the DLP Sensor object. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No spaces. The configuration name cannot be edited once it has been saved.
Match Type	Select the match type: • Any — Data meeting the criteria specified by any one of the dictionaries will be identified as a match. • All — Data meeting the criteria specified by all dictionaries will be identified as a match. The default is **Any**.
Description	Comments about this DLP Sensor object.

5. Click **Save**.
After the DLP Sensor configuration is saved, the Dictionaries section becomes available to configure.

6. Under the **Dictionaries** section, click **Create New** to display the configuration editor.



7. Configure the following Dictionaries settings for the DLP Sensor:

Setting	Description
Status	Enable the **Status** if you intend to apply this sensor.
DLP Dictionary	Select a DLP Dictionary object from the drop-down menu.
Count	Specify the occurrence threshold for the dictionary match. The sensor will be triggered when the dictionary match reaches the specified number of times. Default: 1 Range: 1-255. For example, if the dictionary applies to credit card numbers and the count is set to 4, the sensor will be triggered when credit card number occurs four times in the HTTP request or response.

8. Click **Save**.
   Once the Dictionaries configuration is saved, the editor dialog closes.
9. Click **Save** to update the DLP Sensor configuration.
   Once the DLP Sensor is saved, you can reference it in a DLP Policy.

Configuring a DLP Dictionary object

A DLP dictionary defines the patterns of data. The term "pattern" denotes a set of attributes specific to a given data type. For example, credit card numbers constitute numeric data that follow either the 14-digit or 16-digit patterns associated with credit cards. If the data adheres to these patterns, FortiADC will identify it as a match.

Before you begin:

- You must have a valid FortiGuard DLP service license and have enabled the service on FortiADC. For details, see FortiGuard DLP service on page 705.

Predefined DLP Sensor objects

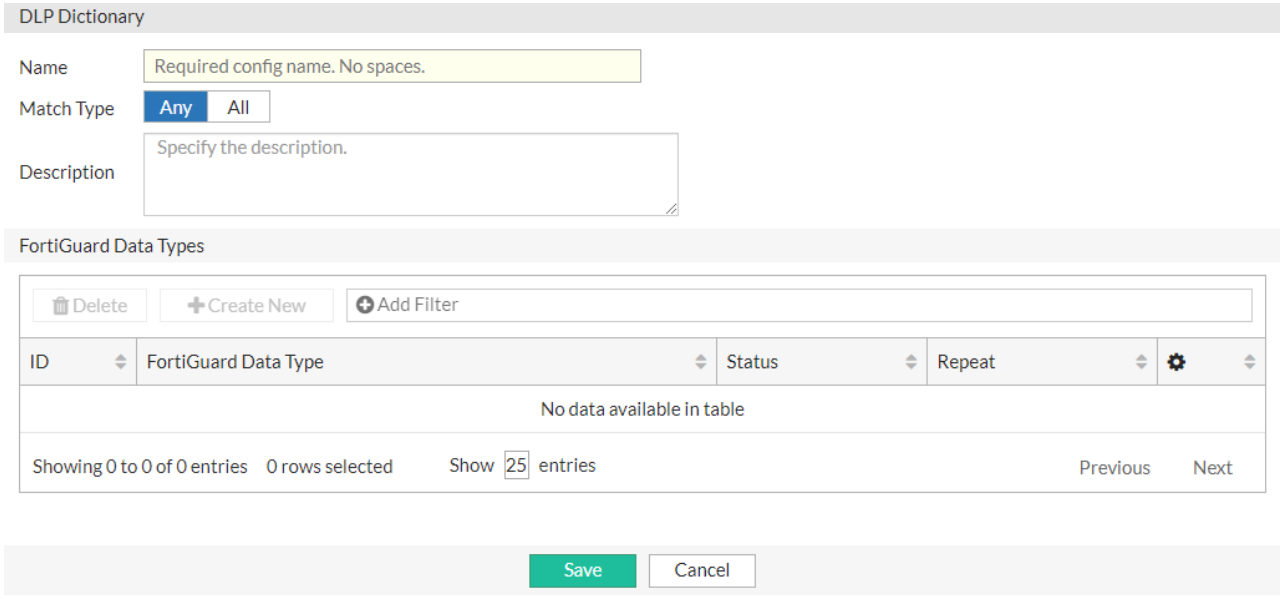
You can use the following predefined DLP Dictionary objects in Data Loss Prevention rules.

Predefined DLP Dictionary object	Match Type	Description
EICAR-TEST-FILE	Any	EICAR Test File for DLP
can-natl_id-pk	Any	
can-natl_id-sin-dict	Any	Canadian SIN Card Number Dictionary
glb-pass-pk	Any	
can-pass-dict	Any	Canadian Passport Dictionary
usa-pass-dict	Any	USA Passport Dictionary
uk-pass-dict	Any	UK Passport Dictionary
aus-pass-dict	Any	Australia Passport Dictionary
fra-pass-dict	Any	France Passport Dictionary
jpn-pass-dict	Any	Japan Passport Dictionary
can-health_service-pk	Any	
can-phin-pk	Any	
can-phin-dict	Any	Canadian Personal Health Identification Number Dictionary
can-health_service-dict	Any	Canadian Health Service Dictionary
glb-cc-pk	Any	
glb-cc-dict	Any	Global Credit Card Dictionary
usa-natl_id-pk	Any	
glb-dl-pk	Any	
can-dl-dict	Any	Canadian Driver's License Dictionary
can-bank_account-pk	Any	
can-bank_account-dict	Any	Canadian Bank Account Dictionary
usa-natl_id-ssn-dict	Any	USA SSN Card Number Dictionary
glb-swift-pk	Any	
source_code-python	Any	Python Source Code Dictionary
source_code-c	Any	C Source Code Dictionary
source_code-java	Any	Java Source Code Dictionary

To configure a DLP Dictionary object:

1. Go to **Web Application Firewall > Data Loss Prevention**.
2. Click the **DLP Dictionary** tab.

3. Click **Create New** to display the configuration editor.

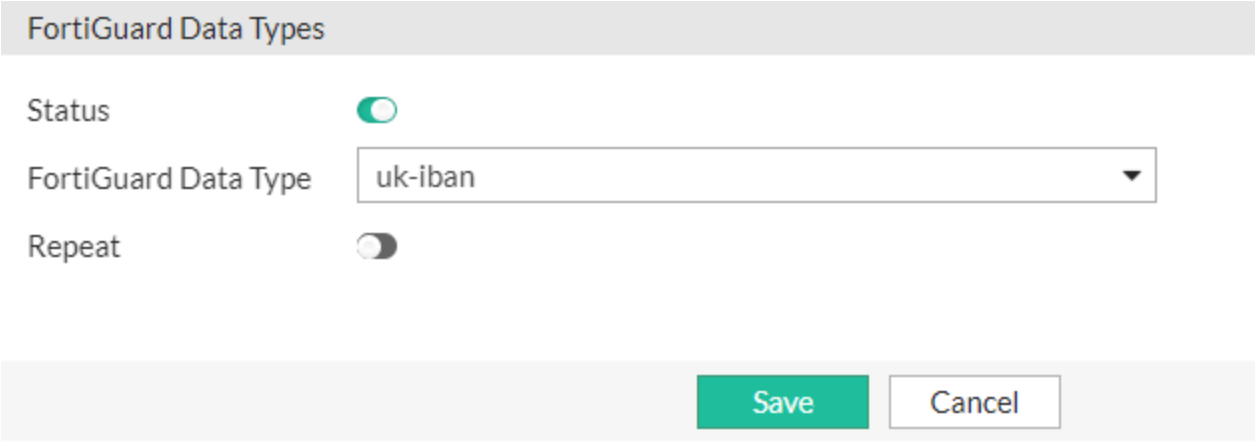


4. Configure the following DLP Dictionary settings:

Setting	Description
Name	Specify a name for the DLP Dictionary object. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No spaces. The configuration name cannot be edited once it has been saved.
Match Type	Select the match type: • Any — Data meeting the criteria specified by any one of the dictionary entries will be identified as a match. • All — Data meeting the criteria specified by all dictionary entries will be identified as a match. The default is **Any**.
Description	Comments about this DLP Dictionary object.

5. Click **Save**.
After the DLP Dictionary configuration is saved, the FortiGuard Data Types section becomes available to configure.

6. Under the **FortiGuard Data Types** section, click **Create New** to display the configuration editor.



7. Configure the following FortiGuard Data Types settings for the DLP Dictionary:

Setting	Description
Status	Enable the **Status** if you intend to apply this data type.
FortiGuard Data Type	Select a FortiGuard Data Type from the drop-down menu.
Repeat	Enable this option if you want to match data exclusively when it appears multiple times. With this option enabled, you can specify the times of occurrence in the DLP Sensor settings.

8. Click **Save**.
Once the FortiGuard Data Types configuration is saved, the editor dialog closes.
9. Click **Save** to update the DLP Sensor configuration.
Once the DLP Dictionary is saved, you can reference it in a DLP Sensor.

Input Validation

The **Web Application Firewall > Input Validation** sub-menu contains features and configurations for Input Validation that enables you to create rules to prevent suspicious HTTP requests by verifying the user input from scan points like URL parameter, HTML form, hidden fields, and upload file.

- Configuring an Input Validation policy on page 718
- Configuring a Parameter Validation rule on page 719
- Configuring a Hidden Field rule on page 721
- Configuring a File Restriction rule on page 722

Configuring an Input Validation policy

An Input Validation policy can prevent suspicious HTTP requests. This function will verify the user input from scan points such as URL parameter, HTML form, hidden fields, and uploaded files. If the input format is incorrect or FortiADC detects

other attacks, the request will be blocked.

Input Validation policies consist of the following components that must have already been configured if they are to be included in the policy:

- Parameter Validation Rule. For details, see Configuring a Parameter Validation rule on page 719.
- Hidden Field Rule. For details, see Configuring a Hidden Field rule on page 721.
- File Restriction Rule. For details, see Configuring a File Restriction rule on page 722.

To configure an Input Validation policy:

1. Go to **Web Application Firewall > Input Validation**.
   The configuration page displays the **Input Validation Policy** tab.
2. Click **Create New** to display the configuration editor.
3. Enter a unique Input Validation policy name. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No space is allowed. The name of the policy cannot be changed after saving initially.
4. Click **Save** to create the Input Validation Policy.
   Once the Input Validation Policy has been saved, the **Parameter Validation Rule**, **Hidden Field Rule**, and **File Restriction Rule** sections are available to configure.
5. Under the **Parameter Validation Rule** section, add Parameter Validation rules to the Input Validation Policy.
   a. Click **Create New** to display the configuration editor.
   b. Select an existing Parameter Validation Rule or create a new configuration.
   c. Click **Save** to submit and exit out of the configuration editor.
6. Under the **Hidden Field Rule** section, add Hidden Field rules to the Input Validation Policy.
   a. Click **Create New** to display the configuration editor.
   b. Select an existing Hidden Field Rule or create a new configuration.
   c. Click **Save** to submit and exit out of the configuration editor.
7. Under the **File Restriction Rule** section, add File Restriction rules to the Input Validation Policy.
   a. Click **Create New** to display the configuration editor.
   b. Select an existing File Restriction Rule or create a new configuration.
   c. Click **Save** to submit and exit out of the configuration editor.
8. Click **Save** to update the Input Validation policy.

After you have configured your Input Validation policy, you can reference it in a WAF Profile.

Configuring a Parameter Validation rule

Define the Parameter Validation rule for Input Validation to determine whether or not parameters are required, and their maximum allowed length. Inputs are typically the `<input>` tags in an HTML form, where the input rules are applicable for visible inputs only, such as buttons and text areas.

The Parameter Validation rule function can do the following:

- Check the HOST by simple string or regular expression matching.
- Check the URL by simple string or regular expression matching.
- 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.

To configure a Parameter Validation rule:

1. Go to **Web Application Firewall > Input Validation**.
2. Click the **Parameter Validation** tab.
3. Click **Create New** to display the configuration editor.
4. Configure the following Parameter Validation Rule settings:

Setting	Description
Name	Enter a unique name for the Parameter Validation Rule. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No space is allowed. Note: Once saved, the name of a Parameter Validation Rule 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	The **Host** option is available if **Host Status** is **enabled**. 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.
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 the action profile that you want to apply. See Configuring WAF Action objects on page 671. 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**.

5. Click **Save**.
   Once the Parameter Validation Rule is saved, the Parameter Validation Rule Element section can be configured.
6. Under the **Parameter Validation Rule Element** section, click **Create New** to display the configuration editor.
7. Configure the following Parameter Validation Rule Element settings:

Setting	Description
Name	Enter a unique name for the Parameter Validation Rule Element. 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.
Use Type Check	Enable/disable to check the data type.

Setting	Description
Argument Type	The **Argument Type** option is available if **Use Type Check** is **enabled**. Select to use predefined data type or customized regular expression.
Data Type	The **Data Type** option is available if **Use Type Check** is **enabled** and **Argument Type** is **Data Type**. Match the string by the predefined data type.
Regular Expression	The **Regular Expression** option is available if **Use Type Check** is **enabled** and **Argument Type** is **Regular Expression**. Match the string by regular expression.

8. Click **Save** to save the Parameter Validation Rule Element configuration and exit from the dialog.
9. Click **Save** to update the Parameter Validation Rule configuration.

After the Parameter Validation rule has been saved, you can include it in an Input Validation Policy.

Configuring a Hidden Field rule

Define the Hidden Field rule for Input Validation to check for hidden parameters from `<input type="hidden">` HTML tags. These hidden parameters are 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.

The Hidden Field rule function can do the following:

- Check the HOST by simple string or regular expression matching.
- Check the URL by simple string or regular expression matching.
- Match the configuration of the fetched URL.

If the conditions are successfully matched, it will execute the specified action.

To configure a Hidden Field rule:

1. Go to **Web Application Firewall > Input Validation**.
2. Click the **Hidden Field** tab.
3. Click **Create New** to display the configuration editor.
4. Configure the following Hidden Field settings:

Setting	Description
Name	Enter a unique name for the Hidden Fields rule. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No space is allowed. Note: Once saved, the name of a Hidden Field rule 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	The **Host** option is available if **Host Status** is **enabled**. 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.

Setting	Description
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 the action profile that you want to apply. See Configuring WAF Action objects on page 671. 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**.

5. Click **Save**.
   Once the Hidden Field configuration is saved, the Post URL and Hidden Fields sections can be configured.
6. Under the **Post URL** section, click **Create New** to display the configuration editor.
7. In the **URL** field, specify the Post URL on which the hidden fields function can work. Click **Save** and exit the configuration dialog.
8. Under the **Hidden Fields** section, click **Create New** to display the configuration editor.



To apply this feature, you must enable Session Management in your protection profile.

9. In the Name field, enter a unique Hidden Fields name. It must match the value of the name in the input type of the HTML request. Click **Save** and exit the configuration dialog.
10. Click **Save** to update the Hidden Field configuration.

After the Hidden Field rule has been saved, you can include it in an Input Validation Policy.

Configuring a File Restriction rule

Configure the **File Restriction Rule** for Input Validation to enforce restrictions on file uploads based on file type and size. The rule performs the following checks:

- Matches the HOST using either simple string or regular expression patterns.
- Matches the URL using either simple string or regular expression patterns.
- Analyzes the uploaded file type and size based on HTTP MIME (Multipurpose Internet Mail Extensions), magic numbers (file signatures), and file extension.

When the defined conditions are met, the system enforces the specified action. Additionally, if multiple files of different types are uploaded in a single HTTP transaction, and one file type violates the rule, the entire transaction will be rejected, resulting in all files being blocked.

File Type Identification in FortiADC

FortiADC employs two methods for file type identification: file type signatures and suffix matching.

File Type Signatures:

FortiADC examines specific attributes of a file to determine its content type by detecting unique signatures, or magic codes, associated with predefined file types based on MIME types and magic numbers (file signatures). If the detected file type matches one specified in the file restriction rule, the system enforces the corresponding action. Supported file type categories include **Audio Files**, **Compressed Files**, **Picture Files**, **Text Files**, and **Video Files**.

Suffix Matching:

FortiADC can also identify files based on their suffix (extension). If the file suffix matches an entry under the **Whole Suffix Files** category in the file restriction rule, the associated action is triggered.

When both file type signature and suffix matching are configured, suffix matching takes precedence. If the file suffix matches, the file restriction rule is applied immediately. If the suffix does not match but the file signature does, the file restriction rule will still be enforced.

For the full list of the supported file types, see Supported File Types on page 724.

To configure a File Restriction rule:

1. Go to **Web Application Firewall > Input Validation**.
2. Click the **File Restriction** tab.
3. Click **Create New** to display the configuration editor.
4. Configure the following File Restriction settings:

Setting	Description
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	The **Host** option is available if **Host Status** is **enabled**. 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.
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 the action profile that you want to apply. See Configuring WAF Action objects on page 671. The default value is Alert.

Setting	Description
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. The default value is 0, and the range is 0-102400 KB.

5. Click **Save**.
   Once the File Restriction configuration is saved, the Upload File Type section can be configured.
6. Under the **Upload File Type** section, click **Create New** to display the configuration editor.
7. In the **File Type** field, select the supported file types for the uploaded file.
8. Click **Save** to update the File Restriction configuration.

After the File Restriction rule has been saved, you can include it in an Input Validation Policy.

Supported File Types

Category	File Type
Audio Files	MP3 MIDI WAVE AVI Apple CoreAudio (.caf) Microsoft Advanced Streaming (.asf) Real Audio File (.ra) Apple Lossless Audio (.m4a) Digital Speech Standard (.dss) Advanced Audio Coding (.aac)
Compressed Files	RAR ZIP TAR 7-ZIP Debian Package Microsoft Cabinet File Unix Archiver File (.ar) Installshield Cabinet Archive Data AIN Archive Data (.ain)

Category	File Type
	BZIP2 Archive (.bz2) WinZIP ZIPX Archive (ZIPx) Gzipped Tape Archive (TGZ) Extensible Archive (XAR)
Picture Files	GIF JPG BMP PNG TIFF/TIF Windows Metafile Format (.wmf) Corel Draw Picture Windows Icon Microsoft Document Image (.mdi) Windows Enhanced Metafile (.emf) Photoshop Image File (.psd) JPEG-2000 Image File Format (.jp2) Multipage PCX Bitmap File (.dcx)
Text Files	PDF XML CHM EXE RTF Windows Help File (.hlp) Windows Mobile Note (.pwi) Windows Registry Text (.reg) SQL Server 2000 Database (.mdf) Java Archive (.jar) Windows Printer Spool File (.shd) Windows Shortcut File (.lnk) Quark Express Document (.qxd) Windows MS Info File (.mof) Microsoft Access Database (.mdb) SPSS Data (.sav) XPS Word (.docx) Word Macro-Enabled (.docm) Word Template (.dotx) Word Macro-Enabled Template (.dotm) Excel (.xlsx)

Category	File Type
	Excel Macro-Enabled (.xlsm) Excel Template (.xltx) Excel Macro-Enabled Template (.xltm) Excel Add-In (.xlam) PPT (.pptx) PPT Macro-Enabled (.pptm) PPT Template (.potx) PPT Macro-Enabled Template (.potm) PPT Add-In (.ppam) PPT Show (.ppsx) PPT Macro-Enabled Show (.ppsm) Visio Drawing (.vsdx) Visio Macro-Enabled Drawing (.vsdm) Visio Stencil (.vssx) Visio Macro-Enabled Stencil (.vssm) Visio Template (.vstx) Visio Macro-Enabled Template (.vstm) VMware Virtual Disk File (.vmdk) RedHat Package Manager file (.rpm) Lotus WordPro document (.lwp) Adobe Encapsulated PostScript file (.eps) Lotus 1-2-3 spreadsheet (.wk) SkinCrafter skin file (.skf) Nero CD Compilation (.nri) TXT Microsoft Office Word (.doc) Microsoft Office Excel (.xls) Microsoft Office PowerPoint (.ppt) Hancom Office Hanword (.hwp) Electronic Publication (.epub) Dynamic link library (.dll) SYS File (.sys) COM File (.com) CMD File (.cmd) Binary File (.bin) Scalable Vector Graphics (.svg) PHP (.php) Perl (.pl) Python (.py) Ruby (.rb)

Category	File Type
	Microsoft Software Installer (.msi) Batch File (.bat) Privacy Enhanced Mail (.pem) x509 certificate (.cer) x509 certificate (.crt)
Video Files	Real Media File (.rm) MPEG v4 3GPP Macromedia Flash Windows Animated Cursor DVD Video Movie File (.vob) MKV
Whole Suffix Files	TXT (.txt) ZIP (.zip) 7-ZIP (.7z) Debian Package (.pkg) Unix Archiver File (.ar) AIN Archive Data (.ain) BZIP2 Archive (.bz2) Gzipped Tape Archive (.tgz) Word (.docx) Word Macro-Enabled (.docm) Word Template (.dotx) Word Macro-Enabled Template (.dotm) Excel (.xlsx) Excel Macro-Enabled (.xlsm) Excel Template (.xltx) Excel Macro-Enabled Template (.xltm) Excel Add-In (.xlam) PPT (.pptx) PPT Macro-Enabled (.pptm) PPT Template (.potx) PPT Macro-Enabled Template (.potm) PPT Add-In (.ppam) PPT Show (.ppsx) PPT Macro-Enabled Show (.ppsm) Visio Drawing (.vsdx) Visio Macro-Enabled Drawing (.vsdm) Visio Stencil (.vssx)

Category	File Type
	Visio Macro-Enabled Stencil (.vssm) Visio Template (.vstx) Visio Macro-Enabled Template (.vstm) PDF (.pdf) XML (.xml) EXE (.exe) Rich Text Format (.rtf) Windows Help File (.hlp) Windows Mobile Note (.pwi) Windows Registry Text (.reg) SQL Server 2000 Database (.mdf) Java Archive (.jar) Windows Printer Spool File (.shd) Window Shortcut File (.lnk) Quark Express Document (.qxd) Windows MS Info File (.mof) Microsoft Access Database (.mdb) SPSS Data (.sav) RedHat Package Manager file (.rpm) VMware Virtual Disk File (.vmdk) Lotus WordPro document (.lwp) Adobe Encapsulated PostScript file (.eps) Lotus 1-2-3 spreadsheet (.wk) SkinCrafter skin file (.skf) Nero CD Compilation (.nri) Microsoft Office Word (.doc) Microsoft Office Excel (.xls) Microsoft Office PowerPoint (.ppt) Hancom Office Hanword (.hwp) PHP (.php) JSP (.jsp) ASPX (.aspx) GIF (.gif) JPG (.jpg) BMP (.bmp) PNG (.png) Microsoft Metafile Format (.wmf) Windows Icon (.icon) Microsoft Document Image (.mdi) Windows Enhanced Metafile (.emf)

Category	File Type
	Photoshop Image File (.psd) JPEG-2000 Image File Format (.jp2) Multipage PCX Bitmap File (.dcx) SQL (.sql) Cascading Style Sheets (.css) ASP (.asp) CSV (.csv) PHP3 (.php3) PHTML (.phtml) Workflow File (.workflow) Scalable Vector Graphics (.svg) MSG (.msg) OpenDocument Spreadsheet (.ods) OpenDocument Text (.odt) Privacy-Enhanced Mail (.pem) Electronic Publication (.epub) Advanced Audio Coding (.aac) Personal Information Exchange (.pfx) Personal Information Exchange (.p12) Microsoft Software Installer (.msi) Batch File (.bat) Dynamic link library (.dll) SYS File (.sys) COM File (.com) CMD File (.cmd) Binary File (.bin) Tab-Separated Values (.tsv) Android Package Kit (.apk) Compressed package file (.xapk) APK set archive (.apks) APKMirror Bundle file (.apkm) Distinguished Encoding Rules (.der)

Access Protection

The **Web Application Firewall > Access Protection** sub-menu contains features and configurations for Access Protection, that include URL Protection, Brute Force Attack Detection, Credential Stuffing Defense.

- Configuring a URL Protection policy on page 730
- Configuring a Brute Force Attack Detection policy on page 731

- Configuring a Credential Stuffing Defense policy on page 732

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 > Access Protection**.
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 730.
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	Select the action profile that you want to apply. See Configuring WAF Action objects on page 671. 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	Select the action profile that you want to apply. See Configuring WAF Action objects on page 671. The default 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.
Exception Name	Select an exception configuration object. Exceptions identify specific hosts or URL patterns that are not subject to processing by this rule.

Configuring a Brute Force Attack Detection policy

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:

1. Go to **Web Application Firewall > Access Protection**.
2. Click the **Brute Force Attack Detection** tab.
3. Click **Create New** to display the configuration editor.
4. 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	Select the action profile that you want to apply. See Configuring WAF Action objects on page 671. 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 a Credential Stuffing Defense policy

Credential Stuffing Defense identifies login attempts using username and password that have been compromised using an always up-to-date feed of stolen credentials. Administrators can configure their supported devices to take various actions if a suspicious login is used including logging, alerts, and blocking.

To configure a Credential Stuffing Defense policy:

1. Go to **Web Application Firewall > Access Protection**.
2. Click the **Credential Stuffing Defense** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the Credential Stuffing Defense configuration.
5. Save the configuration.

Predefined Rules	Description
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 or disable this profile. Default is disable.
Action	Select the action profile that you want to apply. See Configuring WAF Action objects on page 671. The default is Alert.

Predefined Rules	Description
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 use High or Medium.

Note: FortiADC has no built-in Credential Stuffing Defense database. At least one FortiGuard update is required to install the database, otherwise this feature is ineffective. For details, see Configuring FortiGuard service settings on page 270.

CORS Protection

Cross-Origin Resource Sharing (CORS) is a browser mechanism which enables controlled access to resources located outside of a given domain The CORS standard works by adding new HTTP headers that allow servers to describe which origins are permitted to read that information from a web browser. It extends and adds flexibility to the same-origin policy so that websites would not be restricted to accessing resources from the same origin.

However, in the process of enabling information sharing between sites, the significance of CORS configuration may be overlooked and allow for vulnerabilities. One such example is the Cross-Origin Request Site, an OWASP TOP10 Security Misconfiguration vulnerability.

To protect your applications against CORS vulnerabilities, use the CORS Protection feature to ensure that only legitimate CORS requests from allowed web applications can reach your application.

The **Web Application Firewall > CORS Protection** sub-menu includes the following:

- Configuring a CORS Protection Rule on page 733
- Configuring an Allowed Origin List on page 736
- Configuring a CORS Headers List on page 738

Configuring a CORS Protection Rule

The CORS Protection Rule List defines the actions FortiADC may take to protect the Cross-Origin Resource Sharing using the Allowed Origin and optionally, the CORS Headers.

Configuration overview

To enable the CORS protection functionality, you need to configure the following:

- Allowed Origin List — see Configuring an Allowed Origin List on page 736.
- CORS Headers List (optional) — see Configuring a CORS Headers List on page 738.
- CORS Protection Rule List.

After you have configured your CORS Protection, you can add it to your WAF profile configuration under the Input Protection section. For more information, see Configuring a WAF Profile on page 667.

To create and configure the CORS Protection Rule List:

1. Go to **Web Application Firewall > CORS Protection**.
2. Click the **CORS Protection** tab.
3. Click **Create New** to display the configuration editor.
   Configure the following:

Parameter	Description
Name	Enter a unique CORS Protection name. Valid characters should match regular expression `/^[A-Za-z0-9.:_-]*$/`. No space is allowed. **Note**: Once saved, the name of an CORS Protection cannot be changed.
Status	Enable/disable CORS protection. This is disabled by default. **Note**: The CORS Protection Rule List cannot be configured until CORS protection is enabled.

4. Click **Save**.
   The newly created CORS Protection is listed under the **CORS Protection** tab.
5. Locate the newly created CORS Protection on the list and double-click the row or click the **(Edit icon)**.
6. Under **CORS Protection Rule List**, click **Create New** to display the configuration editor.
   Configure the following:

Parameter	Description
Action	Specify the WAF action: • alert • deny • block • silent-block The default action is **block**.
Host Status	Enable/disable to allow this rule to protect a specific domain name or IP address. This is disabled by default.
Host Name	This option appears if **Host Status** is **enabled**. Specify the host name.
Request URL	Specify the request URL as a regular expression. The maximum length is 8192 characters.
Apply to All CORS Traffic	Enable/disable to apply the CORS Protection Rule to all CORS traffic. This is disabled by default. • Disable — The CORS Protection Rule will take effect if all CORS protection parameters matches, including Allowed Origin. • Enable — The CORS Protection Rule will take effect if the Request URL and/or the Host Name (if Host Status is enabled) matches. Once **Apply to All CORS Traffic** is enabled, all options are hidden except **Action**, **Host Status** (**Host Name**), and **Request URL**.
Allowed Origin	Specify the name of the Allowed Origin.

Parameter	Description
	From the drop-down, you may select previously configured Allowed Origin or select **Create New** to create and configure an Allowed Origin directly. For detailed steps, see Configuring an Allowed Origin List on page 736. The allowed origin list ensures only the CORS traffic from the specified applications are allowed.
Insert Allow Credentials	Enable/disable to allow whether the CORS requests from foreign applications can include user credentials. This is disabled by default.
Allowed Credentials	This option appears if **Insert Allow Credentials** is enabled. Select one of the following options: • True • False If the selected **Allowed Origin** is set to `*`, then do not select **True** for **Allowed Credentials**.
Insert Max Age	Enable/disable to specify a maximum time period before the result of the preflight request expires.
Allowed Maximum Age	This option appears if **Insert Max Age** is enabled. Specify the maximum time period in seconds. (Range: 0-86400, default: 0).
Allowed Methods	Enable/disable to allow FortiADC to use the **Methods** specified to verify whether the methods used in the CORS requests are legitimate. This is disabled by default.
Methods	This option appears if **Allowed Methods** is enabled. Specify the method(s): • GET • POST • HEAD • TRACE • CONNECT • DELETE • PUT • PATCH
Allowed Headers	Enable/disable to allow FortiADC to use the **CORS Headers List** to verify whether the headers used in the CORS requests are legitimate. This is disabled by default.
Allowed Headers List	This option appears if **Allowed Headers** is enabled. Specify the name of the CORS Headers List to allow. From the drop-down, you may select previously configured CORS Headers. For detailed steps, see Configuring a CORS Headers List on page 738. FortiADC uses the allowed-headers-list to verify whether the headers used in the CORS requests are legitimate.

Parameter	Description
Exposed Headers	Enable/disable to allow FortiADC to expose the specified headers in the **CORS Headers List** in JavaScript and share with foreign applications. This is disabled by default.
Exposed Headers List	This option appears if **Exposed Headers** is enabled. Specify the name of the CORS Headers List to expose. From the drop-down, you may select previously configured CORS Headers. For detailed steps, see Configuring a CORS Protection Rule on page 733. FortiADC will expose the headers in the exposed-headers-list in JavaScript and share with foreign applications.

7. Click **Save**.

Configuring an Allowed Origin List

The Allowed Origin List specifies the allowed domains using the HTTP response header. The header can contain either a `*` to indicate that all domains are allowed OR a specified domain to indicate the specified allowed domain.

You can create and configure the Allowed Origin List from the **Allowed Origin** tab or as part of the CORS Protection Rule List.



Allowed Origin can only take effect in the CORS Protection rule when the **Apply to All CORS Traffic** is **disabled**. In the CORS Protection Rule List configuration, the Apply to All CORS Traffic option is disabled by default, which then requires you to apply an Allowed Origin List for the CORS Protection rule. If the Allowed Origin List is not applied, the CORS Protection rule would not work as the empty list would not match the condition.

Enabling the Apply to All CORS Traffic option hides the Allowed Origin option, making it inapplicable to the CORS Protection rule.

To create and configure the Allowed Origin List from Allowed Origin tab:

1. Go to **Web Application Firewall > CORS Protection**.
2. Click the **Allowed Origin** tab.
3. Click **Create New** to display the configuration editor.
   Configure the following:

Parameter	Description
Name	Enter a unique Allowed Origin name. Valid characters should match regular expression `/^[A-Za-z0-9.:_-]*$/`. No space is allowed. **Note**: Once saved, the name of an Allowed Origin cannot be changed.

4. Click **Save**.
5. Under **Allowed Origin List**, click **Create New** to display the configuration editor.

Configure the following:

Parameter	Description
Protocol	Select which type of protocols are allowed for the connections between foreign applications and your application. • HTTP • HTTPS • ANY The default is HTTP.
Origin Name	Enter the foreign application's domain name or IP address. Wildcards are supported. (Range: 1-128 characters).
Port	Specify the TCP port number for the CORS connections. (Range: 0-65535; default: 80).
Include Sub Domains	Enable/disable to allow/disallow the Origin Value to match with the domains of its sub level. This is disabled by default.

6. Click **Save**.

To create and configure the Allowed Origin List as part of the CORS Protection Rule List:

1. Go to **Web Application Firewall > CORS Protection**.
2. Click the **CORS Protection** tab.
3. Click **Create New** to display the configuration editor.
   Configure the following:

Parameter	Description
Name	Enter a unique CORS Protection name. Valid characters should match regular expression `/^[A-Za-z0-9.:_-]*$/`. No space is allowed. **Note**: Once saved, the name of an CORS Protection cannot be changed.
Status	Enable/disable CORS protection. This is disabled by default. **Note**: The CORS Protection Rule List cannot be configured until CORS protection is enabled.

4. Click **Save**.
   The newly created CORS Protection is listed under the **CORS Protection** tab.
5. Locate the newly created CORS Protection on the list and double-click the row or click the **(Edit icon)**.
6. Under **CORS Protection Rule List**, click **Create New** to display the configuration editor.
7. In the **Allow Origin** field, select **Create New** from the drop-down.
   The **Allowed Origin** configuration editor is displayed.

8. Configure the following:

Parameter	Description
Name	Enter a unique Allowed Origin name. Valid characters should match regular expression `/^[A-Za-z0-9.:_-]*$/`. No space is allowed. **Note**: Once saved, the name of an Allowed Origin cannot be changed.

9. Click **Save**.
10. Under **Allowed Origin List**, click **Create New** to display the configuration editor.
Configure the following:

Parameter	Description
Protocol	Select which type of protocols are allowed for the connections between foreign applications and your application. • HTTP • HTTPS • ANY The default is HTTP.
Origin Name	Enter the foreign application's domain name or IP address. Wildcards are supported. (Range: 1-128 characters).
Port	Specify the TCP port number for the CORS connections. (Range: 0-65535; default: 80).
Include Sub Domains	Enable/disable to allow/disallow the Origin Value to match with the domains of its sub level. This is disabled by default.

11. Click **Save**.

Configuring a CORS Headers List

The CORS Headers List specifies the HTTP headers that may be "allowed" or "exposed" in the CORS Protection Rule List. If allowed, FortiADC will use the headers list to verify whether the headers used in the CORS requests are legitimate. If exposed, FortiADC will expose the headers in the headers list in JavaScript and share with foreign applications.

The CORS Headers List can be optional as it is only required if **Allowed Headers** or **Exposed Headers** is enabled in the CORS Protection Rule List.

To create and configure the CORS Headers List:

1. Go to **Web Application Firewall > CORS Protection**.
2. Click the **CORS Headers** tab.
3. Click **Create New** to display the configuration editor.
Configure the following:

Parameter	Description
Name	Enter a unique CORS Headers name. Valid characters should match regular expression `/^[A-Za-z0-9.:_-]*$/`. No space is allowed. **Note**: Once saved, the name of a CORS Headers cannot be changed.

4. Click **Save**.
   The newly created CORS Headers is listed under the **CORS Headers** tab.
5. Locate the newly created CORS Headers on the list and double-click the row or click the **(Edit icon)**.
6. Under **CORS Headers List**, click **Create New** to display the configuration editor.
   Configure the following:

Parameter	Description
Header	Specify the HTTP header as a string. (Range: 1-63 characters).

7. Click **Save**.

API Protection

The **Web Application Firewall > Bot Mitigation** sub-menu is organized into the following sections:

- JSON Protection on page 739
- XML Protection on page 742
- OpenAPI Validation on page 746
- API Gateway on page 748
- API Discovery on page 751

JSON Protection

Under the **API Protection** sub-menu, the **JSON Protection** page contains the features that allow you to create rules that enforce security checks that examine client HTTP requests for anomalies in JSON data in HTTP POST operations.

This section includes the following:

- Configuring JSON detection on page 739
- Importing JSON schema on page 741

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 740 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 384 and Configuring a WAF Profile on page 667.

To configure JSON Detection:

1. Go to **Web Application Firewall > API Protection** and select the **JSON Detection** tab.
2. Click **Create New.**
3. Complete the configuration as described in JSON Detection on page 740.
4. Click **Save**.

JSON Detection

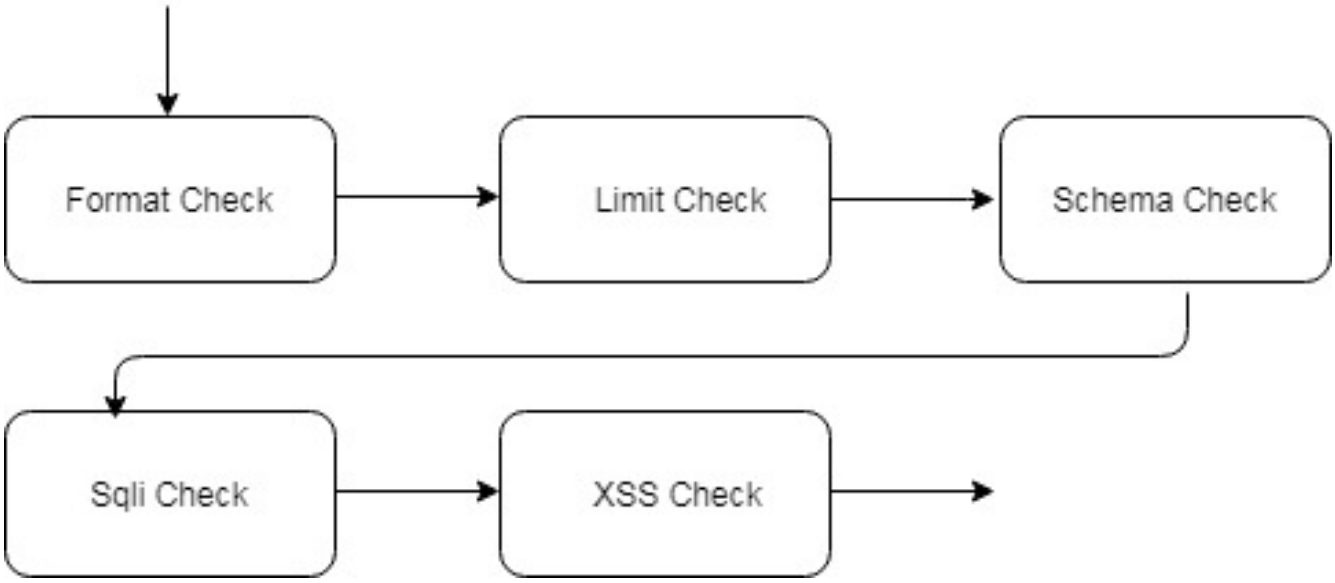
Settings	Guidelines
Name	Enter the name of the JSON Detection profile. You will use the name to select the JSON Detection profile in WAF profiles. No spaces.
JSON Format Checks	Enable to configure security checks for incoming HTTP requests to determine whether they are well-formed. You can set FortiADC 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

Settings	Guidelines
	• Limit Max String
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. *Available only when **JSON Limit Checks** is enabled.*
Limit Max Depth	Limits the maximum depth in a JSON value. The default value is 16. The valid range is 0–4096. *Available only when **JSON Limit Checks** is enabled.*
Limit Max Object Member	Limits the number of members in a JSON object. The default value is 64. The valid range is 0–4096. *Available only when **JSON Limit Checks** is enabled.*
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. *Available only when **JSON Limit Checks** is enabled.*
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	Select the action profile that you want to apply. See Configuring WAF Action objects on page 671. The default is Alert.
Exception Name	Optional. Select the exception profile that you want to apply to the JSON Detection profile. See Configuring WAF Exception objects on page 673.

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 > API Protection** 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**.

XML Protection

Under the **API Protection** sub-menu, the **XML Protection** page contains the features that allow you to create rules that examine client requests for anomalies in XML code.

This section includes the following:

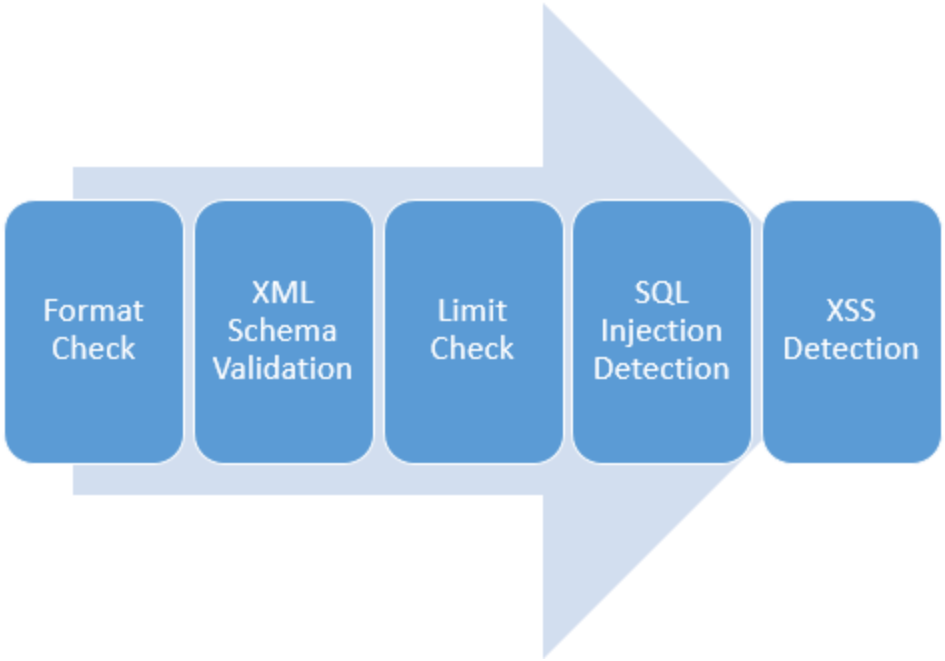
- Configuring XML Detection on page 743
- Importing XML schema on page 745
- Uploading WSDL files on page 746

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 743 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 384 and Configuring a WAF Profile on page 667.

To configure XML Detection:

1. Go to **Web Application Firewall > API Protection** and select the **XML Detection** tab.
2. Click **Create New.**

3. Complete the configuration as described in XML Detection on page 744.
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.
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 745
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. *Available only when **XML Limit Checks** is enabled.*
Max Attribute	Limits the maximum length of each attribute name. The default value is 128. The valid range is

Settings	Guidelines
Name Length	1–2048. *Available only when **XML Limit Checks** is enabled.*
Max Attribute Value Length	Limits the maximum length of each attribute value. The default value is 128. The valid range is 1–2048. *Available only when **XML Limit Checks** is enabled.*
Max Cdata Length	Limits the length of the CDATA section for each element. The default value is 65535. The valid range is 1–65535. *Available only when **XML Limit Checks** is enabled.*
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. *Available only when **XML Limit Checks** is enabled.*
Max Element Depth	Limits the maximum number of nested levels in each element. The default value is 256. The valid range is 1–65535. *Available only when **XML Limit Checks** is enabled.*
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. *Available only when **XML Limit Checks** is enabled.*
Max Namespace	Limits the number of namespace declarations in the XML document. The default value is 16. The valid range is 0–256. *Available only when **XML Limit Checks** is enabled.*
Max Namespace URL Length	Limits the URL length for each namespace declaration. The default value is 256. The valid range is 0–1024. *Available only when **XML Limit Checks** is enabled.*
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	Select the action profile that you want to apply. See Configuring WAF Action objects on page 671. The default is Alert.
Exception Name	Optional. Select the exception profile that you want to apply to the XML Detection profile. See Configuring WAF Exception objects on page 673.

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 > API Protection** 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> API Protection**.
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**.

OpenAPI Validation

Under the **API Protection** sub-menu, the **OpenAPI Validation** page contains the features that allow you to create rules through defining a standard, language-agnostic interface to RESTful APIs, which allows both humans and computers to discover and understand the capabilities of the service without access to source code, documentation, or through network traffic inspection.

This section includes the following:

- Configuring OpenAPI Detection on page 747
- Importing OpenAPI schema on page 748

Configuring OpenAPI Detection

The OpenAPI Specification (OAS) defines a standard, language-agnostic interface to RESTful APIs, which allows both humans and computers to discover and understand the capabilities of the service without access to source code, documentation, or through network traffic inspection. When properly defined, you can understand and interact with the remote service with a minimal amount of implementation logic.

FortiADC can parse the OpenAPI description file and provide additional security to APIs by making sure that access is based on the definitions described in the OpenAPI file.

Note: FortiADC supports OpenAPI 3.0.

To configure OpenAPI Detection:

1. Go to **Web Application Firewall > OpenAPI Validation**.
2. Click the **OpenAPI Detection** tab.
3. Click **Create New** to display the configuration editor and set up the configuration.
4. Save the configuration.

API Detection Configuration

Settings	Guidelines
Name	Configure the name. Valid characters are A-Z, a-z, 0-9, _, and -. Whitespaces not allowed. Note: Once saved, the name cannot be changed.
OpenAPI Schema Check	Before enabling OpenAPI Schema Check, you must upload an OpenAPI schema file to check whether OpenAPI content is permitted. Enable to use OpenAPI schema to validate OpenAPI content. See Importing OpenAPI schema on page 748.
OpenAPI Schema	Select the OpenAPI schema file that you want to use to check whether OpenAPI content is valid.
Action	Select the action profile that you want to apply. See Configuring WAF Action objects on page 671. The default 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 the severity level FortiADC uses when using Input Validation: • Low • Medium • High 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.

Importing OpenAPI schema

An OpenAPI schema file defines or describes the API including information like the API URL, parameter names in the URL, type of data parameters should have (string, integer, etc), where parameters are submitted (URL, header, body, etc.), and so on. For more information about OpenAPI files, see https://github.com/OAI/OpenAPI-Specification.

Once you upload the valid OpenAPI schema file, FortiADC will parse the file and then block requests that do not match the definitions in the file.

Before you begin, you must:

l Prepare a trusted OpenAPI schema file in YAML or JSON format that you can import to FortiADC. Acceptable file types are .tar, .tar.gz, and .zip.

To import an OpenAPI schema file:

1. Go to **Web Application Firewall > OpenAPI Validation**.
2. Click the **OpenAPI Schema** tab.
3. Click **Create New**.
4. Enter the name of the OpenAPI schema configuration. You will use the name to select the schema file in OpenAPI Detection profiles.
5. Click **Choose File** and select the OpenAPI schema file that you want to import.
6. Click **Save**.

API Gateway

Under the **API Protection** sub-menu, the **API Gateway** allows you to create policies that include an API management tool that sits between a client and a collection of backend services. It acts as a reverse proxy to accept all API calls and return the appropriate result.

This section includes the following:

- Configuring an API Gateway policy on page 748
- Configuring an API Gateway User on page 749
- Configuring an API Gateway Rule on page 750

Configuring an API Gateway policy

An API gateway is an API management tool that sits between a client and a collection of backend services. It acts as a reverse proxy to accept all API calls and return the appropriate result.

API gateway on FortiADC provides the following functions:

- API user management
- API key verification
- API access control
- Rate limit control
- Attach HTTP Header in API call



To configure an API Gateway Policy:

1. Go to **Web Application Firewall > API Gateway**.
2. Click the **API Gateway Policy** tab.
3. Click **Create New** to display the configuration editor and set up the configuration.
4. Save the configuration.

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. Whitespaces are not allowed. After you initially save the configuration, you cannot edit the name.
Rule Name	Specify one or more rules created in API Gateway Rule to be used in policy. The rules will be checked one by one from top to bottom until URL in request is matched to the Full URL Pattern in a rule.

Configuring an API Gateway User

1. Go to **Web Application Firewall > API Gateway**.
2. Click the **API Gateway User** tab.
3. Click **Create New** to display the configuration editor and set up the configuration.
4. Save the configuration.

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. Whitespaces are not allowed. After you initially save the configuration, you cannot edit the name.
Comments	(Optional) Enter a description or comments for the user.
UUID	Non-editable. Automatically generated when the user is created.
API Key	Non-editable. Automatically generated when the user is created.
Restricted Access IPs	Restrict this API key so that it may only be used from the specified IP addresses.
Restrict HTTP Referers	Restrict this API key so that it may only be used when the specified URLs are present in the Referer HTTP header. This can be used to prevent an API key from being reused on other client-side web applications that don't match this URL. Only full URLs that begin with http:// or https:// are supported.

Configuring an API Gateway Rule

1. Go to **Web Application Firewall > API Gateway**.
2. Click the **API Gateway Rule** tab.
3. Click **Create New** to display the configuration editor and set up the configuration.
4. Save the configuration.

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. Whitespaces are not allowed. After you initially save the configuration, you cannot edit the name.
Host Status	Enable/Disable for applying this rule only to HTTP requests for specific web hosts.
Host	Select the name of a protected host that the Host: field of an HTTP request must be in to match the API gateway rule. This option is available only if Host Status is enabled.
Full URL Pattern	Matching string. Regular expressions are supported.
Method	Select one or more HTTP methods are allowed when access the API.
API Key Verification	When a user makes an API request, the API key will be included in the HTTP header or parameter. FortiWeb obtains the API key from the request. When this option is enabled, FortiWeb verifies the key to check whether the key belongs to an valid API user.
API Key Carried In	Indicate where to find the API key in HTTP request: • HTTP Parameter • HTTP Header Available only when API Key Verification is enabled.
HTTP Header Name	Enter the header filed name of the API key.
HTTP Parameter Name	Enter the parameter name of the API key.
Rate Limit Status	Enable/Disable to do rate limit for API calls.
Rate Limit Requests	Sets the condition for the limit of the number of API requests received. If the number of requests received within the time frame (set in Rate Limit Period), this condition is fulfilled.
Rate Limit Period	Sets the time spent during which to count how many times a request is received.
Action	Select the action profile that you want to apply. See Configuring WAF Action objects on page 671. The default 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.

Settings	Guidelines
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.
User	Specify one or more users created in API Gateway User to define which users have the permission to access the API.
Attach HTTP Header	Insert specific header lines into HTTP header. Need to specify the fieldname and value is seach entry.

API Discovery

Under the **API Protection** sub-menu, the **API Discovery** page contains the features that allow you to protect your application against malicious APIs through discovered API endpoints.

Through API Discovery policies, FortiADC is enabled to automatically discover external API endpoints from HTTP/HTTPS requests and responses that have passed through API validity checks, wherein the API is parsed for information including the Host, Paths, parameters and their schemas from query requests or entity bodies, as well as classify parameters that match PII (Personal Identifiable Information) signatures. API Discovery also supports manually imported OAS files compliant with OpenAPI 3.0 and Swagger 2.0 standard to parse and discover as internal API endpoints that can also be matched by incoming API requests or responses.

This section includes the following:

- Configuring API Discovery on page 751
- Understanding the API endpoints database in API View on page 754

Configuring API Discovery

API Discovery policies allow FortiADC to automatically discover external API endpoints from HTTP/HTTPS requests and responses that have passed through API validity checks, wherein the API is parsed for information including the Host, Paths, parameters and their schemas from query requests or entity bodies, as well as classify parameters that match PII (Personal Identifiable Information) signatures. API Discovery also supports manually imported OAS files compliant with OpenAPI 3.0 and Swagger 2.0 standard to parse and discover as internal API endpoints that can also be matched by incoming API requests or responses. The discovered external and internal API endpoints can then be directly applied in API security rules based on the Host, Path, and request rate. Once the API requests and responses pass the API validity check that matches the rule, the specified security action will be triggered to protect against the malicious APIs.



API Discovery is based on VDOMs, where API endpoints are discovered and stored per VDOM. The total API endpoints database size for each VDOM is 1 G.

API Discovery is based on VDOMs, where API endpoints are discovered and stored per VDOM.

Basic API Discovery workflow

API Discovery policies depend entirely on internal or external API endpoints to function. For external API endpoints to be discovered, the API Discovery policy must be referenced in a WAF Profile used in an active virtual server. Follow the basic workflow below to deploy API Discovery.

Basic steps:

1. Create and enable an API Discovery policy to enable automatic discovery for external API endpoints. For details, see Configuring automatic discovery for external API endpoints on page 752.
2. Reference the API Discovery policy in a WAF Profile. Once the WAF Profile is used in a virtual server, external API endpoints will be discovered from HTTP/HTTPS requests and responses. For details, see Configuring a WAF Profile on page 667.
3. Optionally, import an OAS file to parse and discovery as internal API endpoints. For details, see Manually import internal API endpoints on page 753.
4. Optionally, check and refresh the API View page for automatically discovered API endpoints.
5. Optionally, check the **API Threats** page under **FortiView > Data Analytics** to view all discovered API endpoints that have triggered attacks.

Configuring automatic discovery for external API endpoints

External API endpoint discovery is automatically activated when the virtual server referencing the API Discovery policy receives API responses/requests. The API Discovery policy will automatically begin validating APIs and parsing endpoints to build your API endpoints database which can then be viewed in the API View page. From the API Discovery policy, you can add API Security rules to trigger alerts and actions against APIs identified as malicious through API Discovery.

The maximum number of API Discovery policies is 256.

To configure an API Discovery policy:

1. Go to **Web Application Firewall > API Discovery**.
   The configuration page displays the **API Discovery** tab.
2. Click **Create New** to display the configuration editor.
3. Configure the following settings:

Setting	Description
Name	Enter a unique API Discovery name. Valid characters should match regular expression `/^[A-Za-z0-9.:_-]*$/`. No space is allowed. **Note**: Once saved, the name of an Allowed Origin cannot be changed.
API Discovery	Enable or disable the API Discovery configuration.

4. Click **Save**.
   Once the API Discovery configuration is saved, the API Security Rule section can be configured.
5. Under the API Security Rule section, click **Create New** to display the configuration editor.
   The maximum number of API security rules under one API Discovery policy is 256.
6. Configure the following API Security Rule settings:

Setting	Description
Host	Specify the HTTP Host header. This is required. Maximum length is 255 characters. Example: 192.168.0.253, [2001:1234::a41:6e]:8443, or demo.fortinet.com.

Setting	Description
	Once the API Discovery policy is activated, the policy matches only if the Host header matches this value. Complete, exact matching is required. For example, www.example.com matches www.example.com but not www.example.com.hk.
Path	Specify the API resource path. Text string and simple regex is supported. Example: /login. Begin with '/'.
API Request Rate	Specify the allowable requests per second. Default: 0. Range: 0 - 100000000. **Note**: 0 means there is no limit.
Action	Specify a WAF action object to apply when a bot is detected. See Configuring WAF Action objects on page 671. The default action is alert.
Severity	Select the event severity to log when a bot is detected: • High — Log as high severity events. • Medium — Log as a medium severity events. • Low — Log as low severity events. The default is Low.

7. Click **Save**.
   Once the API Security Rule configuration is saved, the dialog closes.
8. Click **Save** to update the API Discovery configuration with the API Security Rule changes.

After you have completed the API Discovery policy configuration, reference it in a WAF Profile that will be used in a virtual server.



If a previously enabled API Discovery policy is disabled, new API endpoints will not be discovered. However, the API security rules associated with that API Discovery policy will continue to work.

Manually import internal API endpoints

API Discovery supports manually imported OAS files compliant with OpenAPI 3.0 and Swagger 2.0 standard to parse and discover as internal API endpoints that can also be matched by incoming API requests or responses.

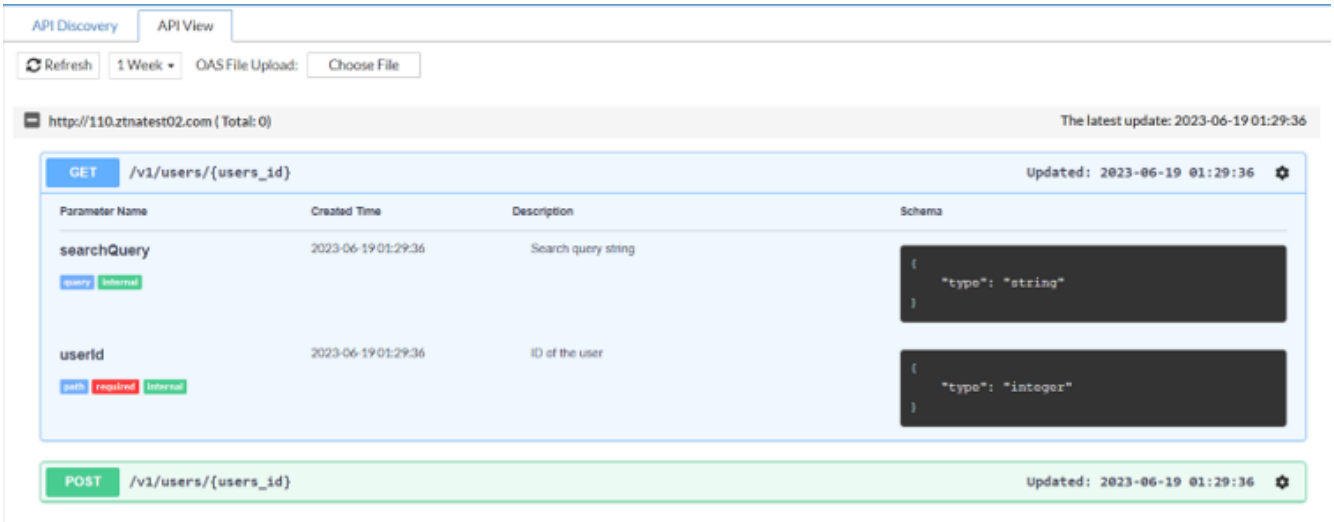
JSON and YAML file types are supported.

To import an OAS file:

1. Go to **Web Application Firewall > API Discovery**.
2. Click the **API View** tab.
3. Click **Choose File** to display the file upload dialog.
4. Locate and select the OAS file to import.
   After the OAS file is successfully imported, the file contents are parsed into internal API endpoints that can be viewed on the API View page.

Understanding the API endpoints database in API View

From the **API View** tab, you can view your database of discovered internal and external API endpoints. You can refresh the page to view newly discovered external API endpoints.



Each API endpoint is parsed for the following key elements:

API Information	Description
Host	The name of the protected host that the Host: field of an HTTP request must be in to match the API Discovery policy. Example: 110.ztnatest02.com **Note**: Maximum supported length of Host is 255 characters.
Base_URL	The base_url consists of the prefix HTTP/HTTPS and the Host. Example: http://110.ztnatest02.com **Note**: Maximum supported length of Base_URL is 255 characters.
Path	The resource path of the API. Example: /api/test
API Endpoint	The API endpoint is organized and displayed as the Path with HTTP method. Example: GET: /api/test
Parameter	Parsed from the HTTP query request or entity body. Example: path, required, internal

API View endpoint organization

Each API endpoint entry is identified and stored by its HTTP method and path. However, each entry is maintained per method, so in cases where endpoints share the same path but use different HTTP methods, a separate entry will be maintained per HTTP method.

API endpoints will be displayed in a simplified view when four or more requests that have a different integer as the last sub-path are discovered. In which case, the endpoints will be simplified and reorganized as a single endpoint template with the following specifications:

- The Path of the endpoint will be similar to `/v1/users/{users_id}`.
- An extra parameter will be generated, named “users_id” with a label “required”.

Supported API types and HTTP methods for API Discovery

- RESTful API — GET/POST/PUT with XML/JSON, or DELETE
- SOAP API — POST with SOAP+XML
- OAS file types — JSON or YAML

Prerequisites for APIs to be discovered

- The **return code** must be **200-299**.
- The HTTP headers and corresponding entity body must have the following in the request/response:

HTTP Method	Requirement
GET	“Accept: application/json”, “Accept: application/xml” or “Accept: */*” in the request, and corresponding “Content-Type:” header with entity body in the response.
PUT/POST	“Content-Type: application/json”, “Accept: application/xml” or “Accept: application/soap+xml” with corresponding entity body in the request.
DELETE	No requirement; all DELETE requests can be discovered.

- For RESTful APIs: the last sub-path **should not** include the dot (**.**) character. SOAP API supports including the dot (**.**) in the last sub-path.
- For OpenAPI 3.0: servers and paths are required in the OAS file.
- For Swagger 2.0: host, paths & schemes are required in the OAS file.

API endpoint parameters

The API Discovery engine parses through raw API data to extract information that make up an API endpoint. Parameters can be parsed from the key elements of an API, which refers to HTTP/HTTPS, Host, Path, and Parameters. Further information can be parsed from the classification or labels of parameters, PII and schema.

Parameters can be parsed from the following API information:

Source	Parsed parameters
HTTP query request	Parameters can be parsed from the HTTP query request. For example: `/query?`**`username`**`=admin&`**`password`**`=test` • username = admin • password = test

Source	Parsed parameters
Elements or keys in the entity body	For JSON API type, API Discovery parses up to 1 layer. For XML/SOAP API type, API Discovery parses up to 2 or 3 layers: • If sub-elements have the same name, the schema of the elements will be recognized as an "array", and 3 layers of the nested parameters will be parsed. • Otherwise, the parameters will be treated as an object, and only 2 layers of the nested parameters will be parsed.
Classification or labels of parameters	• Position — requestBody, responseBody, query, path (for external API endpoints) • Position (“in”) — query, path, header or cookie (for internal API endpoints) • Required — the reorganized parameter `{api_id}` will be labeled as “required” (for external API endpoints); labeled if "required“ is true in files (for internal API endpoints) • Source — **external** for API endpoints parsed from Automatic discovery, and **internal** for imported OAS files • Description — added automatically for the template parameter (for external API endpoints); defined in the file (for internal API endpoints) **Note**: Maximum supported description length is 2048 characters.
PII (Personal Identifiable Information)	• Scan name — Username, Password, Token • Scan value — Credit Card Number, Email, URL, IP Address, GUID • Scan name & value — CVV
Schema of parameters	String, number, integer, Boolean, object, array

Bot Mitigation

The **Web Application Firewall > Bot Mitigation** sub-menu includes the following:

- Configuring a Bot Detection policy on page 756
- Configuring a Threshold Based Detection policy on page 762
- Configuring a Biometrics Based Detection policy on page 767
- Configuring a Fingerprint Based Detection policy on page 769

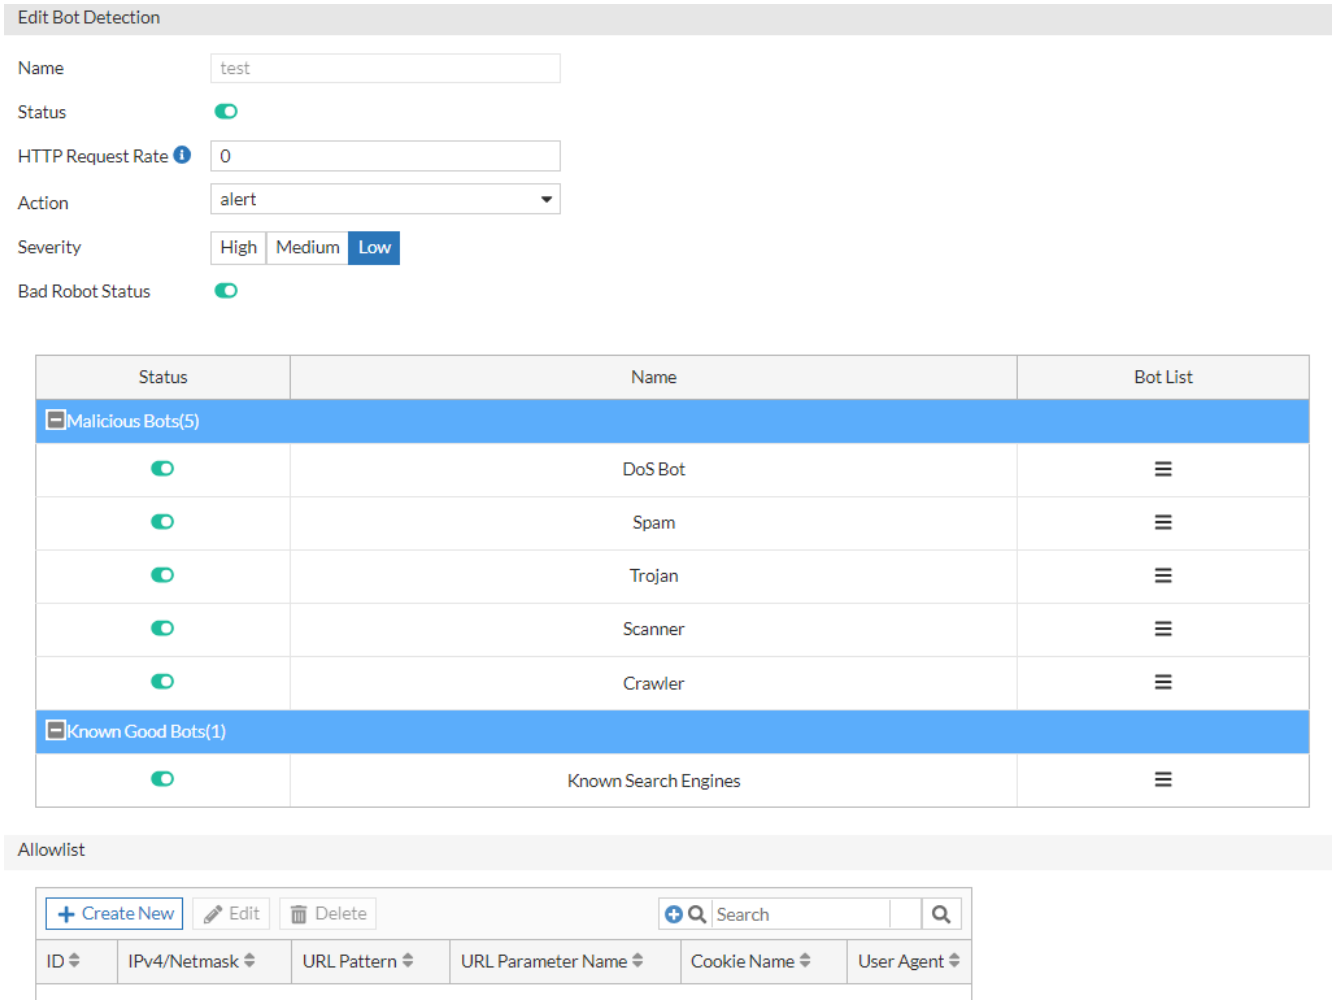
Configuring a Bot Detection policy

Bot Detection policies employ signature analysis and behavioral tracking to identify client traffic likely generated by automated bots rather than genuine human users. Legitimate bots, such as search engine crawlers, are classified as "good bots" because they perform essential search indexing operations, which can increase the visibility of your site to legitimate users.

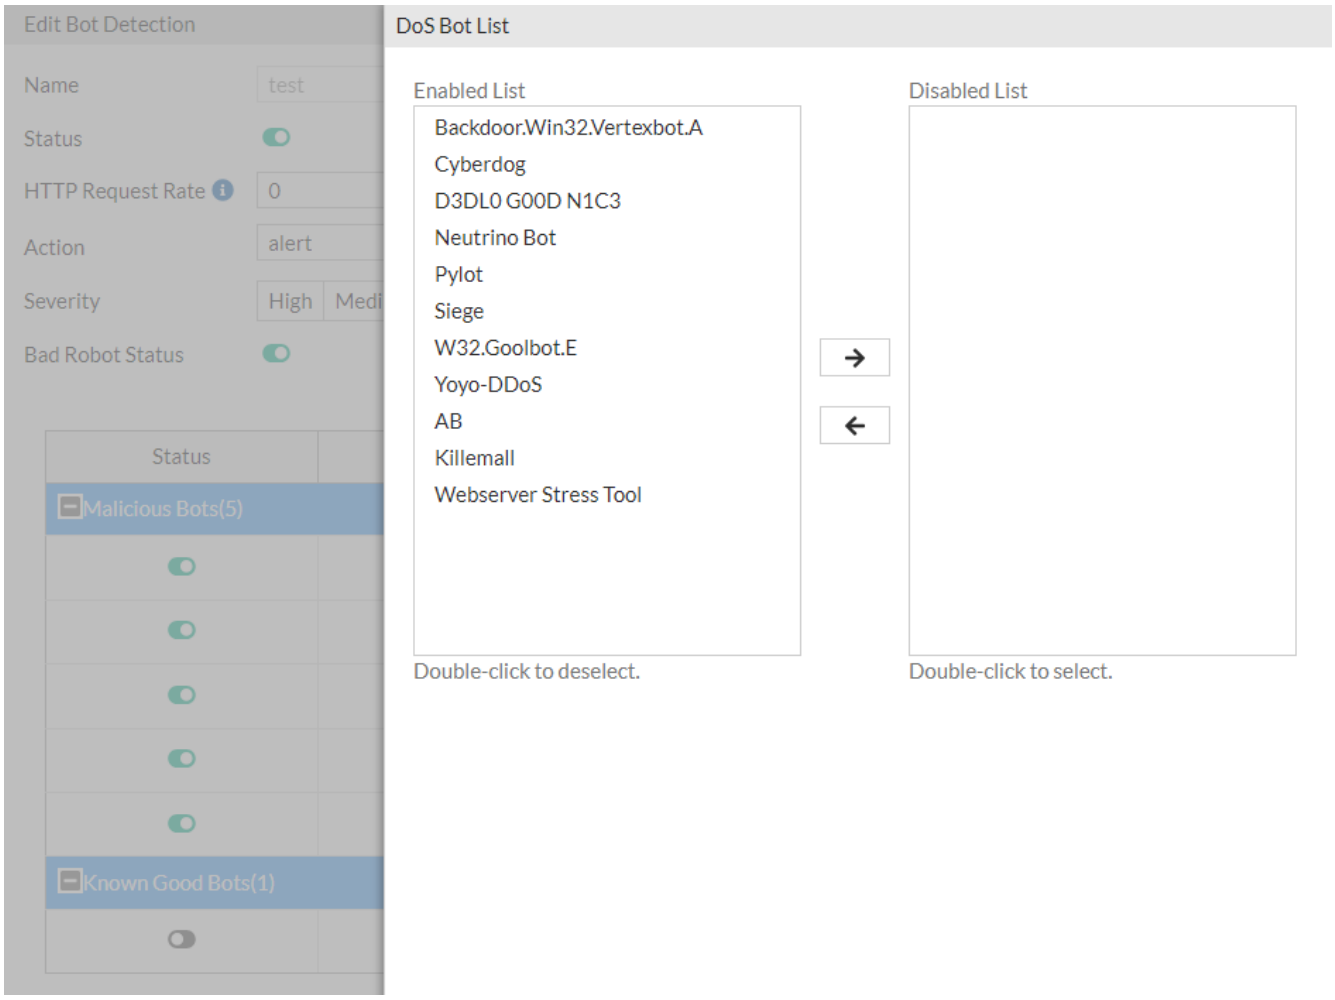
Conversely, "bad bots" are known to generate malicious traffic that can compromise site availability and integrity. Examples of such activities include Distributed Denial of Service (DDoS) attacks and content scraping. To mitigate these

threats, it is crucial to deploy effective bot mitigation strategies, such as IP reputation analysis, rate limiting, anomaly detection algorithms, and CAPTCHA challenges, to identify and block these harmful bots in real-time.

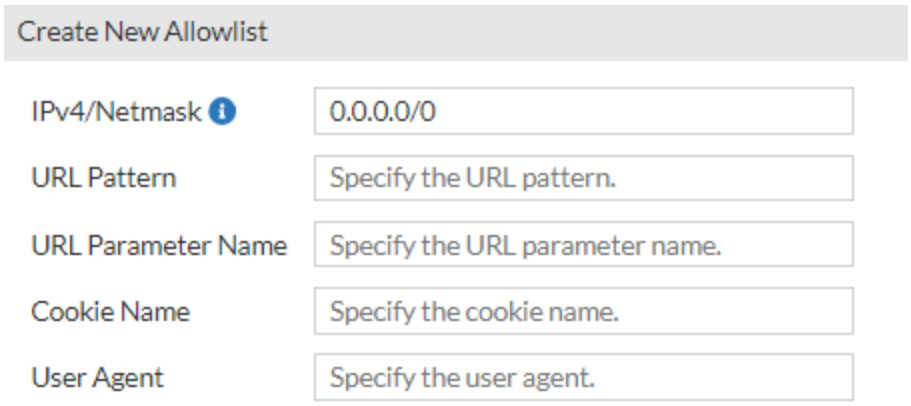
The FortiADC Bot Detection policy allows you to monitor and protect against suspected bot traffic using classifications from the Web Application Firewall (WAF) Signature database. Bots are categorized into five types of Malicious Bots and one type of Known Good Bots. You can customize the detection settings by selecting specific bots within these categories to include or exclude from monitoring and mitigation actions.



From the Bot List column, you can click the ☰ (collapsed menu icon) to enable or disable specific bots for detection.



For more granular control, you can configure an allowlist to specify IP addresses that should be exempt from bot detection, even if they match bot signatures. For instance, if a user utilizes a scanning tool such as DirBuster, which generates HTTP requests that would typically be classified as Malicious Bot activity, adding the user's source IP address to the allowlist ensures that this traffic is not flagged or blocked by the bot detection system. This allows for legitimate testing and scanning activities to proceed without interference from bot mitigation measures.



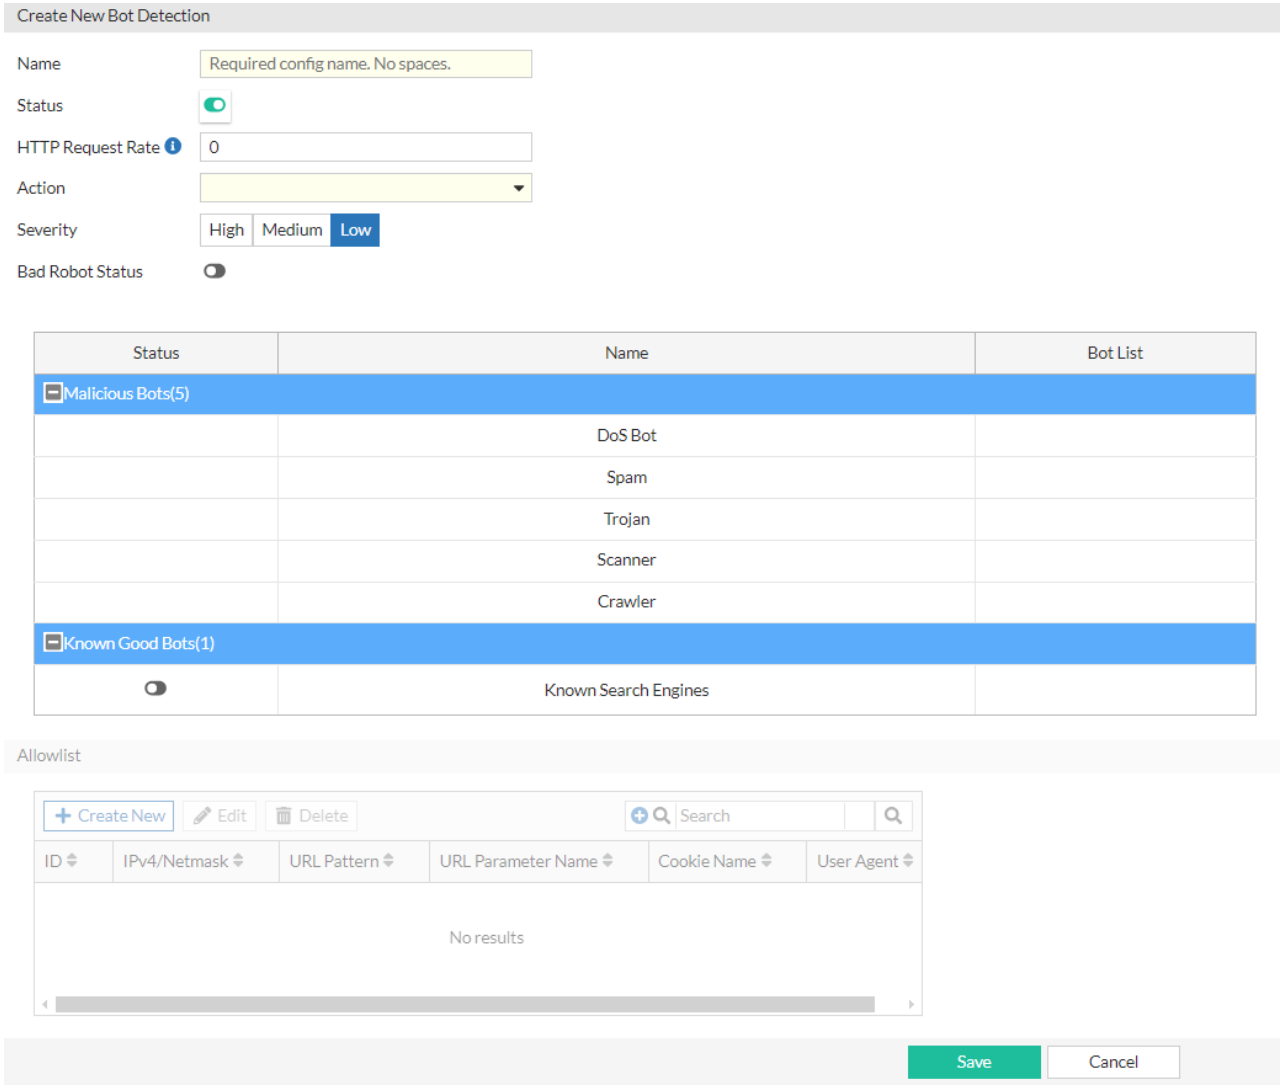
After you have configured Bot Detection policies, you can select them in the WAF Profile. Once you have attached the WAF Profile with Bot Detection to a virtual server, you can monitor the real-time data analytics from FortiView.

Before you begin:

- You must configure the connection to FortiGuard so the system can receive periodic WAF Signature Database updates, including Known Good Bot and Malicious Bot signatures and lists. See Configuring FortiGuard service settings.
- You must have Read-Write permission for Security settings.

To configure a Bot Detection policy:

1. Go to **Web Application Firewall > Bot Detection**.
2. Click **Create New** to display the configuration editor.
3. Enable the **Status** to view the Bot Detection configuration options.

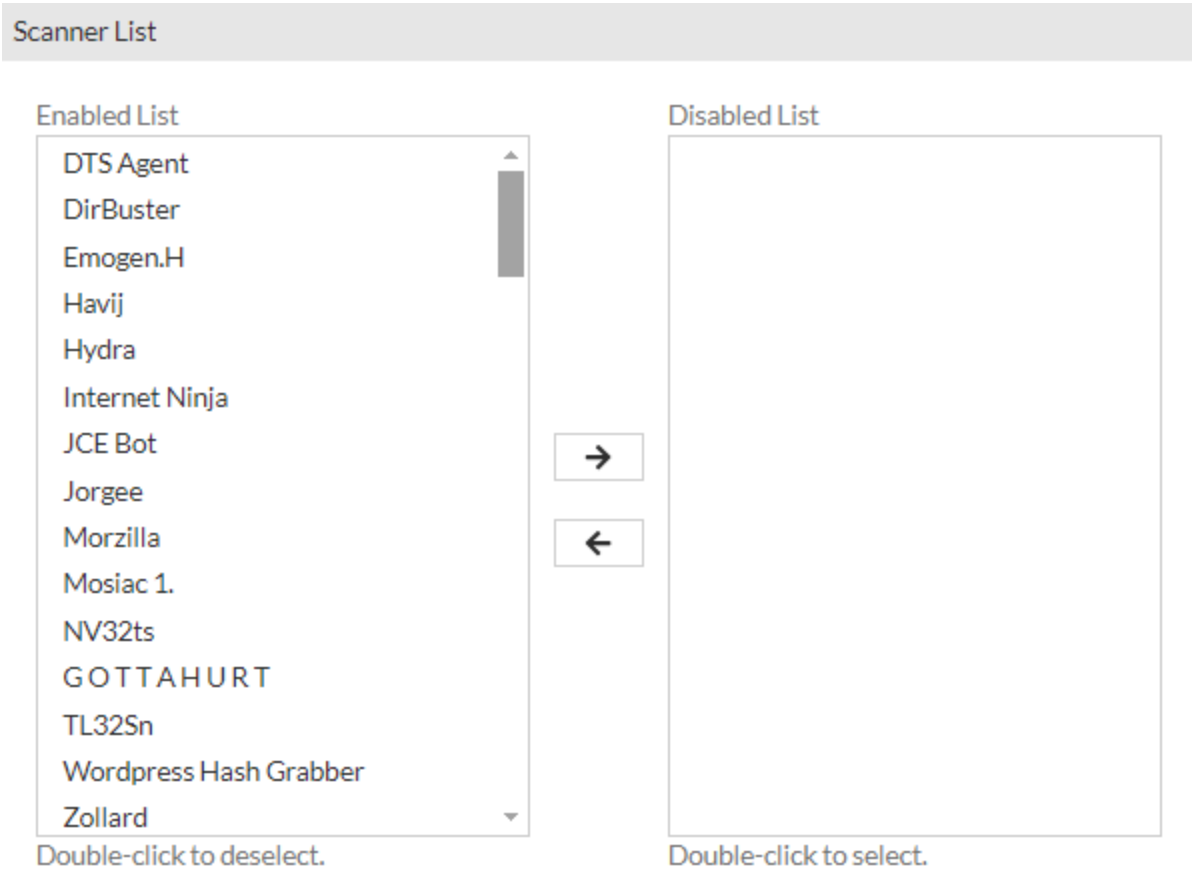


4. Once the Status is enabled, configure the following settings:

Setting	Description
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 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	Select the action profile that you want to apply. See Configuring WAF Action objects on page 671. 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.
Bad Robot Status	Enable or disable detection for Malicious Bots. When enabled, a Malicious Bots blocklist is generated from WAF Signature updates from FortiGuard.
Malicious Bots	The **Malicious Bots** statuses are available if **Bad Robot Status** is **enabled**. You can enable or disable any of the five Malicious Bots categories: • DoS Bot • Spam • Trojan • Scanner • Crawler All Malicious Bots categories are enabled by default. **Note**: The Bot List can be configured once the initial Bot Detection policy configuration is saved.
Known Good Bots	Enable or disable detection for Known Good Bots. When enabled, a Known Good Bots allowlist is generated from WAF Signature updates from FortiGuard. **Note**: The Bot List can be configured once the initial Bot Detection policy configuration is saved.

5. Click **Save**.
   Once the Bot Detection policy is saved initially, the Bot List for each enabled bot categories and the Allowlist becomes configurable.
6. Optionally, you can configure Bot Lists to include or exclude specific bots for detection.
   a. Under the Bot List column of each bot category, click the ☰ (collapsed menu icon) to display the configuration editor.
   b. By default, all bots are enabled. You can select and move each item to be in the **Enabled List** or **Disabled List**.

The example below shows the Bot List for the Scanner category of Malicious Bots.



c. Click **Save** to commit the changes and exit the dialog.

7. Optionally, you can configure an allowlist to specify IP addresses that should be exempt from bot detection, even if they match bot signatures.
   a. Under the Allowlist section, click **Create New** to display the configuration editor.
   b. Configure the following Bot Detection Allowlist settings:

Setting	Description
IPv4/Netmask	Matching subnet (CIDR format: 0.0.0.0/0).
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.

   c. Click **Save** to commit the changes and exit the dialog.
8. Click **Save** to commit all changes to the Bot Detection policy.
   Once the Bot Detection policy is saved, it will be listed on the Bot Detection page. You can now reference this Bot Detection policy in a WAF Profile.

Configuring a Threshold Based Detection policy

Using Threshold Based Detection policies, FortiADC can determine whether requests are generated by robots instead of a human by detecting suspicious behavior patterns that exceed the normal threshold defined in the policy. Threshold Based Detection rules are defined by the number of times a type of behavior is allowed to occur within a specified amount of time. Once the number of occurrence exceeds the defined threshold value, an action is triggered in response to detecting the suspicious behavior.

FortiADC supports the following three types of Threshold Based Detection:

- Crawler Detection — Detects web crawlers that are usually used to map out your application structure by monitoring the frequency of HTTP response codes. If the occurrence of a specified HTTP response code exceeds the allowable threshold in the specified time frame, FortiADC will execute the relevant action for the traffic.
- Content Detection — Detects malicious tools that try to download large amounts of content such as text/HTML and application/ XML from your website by monitoring the frequency of download activities. If the occurrence of the download activity exceeds the allowable threshold within the specified time frame, FortiADC will execute the relevant action for the traffic.
- Attack Detection — Detects suspicious attack behavior patterns indicative of a bot attack by monitoring the frequency of attacks detected in specific WAF Attack modules. If the occurrence of specific attacks exceeds the allowable threshold within the specified time frame, FortiADC will execute the relevant action for the traffic.

FortiADC offers Predefined Threshold Based Detection policy configurations on page 765 that can be applied as is or used as a template for customization.

After you have configured Threshold Based Detection policies, you can select them in WAF profiles.

Before you begin:

- You must have Read-Write permission for Security settings.

To configure a Threshold Based Detection policy:

1. Go to **Web Application Firewall > Threshold Based Detection**.
2. In the **Threshold Based Detection** tab, click **Create New** to display the configuration editor.
3. Configure the following Biometrics Based Detection settings:

Setting	Description
Name	Specify a name for the Threshold Based Detection rule. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No spaces. The configuration name cannot be edited once it has been saved.
Comments	Optionally, enter comments about the Threshold Based Detection policy.
Crawler Detection	
Crawler Status	Enable/Disable Crawler Detection. This is disabled by default.
Response Code	Specify the 3 digit HTTP response code(s) to

Setting	Description
	check. Enter as a single code (e.g. 403), multiple codes (e.g. 403,404), or as a range (e.g. 500-503). Range: 100-599.
Crawler Action	Select the action profile to apply when a web crawler bot is detected. See Configuring WAF Action objects on page 671. The default action is alert.
Crawler Severity	Select the event severity to log when a web crawler bot is detected: • High — Log as high severity events. • Medium — Log as a medium severity events. • Low — Log as low severity events. The default is low.
Crawler Occurrence Limit	Specify the maximum number of responses that can be received from the specified **Response Code** within the time frame (set in **Crawler Occurrence Within**). If the limit is exceeded, the specified **Crawler Action** will be triggered. Default: 100, Range: 1-100000.
Crawler Occurrence Within	Specify the time span during which to count how many times a response is received from the specified **Response Code**. Default: 60 seconds, Range: 1-600 seconds.
Content Detection	
Content Scraping Status	Enable/disable Content Detection. This is disabled by default.
Content Type	Select one or more content type to monitor for content scraping: • Text/HTML • Text/Plain • Text/XML • Application/XML • Application/Soap+XML • Application/JSON
Content Action	Select the action profile to apply when a content scraping bot is detected. See Configuring WAF Action objects on page 671. The default action is alert.
Content Severity	Select the event severity to log when a content scraping bot is detected:

Setting	Description
	• High — Log as high severity events. • Medium — Log as a medium severity events. • Low — Log as low severity events. The default is low.
Content Occurrence Limit	Specify the maximum number of responses that can be received from the specified **Content Type** within the time frame (set in **Content Occurrence Within**). If the limit is exceeded, the specified **Content Action** will be triggered. Default: 100, Range: 1-100000.
Content Occurrence Within	Specify the time span during which to count how many times a response is received from the specified **Content Type**. Default: 60 seconds, Range: 1-600 seconds.
Attack Detection	
Attack Detection Status	Enable/disable Attack Detection. This is disabled by default.
Attack Modules	Select one or more attack modules to monitor for bot attacks: • Web Attack Signature • Input Validation • Brute Force Attack Detection • URL Protection • HTTP Protocol Constraint • Credential Stuffing Defense Click **Advanced** to expand the selection list: • Data Leak Prevention • SQL/XSS Injection Detection • Cookie Security • CSRF Protection • CORS Protection • JSON Validation • OpenAPI Validation • XML Protection • API Gateway
Attack Action	Select the action profile to apply when a bot attack is detected. See Configuring WAF Action objects on page 671. The default action is alert.
Attack Severity	Select the event severity to log when a bot attack is detected:

Setting	Description
	• High — Log as high severity events. • Medium — Log as a medium severity events. • Low — Log as low severity events. The default is Low.
Attack Occurrence Limit	Specify the maximum number of responses that can be received from the specified **Attack Module** within the time frame (set in **Attack Occurrence Within**). If the limit is exceeded, the specified **Attack Action** will be triggered. Default: 100, Range: 1-100000.
Attack Occurrence Within	Specify the time span during which to count how many times a response is received from the specified **Attack Module**. Default: 60 seconds, Range: 1-600 seconds.

4. Click **Save**.
   The newly configured Threshold Based Detection policy is added to the **Threshold Based Detection** page.

Predefined Threshold Based Detection policy configurations

You can apply any of the predefined Threshold Based Detection policies in WAF profiles or you can clone a predefined configuration to use as a template to define your own policy.

Name	Comments	Predefined settings
Bot_Detect	Detect suspicious bot with CAPTCHA action	**Crawler Status — Enabled** Response Code — 403,404 Crawler Action — captcha Crawler Severity — Medium Crawler Occurrence Limit — 100 Crawler Occurrence Within — 60 (seconds) **Content Scraping Status — Enabled** Content Type — Text/HTML, Text/Plain, Text/XML, Application/XML, Application/Soap+XML, Application/JSON Content Action — captcha Content Severity — Medium Content Occurrence Limit — 100 Content Occurrence Within — 60 (seconds) **Attack Detection Status — Enabled** Attack Modules — Web Attack Signature, Input Validation, Brute Force Attack Detection, URL Protection, HTTP Protocol Constraint, Credential Stuffing Defense Attack Action — captcha

Name	Comments	Predefined settings
		Attack Severity — Medium Attack Occurrence Limit — 100 Attack Occurrence Within — 60 (seconds)
Content_Scraping_Detect	Monitor the frequency of illegal content scraping with ALERT action	**Crawler Status — Disabled** **Content Scraping Status — Enabled** Content Type — Text/HTML, Text/Plain, Text/XML, Application/XML, Application/Soap+XML, Application/JSON Content Action — alert Content Severity — Low Content Occurrence Limit — 100 Content Occurrence Within — 60 (seconds) **Attack Detection Status — Disabled**
Crawler_Detect	Monitor the frequency of 403 and 404 response codes with ALERT action	**Crawler Status — Enabled** Response Code — 403,404 Crawler Action — alert Crawler Severity — Low Crawler Occurrence Limit — 100 Crawler Occurrence Within — 60 (seconds) **Content Scraping Status — Disabled** **Attack Detection Status — Disabled**
High-Level-Security	Block all suspicious threshold violations	**Crawler Status — Enabled** Response Code — 403,404 Crawler Action — deny Crawler Severity — High Crawler Occurrence Limit — 100 Crawler Occurrence Within — 60 (seconds) **Content Scraping Status — Enabled** Content Type — Text/HTML, Text/Plain, Text/XML, Application/XML, Application/Soap+XML, Application/JSON Content Action — deny Content Severity — High Content Occurrence Limit — 100 Content Occurrence Within — 60 (seconds) **Attack Detection Status — Enabled** Attack Modules — Web Attack Signature, Input Validation, Brute Force Attack Detection, URL Protection, HTTP Protocol Constraint, Credential Stuffing Defense • Advanced — Data Leak Prevention, SQL/XSS

Name	Comments	Predefined settings
		Injection Detection, Cookie Security, CSRF Protection, CORS Protection, JSON Validation, OpenAPI Validation, XML Protection, API Gateway Attack Action — deny Attack Severity — High Attack Occurrence Limit — 100 Attack Occurrence Within — 60 (seconds)
Illegal_User_Detect	Detect illegal user with CAPTCHA action	**Crawler Status — Disabled** **Content Scraping Status — Disabled** **Attack Detection Status — Enabled** Attack Modules — Brute Force Attack Detection, Credential Stuffing Defense Attack Action — captcha Attack Severity — Medium Attack Occurrence Limit — 100 Attack Occurrence Within — 60 (seconds)
Vulnerability_Scan	Monitor the frequency of web attack signature violations with CAPTCHA action	**Crawler Status — Disabled** **Content Scraping Status — Disabled** **Attack Detection Status — Enabled** Attack Modules — Web Attack Signature Attack Action — captcha Attack Severity — Medium Attack Occurrence Limit — 100 Attack Occurrence Within — 60 (seconds)

Configuring a Biometrics Based Detection policy

Using Biometrics Based Detection policies, FortiADC can determine whether requests are generated by robots instead of a human by checking client events within a specified period. With JavaScript enabled on the client browser, FortiADC can collect behavioral biometrics (such as mouse movement, keyboard, screen touch, and scroll) and monitor as client events for a specified period. FortiADC can then determine whether the behavioral biometrics from the request is indicative of a bot or a human.

After you have configured Biometrics Based Detection policies, you can select them in WAF profiles.

Before you begin:

- You must have Read-Write permission for Security settings.

To configure a Biometrics Based Detection policy:

1. Go to **Web Application Firewall > Biometrics Based Detection**.
2. In the **Biometrics Based Detection** tab, click **Create New** to display the configuration editor.

3. Configure the following Biometrics Based Detection settings:

Setting	Description
Name	Specify a name for the Biometrics Based Detection rule. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No spaces. The configuration name cannot be edited once it has been saved.
Ignore JS Check	Enable/disable redirect to a warning page to enable JavaScript. This is disabled by default. • Disable — FortiADC will check if JavaScript is enabled on the client browser. If JavaScript is not enabled, then FortiADC will redirect to a warning page to let the user enable JavaScript. If the client does not enable JavaScript within 10 seconds, the traffic may be recognized as a bad bot. • Enable — FortiADC will not check if JavaScript is enabled on the client browser. If JavaScript is enabled on the client browser, events can be collected normally and FortiADC can determine if it is a bot or not. But if JavaScript is disabled on the client browser, the client will be recognized as a bot after the Event Collection Time, since events cannot be collected by FortiADC.
Monitor Client Events	Select one or more client events to monitor: • Mouse Movement • Click • Keyboard • Screen Touch • Scroll By default, Mouse Movement, Click, and Keyboard are preselected. If the configuration is saved with no Monitor Client Events selected, it will default to the preselected client events.
Event Collection Time	Specify for how long the events will be collected from the client. Default: 60 Range: 10-3600 seconds.
Bot Effective Time	Specify the time interval before FortiADC tests and verifies a bot again, once a bot has been detected. Default: 5 Range: 1-60 minute(s).
JS Request URL	Specify the URL to use to insert JavaScript code to the client browser. Default: `/fadc_client/default_index.js`.
Action	Specify a WAF action object to apply when a bot is detected. See Configuring WAF Action objects on page 671. The default action is alert.
Severity	Select the event severity to log when a bot is detected: • 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.

4. Click **Save**.
   Once the configuration is saved, the **URL List** becomes configurable. The Biometrics Based Detection policy will be applied to the request URLs in the URL List.
5. Under the **URL List** section, click **Create New** to display the configuration editor.
6. Configure the following URL List settings:

Setting	Description
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	The **Host** option is available if **Host Status** is enabled. 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, `www.example.com` matches `www.example.com` but not `www.example.com.hk`.
Request URL	The literal URL, such as `/index.php`, or a regular expression, such as `^/*.php` that the HTTP request must contain in order to match the rule. Multiple URLs are supported.

7. Click **Save**.
   Once the URL List configuration is saved, you are returned to the Biometrics Based Detection configuration editor.
8. Click **Save** again to apply the newly created URL List configuration to the Biometrics Based Detection configuration.

Configuring a Fingerprint Based Detection policy

Fingerprint Based Detection detects the client's fingerprint using multiple characteristics to ascertain whether an access request originates from a human or a bot. Using the Fingerprint Based Detection policy, FortiADC collects and monitors the behavioral fingerprint information, such as WebDriver, WindowsProperties and MimeTypesConsistent, with JavaScript enabled on the client browser. The collected information forms the client fingerprint that is then checked against known automation tools and frameworks to determine whether the requests are generated by automation tools such as Headless Chrome, Selenium or Electron.

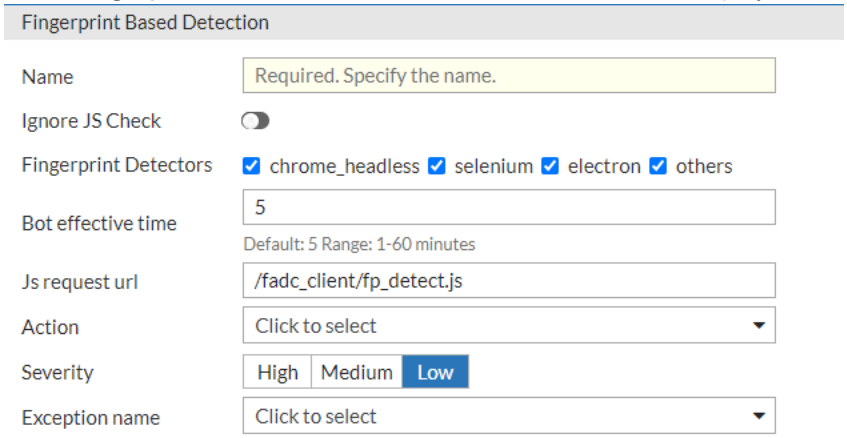
After you have configured Fingerprint Based Detection policies, you can select them in WAF profiles.

Before you begin:

- You must have Read-Write permission for Security settings.

To configure a Fingerprint Based Detection policy:

1. Go to **Web Application Firewall > Fingerprint Based Detection**.
2. In the **Fingerprint Based Detection** tab, click **Create New** to display the configuration editor.



3. Configure the following Fingerprint Based Detection settings:

Setting	Description
Name	Specify a name for the Fingerprint Based Detection rule. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No spaces. The configuration name cannot be edited once it has been saved.
Ignore JS Check	Enable/disable redirect to a warning page to enable JavaScript. This is disabled by default. • Disable — FortiADC will check if JavaScript is enabled on the client browser. If JavaScript is not enabled, then FortiADC will redirect to a warning page to let the user enable JavaScript. If the client does not enable JavaScript, Fingerprint Based Detection cannot be activated. • Enable — FortiADC will not check if JavaScript is enabled on the client browser. If JavaScript is enabled on the client browser, behavioral fingerprint information can be collected normally and FortiADC can determine if it is a bot or not. But if JavaScript is disabled on the client browser, Fingerprint Based Detection cannot be activated.
Fingerprint Detectors	Select one or more fingerprint detectors: • chrome_headless — Headless Chrome • selenium — Selenium • electron — Electron • others — includes SlimerJS and CefSharp All fingerprint detectors are preselected by default. If the configuration is saved with no Fingerprint Detectors selected, it will default to the preselected.
Bot Effective Time	Specify the time interval before FortiADC tests and verifies a bot again, once a bot has been detected. Default: 5 Range: 1-60 minute(s).
JS Request URL	Specify the URL to use to insert JavaScript code to the client browser. Default: `/fadc_client/fp_detect.js`.

Setting	Description
Action	Specify a WAF action object to apply when a bot is detected. See Configuring WAF Action objects on page 671. The default action is alert.
Severity	Select the event severity to log when a bot is detected: • 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.

4. Click **Save**.
   Once the configuration is saved, the **URL List** becomes configurable. The Fingerprint Based Detection policy will be applied to the request URLs in the URL List.
5. Under the **URL List** section, click **Create New** to display the configuration editor.
6. Configure the following URL List settings:

Setting	Description
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	The **Host** option is available if **Host Status** is enabled. 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, `www.example.com` matches `www.example.com` but not `www.example.com.hk`.
Request URL	The literal URL, such as `/index.php`, or a regular expression, such as `^/*.php` that the HTTP request must contain in order to match the rule. Multiple URLs are supported.

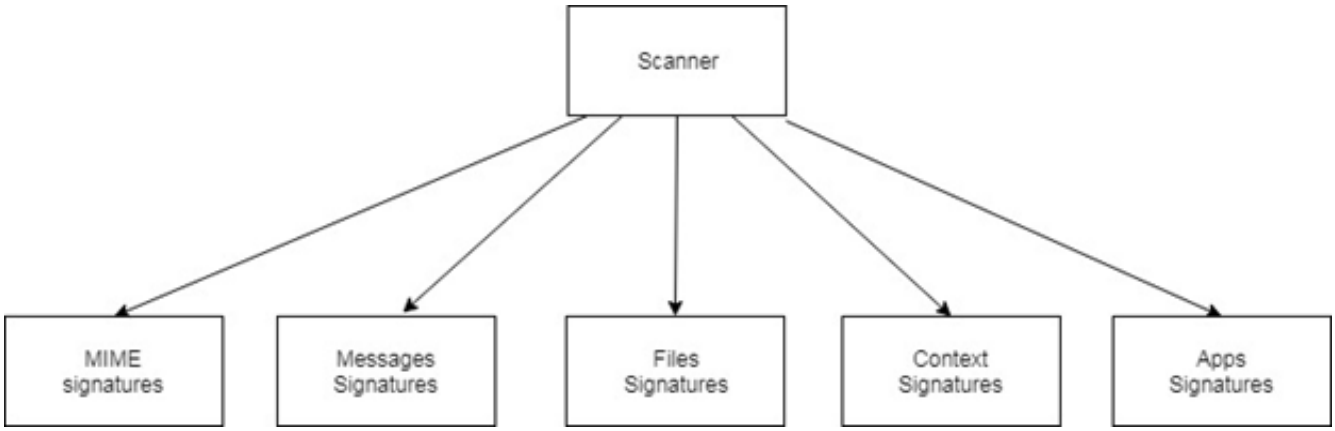
7. Click **Save**.
   Once the URL List configuration is saved, you are returned to the Fingerprint Based Detection configuration editor.
8. Click **Save** again to apply the newly created URL List configuration to the Fingerprint Based Detection configuration.

Once the Fingerprint Based Detection policy is saved, you can apply it to the WAF profile.

Web Vulnerability Scanner

From the **Web Application Firewall > Web Vulnerability Scanner** sub-menu, you can access 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 signature**s 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 signature**s 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.

A report will be generated after a web vulnerability scan is completed. FortiADC will generate a WAF profile based on the results of the scan report. For example, if the scan report detects an SQL injection vulnerability, a WAF profile containing SQL/XSS Injection Detection settings will be generated and attached to the VIP to protect servers behind VS.

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 Scan History on page 776

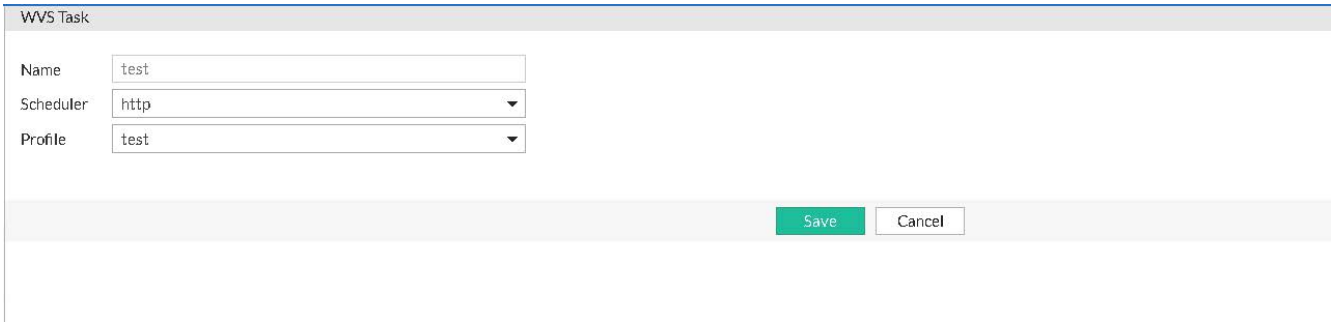
Notes

- 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 primary can start the scanning; it will be triggered only if it is primary.
- **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 Schedule Group on page 368
Profile	Select a profile. Profiles are configured under **WVS Profile**, the tab to the right of WVS task. See Configuring a WVS Profile on page 774

WVS Task Dialogue

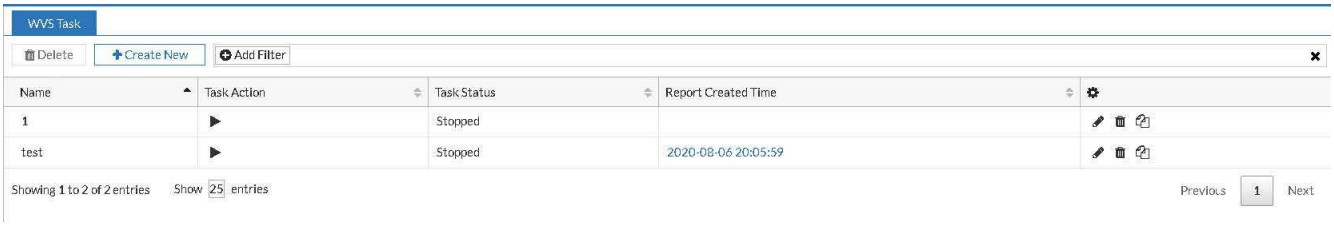


WVS Task dashboard

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.

Settings	Guidelines
	Diamond—Task Status "Stopped."
Report Created Time	Time the task was created.
	Edit Delete Clone

Run/Stop WVS-task



Scan Profile

Under the Web Vulnerability Scanner sub-menu, the Scan Profile page contains the features that allow you to configure WVS settings:

- Configuring a WVS Profile on page 774
- Configuring a WVS Login on page 775
- Configuring WVS Exceptions on page 776

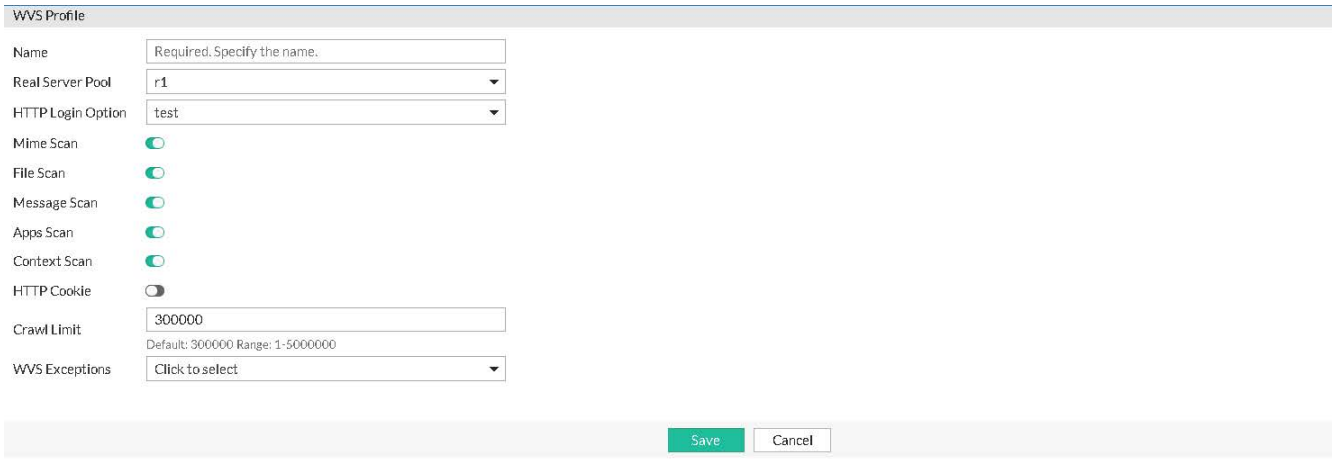
Configuring a WVS Profile

Creates a WVS Profile that can be selected in Web Vulnerability Scanner on page 771It gives you the option to select which types of scans you want.

To configure a 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 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 552
HTTP Login Option	Select an HTTP Login Option. Configure a login option in Configuring a WVS Login on page 775
Mime Scan File Scan Message Scan Apps Scan Context Scan	You have five scan options. See Web Vulnerability Scanner on page 771
HTTP Cookie	Enable HTTP Cookie in Web Vulnerability Scanner profile.
Crawl Limit	Specify a crawl limit.
WVS Exceptions	Specify a WVS Exception. See Configuring WVS Exceptions on page 776.

Configuring a WVS Login

To configure a WVS Login:

1. Go to **Web Application Firewall** > **Web Vulnerability Scanner** > **Scan Profile**. Select 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. 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.

Configuring WVS Exceptions

Creates a WVS Exception that can be selected in Web Vulnerability Scanner on page 771

To configure a WVS Exception:

1. Go to **Web Application Firewall** > **Web Vulnerability Scanner** > **Scan Profile**
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.

Scan History

After the scan is completed, the report is displayed on **Scan History** page.

WVS Report

Settings	Guidelines
Name	Name of the task.
Created Time	Time the report was created.
	On the far right. • Download the report. It can be imported in **Scan Integration**. See Scan Integration on page 778. • Delete the report. • Preview the report. • Generate Policy. See Generating Automatic Policy.
Add Filter	Add Filter—Sort by Name, Created Time

Generating Automatic Policy

By analyzing the scan results in the imported report, FortiADC automatically generates a WAF profile to prevent the reported attacks. In the Automatic Policy, you will required to specify the name of the generated WAF profile and the actions to be taken upon the attacks.

To edit the Automatic Policy:

1. In **Web Application Firewall** > **Web Vulnerability Scanner** > **Scan History**, find the desired report in the table, click the  Icon to edit the Automatic Policy.
2. Configure the following settings.

Generate Policies Automatically	If enabled, FortiADC will automatically generate an Automatic Policy and update it every time it runs the web vulnerability scan.
Merge the Report to Existing Profile	If disabled, FortiADC generates a new WAF profile based on the scan results. If enabled, the WAF settings based on the scan results will be merged to an existing WAF profile. If there are conflict settings, the new ones will overwrite the existing ones.
Profile Name	Enter a name for the newly generated WAF profile, and select an existing WAF profile.
Action - High	Select the action that FortiADC will take if High severity attacks are detected.
Action - Medium	Select the action that FortiADC will take if Medium severity attacks are detected.
Action - Low	Select the action that FortiADC will take if Low severity attacks are detected.

3. Click **Save**.

After the auto policy is saved, you can view it on **Scan Integration**. See Scan Integration.

Scan Integration

FortiADC generates a WAF profile based on the results of the scan report. For example, if the scan report detects an SQL injection vulnerability, a WAF profile containing SQL/XSS Injection Detection settings will be generated and attached to the VIP to protect servers behind VS.

The Automatic Policy contains the automatically generated WAF profile and specify the actions to be taken on the attacks. It is displayed on **Scan Integration** page.

You can also manually generate an Automatic Policy by importing a scan report . FortiADC supports scan reports from the following products:

- Acunetix
- IBM AppScan Standard
- WhiteHat
- HP WebInspect
- Qualys
- Telefonica FAAST
- ImmuniWeb
- FortiWeb
- FortiADC

To import a scan report:

1. Go to **Web Application Firewall** > **Web Vulnerability Scanner** > **Scan Integration**.
2. Click **Scanner File Import**.

3. Configure the following settings.

Scanner Type	Select the type of scanner report you want to import. • Acunetix • IBM AppScan Standard • WhiteHat • HP WebInspect • Qualys • Telefonica FAAST • ImmuniWeb • FortiWeb Scanner • FortiADC Scanner Some types of reports have specific requirements. For details, see WhiteHat Sentinel scanner report requirements, Telefónica FAAST scanner report requirements, and HP WebInspect scanner report requirements.
Upload File	Upload the scanner report file.
Generate Policies Automatically	This is by default enabled. If disabled, FortiADC will not generate an Automatic Policy the next time it runs Web Vulnerability Scan.
Merge the Report to Existing Profile	If disabled, FortiADC generates a new WAF profile based on the scan results. If enabled, the WAF settings based on the scan results will be merged to an existing WAF profile. If there are conflict settings, the new ones will overwrite the existing ones.
Profile Name	Enter a name for the newly generated WAF profile, and select an existing WAF profile.
Action - High	Select the action that FortiADC will take if High severity attacks are detected.
Action - Medium	Select the action that FortiADC will take if Medium severity attacks are detected.
Action - Low	Select the action that FortiADC will take if Low severity attacks are detected.

4. Click **Save**.

WhiteHat Sentinel scanner report requirements

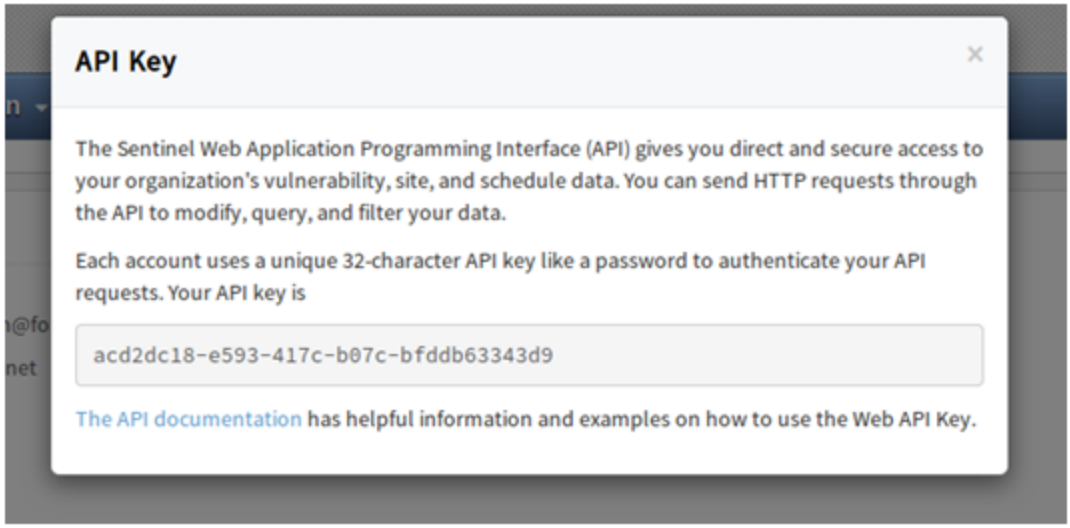
To allow [[[Undefined variable FortiWebVariables.FortiWeb]]] to generate rules using a WhiteHat Sentinel scanner report, ensure that the parameters “display_vulnerabilities” and “display_description” are enabled when you run the scan.

You can upload a WhiteHat Sentinel scanner report using either a report file you have downloaded manually or directly import the file from the WhiteHat portal using the RESTful API. Importing a scanner file from the WhiteHat portal requires the API key and application name that WhiteHat provides.

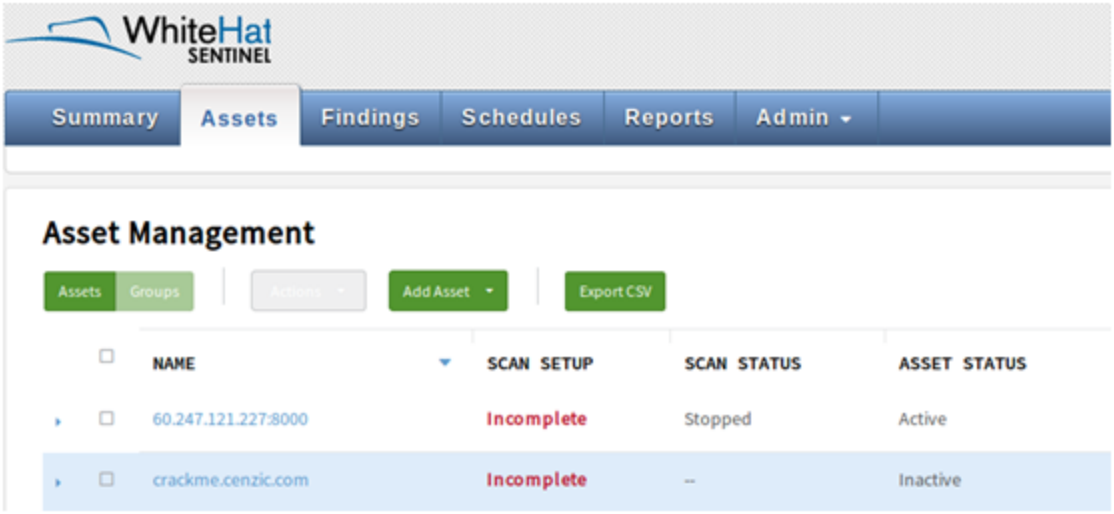
To retrieve the WhiteHat API key and application name

1. Go to the following location and log in: https://source.whitehatsec.com/summary.html#dashboard
2. In the top right corner, click **My Profile**.

3. Click View My API Key and enter your password.
   Your API key is displayed. For example:



4. To view the application name, navigate to the Assets tab. The application name is the NAME value. For example:

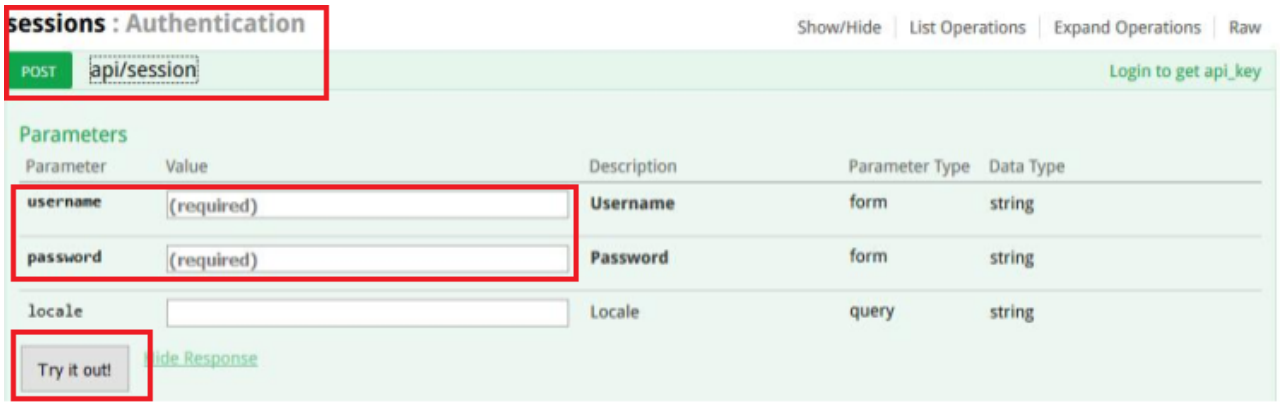


Telefónica FAAST scanner report requirements

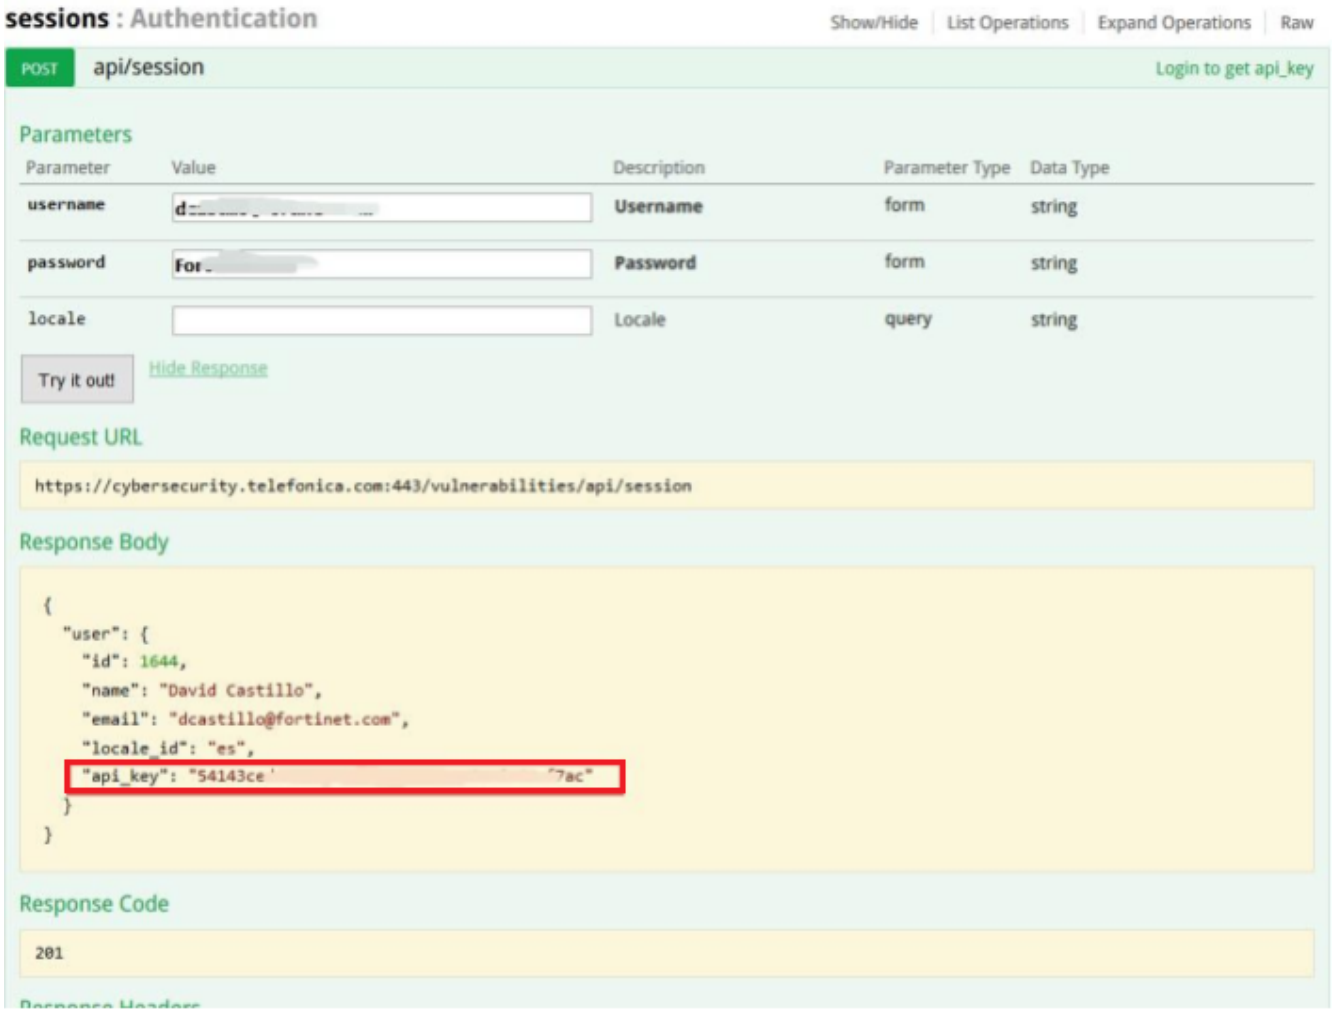
You can upload a Telefónica FAAST scanner report using either a report file you have downloaded manually or directly import the file from the Telefónica FAAST portal using the RESTful API. Importing a scanner file from the Telefónica FAAST portal requires the API key that Telefónica FAAST provides. One Telefónica FAAST scanner account can apply for an API key.

To apply for a Telefónica FAAST API key

1. Go to the following location and log in:
   https://cybersecurity.telefonica.com/vulnerabilities/es/api_docs
2. In the **session : Authentication** page, please select **POST > api/session** for the method, and fill in the blanks for **username** and **password**. Then click **Try it out**.



3. The API key will be gave in the **Response Body** if the username and password are authorized.



HP WebInspect scanner report requirements

To generate rules from HP WebInspect, when you export the report, for the **Details** option, select either **Full** or **Vulnerabilities**.

Web Anti-Defacement

The Web Anti-Defacement feature examines a website's files for changes at specified time intervals. If it detects a change that could indicate a defacement attack, it will notify you and quickly react by automatically restoring the website contents to the previous backup.

To configure a Web Anti-Defacement policy:

1. Go to **Web Application Firewall> Web Anti-Defacement**.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration.
4. Click **Test Connection** to test the connection between the FortiADC and the web server.
5. Save the 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.
Monitor	Enable/Disable to monitor the website's files for changes, and to download backup revisions for reverting the website to its previous revision.
Host Name/IP Address	Type the IP address or FQDN of the web server.
Connection Type	Select which protocol to use when connecting to the website in order to monitor its contents and download website backups. • FTP • SSH
Port	Enter the TCP port number on which the website's real server listens. The standard port number for FTP is 21; the standard port number for SSH is 22. The valid range is 1 to 65535.
Folder of Web Site	Type the path to the website's folder, such as public_html or wwwroot, on the real server. The path is relative to the initial location when logging in with the user name that you specify in Username.
Username	Enter the user name that the FortiADC will use to log in to the website's real server.
Password	Enter the password for the username you entered
Monitor Interval for Root Folder	Enter the time interval in seconds between each monitoring connection from the FortiADC to the web server. During this connection, the FortiADC examines Folder of Web Site(but not its subfolders) to see if any files have changed by comparing the files with the latest backup. If it detects any file changes, FortiADC will download a new backup revision. If you have enabled Restore in Automatic Action, FortiADC will revert the files to their previous version. The valid range is 1 to 86400 seconds and default value is 600 seconds.

Settings	Guidelines
Moniter Interval for Other Folder	Enter the time interval in seconds between each monitoring connection from the FortiADC to the web server. During this connection, the FortiADC examines subfolders to see if any files have been changed by comparing the files with the latest backup. If it detects any file changes, the FortiADC will download a new backup revision. If you have enabled Restore in Automatic Action, FortiADC will revert the files to their previous version. The valid range is 1 to 86400 seconds and default value is 600 seconds.
Skip Files Larger Than	Type a file size limit in kilobytes (KB) to indicate which files will be included in the website backup. Files exceeding this size will not be backed up. The valid range is 1 to 102400 KB and the default file size limit is 10240 KB. Note: Backing up large files can impact performance.
Skip Files with these Extensions	Type zero or more file extensions, such as iso, avi, to exclude from the website backup. Separate each file extension with a comma. Note: Backing up large files, such as video and audio, can impact performance.
Automatic Action	Select to decide which action will be excuted when the FortiADC detects file changes. • Disable - Accept changes and record the change in “Total Changed” table when FortiADC detects that the web site has been changed. You can manually restore the web site to a previous revision. • Acknowledge - Automatically accept changes to the web site when FortiADC detects that the web site has been changed • Restore - Enable to automatically restore the web site to the previous revision number when FortiADC detects that the website has been changed.

Accepting or reverting changed files

The anti-defacement feature maintains a list of files that have changed for each website it monitors. You can use this list to review, accept, and revert the changes.

To restore all the website files, use Automatic Action - Restore.

Alternatively, to automatically acknowledge all changes to files (for example, if you are updating the website), use Automatic Action - Acknowledge.

To accept or revert changed files:

1. Go to **Web Application Firewall > Web Anti-Defacement**. For the appropriate website, click the value in the **Total Changed** column.
2. Do one of the following:
   a. Select an item in the list, and then click the Acknowledge icon to accept the individual change. FortiADC clears the item from the list.
   b. Select an item in the list, and then click the Revert to icon. In the list of previous versions, click the Revert to this version icon for the version to revert to. FortiADC adds this revert action as a new version in the list.

WAF Adaptive Learning

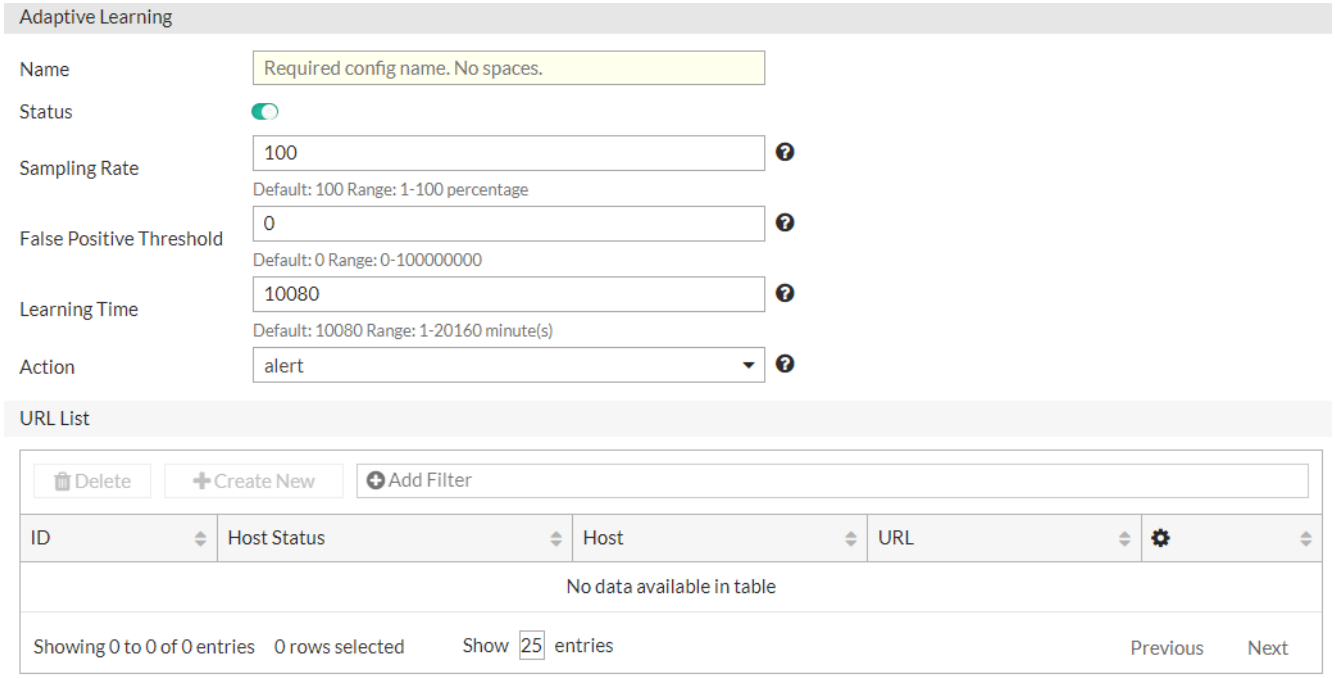
The new FortiADC WAF Adaptive Learning feature dynamically generates tailored recommendations to refine protection policies, enabling rapid and effective WAF configuration to counter threats. By continuously performing deep packet inspection and traffic analysis, the Adaptive Learning engine constructs comprehensive datasets from traffic patterns, allowing it to autonomously identify and mitigate threats with precision. Using these insights, the Adaptive Learning engine generates actionable recommendations for WAF policy optimization, ensuring security measures align with the unique traffic characteristics of each web application. This enhances both security posture and operational performance, allowing for real-time adjustment and fine-tuning of security measures.

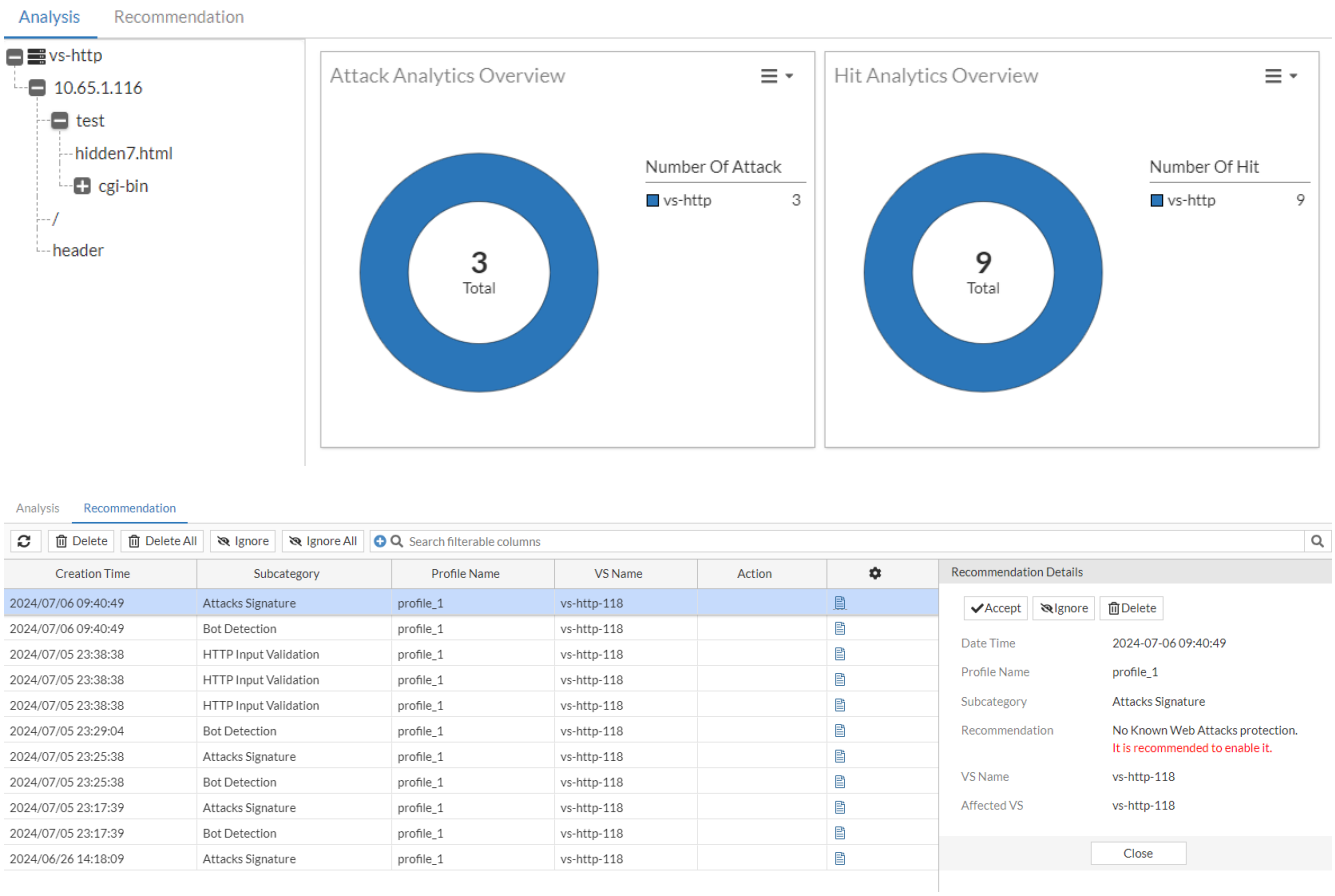
Note: The WAF Adaptive Learning feature requires the WAF Signature license or Application Security bundle license. If you do not already have a valid license, FortiADC offers a 30-day trial license to explore the WAF Adaptive Learning functionality. The trial license is activated automatically upon upgrading to FortiADC 7.6.0.

Here are some practical examples of how users can apply the WAF Adaptive Learning functionality:

- Create new WAF policies as part of a new virtual server setup — users can utilize the Adaptive Learning recommendation functionality to automatically generate and apply basic WAF policies for a new WAF Profile.
- Enhance the security of existing WAF policies — users can apply the Adaptive Learning recommendations that identify application traffic pattern changes to adjust policy settings accordingly.
- Identify false positive triggers that unnecessarily trigger WAF policy violations — users can leverage the Adaptive Learning engine's functionality to identify false positive triggers in WAF policies to adjust policy settings to exclude these false triggers.

The WAF Adaptive Learning engine features two key components: the Adaptive Learning policy that defines the dataset, and the Adaptive Learning statistics that include the analytical results and actionable recommendations. To learn more details about how to configure the Adaptive Learning policies and view the analysis results, see Configuring an Adaptive Learning policy on page 787 and Viewing Adaptive Learning Analysis and Recommendation on page 793.





In FortiADC 7.6.0, the Adaptive Learning feature supports the following WAF modules:

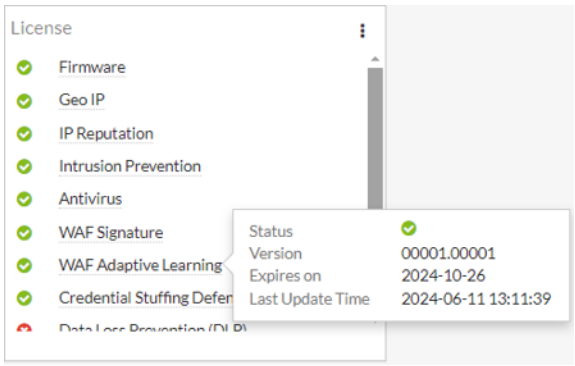
- Web Attack Signature
- Bot Detection
- Input Validation (Parameter validation and Hidden field validation)
- JSON Protection
- XML Protection

More WAF features will be supported in Adaptive Learning in future releases.

Licensing requirements

The WAF Signature license or Application Security bundle license (which covers Web Security Signatures and Adaptive Learning) is required to use the WAF Adaptive Learning feature.

If there is no valid WAF Signature license, a 30-day trial for the Adaptive Learning feature is available. This 30-day trial will be automatically activated upon the first boot up in FortiADC 7.6.0. Once the 30-day trial expires, all Adaptive Learning related configurations will become hidden in the GUI and the existing Adaptive Learning policy will stop working if no valid WAF Signature license is in place.



		Network Security	Application Security	AI Security
Network Security	IP Reputation and GeoIP	●	●	●
	Advanced Malware Protection	●	●	●
	Intrusion Prevention Service	●	●	●
Application Security	WAF Signatures and Adaptive Learning		●	●
	Credential Stuffing Defense		●	●
	Sandbox Cloud		●	●
	FortiGuard DLP		●	●
AI Security	Threat Analytics			●
	Advanced Bot Protection			●

Basic workflow:

1. Configure an Adaptive Learning policy. For detailed steps, see Configuring an Adaptive Learning policy on page 787.
2. Apply the Adaptive Learning policy in a WAF Profile.
3. Apply the WAF Profile to a virtual server.
4. Monitor Adaptive Learning statistics and take action according to Adaptive Learning recommendations. For details, see Viewing Adaptive Learning Analysis and Recommendation on page 793.



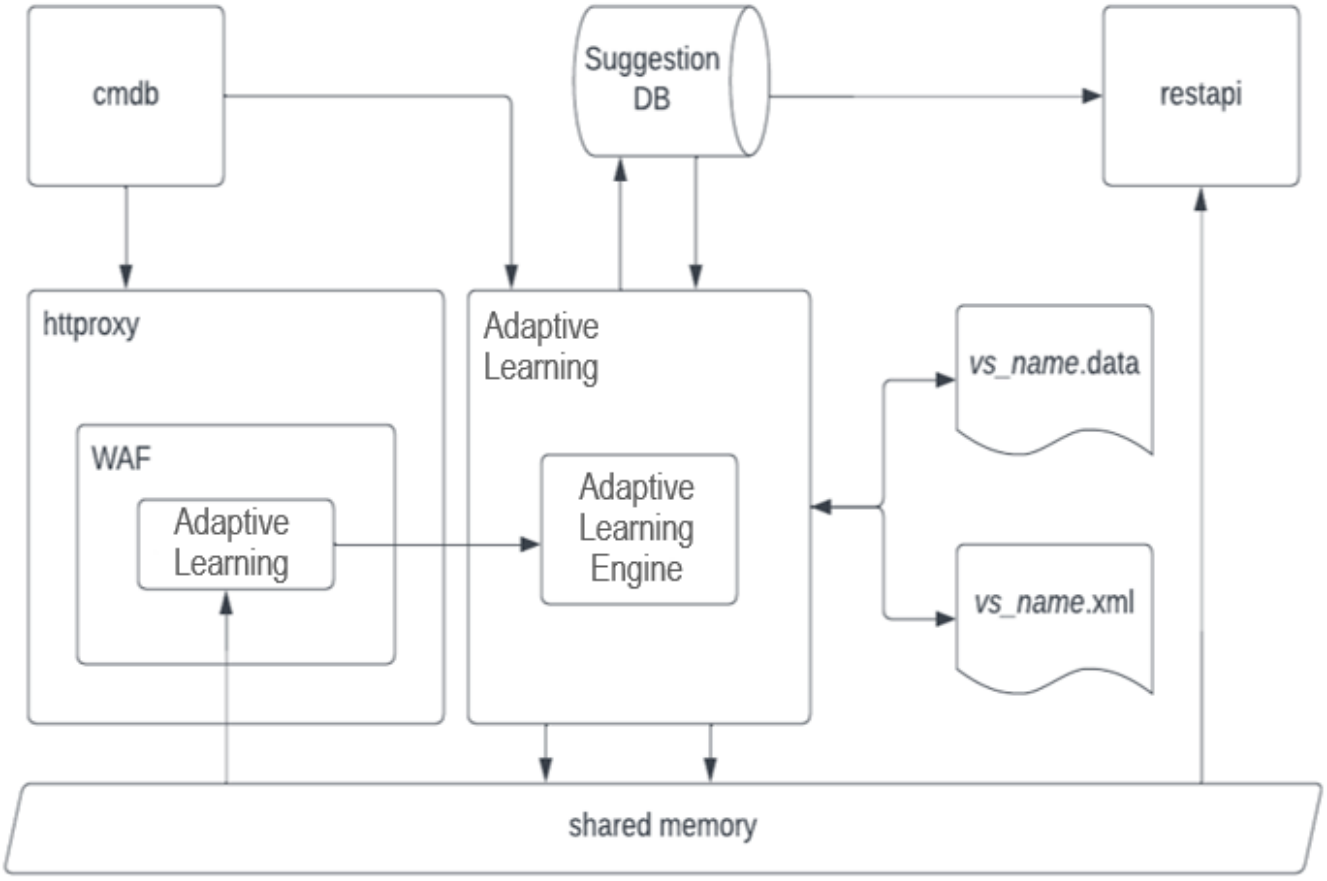
Be aware of several limitations in the Adaptive Learning feature implementation in version 7.6.0:

- Adaptive Learning statistics are not synchronized to HA peer nodes.
- Exceptions is currently not well supported — For example, when the incoming request matches the exception configured in the JSON/XML Protection policy or the allowlist of a Bot Detection policy, Adaptive Learning will still generate recommendations. However, if an Exception is configured in the WAF Profile, any traffic matched will be ignored by all WAF processing, including Adaptive Learning.
- For HTML Parameters or Hidden Field learning and recommendations, a GET request must be sent before sending a POST request; otherwise, the parameters may not be learned correctly. For the same virtual server, a specific parameter or hidden field will only support one URL, otherwise the rules sorted lower in the policy will not be updated correctly after accepting the recommendation.

Configuring an Adaptive Learning policy

To enable the FortiADC Adaptive Learning engine for continuous deep analysis of incoming traffic, you can configure an Adaptive Learning policy or use one of the three predefined configurations offered by FortiADC. By applying this to a WAF profile, the engine samples traffic at a defined rate based on the configured WAF policies. This traffic sampling allows the engine to dynamically learn and adapt to various web application elements, such as hostnames, URLs, query parameters, hidden fields, cookies, and file types. This advanced learning capability ensures that the WAF policies are precisely tuned to specific traffic characteristics, enhancing both security and performance.

The Adaptive Learning workflow begins with the httproxy module, where it receives HTTP requests and responses and then forwards them to the WAF Adaptive Learning module for inspection. The module filters the traffic, extracting critical elements such as URLs, source IP addresses, HTTP methods, parameters, cookies, and status codes. This extracted data is then dispatched to the Adaptive Learning engine for further analysis. Traffic is filtered based on a predefined sampling rate, and URLs with specific suffixes (e.g., .jpg, .ico, .mp3, .css) are excluded from the analysis.



The Adaptive Learning engine performs continuous analysis on the filtered traffic received from httproxy, storing statistical information in its memory and periodically flushing it to local database files. If the learning results demonstrate sufficient stability within the defined Learning Time, the Adaptive Learning engine can generate and propose specific recommendations.

In FortiADC 7.6.0, the Adaptive Learning feature supports the following WAF modules:

- Web Attack Signature
- Bot Detection
- Input Validation (Parameter validation and Hidden field validation)
- JSON Protection
- XML Protection



Consider the following module-specific guidelines and best practices to optimize the Adaptive Learning engine's analysis.

To learn the HTML parameters or hidden field correctly:

- Ensure a decompression policy is set, as it is typically required for proper parsing of HTML content.
- A GET request must be made before a POST request to ensure parameters are correctly recognized.

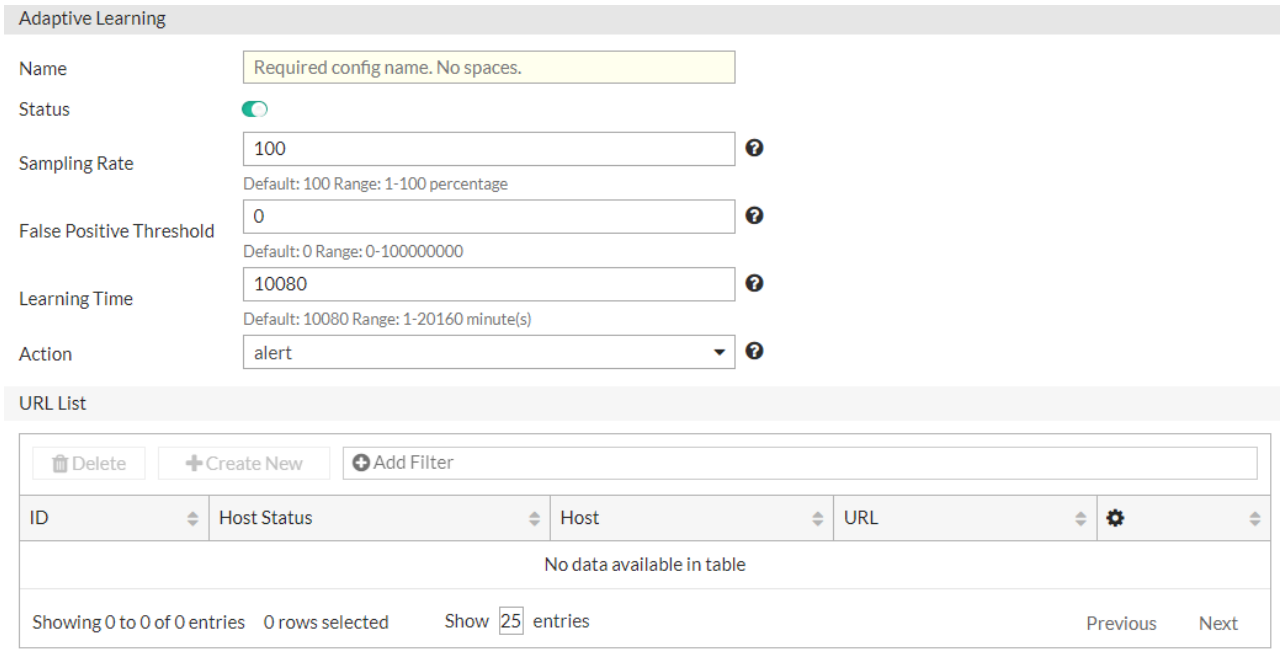
To learn JSON/XML correctly:

- A Decompression policy is required to correctly parse JSON/XML if the POST body is compressed with the HTTP header "Content-Encoding: gzip".
- The received POST request should include an HTTP header with either "Content-Type: application/json" or "Content-Type: application/xml".
- The JSON/XML content should adhere to the correct format.

Adaptive Learning policies are configured per VDOM, and will take effect only after it is applied to a WAF Profile. Once this WAF Profile is attached to a virtual server, the Adaptive Learning engine will begin to sample the incoming traffic to generate statistics and recommendations.

To configure an Adaptive Learning policy:

1. Go to **Web Application Firewall > Adaptive Learning**.
2. Click **Create New** to display the configuration editor.
3. Enable **Status** to view the Adaptive Learning configuration parameters.



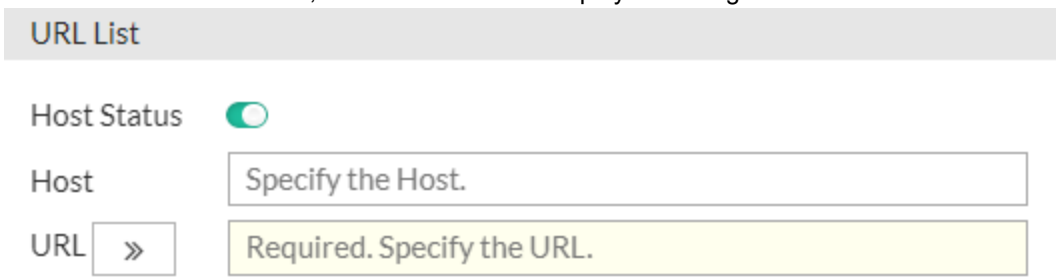
4. Configure the following Adaptive Learning policy settings:

Setting	Description
Name	Specify a unique name for the Adaptive Learning configuration object. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No space is allowed. Once saved, the name of the Adaptive Learning configuration cannot be changed.
Sampling Rate	Specify the percentage of received requests and their responses that will be sampled. For example, if the sampling rate is 50%, then for every 100 requests, the first 50 requests will be sampled. The default is 100, and the acceptable range is 1-100.

False Positive Threshold	Specify the threshold at which triggered events should be considered a false positive.
	In scenarios when requests that trigger a WAF policy violation are received from multiple different sources within a certain time period, the False Positive Threshold can be set to allow the Adaptive Learning engine to identify these triggered events as false positives and recommend adjustments to the WAF policy.
	The default is 0, and the acceptable range is 0-100000000.
	For example:
	`False Positive Threshold - 2`
	`Learning Time - 10`
	`WAF policy - WAF Signature Profile`
	When requests trigger a specific WAF signature ID violation are received from 2 different clients within the 10 minute Learning Time, then Adaptive Learning will generate a recommendation to disable the specific signature ID avoid triggering false positive results.

Setting	Description
Learning Time	Specify the Learning Time period in minutes. The default is 10080 minutes, and the acceptable range is 1-20160 minute(s). Adaptive Learning will only generate recommendations if the analysis (or "learning") results are "stable" within the specified time period. For the learning results to be stable, the Adaptive Learning engine must not detect any drastic flux in request rates, parameter lengths or types, or longer JSON/XML element names of values, among other configurable limit checks that are configurable in the policies.
Action	Set the action to take after you accept the recommendation for the WAF policy from Adaptive Learning. • alert — WAF policies will allow the traffic to pass and log the event. • deny — WAF policies will the drop current attack session by HTTP 403 message, and log the event. • block — WAF policies will drop the current attack session by HTTP 403 message and block the attacker (according the attacker's IP address) for 1 hour, and log the event. • silent-deny — WAF policies will drop the current attack session by HTTP 403 message, without logging the event. • captcha — WAF policies will allow the traffic to pass if the client successfully fulfills the CAPTCHA request, and log the event. The default action is alert.

5. Click **Save** to save the Adaptive Learning settings.
Once the configuration is saved, the **URL List** becomes configurable. The Adaptive Learning policy will be applied to the request URLs in the URL List.
6. Under the **URL List** section, click **Create New** to display the configuration editor.



7. Configure the following URL List settings:

Setting	Description
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	The **Host** option is available if **Host Status** is enabled.

Setting	Description
	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, `www.example.com` matches `www.example.com` but not `www.example.com.hk`. The maximum character length is 256.
Request URL	The literal URL, such as `/index.php`, or a regular expression, such as `^/*.php` that the HTTP request must contain in order to match the rule. Multiple URLs are supported. The maximum character length is 1024.

8. Click **Save**.
   Once the URL List configuration is saved, you are returned to the Adaptive Learning configuration editor.
9. Click **Save** again to apply the newly created URL List configuration to the Adaptive Learning configuration.

Once the Adaptive Learning policy is saved, you can apply it to the WAF profile.

Predefined Adaptive Learning policy configurations:

FortiADC offers three predefined Adaptive Learning policies you can apply directly in the WAF Profile or you can clone to use as a template to define your own policy. Please note that these predefined configurations are read-only and cannot be modified directly.

Predefined policy	Parameter	Setting
Fast_Learning	Sampling Rate	100
	False Positive Threshold	0
	Learning Time	1440
	Action	alert
	URL List	
	Host Status	disable
	URL	/
Medium_Learning	Sampling Rate	80
	False Positive Threshold	0
	Learning Time	10080
	Action	alert
	URL List	
	Host Status	disable
	URL	/

Predefined policy	Parameter	Setting
Slow_Learning	Sampling Rate	50
	False Positive Threshold	0
	Learning Time	20160
	Action	alert
	URL List	
	Host Status	disable
	URL	/

Viewing Adaptive Learning Analysis and Recommendation

From the **Adaptive Learning View** page, you can access statistical outputs derived from the Adaptive Learning engine's continuous, deep analysis of comprehensive datasets constructed from incoming traffic samples. The Adaptive Learning statistics are organized into two tabs: Analysis on page 793 and Recommendation on page 798.

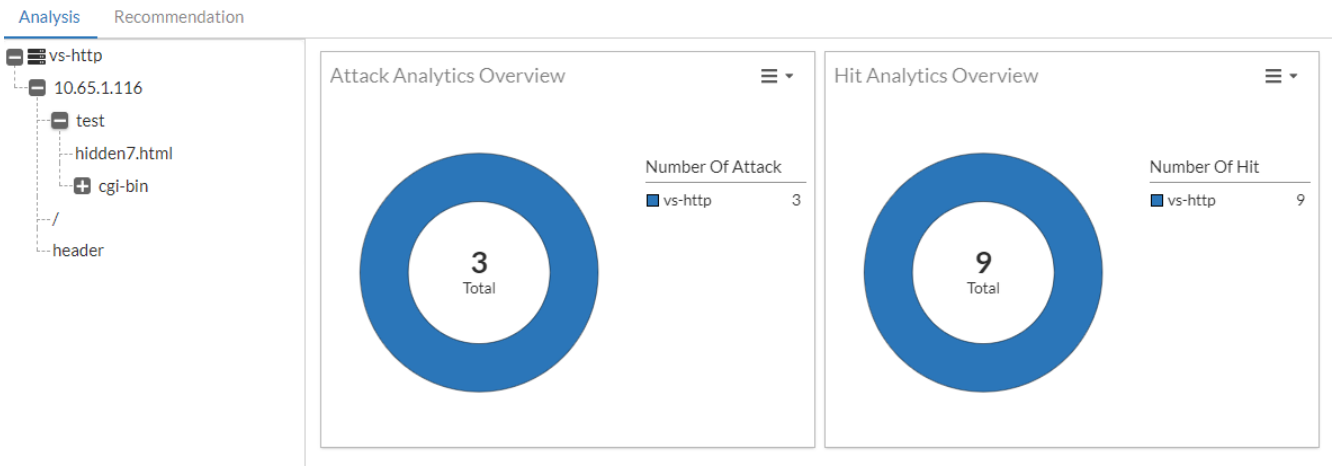
Analysis

The Analysis page provides graphical representations of the statistical outputs generated from the WAF Adaptive Learning engine. This page features a navigation tree that allows you to drill down to various levels of data. Each view presents a dashboard composed of various widgets that can be moved and resized. Statistics may be presented as tables or circle graphs that break-down the percentage against the total count.

Adaptive Learning statistics data are saved to the local database, ensuring persistence across reboots and upgrades. Additionally, if the interval between two requests received under the same virtual server exceeds one minute, the learned statistics will be written to the database file.

Overviews

The Overview dashboard serves as the landing page of the Analysis tab. Here, you can view the Attack Analytics Overview and Hit Analytics Overview graphs that provide a comprehensive overview of the statistics collected from all virtual servers with Adaptive Learning enabled.

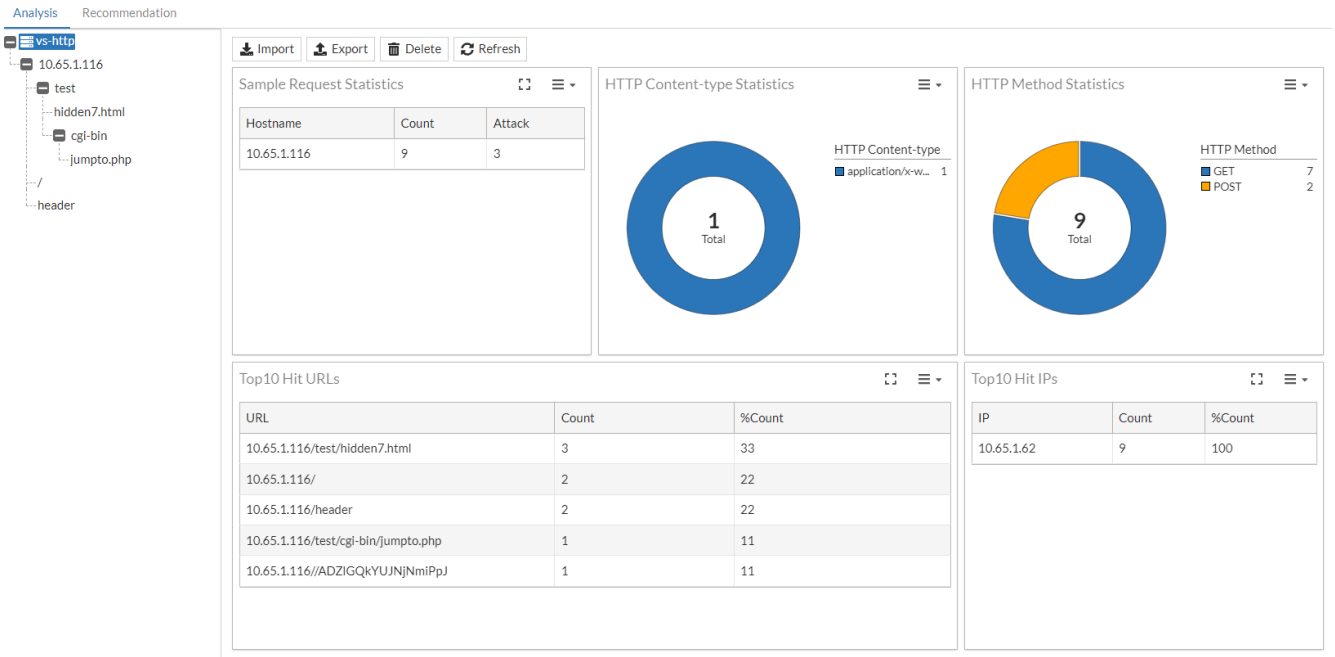


Virtual Server Statistics

At the virtual server level, various graphs and tables consolidate the Adaptive Learning analysis statistics from the virtual server traffic.

You can perform the following actions on the virtual server statistics:

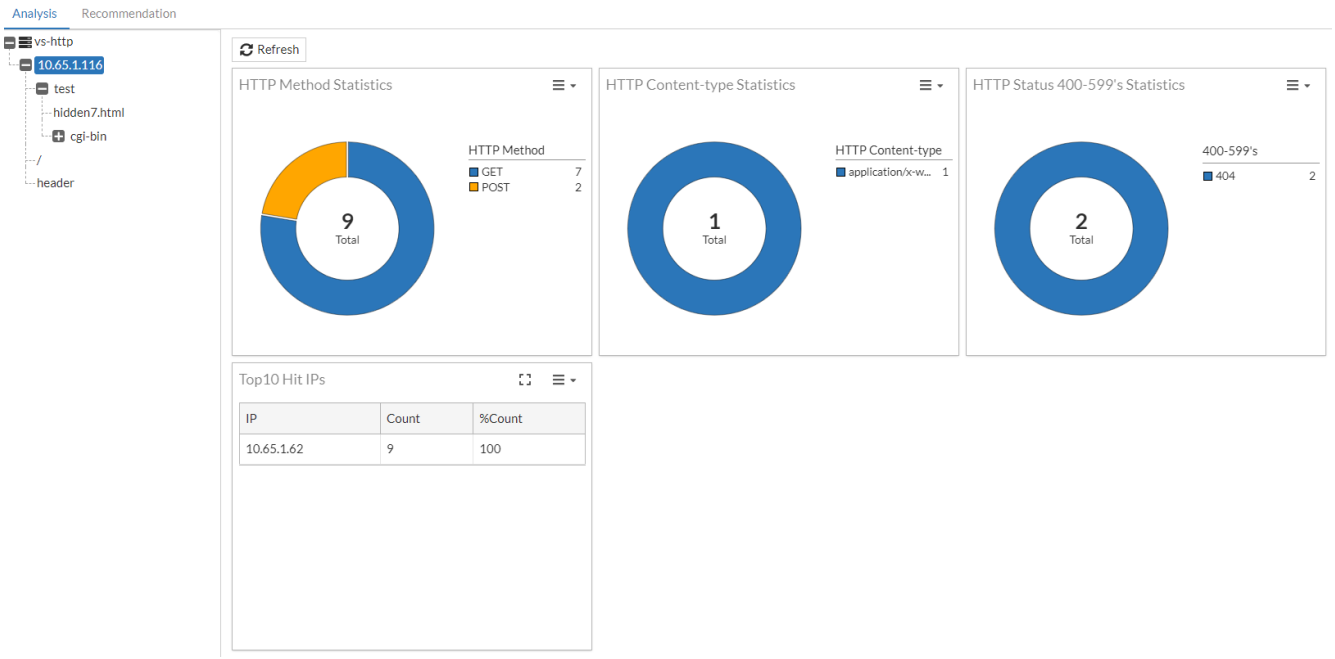
- Import — You may import a .zip file of statistics exported from another virtual server you wish to migrate.
- Export — The virtual server URL tree (.xml) and associated statistics (.bin) will be compressed into a .zip file and downloaded automatically. The .bin file statistics for each URL includes the parameters, hidden fields and other information that are used to determine if a recommendation should be generated.
- Delete — Both the URL trees and associated statistics will be cleared. All subsequent recommendations will be generated based on new traffic sampling.
- Refresh — Manually refresh the statistics for the current page.



Statistic	Description
Count	The number of requests sent to Adaptive Learning for processing.
Attack	The number of requests that are denied by WAF policies.
%Count	The percentage of the current requests relative to the total number of requests.

Host IP Statistics

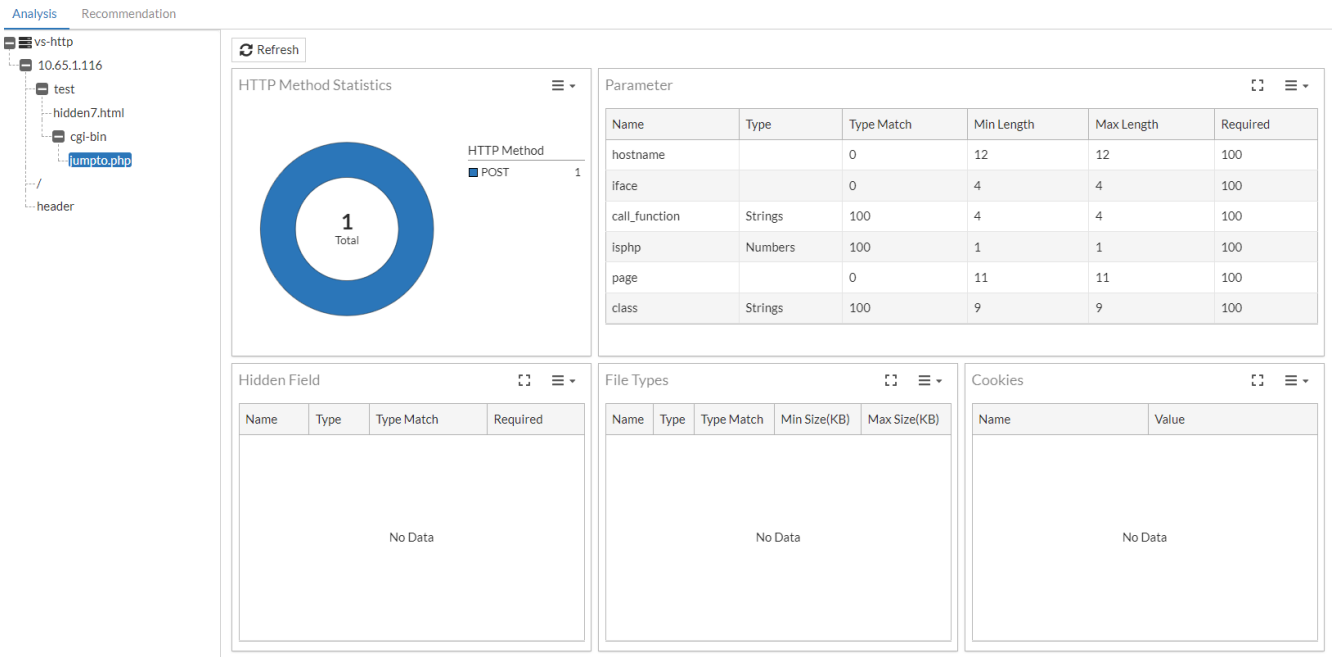
At the Host IP level, various graphs and tables consolidate the Adaptive Learning analysis statistics from each Host IP address.



Statistic	Description
Count	The number of requests sent to Adaptive Learning for processing.
%Count	The percentage of the current requests relative to the total number of requests.

URL Statistics

At the URL level, critical data about the URL is presented in a graph and several tables.



HTTP Method Statistics

The HTTP Method Statistics graph breaks down the types of methods that can be analyzed from this URL.

Parameter

The Parameter table captures the statistics used to generate the parameter validation policies.

Note: As HTML content is usually compressed in transferring, for most of the time, decompression should be enabled on the VS to correctly parse the HTML content.

Column	Description
Name	Parameters are parsed from the `<input>` fields within `<form>` elements or from query parameters in a URL in HTML. Parameters extracted from HTML forms are processed by the Adaptive Learning engine and used to generate recommendations. Parameters initially learned from query requests are displayed in the statistics table but do not generate recommendations. However, these parameters can be utilized to adjust those already learned from HTML forms.
Type	The type is parsed by the WAF's built-in DLP sensitive data-type database based on parameter values from the latest POST request, rather than the input type defined in the HTML source file. The input type parsed from the HTML form determines if a parameter should be learned or ignored, while the type displayed in the statistics table is parsed by the DLP database. • Parameter types ignored by Adaptive Learning: range, color, checkbox, radio, submit, reset, button, image. • Supported data types: Both predefined and customized data types are supported. If the type cannot be recognized by the Sensitive Data-Type

Column	Description
	(SDT) database, it will be shown as empty. Though the type is parsed from the POST request, a prior GET request is necessary to identify the parameter; otherwise, the parameter may not be recognized correctly.
Type Match	The percentage of the current/latest type among all received requests is calculated. Whenever the type changes, the count for the previous type is reset to 0.
Min Length	The minimum length of the value ever received for the current parameter.
Max Length	The maximum length of the value ever received for the current parameter.
Required	The percentage of times the parameter has a non-empty value out of the total occurrences. For example, if the parameter "fruit" is received 10 times and 3 times the value is empty, then the required percentage is 70%. This means that 70% of the times the parameter "fruit" is received, it has a non-empty value.

Hidden Field

The Hidden Field table captures the statistics used to generate recommendation for hidden field validation policies.

Column	Description
Name	Parameters with an input type "hidden" in the HTML form will be recognized as Hidden Field. Parameters are parsed from the `<input>` fields within `<form>` elements or from query parameters in a URL in HTML. Parameters extracted from HTML forms are processed by the Adaptive Learning engine and used to generate recommendations. Parameters initially learned from query requests are displayed in the statistics table but do not generate recommendations. However, these parameters can be utilized to adjust those already learned from HTML forms.
Type	The type is parsed by the WAF's built-in DLP sensitive data-type database based on parameter values from the latest POST request, rather than the input type defined in the HTML source file. The input type parsed from the HTML form determines if a parameter should be learned or ignored, while the type displayed in the statistics table is parsed by the DLP database. • Parameter types ignored by Adaptive Learning: range, color, checkbox, radio, submit, reset, button, image. • Supported data types: Both predefined and customized data types are supported. If the type cannot be recognized by the Sensitive Data-Type (SDT) database, it will be shown as empty. Though the type is parsed from the POST request, a prior GET request is necessary to identify the parameter; otherwise, the parameter may not be recognized correctly.
Type Match	The percentage of the current/latest type among all received requests is

Column	Description
	calculated. Whenever the type changes, the count for the previous type is reset to 0.
Required	The percentage of times the parameter has a non-empty value out of the total occurrences. For example, if the parameter "fruit" is received 10 times and 3 times the value is empty, then the required percentage is 70%. This means that 70% of the times the parameter "fruit" is received, it has a non-empty value.

File Types

The File Types table captures the statistics of the files uploaded through the requests.

Note: Currently, these statistics are not actively used to generate recommendations, however, these File Type statistics may be used to generate File Restriction policies in the future.

Column	Description
Name	Parsed from the multipart or form-data parts of the file.
Type	The type of the uploaded file is parsed by the WAF's built-in file-type database, which is used for file restriction policies. This is independent of the type defined in the multipart/form-data headers.
Type Match	The percentage of the current/latest type among all received requests is calculated. Whenever the type changes, the count for the previous type is reset to 0.
Min Size (KB)	The minimum size ever uploaded of the same file name.
Max Size (KB)	The maximum size ever uploaded of the same file name.

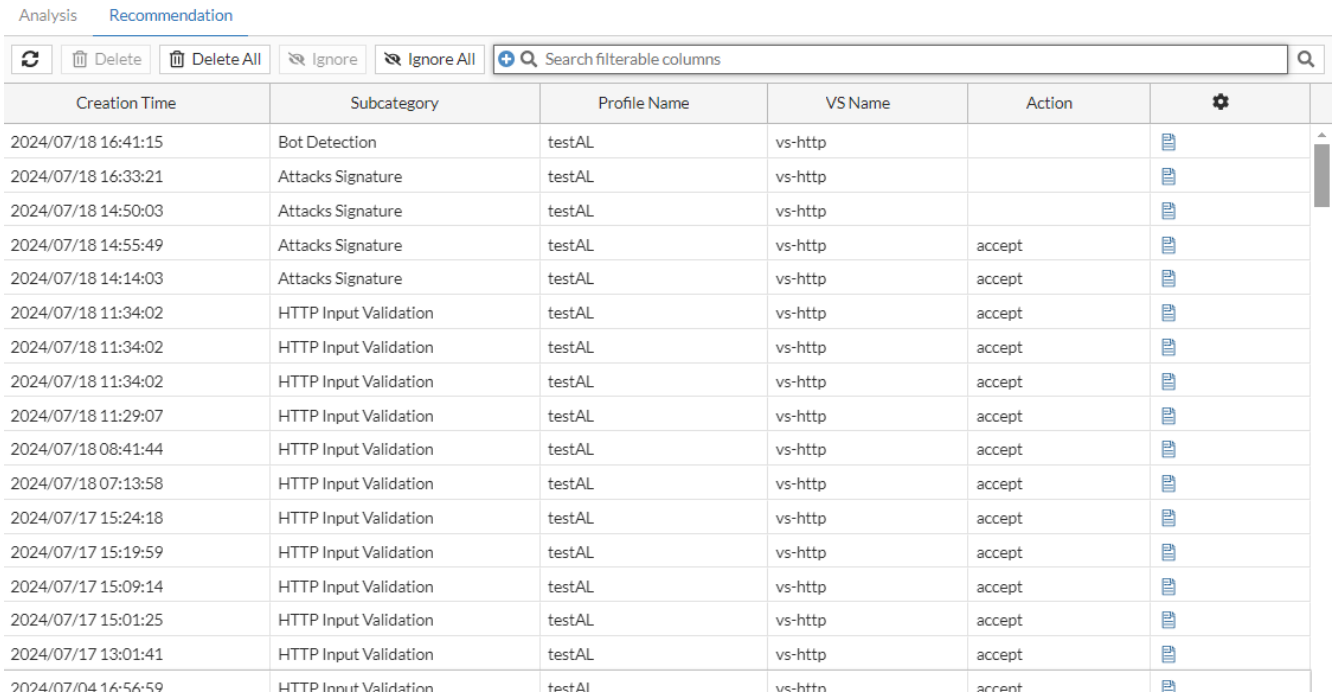
Cookies

The Cookies table captures the statistics of cookies through the requests. This statistic is not used to generate recommendations.

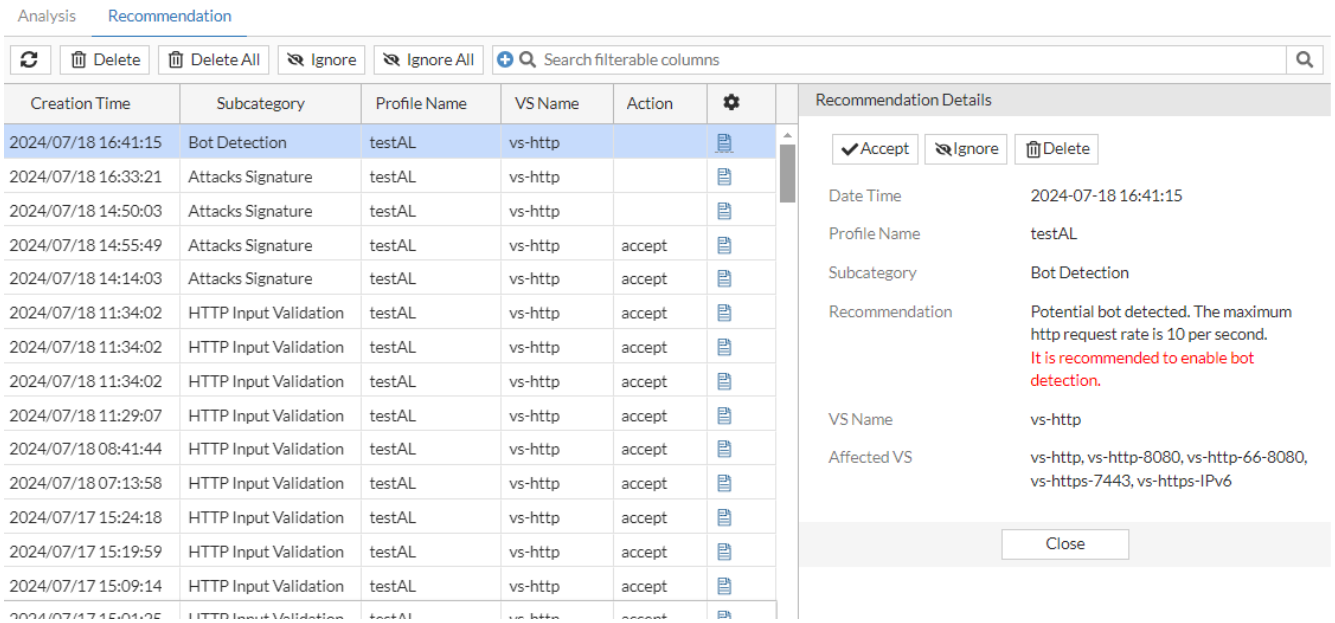
Column	Description
Name	Parsed from the cookie pairs from request the request header "Cookie: **name**=value".
Value	Parsed from the cookie pairs from request the request header "Cookie: name=**value**".

Recommendation

From the Recommendation page, you can view all the recommendations generated from the deep analysis performed by the Adaptive Learning engine.



To view the Recommendation Details, select the entry and click (page icon) to display the **Recommendation Details** dialog.



Date Time	Records the time when the recommendation was generated.
Profile Name	The WAF Profile name.
Subcategory	The WAF module to which the recommendation applies.
Recommendation	The rationale behind the recommendation and the specific action it advises.

VS Name	The name of the virtual server where the WAF Profile is applied.
Affected VS	Lists the virtual servers that will be affected by the action taken for this recommendation.

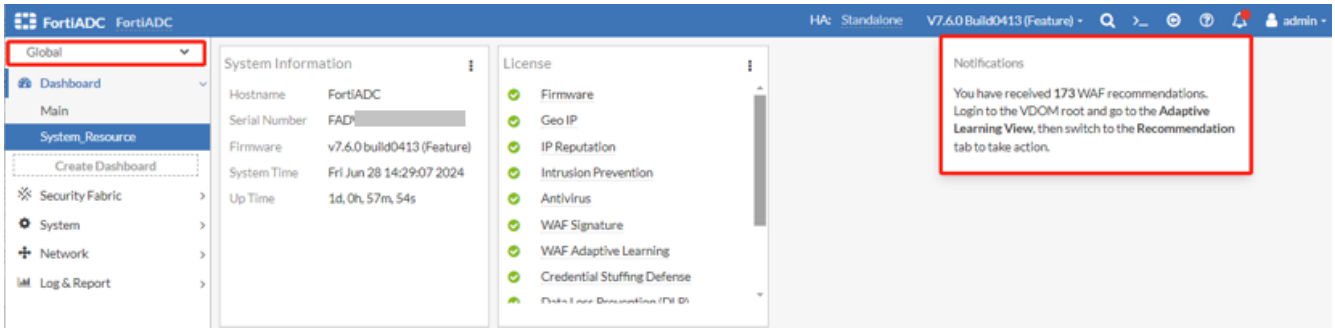
There are three actions you can take to address this recommendation:

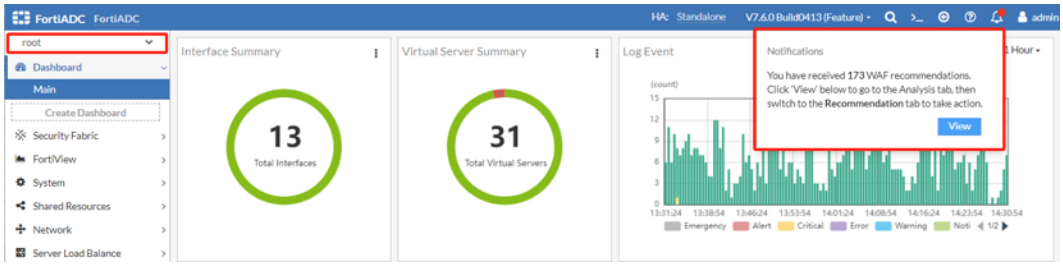
- Accept — The corresponding policy will be automatically generated or adjusted. If the current policy is read-only (a predefined policy), it will be replaced with a cloned version and adjusted with the recommended settings.
- Ignore — This recommendation will be ignored, but will still be kept for reference. The corresponding notification will be considered "read", decreasing the notification count.
- Delete — This recommendation will be deleted and removed from the recommendation list. The corresponding notification will be considered "read", decreasing the notification count.

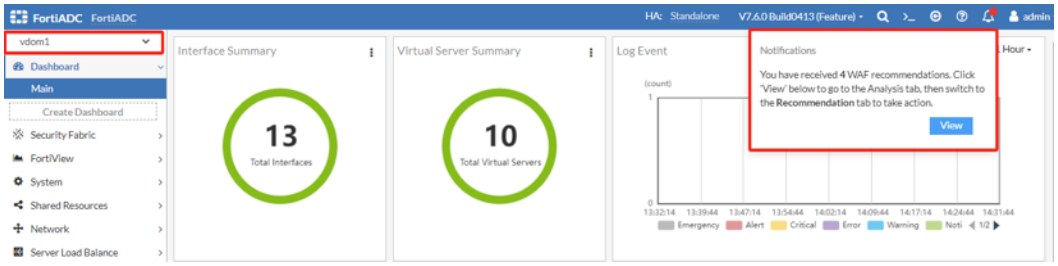
For practical demonstrations of how to use Adaptive Learning recommendations, see Practical Applications of Adaptive Learning on page 801.

Recommendation Notifications

Adaptive Learning generates notifications for each unresolved recommendation. The displayed recommendation count differs based on the context in which the notification is viewed, either from the Global VDOM or specific VDOMs (root or non-root). In the Global VDOM, the recommendation count aggregates only those generated within the root VDOM. Conversely, when viewed from individual VDOMs, the count reflects recommendations generated exclusively within the respective VDOM.







Practical Applications of Adaptive Learning

This section describes examples of how to practically apply the Adaptive Learning functionality for various WAF modules.

- Web Attack Signature examples on page 801
- Bot Detection examples on page 802
- HTML Parameter Validation examples on page 803
- JSON Protection examples on page 804
- XML Protection examples on page 807

Web Attack Signature examples

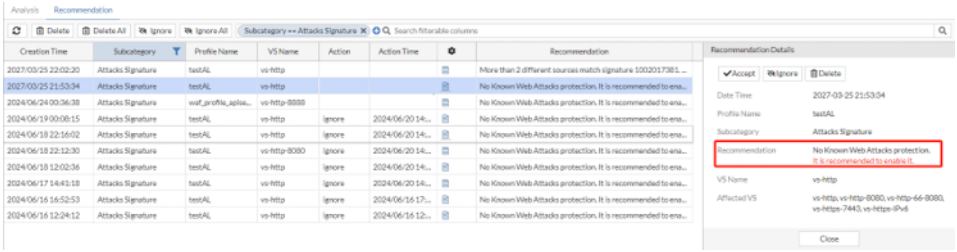
Triggering recommendations to set a new profile

If no WAF signature profile is configured in the WAF Profile, Adaptive Learning will recommend setting one upon receiving a request. Upon accepting the recommendation, a new WAF signature profile will be automatically created and attached to the WAF Profile.

This new WAF signature profile recommended by Adaptive Learning enables the following signatures:

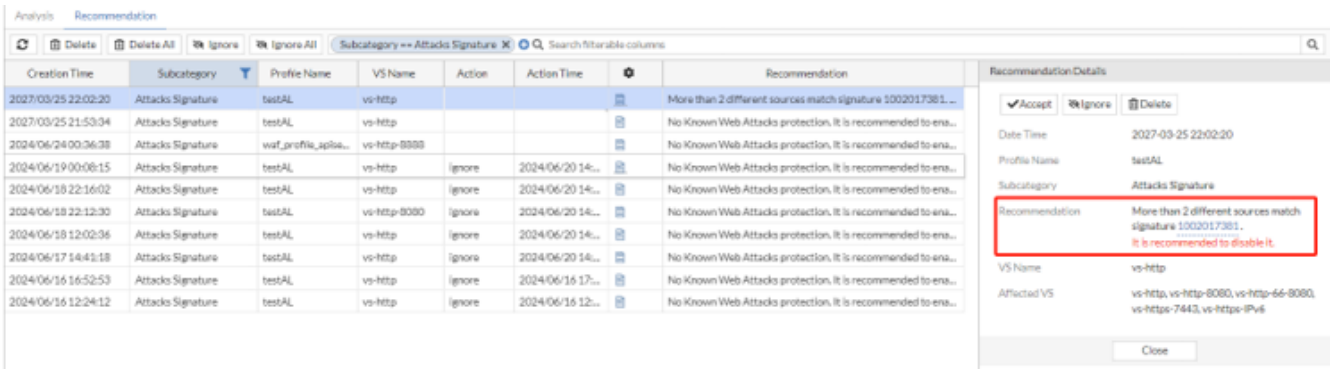
- Bad Robot
- Credit Card Detection
- Cross Site Scripting
- Generic Attacks
- Information Disclosure
- Known Exploits
- SQL Injection
- Trojans

The WAF action will be the same as the one set in the Adaptive Learning policy.



Triggering a False Positive recommendation to disable a signature ID

If requests triggering a specific signature ID violation exceed the client count set in the False Positive Threshold within the defined Learning Time, Adaptive Learning will recommend disabling that signature ID. Upon accepting the recommendation, the specific signature ID will be disabled in the WAF signature profile. If the signature profile in use is a predefined configuration, Adaptive Learning will automatically replace it with a cloned version where the affected signature ID is disabled.

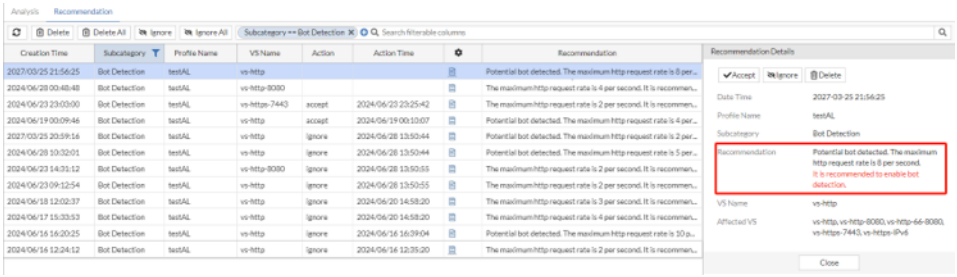


Bot Detection examples

Triggering recommendations to set a new policy

If no Bot Detection policy is configured in the WAF profile, the system injects JavaScript into the HTML file before delivering it to the client. The system then monitors for a subsequent request indicating that the JavaScript has been executed, verifying the client's support for JavaScript. If the client fails to respond, it is flagged as a potential bot, and a recommendation to configure a Bot Detection Policy is generated.

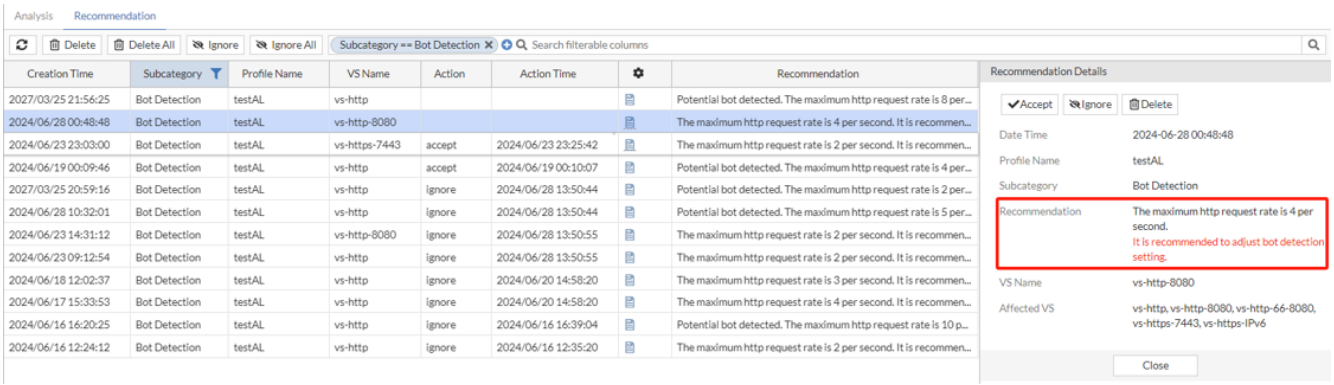
Upon accepting the recommendation, a Bot Detection policy will be created and attached to the WAF profile. This policy will enable detection for all Malicious Bots and Known Good Bots. The HTTP request rate will be set to match the current traffic sampling rate in the Adaptive Learning policy.



Triggering recommendations to adjust the request rate

When a Bot Detection policy is enabled, if the maximum HTTP request rate differs from the configured policy rate, Adaptive Learning will recommend to adjust the request rate.

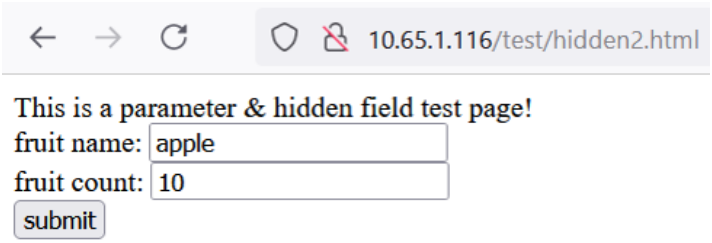
Upon accepting the recommendation, the HTTP request rate in the Bot Detection policy will be updated.



Triggering a False Positive recommendation to disable a malicious bot

When the Bad Robot Status option is enabled in the Bot Detection policy, if requests that trigger the bad bot detection come from more source IPs than the number specified in the False Positive Threshold, Adaptive Learning will recommend to disable the Bad Robot Status.

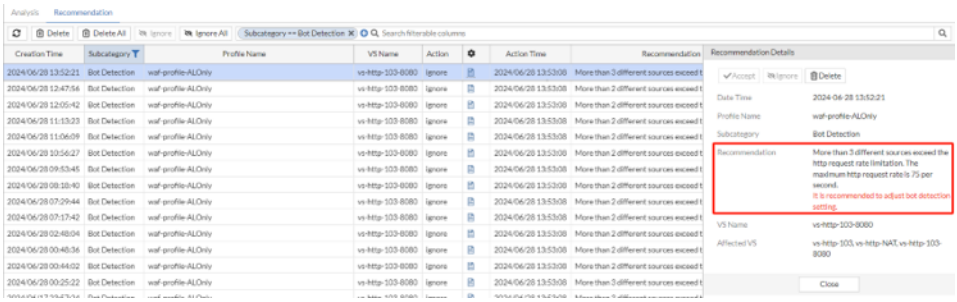
Upon accepting the recommendation, the Bad Robot Status in the Bot Detection policy will be disabled.



Triggering a False Positive recommendation to adjust the request rate

When the HTTP request rate is set, if the maximum request rate from more source IPs than the number specified in the False Positive Threshold exceeds the configured request rate within the defined Learning Time, Adaptive Learning will recommend to adjust the request rate.

Upon accepting the recommendation, the HTTP request rate in the Bot Detection policy will be updated.



HTML Parameter Validation examples

Triggering recommendations to set a new policy

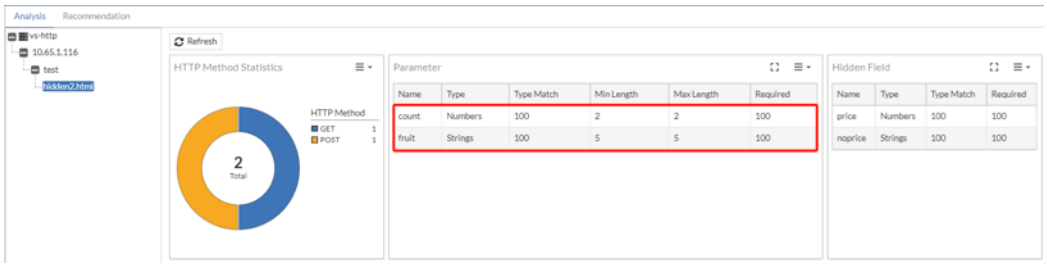
Adaptive Learning will recommend to set a Parameter Validation policy if a parameter is learned while there is no policy is set. The following conditions are considered:

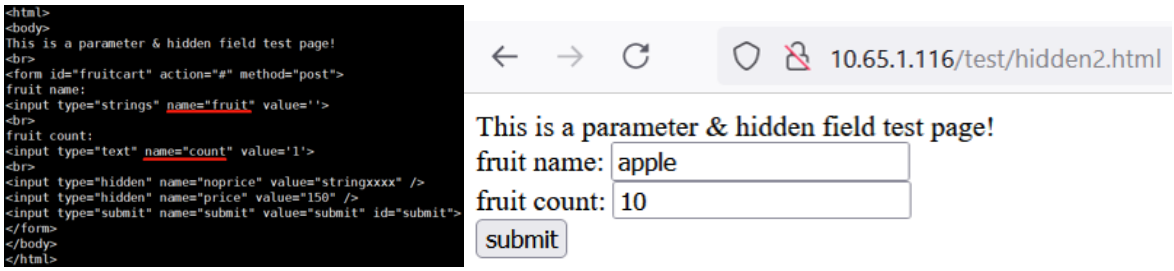
- No Input Validation policy is set in the WAF Profile.
- No Parameter Validation policy is included in the Input Validation policy in use.
- No Parameter Validation rule is set in the Parameter Validation policy in use.

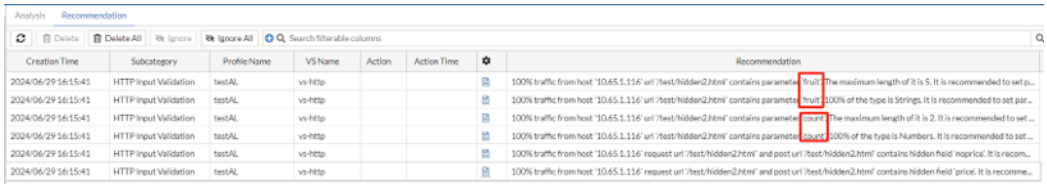
Typically, two recommendations will be generated for one parameter: one to set the Max Length, and another one to set the Data Type.

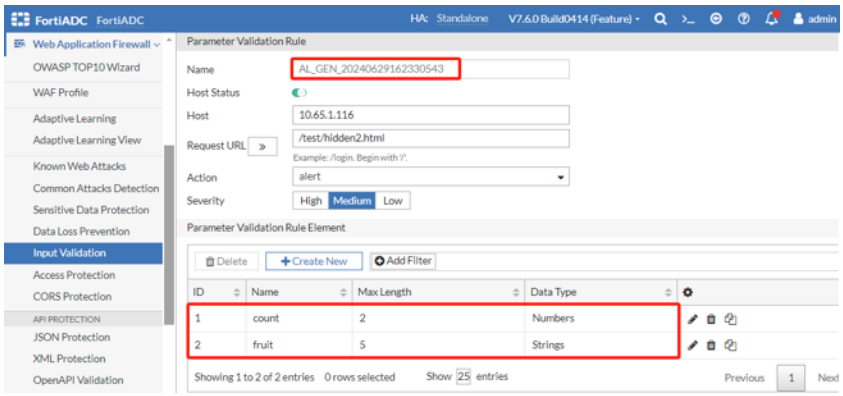
Upon accepting the recommendations, an Input Validation policy, a Parameter Validation policy or rule for the learned parameters will be created and attached to the WAF Profile. The policies will include the following specifications:

- The Request URL is set as the GET URL.
- The Max Length and Data Type are set according to the learned parameter.









JSON Protection examples

Triggering recommendations to set a new policy

Adaptive Learning will recommend to set a JSON Protection policy when JSON data is detected in the HTTP POST while there is not policy set.

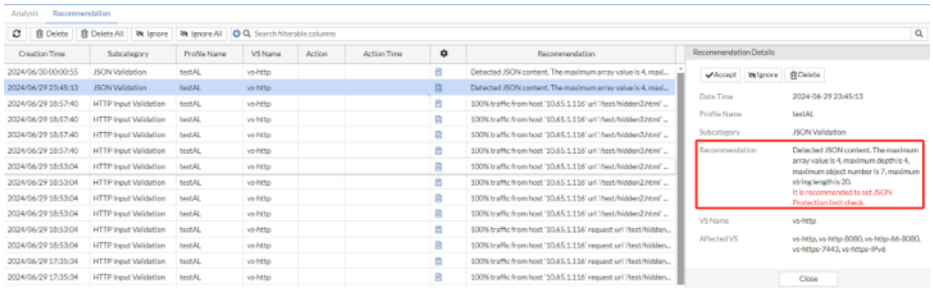
Upon accepting the recommendation, a JSON Protection policy will be created and attached to the WAF profile.

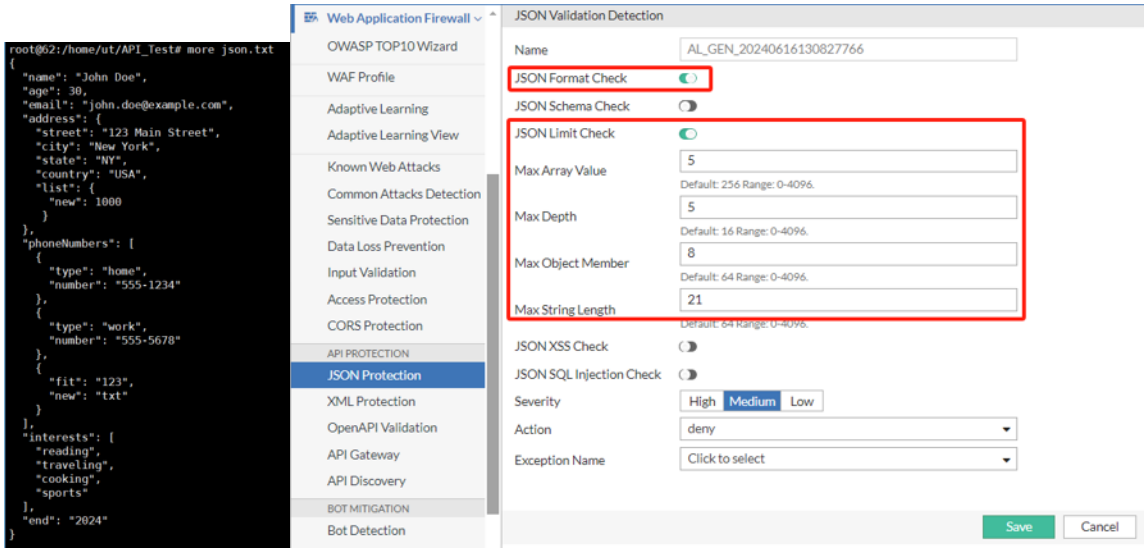
This new JSON Protection policy recommended by Adaptive Learning sets the following parameters:

- JSON Format Check is enabled.
- JSON Limit Check is enabled.
- Max Array Value, Max Depth, Max Object Member, Max String Length is set as the actual values parsed from the JSON data plus one. For example, the Array Value, Depth, Object Member, String Length parsed from the received

JSON data are 4, 4, 7, 20, then the Max Array Value, Depth, Max Object Member, Max String Length are set as 5, 5, 8, 21 respectively.

- Other parameters are kept as default values.
- The WAF action will be the same as the one set in the Adaptive Learning policy.



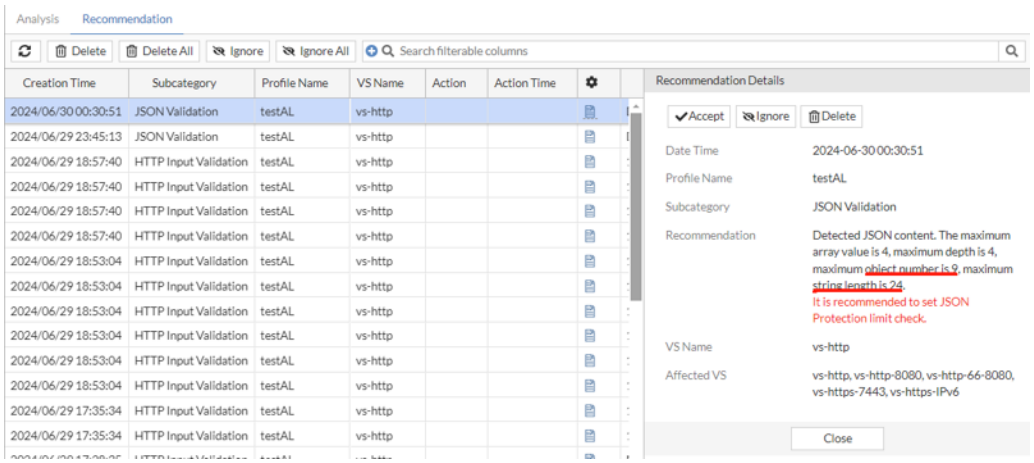


Triggering recommendations to adjust a policy

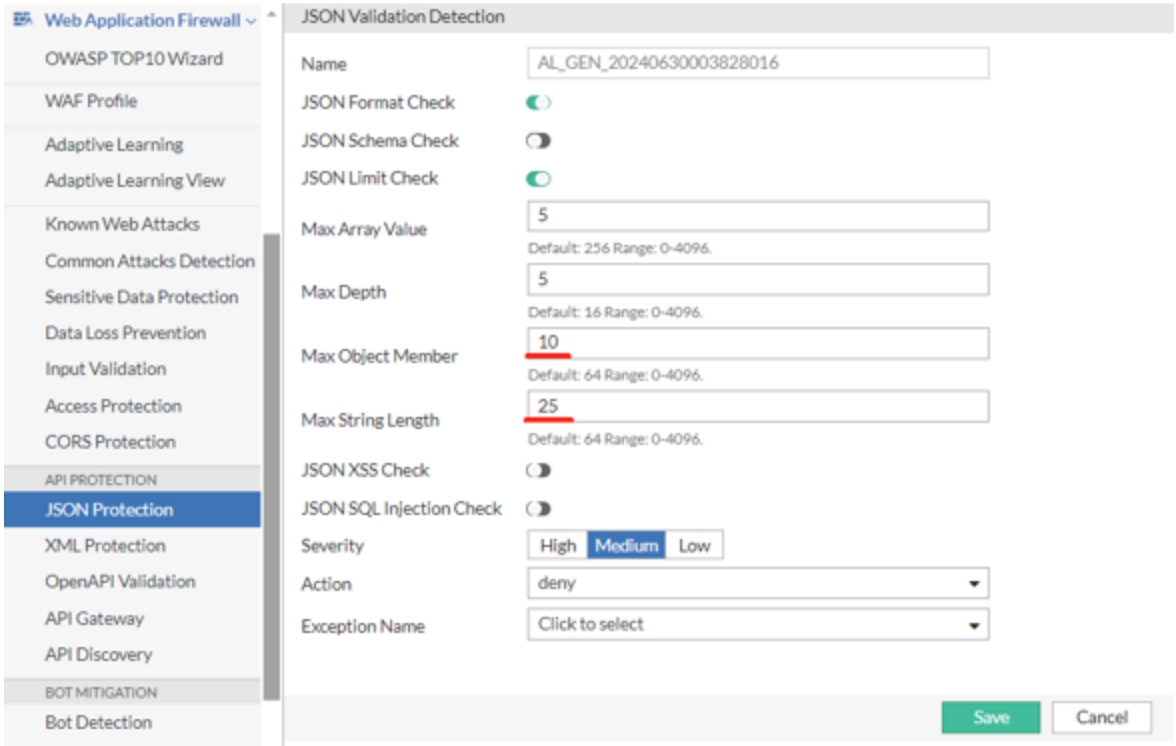
If any pattern parsed from the newly received JSON data exceeds the limits configured in the JSON protection policy, Adaptive Learning will recommend to adjust the JSON limits.

Upon accepting the recommendation, the corresponding limits in the policy will be updated:

- Max Array Value, Max Object Member, Max String Length will be set as the actual values parsed from the JSON data plus one.
- Other parameters will be kept without any changes.


Analysis
Recommendation
Delete
Delete All
Ignore
Ignore All
Search filterable columns
Creation Time
Subcategory
Profile Name
VS Name
Action
Action Time
2024/06/30 00:30:51 JSON Validation testAL vs-http
2024/06/29 23:45:13 JSON Validation testAL vs-http
2024/06/29 18:57:40 HTTP Input Validation testAL vs-http
2024/06/29 18:57:40 HTTP Input Validation testAL vs-http
2024/06/29 18:57:40 HTTP Input Validation testAL vs-http
2024/06/29 18:57:40 HTTP Input Validation testAL vs-http
2024/06/29 18:53:04 HTTP Input Validation testAL vs-http
2024/06/29 18:53:04 HTTP Input Validation testAL vs-http
2024/06/29 18:53:04 HTTP Input Validation testAL vs-http
2024/06/29 18:53:04 HTTP Input Validation testAL vs-http
2024/06/29 18:53:04 HTTP Input Validation testAL vs-http
2024/06/29 18:53:04 HTTP Input Validation testAL vs-http
2024/06/29 17:35:34 HTTP Input Validation testAL vs-http
2024/06/29 17:35:34 HTTP Input Validation testAL vs-http
Recommendation Details
Accept
Ignore
Delete
Date Time 2024-06-30 00:30:51
Profile Name testAL
Subcategory JSON Validation
Recommendation Detected JSON content. The maximum array value is 4, maximum depth is 4, maximum object number is 9, maximum string length is 24.
It is recommended to set JSON Protection limit check.
VS Name vs-http
Affected VS vs-http, vs-http-8080, vs-http-66-8080, vs-https-7443, vs-https-IPv6
Close

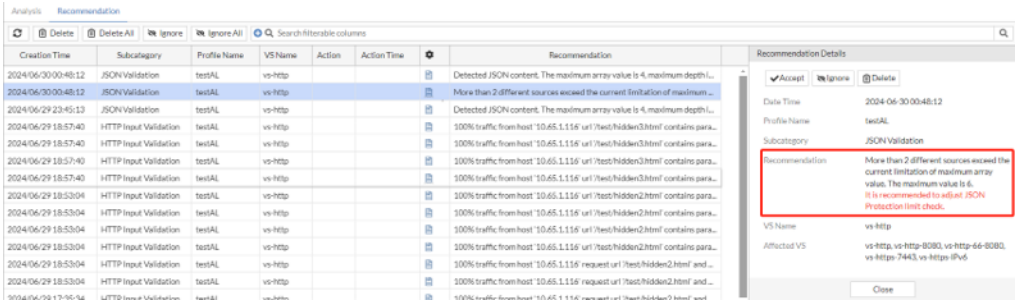

root@62:/home/ut/API_Test# more json2.txt
{
"name": "John Doe",
"age": 30,
"email": "john.doe2024@example.com",
"address": {
"street": "123 Main Street",
"city": "New York",
"state": "NY",
"country": "USA"
},
"phoneNumbers": [
{
"type": "home",
"number": "555-1234"
},
{
"type": "work",
"number": "555-5678"
}
],
"interests": [
"reading",
"traveling",
"cooking"
],
"xxxxxx": "llllll",
"yyyyyy": "mmmmm",
"zzzzzz": "nnnnnn"
}

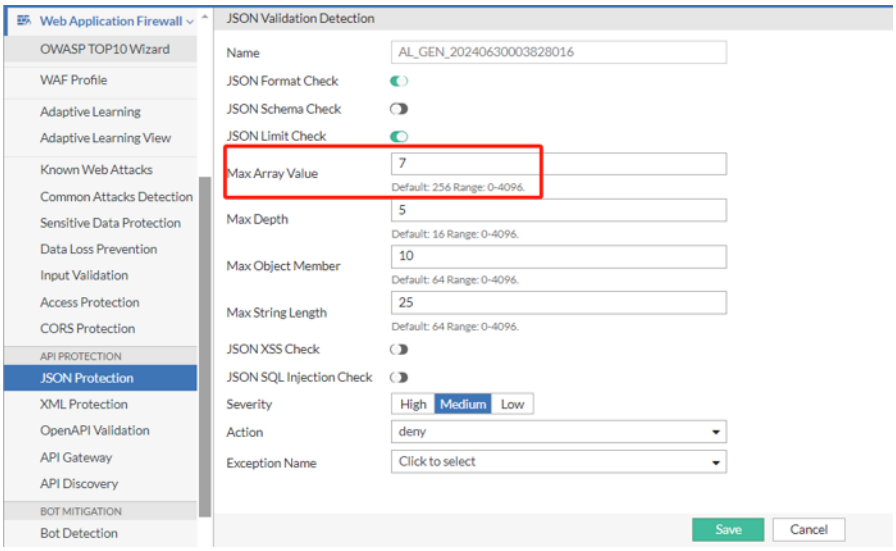

Web Application Firewall
OWASP TOP10 Wizard
WAF Profile
Adaptive Learning
Adaptive Learning View
Known Web Attacks
Common Attacks Detection
Sensitive Data Protection
Data Loss Prevention
Input Validation
Access Protection
CORS Protection
API PROTECTION
JSON Protection
XML Protection
OpenAPI Validation
API Gateway
API Discovery
BOT MITIGATION
Bot Detection
JSON Validation Detection
Name AL_GEN_20240630003828016
JSON Format Check
JSON Schema Check
JSON Limit Check
Max Array Value 5
Default: 256 Range: 0-4096.
Max Depth 5
Default: 16 Range: 0-4096.
Max Object Member 10
Default: 64 Range: 0-4096.
Max String Length 25
Default: 64 Range: 0-4096.
JSON XSS Check
JSON SQL Injection Check
Severity High Medium Low
Action deny
Exception Name Click to select
Save
Cancel

Triggering a False Positive recommendation to adjust a policy

When a JSON protection policy is configured and the number of source IPs transmitting JSON data with elements exceeding the configured limits surpasses the False Positive Threshold, Adaptive Learning will recommend to adjust the JSON limits.

Upon accepting the recommendation, the corresponding limits in the policy will be updated accordingly.





XML Protection examples

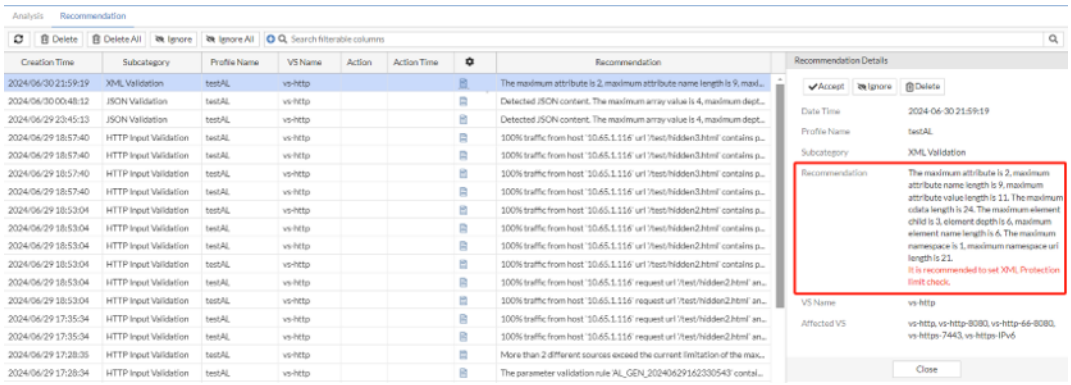
Triggering recommendations to set a new policy

Adaptive Learning will recommend to set a XML Protection policy when XML data is detected in the HTTP POST while there is not policy set.

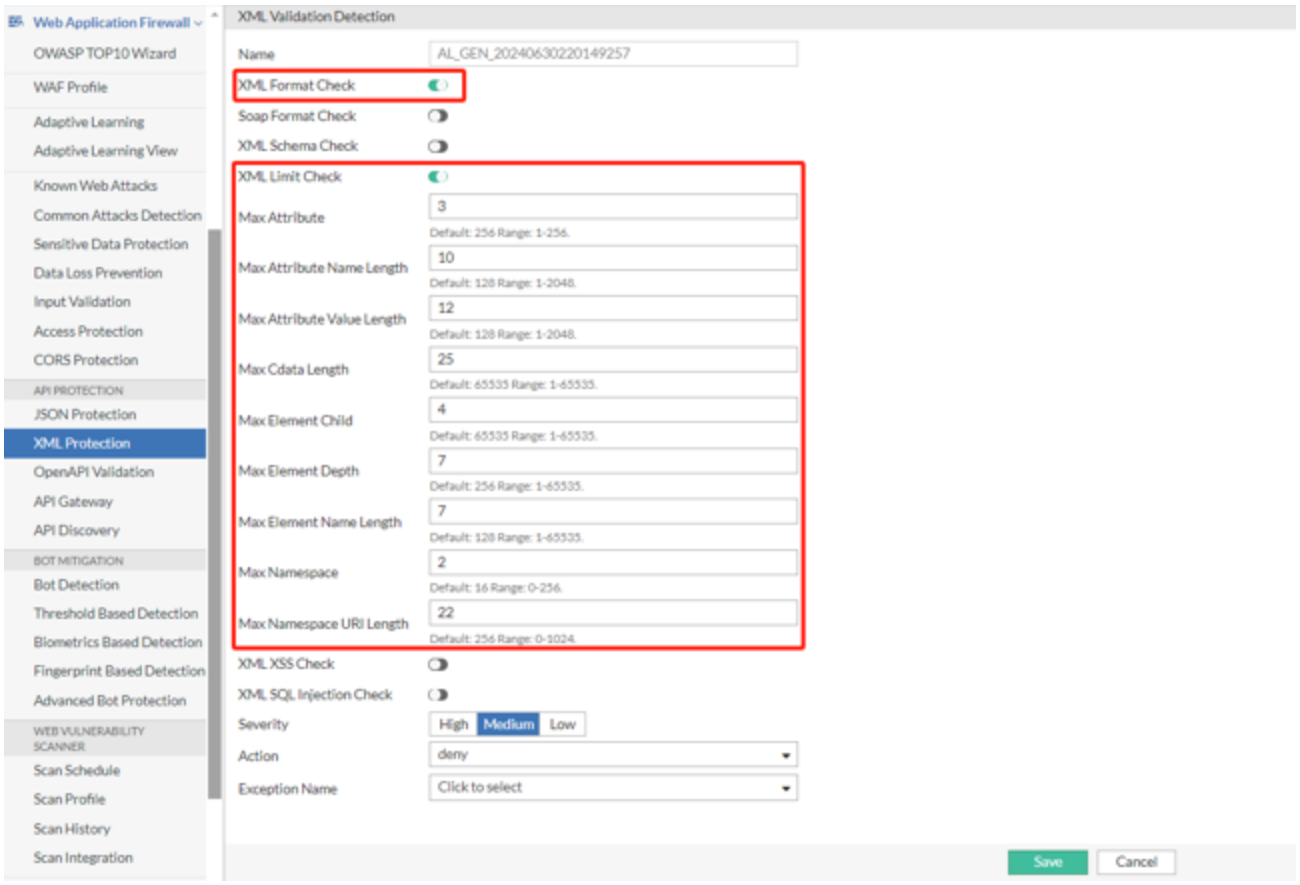
Upon accepting the recommendation, a XML Protection policy will be created and attached to the WAF profile.

This new XML Protection policy recommended by Adaptive Learning sets the following parameters:

- XML Format Check is enabled.
- XML Limit Check is enabled.
- Max Attribute, Attribute Name Length, Attribute Value Length (namespace value length is excluded here), Cdata Length, Element Child, Element Depth, Element Name Length, Max Namespace, Namespace URI Length plus one.
- Other parameters are kept as default values.
- The WAF action will be the same as the one set in the Adaptive Learning policy.




```
root@62:/home/ut/API_Test# more xml_full.txt
<?xml version="1.0"?>
<books xmlns:ns="http://example.com/ns">
  <book id="book3" available="true">
    <name><![CDATA[A Little Life]]></name>
    <reader type="fiction">
      <uname id="1">
        <tip>no</tip>
      </uname>
      <uname id="2">Tom</uname>
    </reader>
  </book>
  <book id="book4" available="false">
    <name>Darkmans</name>
    <reader type="fiction">
      <uname id="3">Jerry Fan</uname>
      <uname id="4">Nancy Liu</uname>
    </reader>
  </book>
  <ns:book id="book5" available="true">
    <ns:name><![CDATA[Sapiens: A Brief History]]></ns:name>
    <ns:reader type="non-fiction">
      <ns:uname id="6">Eva Green</ns:uname>
      <ns:uname id="7">Sam Harris</ns:uname>
    </ns:reader>
  </ns:book>
</books>
```

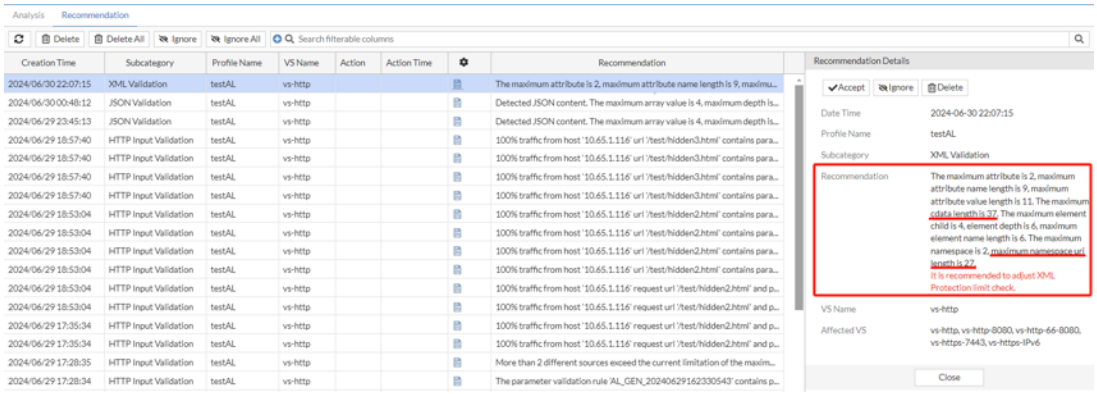


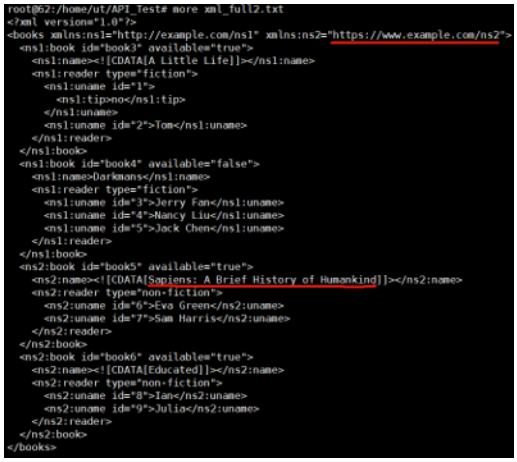
Triggering recommendations to adjust a policy

If any pattern parsed from the newly received XML data exceeds the limits configured in the XML protection policy, Adaptive Learning will recommend to adjust the XML limits.

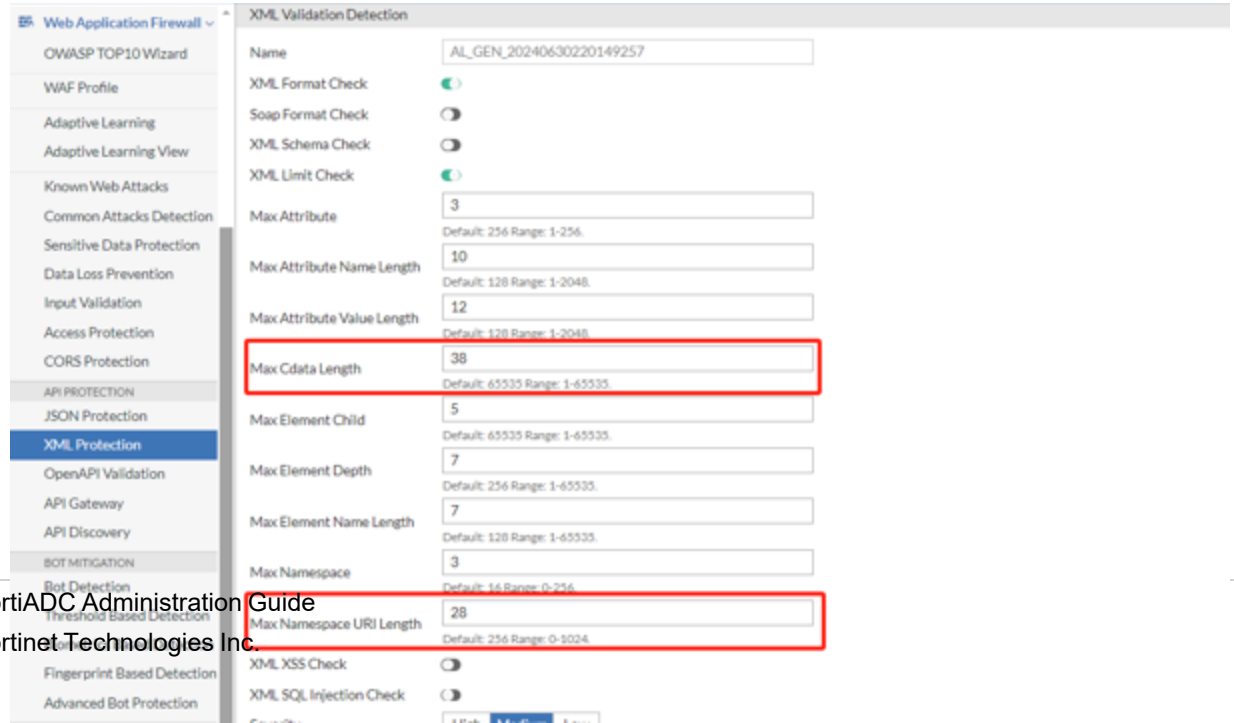
Upon accepting the recommendation, the corresponding limits in the policy will be updated:

- Max Attribute, Attribute Name Length, Attribute Value Length (namespace value length is excluded here), Cdata Length, Element Child, Element Depth, Element Name Length, Max Namespace, Namespace URI Length.
- Other parameters will be kept without any changes.





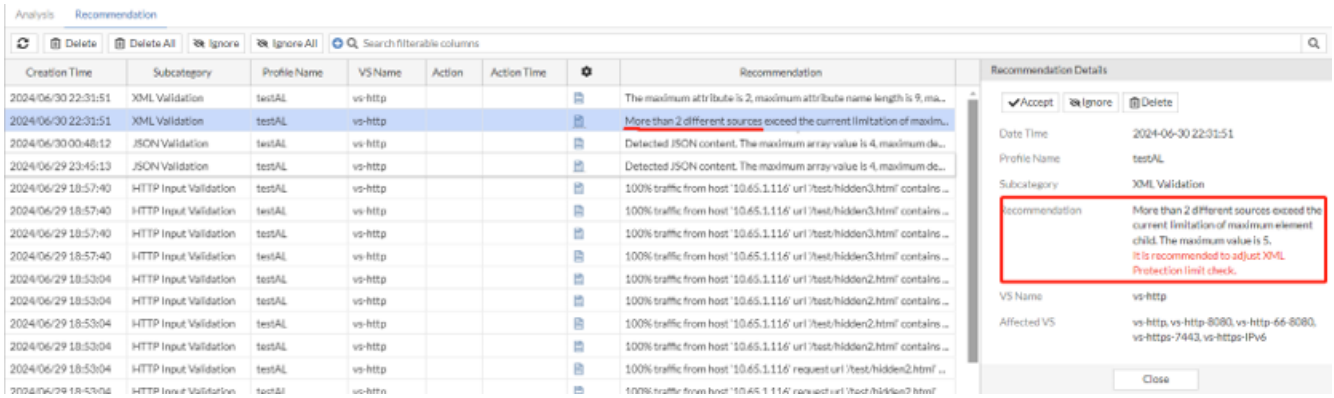
```
root@62:/home/ut/API_Test# more xml_full2.txt
<?xml version="1.0"?>
<books xmlns:ns1="http://example.com/ns1" xmlns:ns2="https://www.example.com/ns2">
  <ns1:book id="book3" available="true">
    <ns1:name><![CDATA[A Little Life]]></ns1:name>
    <ns1:reader type="fiction">
      <ns1:uname id="1">
        <ns1:tip>no</ns1:tip>
      </ns1:uname>
      <ns1:uname id="2">Tom</ns1:uname>
    </ns1:reader>
  </ns1:book>
  <ns1:book id="book4" available="false">
    <ns1:name>Darkmans</ns1:name>
    <ns1:reader type="fiction">
      <ns1:uname id="3">Jerry Fan</ns1:uname>
      <ns1:uname id="4">Nancy Liu</ns1:uname>
      <ns1:uname id="5">Jack Chen</ns1:uname>
    </ns1:reader>
  </ns1:book>
  <ns2:book id="book5" available="true">
    <ns2:name><![CDATA[Sapiens: A Brief History of Humankind]]></ns2:name>
    <ns2:reader type="non-fiction">
      <ns2:uname id="6">Eva Green</ns2:uname>
      <ns2:uname id="7">Sam Harris</ns2:uname>
    </ns2:reader>
  </ns2:book>
  <ns2:book id="book6" available="true">
    <ns2:name><![CDATA[Educated]]></ns2:name>
    <ns2:reader type="non-fiction">
      <ns2:uname id="8">Ian</ns2:uname>
      <ns2:uname id="9">Julia</ns2:uname>
    </ns2:reader>
  </ns2:book>
</books>
```

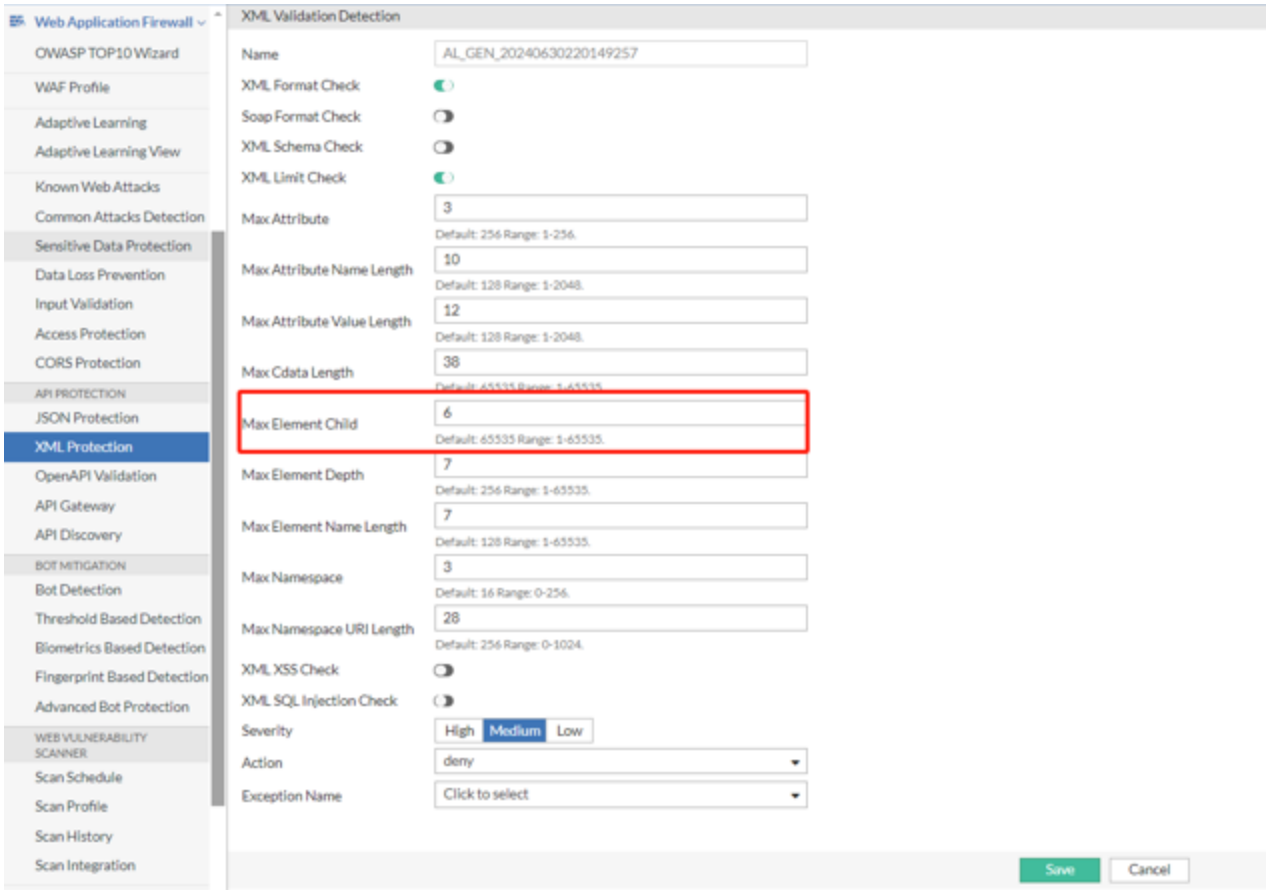


Triggering a False Positive recommendation to adjust a policy

When an XML protection policy is configured and the number of source IPs transmitting XML data with elements exceeding the configured limits surpasses the False Positive Threshold, Adaptive Learning will recommend to adjust the XML limits.

Upon accepting the recommendation, the corresponding limits in the policy will be updated accordingly.





Adaptive Learning troubleshooting and debugging

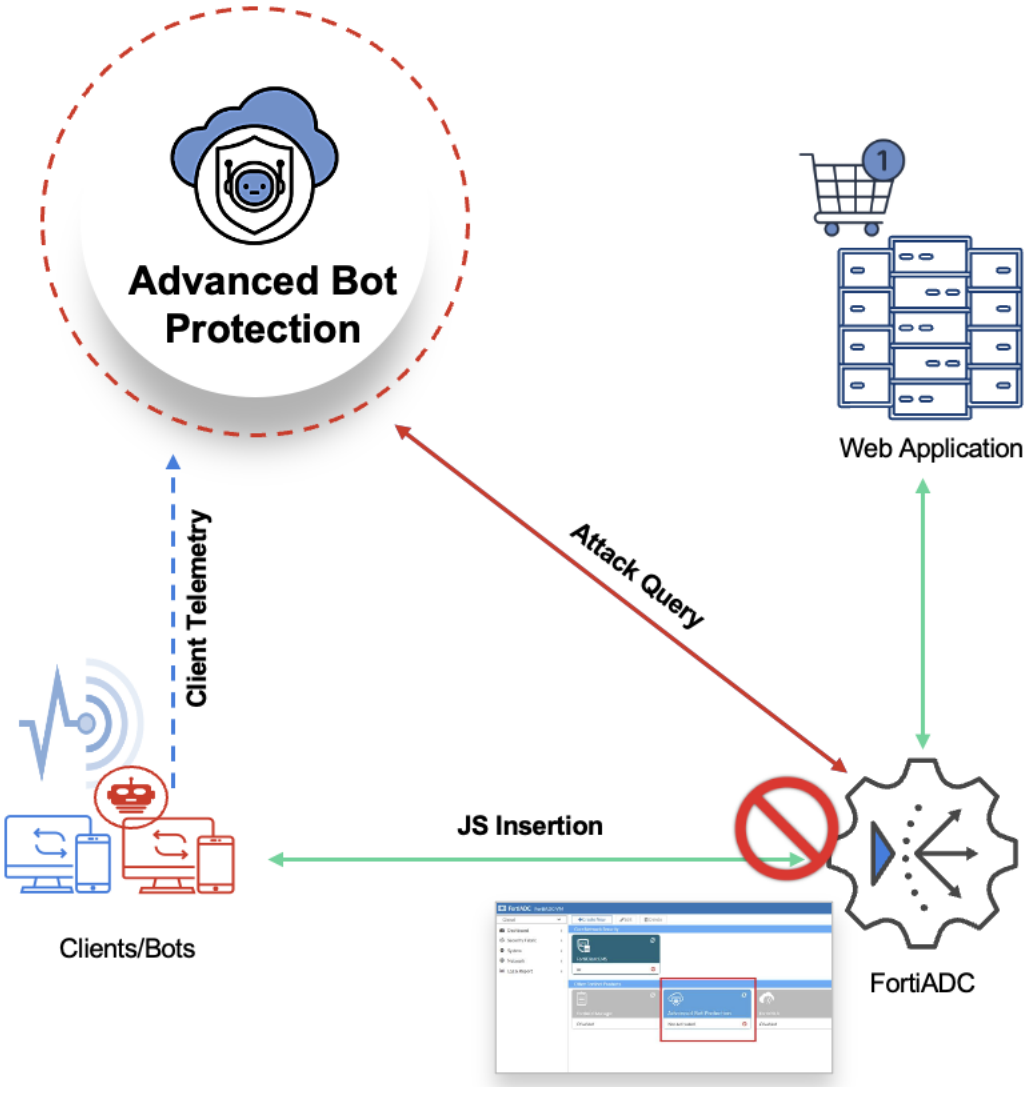
The following tools are available to troubleshoot and debug Adaptive Learning issues.

CLI commands to view debug logs relating to Adaptive Learning

Command	Guidelines
`diagnose debug module autolearn all | error | info | conf | event | stat | show` `diagnose debug enable`	To view the debug information for Adaptive Learning statistics collection, processing and recommendation generation.
`diagnose debug module waf adaptive-learning`	To view the debug information for traffic processed by httproxy, before delivering to Adaptive Learning.

Advanced Bot Protection (ABP)

FortiGuard ABP (Advanced Bot Protection) is a Fortinet SaaS advanced bot mitigation solution designed to detect and mitigate sophisticated bots that may be used to conduct malicious attacks on online applications, including data scraping, credential stuffing, DDoS attacks, and fraudulent activities. To safeguard your digital assets, websites, and applications, FortiGuard Advanced Bot Protection employs advanced deep learning algorithms and behavior analysis to identify and block suspicious activities. It analyzes user behavior patterns, device fingerprints, and more to distinguish between genuine users and malicious bots.



The FortiGuard ABP integration with FortiADC works by using client information collected by JavaScript insertion, which allow the client and FortiADC (via Fabric connector) to communicate with the Advanced Bot Protection Cloud for data telemetry information (such as headers and device fingerprinting). Once FortiADC is connected with FortiGuard ABP, you can configure an Advanced Bot Protection policy to apply to Web Application Firewall (WAF) profiles and virtual servers. FortiADC reports the telemetry data to FortiGuard ABP which then inspects the HTTP request to determine if

the client is human or a bot and sends instructions back to FortiADC to initiate an action against the request (such as block, CAPTCHA, or allow).

Prerequisites

Before you begin to integrate FortiGuard Advanced Bot Protection and configure Advanced Bot Protection on FortiADC, you must have the following:

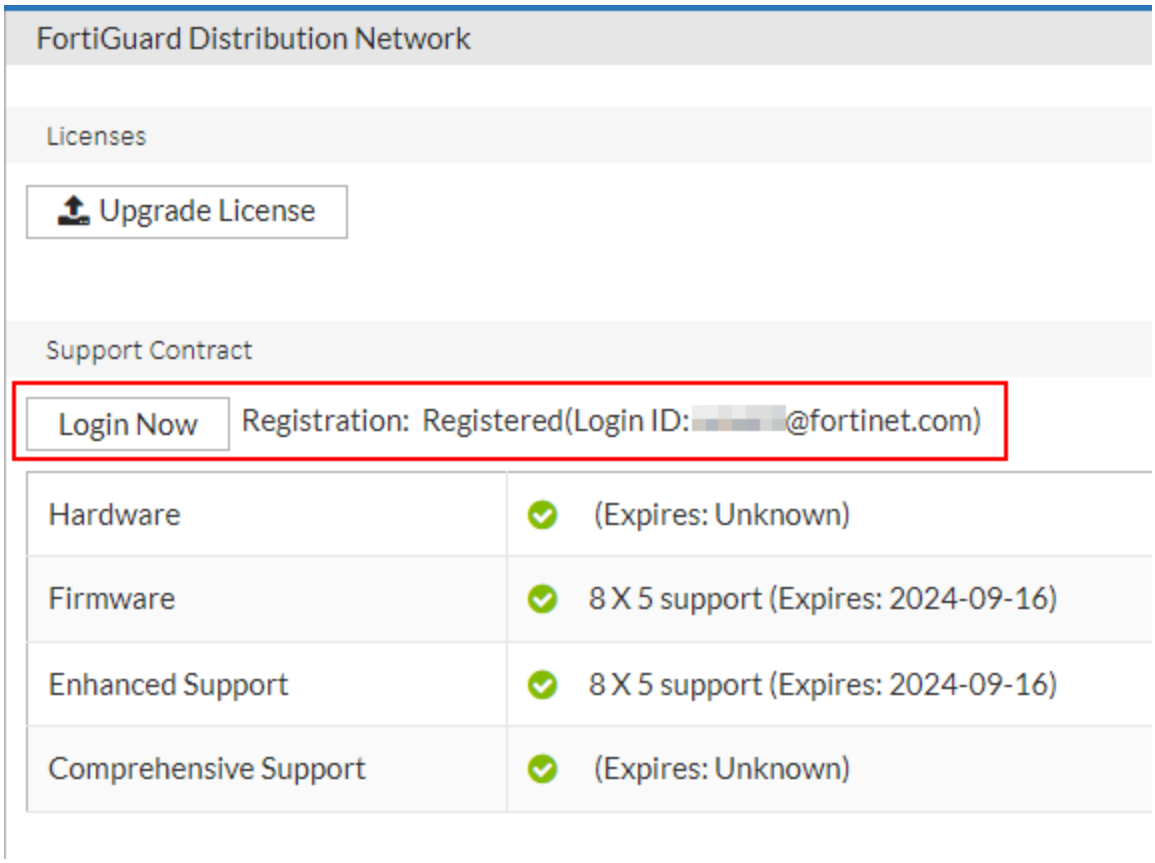
Valid FortiGuard Advanced Bot Protection license

FortiGuard ABP is available on a Standalone license which is a Fortinet Support account based license that is verified by the FortiGuard ABP User Portal instead of through FortiGuard. For more information, login to https://fortiabp.forticloud.com/.

Ensure the account used to register for the FortiGuard ABP license matches the account information from your Fortinet Support Contract. Otherwise, the FortiADC will not be able to connect to FortiGuard ABP.

FortiGuard ABP and Fortinet Support utilize a common account management system, allowing you to log into FortiGuard ABP directly using the credentials you have registered in your Fortinet Support contract. Upon logging into FortiGuard ABP with your account, the system will automatically check the validity of your FortiGuard ABP license. This same authentication process is applied when FortiADC sends requests to the FortiGuard ABP system to verify the account's validity to determine whether you can successfully enable FortiGuard ABP in FortiADC.

Note: The Fortinet Support Contract information can be found in **System > FortiGuard**, under the **Support Contract** section.



Supported FortiADC hardware, VM, or cloud platform

- Hardware — All FortiADC hardware models support FortiGuard ABP.
- VM — All FortiADC VM environments support FortiGuard ABP.
- Cloud platform — FortiADC instances on cloud platforms with BYOL support FortiGuard ABP (PAYG FortiADC instances do not support FortiGuard ABP).

Supported modes in FortiADC

Server Load Balancing:

- Layer 7 HTTP/HTTPS/HTTP2
- Layer 2 HTTP
- SSLi
- WCCP



Decompression must be enabled to support JavaScript insertion, which is critical to FortiGuard Advanced Bot Protection functionality.

Ensure a Decompression policy is enabled in your HTTP/HTTPS Application Profile to allow JavaScript to be inserted into compressed HTTP/HTTPS web content. To increase performance, most websites utilize HTTP compression to reduce the size of transmitted data. This compressed HTTP/HTTPS content must be decompressed to allow FortiADC to insert the required JavaScript tag to the HTML. If the real server response is not compressed, then decompression is unnecessary. However, if the real server response is compressed then decompression must be enabled, otherwise the JavaScript will fail to insert.

Virtual Domain:

- ABP global setting is configured in Global view.
- ABP policies are configured per VDOM, and functions at the VDOM level.
- ABP entries are shared by all VDOMs while their status are identified by each VDOM.

High Availability:

- All HA modes supported: Active-Passive, Active-Active, Active-Active-VRRP.
- ABP global settings and policies are synchronized.
- ABP entries are **not** synchronized.
- ABP global status is **not** synchronized. If the secondary node fails to connect to the FortiGuard ABP portal, its status will not be **Connected**.

Basic configuration

To deploy Advanced Bot Protection on FortiADC, follow the workflow below:

1. On FortiADC, enable the Advanced Bot Protection Fabric Connector to connect FortiADC to the FortiGuard Advanced Bot Protection server. For details, see Enabling the Advanced Bot Protection connector on page 815.
2. On the FortiGuard Advanced Bot Protection User Portal, obtain the Application ID of an existing FortiGuard ABP Application or create a new configuration. This Application ID is required to connect the FortiGuard ABP Application

to the FortiADC Advanced Bot Protection policy. For details, see Obtaining the Application ID from the FortiGuard ABP User Portal on page 816.

3. Configure a FortiADC Advanced Bot Protection policy. For details, see Configuring an Advanced Bot Protection policy on page 821.
4. Apply the Advanced Bot Protection policy to a WAF profile. For details, see Configuring a WAF Profile on page 667.
5. Apply the WAF profile that references the Advanced Bot Protection policy to a virtual server to activate Advanced Bot Protection for server load balancing.

Enabling the Advanced Bot Protection connector

After you have registered your FortiADC as a connector device in FortiGuard ABP, you can connect your FortiADC to the FortiGuard ABP server via the Fabric connector.

The Advanced Bot Protection connector can only be enabled and disabled, no further configuration is required. Once enabled and successfully connected, the Advanced Bot Protection module becomes available under the Web Application Firewall menu in the GUI.

Before you begin:

- You must have a valid FortiGuard Advanced Bot Protection license.
- You must have registered your FortiADC as a connector device in FortiGuard ABP. For details, see Obtaining the Application ID from the FortiGuard ABP User Portal on page 816.

To enable the Advanced Bot Protection connector:

1. Go to **Security Fabric > Fabric Connectors**.
2. Under **Other Fortinet Products** section, locate the **Advanced Bot Protection** connector.
3. **Enable** the Advanced Bot Protection connector. Once the connector is enabled, the connection status will display. The Advanced Bot Protection connector is ready when the status is **Connected**.



4. The and icons indicate whether the Advanced Bot Protection connector has successfully connected to the FortiGuard ABP server. Hover over the Advanced Bot Protection connector to see the status details. The table below lists the possible connection statuses for the Advanced Bot Protection connector.

Icon	ABP connector status	Guidelines
	Connected	The FortiADC is successfully connected to the FortiGuard ABP

Icon	ABP connector status	Guidelines
		server.
	Account license invalid	The FortiGuard ABP license is not valid. Please verify your license details or contact Fortinet Support.
	Couldn't connect to server	Unable to connect to the FortiGuard ABP server. Please check your network settings.
	Couldn't resolve hostname	Unable to resolve the hostname of the FortiGuard ABP server. Please check your network settings.
	No available SN cert	The device does not have an available SN certificate. Please check your local certificate.
	No available CA cert	The device does not have an available CA certificate. Please check your CA certificate.
	Problem with the local certificate	An error occurred with the remote server certificate. Please check your local certificate.
	SSL peer certificate or SSH remote key was not OK	An error occurred with the remote server certificate involving the SSL peer certificate or SSH remote key. Please check your local certificate.

Once the Advanced Bot Protection fabric connector is successfully connected, the Advanced Bot Protection module becomes available under the Web Application Firewall menu in the GUI.



If you want to disable the Advanced Bot Protection connector, the following settings will be impacted:

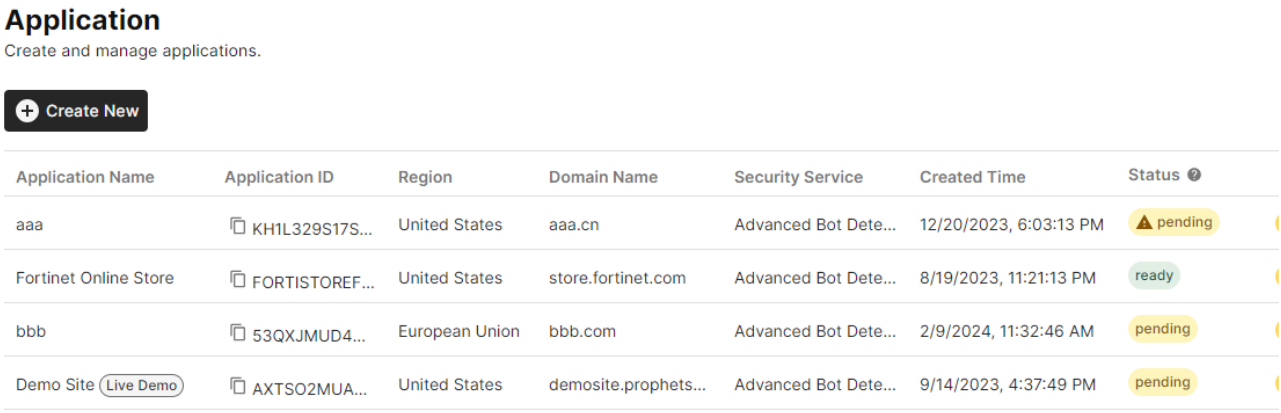
- The Advanced Bot Protection module under Web Application Firewall will be hidden in the GUI.
- Existing ABP policies will be hidden in the GUI. These ABP policies can be seen and can be deleted through CLI only.
- In the WAF Profile configuration, the Advanced Bot Protection option will not be available in the GUI. It can only be set through CLI.

Obtaining the Application ID from the FortiGuard ABP User Portal

The FortiADC Advanced Bot Protection policy binds to the FortiGuard ABP Application through its assigned Application ID to initiate bot detection and protection for the Application. Before you can configure an Advanced Bot Protection policy, you will need to obtain the Application ID either from an existing Application or create a new configuration in the FortiGuard ABP User Portal.

When an Application is created, the Application ID is automatically assigned. You can immediately begin to configure Advanced Bot Protection policies on FortiADC using this Application ID, however, the Advanced Bot Protection functionality will not activate until the FortiGuard ABP Application is fully analyzed and Pre-Provisioned to protect the Application. Pre-Provisioning is required to identify all URLs that should be protected in your Application domain (such as login URLs), and the locations to which JavaScript need to be inserted to collect client information. Without these resources, FortiADC will not be able to insert the necessary JavaScript for bot detection.

Pre-Provisioning is triggered upon creating the Application, and requires 2 to 3 days to complete. During this process, your FortiGuard ABP Application will be in **pending** status until Pre-Provisioning is complete. When the Application status is **ready**, Advanced Bot Protection can be activated in your FortiADC.



Before you begin:

- You must have a valid FortiGuard Advanced Bot Protection license.
- You must have access to the FortiGuard Advanced Bot Protection User Portal.

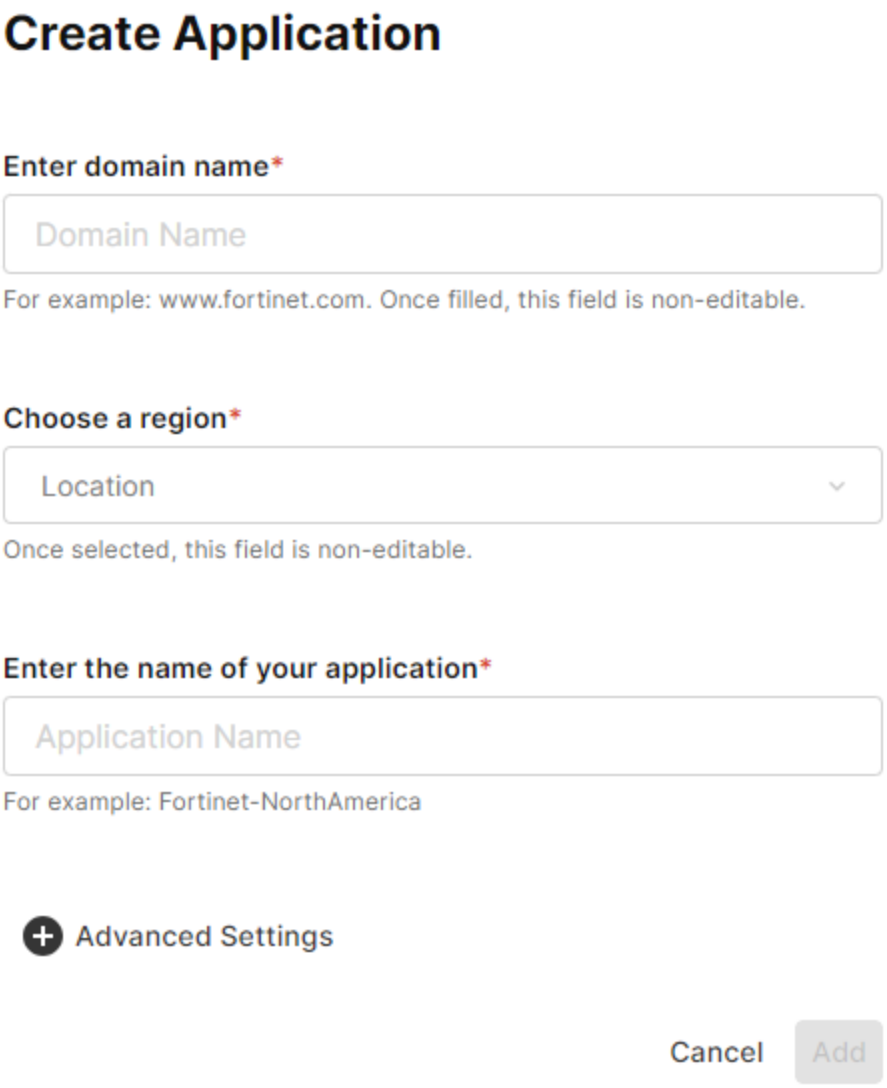
To obtain the Application ID from the FortiGuard ABP User Portal:

1. Login to FortiGuard Advanced Bot Protection (https://fortiabp.forticloud.com/).
2. Navigate to **Application**.
3. On the Application page, locate the entry for your FortiADC FortiGuard ABP Application.
4. Under the **Application ID** column, copy the ID.

To add your application in FortiGuard ABP:

1. Login to FortiGuard Advanced Bot Protection.
2. Navigate to **Application**.

3. Click **Create New** to display the configuration editor.

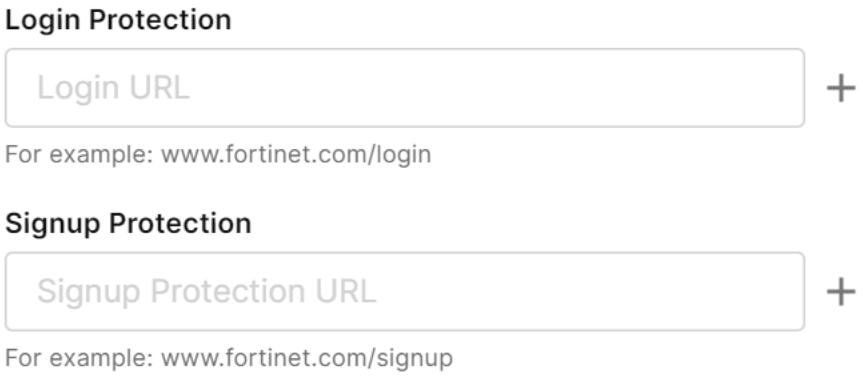


4. In the Create Application page, configure the following:

5.

Setting	Description
Domain Name	Enter the domain name of your application. Once this application is added, the domain name cannot be edited.
Location	Select the location for your application servers. FortiGuard Advanced Bot Protection is hosted in both the EU and US regions of Google Cloud. Opting for a region near your application server can significantly decrease network latency when FortiGuard Advanced Bot Protection processes your traffic. Once this application is added, the location cannot be edited.
Application Name	Provide a distinctive name for your application to facilitate easy identification.

6. Optionally, click **Advanced Settings** to enter the login URL of your application.



   This setting is optional as your domain will be analyzed in the Pre-Provisioning process and the login URLs/Signup Protection URLs can be automatically identified during this phase. However, you can manually include additional login URLs and Signup Protection URLs to ensure they are not overlooked in the Pre-Provisioning process.
7. Click **Add**.
   Your application will be listed in the Application page and automatically assigned an Application ID. The Pre-Provisioning process is initiated once the application is successfully added.
   This application will be in **pending** status until Pre-Provisioning is complete. Once the Application status is **ready**, you can configure the Advanced Bot Protection policies in your FortiADC.

FortiGuard ABP Pre-Provisioning

Once you have added your FortiGuard ABP Application, the Pre-Provisioning process begins. This triggers a request to the Professional Engagement Team (PET) to analyze your application details and to identify URLs to protect and insert the required JavaScript to enable the Advanced Bot Protection functionality. This process is currently conducted internally by the PET and requires 2 to 3 days to complete. If you wish to modify or add entries, please submit a request with Fortinet Support.

When Pre-Provisioning is complete, you can access the information on the protected URL entries and the locations where JavaScript is inserted for bot protection from the **Pre-Provisioning** page under the Configurations menu.

Pre-Provisioning

Pre-generated configuration for connector device.

Application ID

For more detailed information about the integration in **FortiWeb**, please refer to the document library.

For more detailed information about the integration in **FortiADC**, please refer to the document library.

Insert JS Entries

Domain	URL	Location
demosite.prophetsec.net	/accounts/login/	head
demosite.prophetsec.net	/	head
demosite.prophetsec.net	/accounts/signup/	head
demosite.prophetsec.net	/payment/stripe/	head
demosite.prophetsec.net	/payment/paypal/	head

Protect Entries

Domain	URL	Method
demosite.prophetsec.net	/accounts/login/	get, post, put, delete
demosite.prophetsec.net	/	get, post, put, delete
demosite.prophetsec.net	/accounts/signup/	get, post, put, delete
demosite.prophetsec.net	/payment/stripe/	get, post, put, delete

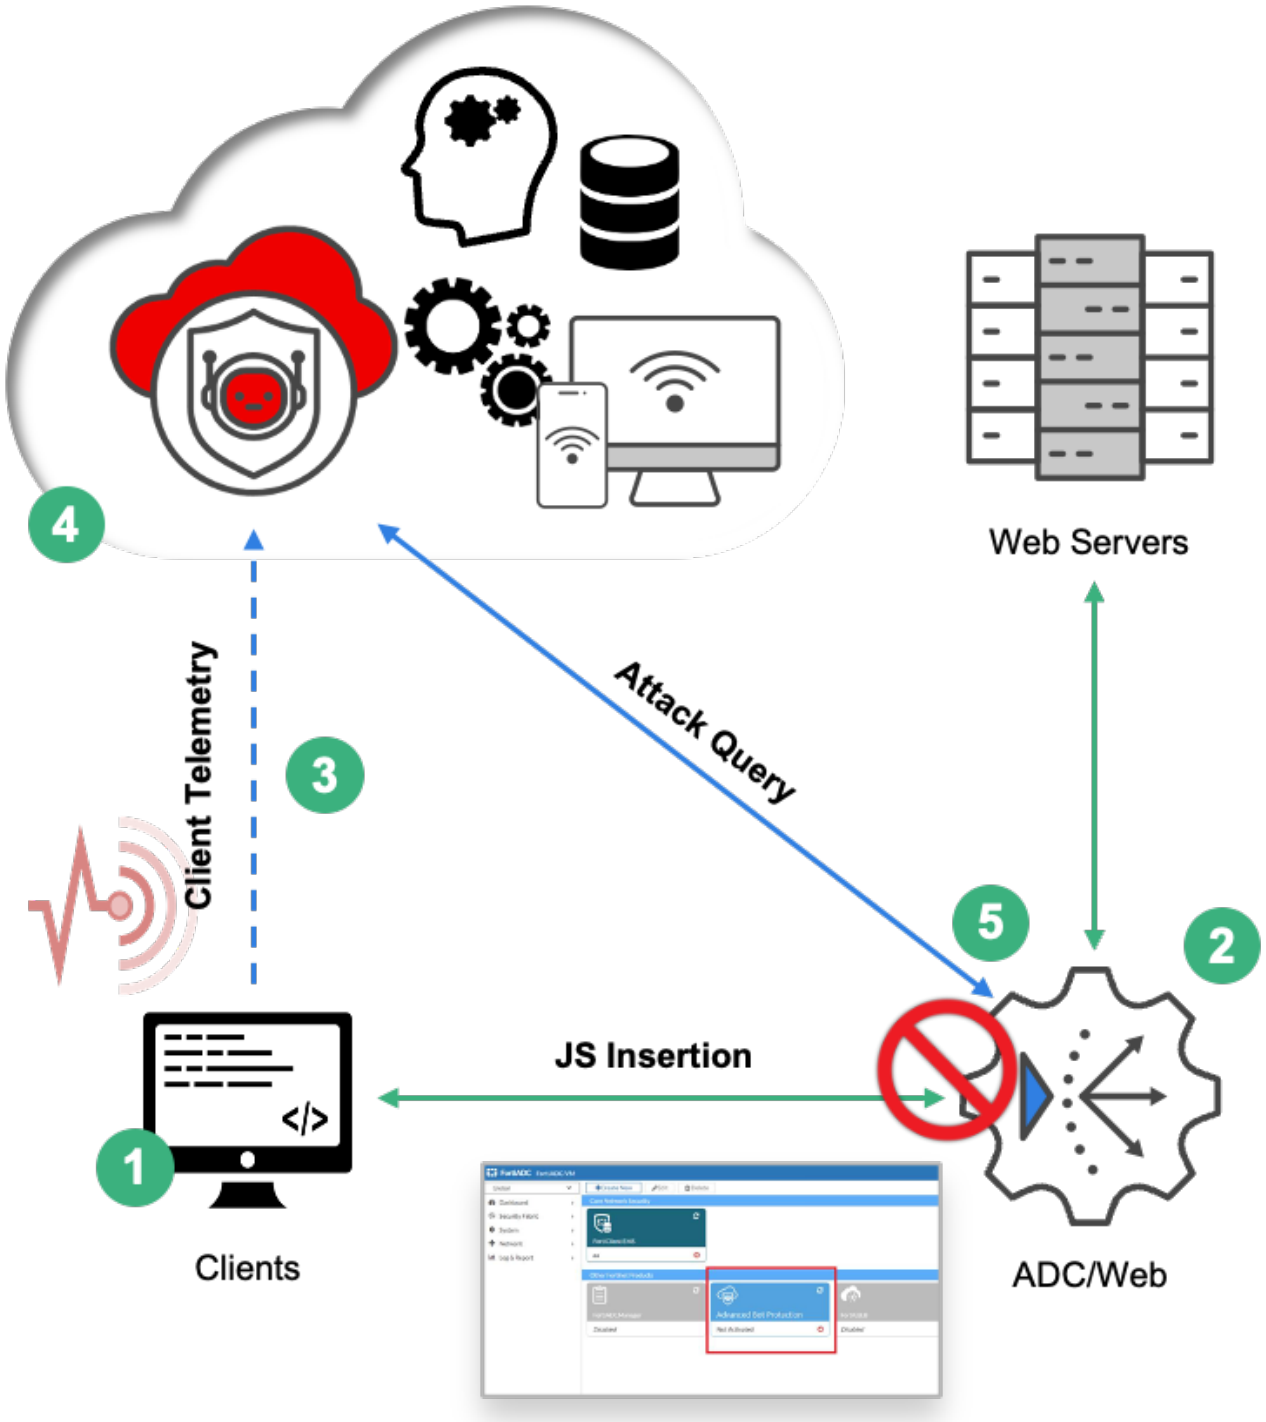
Configuring an Advanced Bot Protection policy

Once your FortiADC is successfully connected to the FortiGuard Advanced Bot Protection server via the Advanced Bot Protection Fabric Connector, you can configure an Advanced Bot Protection policy for your virtual server to protect your online applications from malicious bots and automated attacks. By incorporating FortiGuard ABP into FortiADC's server policy, client traffic will be directed to the FortiGuard ABP service deployed on Google Cloud where it will be analyzed to identify any malicious bot behavior and initiate appropriate actions in response.

FortiGuard ABP features a multi-dimensional deep learning engine that learns and tracks bot attacks over time, delivering the highest possible accuracy of classification between humans, and good and bad bots.

FortiGuard ABP protects against a wide range of threats, including the following:

- Data harvesting
- Credential stuffing attacks
- Account takeover attempts
- DDoS attacks



No.	Description
1	User request reaches FortiADC (as Reverse Proxy).
2	FortiADC inserts a JavaScript to the HTTP/S response for telemetric information.

No.	Description
3	The client and FortiADC (via Fabric connector) share telemetry data (such as IP, headers, and device fingerprinting) with the Advanced Bot Protection engine.
4	Using Deep Learning, FortiGuard ABP determines if the client is a human or a bot.
5	FortiGuard ABP sends instructions back to FortiADC to initiate an action against the request (such as block, CAPTCHA, or allow).

After you have configured the Advanced Bot Protection policy, you can reference it in a WAF Profile and apply it in a virtual server policy. However, the Advanced Bot Protection policy does not activate until the FortiGuard ABP Application is fully analyzed and Pre-Provisioned to protect the Application.

Pre-Provisioning is required to identify all URLs that should be protected in your Application domain (such as login URLs), and the locations to which JavaScript need to be inserted to collect client information. Without these resources, FortiADC will not be able to insert the necessary JavaScript for bot detection.

Pre-Provisioning is triggered upon creating the Application, and requires 2 to 3 days to complete. During this process, your FortiGuard ABP Application will be in **pending** status until Pre-Provisioning is complete. When the Application status is **ready**, Advanced Bot Protection can be activated in your FortiADC.



When Advanced Bot Protection is initially activated, it is recommended to set the WAF action that would allow you to observe and log any events detected by the FortiGuard Advanced Bot Protection, instead of immediately setting to block.

FortiGuard Advanced Bot Protection uses a multidimensional deep learning engine to learn and track bot attacks over time by using sophisticated AI model training. As FortiGuard ABP builds its training model, it will continue to improve and refine its bot detection capabilities. However, this may mean triggering false positives in the initial stages of the AI model training.

Before you begin:

- You must have enabled and successfully connected the Advanced Bot Protection connector on **Security Fabric > Fabric Connectors**.
  The Advanced Bot Protection module under Web Application Firewall is available only after the Advanced Bot Protection connector is enabled.
- You must have Read-Write permission for Security settings.
- You must have access to the FortiGuard Advanced Bot Protection User Portal to obtain the Application ID from an existing Application or create a new configuration. For more information, see Obtaining the Application ID from the FortiGuard ABP User Portal on page 816.

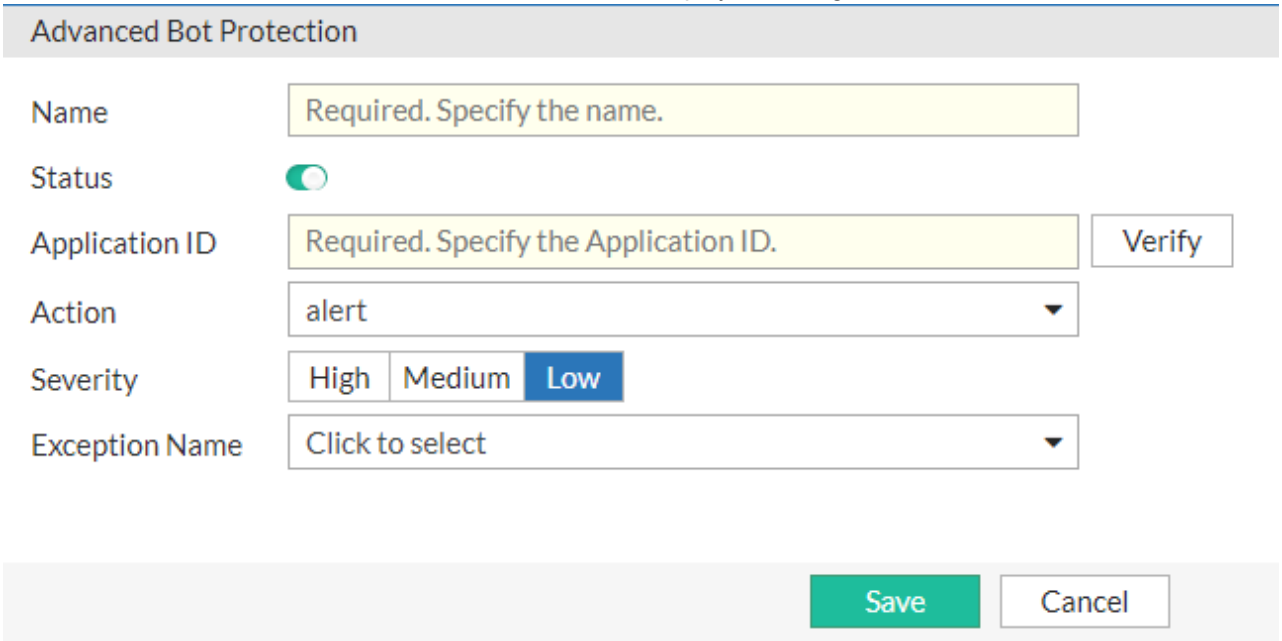


Decompression must be enabled to support JavaScript insertion, which is critical to FortiGuard Advanced Bot Protection functionality.

Ensure a Decompression policy is enabled in your HTTP/HTTPS Application Profile to allow JavaScript to be inserted into compressed HTTP/HTTPS web content. To increase performance, most websites utilize HTTP compression to reduce the size of transmitted data. This compressed HTTP/HTTPS content must be decompressed to allow FortiADC to insert the required JavaScript tag to the HTML. If the real server response is not compressed, then decompression is unnecessary. However, if the real server response is compressed then decompression must be enabled, otherwise the JavaScript will fail to insert.

To configure an Advanced Bot Protection policy:

1. Go to **Web Application Firewall > Advanced Bot Protection**.
2. In the **Advanced Bot Protection** tab, click **Create New** to display the configuration editor.

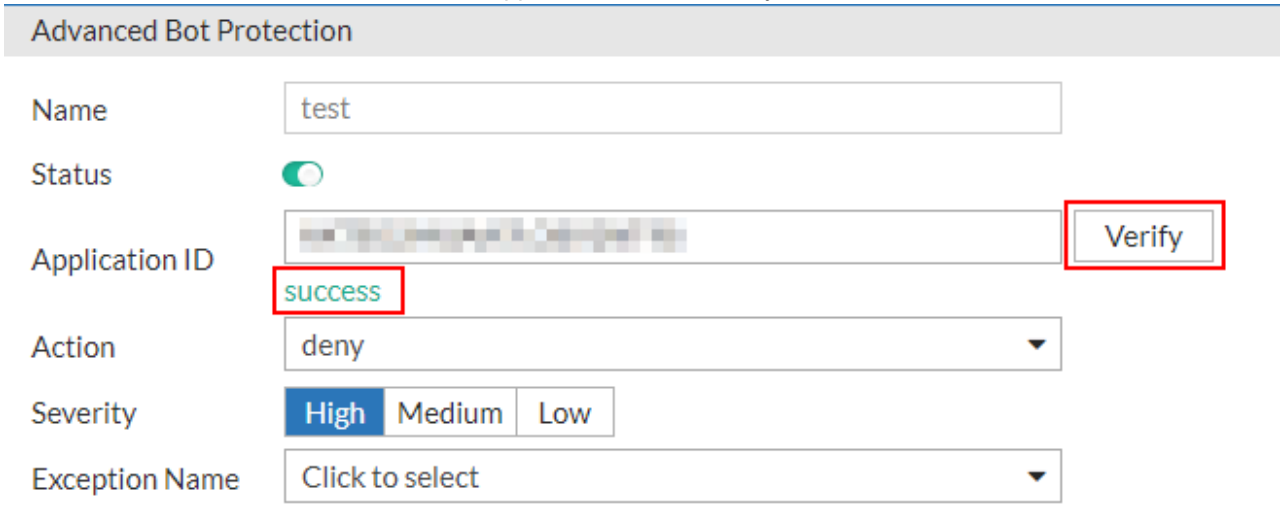


3. Configure the following Advanced Bot Protection settings:

Setting	Description
Name	Specify a name for the Advanced Bot Protection policy. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No spaces. The configuration name cannot be edited once it has been saved.
Status	Enable/disable the status of this Advanced Bot Protection policy. **Status** must be **enabled** to display Advanced Bot Protection configuration options.
Application ID	Specify the Application ID assigned to your FortiGuard ABP Application. The Application ID is used to bind this Advanced Bot Protection policy to the FortiGuard ABP Application. For steps on how to obtain the Application ID from the FortiGuard ABP User Portal, see Obtaining the Application ID from the FortiGuard ABP User Portal on page 816.
Action	Specify a WAF action object to apply when a bot is detected. You can specify a predefined or user-defined WAF action profile. (See Configuring WAF Action objects on page 671.) Predefined WAF actions: • alert — WAF policies will allow the traffic to pass and log the event. • block — WAF policies will drop the current attack session by HTTP 403 message and block the attacker (according the attacker's IP address) for 1 hour, and log the event.

Setting	Description
	• captcha — WAF policies will allow the traffic to pass if the client successfully fulfills the CAPTCHA request, and log the event. • deny — WAF policies will the drop current attack session by HTTP 403 message, and log the event. • silent-deny — WAF policies will drop the current attack session by HTTP 403 message, without logging the event. The default action is alert.
Severity	Select the event severity to log when a bot is detected: • 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.

4. Optionally, click **Verify** to validate the inputted Application ID against multiple parameters to ensure the connection between FortiADC and the FortiGuard ABP Application is successfully established.



For more information, see Verifying the Application ID and connection status to FortiGuard ABP on page 826.

5. Click **Save**.
Once the Advanced Bot Protection policy is saved, you can reference it in a WAF Profile configuration.



It is strongly recommended to verify the Application ID and FortiGuard ABP server connection prior to completing the Advanced Bot Protection policy configuration. Even though this is an optional step, it is helpful to diagnose any potential issues and apply fixes early.

When the Advanced Bot Protection policy is created, an internal verification is automatically conducted to verify the status of the Application ID and FortiGuard ABP server connection. If the Application ID is not valid, or any other validation parameters has failed, the Advanced Bot Protection policy will fail to function and the system will log the failure to send the ABP policy.

Verifying the Application ID and connection status to FortiGuard ABP

The Advanced Bot Protection policy can only function if the Application ID is valid and the connection to the FortiGuard ABP server is successfully established. FortiADC validates multiple parameters, including if the Application ID is available, the FortiGuard ABP server connectivity, if the FortiGuard ABP license is valid. There are two types of messages as differentiated by text color: green text indicates a positive status where all required parameters are validated successfully; and red text that indicate one or more parameters did not pass validation.

The following table describes some common verification results.

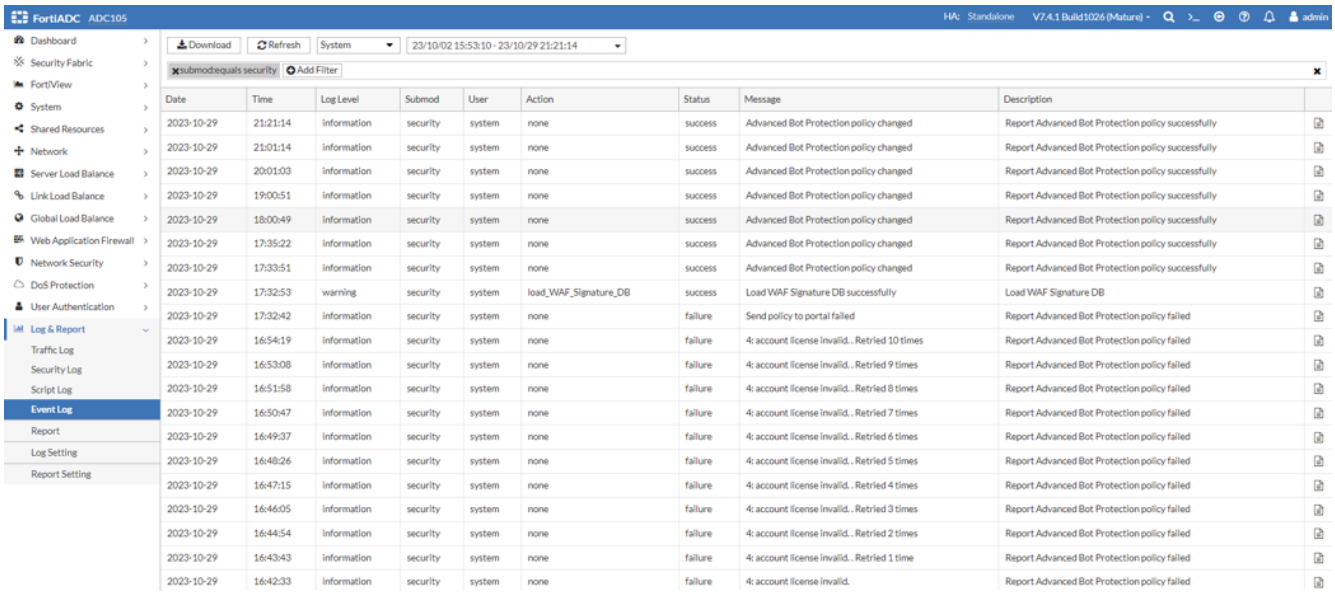
Verification status message	Guidelines
Success (green)	All required parameters pass validation; application ID is available, FortiGuard ABP server certificate is valid, network connectivity is good, etc.
Application not found (red)	The Application ID does not exist. This could be an input error.
Account license invalid (red)	The FortiGuard ABP license is not valid. Please verify your license details or contact Fortinet Support.
Couldn't connect to server (red)	Unable to connect to the FortiGuard ABP server. Please check your network settings.
Couldn't resolve hostname (red)	Unable to resolve the hostname of the FortiGuard ABP server. Please check your network settings.
No available SN cert (red)	The device does not have an available SN certificate. Please check your local certificate.
No available CA cert (red)	The device does not have an available CA certificate. Please check your CA certificate.
Problem with the local certificate	An error occurred with the remote server certificate. Please check your local certificate.
SSL peer certificate or SSH remote key was not OK	An error occurred with the remote server certificate involving the SSL peer certificate or SSH remote key. Please check your local certificate.

Advanced Bot Protection troubleshooting and debugging

The following tools are available to troubleshoot and debug Advanced Bot Protection issues.

Event logs that record the interactions between FortiADC and FortiGuard Advanced Bot Protection when system events occur

You can check the system event logs to see the communication between FortiADC and FortiGuard ABP.



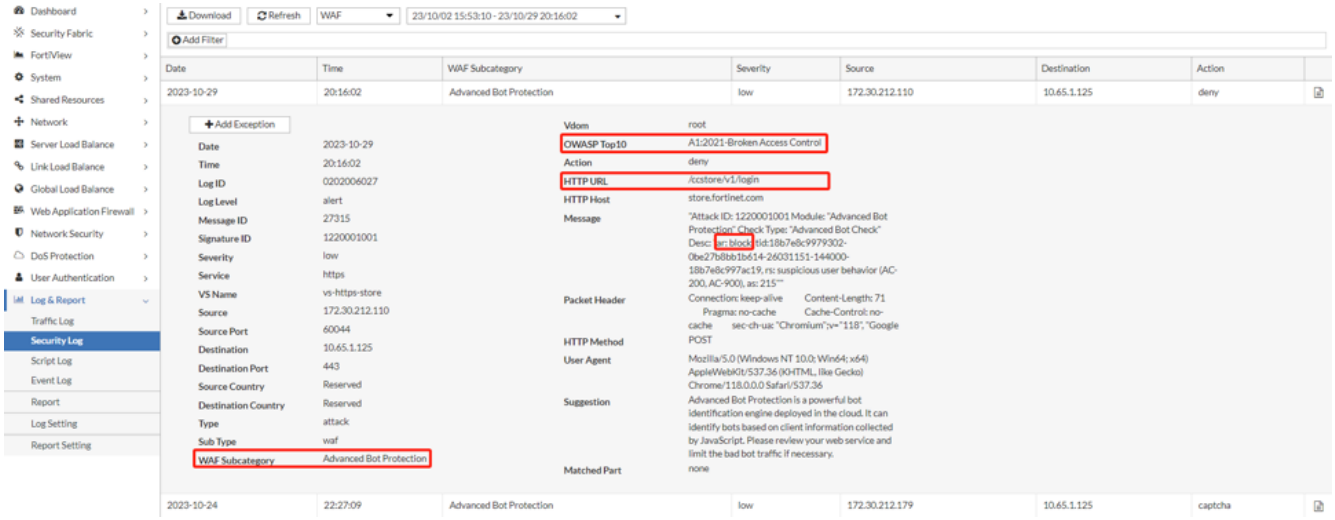
The following system events are logged:

- Configuration changes to the Advanced Bot Protection policy.
- When FortiADC has successfully reports Advanced Bot Protection policy configurations to FortiGuard ABP.
- When FortiADC fails and retries to report Advanced Bot Protection policy configurations to FortiGuard ABP.

Attack Logs that record interactions between FortiADC and FortiGuard ABP when suspicious behavior is detected

You can check the attack logs to see the WAF action triggered by the response received from FortiGuard ABP.

In the example below, the attack logged a block WAF action that was triggered by FortiGuard ABP detecting suspicious user behavior.



Some interactions may not trigger an attack log.

- When the WAF action for the Advanced Bot Protection policy is set to **allow**, triggered attacks will not be logged.
- When FortiADC receives an "empty" response from FortiGuard ABP.

- When FortiADC receives a 401 Authentication error (or another error) from FortiGuard ABP.
- When FortiADC does not receive a response from FortiGuard ABP until timeout (10 seconds).

CLI commands to view debug logs relating to Advanced Bot Protection

Command	Guidelines
`diagnose debug module wafmonitor all` `diagnose debug enable`	To view the debug information for interactions between the wafmonitor daemon and the FortiGuard ABP. Interactions include fetching the Advanced Bot Protection policy entries from FortiGuard ABP and reporting FortiADC configurations to FortiGuard ABP. **Note**: The `diagnose debug module wafmonitor all` command prints the debug information for all WAF modules that use the daemon wafmonitor. You will need to manually check for the debug information regarding the Advanced Bot Protection module.
`diagnose debug module waf advanced_bot`	To view the debug information for traffic processed by the Advanced Bot Protection policy.
`diagnose debug module framework_http`	To view the debug information for HTTP packet processing by WAF modules. **Note**: The `diagnose debug module framework_http` command prints the debug information for all WAF modules. You will need to manually check for the debug information regarding the Advanced Bot Protection module.

Network Security

The **Network Security** menu contain features and configurations that allow you to protect your network against malicious attacks.

This section is organized into the following sub-menu topics:

- Firewall on page 829
- Intrusion Prevention on page 836
- AntiVirus on page 841
- IP Reputation on page 847
- Geo IP Protection on page 850
- Zero Trust Network Access (ZTNA) on page 854

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.
- AntiVirus—Drop traffic that matches in FortiSandbox's Malware Signature Database.
- 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.

Firewall

The **Network Security > Firewall** sub-menu contains features and configurations that allow you to define your Firewall and Connection Limit policies.

- Configuring an IPv4 firewall policy on page 830
- Configuring an IPv6 firewall policy on page 832
- Configuring an IPv4 connection limit policy on page 834
- Configuring an IPv6 connection limit policy on page 835

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**.
2. Click the **IPv4 Firewall Policy** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in Firewall policy configuration on page 830.
5. Save the configuration.
6. 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 *if your FortiADC is configured for link load-balancing and/or traffic routing*. In both cases, the system will use this interface to forward traffic to its destination.

Settings	Guidelines
	Note: You *MUST* 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 Type	Select the source type to use to form the matching tuple. • Address • Address Group • External Resource
Source	The **Source** option is available if the **Source Type** is **Address**. Specify the Address object to use as the source.
Source Address Group	The **Source Address Group** option is available if the **Source Type** is **Address Group**. Specify the Address Group object to use as the source.
Source External Address Group	The **Source External Address Group** option is available if the **Source Type** is **External Resource**. Specify the external IP address list imported through the IP Address connector to use as the source. For details, see IP Address Connector on page 106.
Destination Type	Select the destination type to use to form the matching tuple. • Address • Address Group • External Resource
Destination	The **Destination** option is available if the **Destination Type** is **Address**. Specify the Address object to use as the destination.
Destination Address Group	The **Destination Address Group** option is available if the **Destination Type** is **Address Group**. Specify the Address Group object to use as the destination.
Destination External Address Group	The **Destination External Address Group** option is available if the **Destination Type** is **External Resource**. Specify the external IP address list imported through the IP Address connector to use as the destination. For details, see IP Address Connector on page 106.
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.
Deny Log	The **Deny Log** option is available is the **Action** is **Deny**. Enable/disable logging for denied action.

Settings	Guidelines
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**.
2. Click the **IPv6 Firewall Policy** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in Firewall policy configuration on page 832.
5. Save the configuration.
6. 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.

Settings	Guidelines
	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 *if your FortiADC is configured for link load-balancing and/or traffic routing*. In both cases, the system will use this interface to forward traffic to its destination. **Note**: You *MUST* 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 Type	Select the source type to use to form the matching tuple. • Address • Address Group • External Resource
Source	The **Source** option is available if the **Source Type** is **Address**. Specify the Address object to use as the source.
Source Address Group	The **Source Address Group** option is available if the **Source Type** is **Address Group**. Specify the Address Group object to use as the source.
Source External Address Group	The **Source External Address Group** option is available if the **Source Type** is **External Resource**. Specify the external IP address list imported through the IP Address connector to use as the source. For details, see IP Address Connector on page 106.
Destination Type	Select the destination type to use to form the matching tuple. • Address • Address Group • External Resource
Destination	The **Destination** option is available if the **Destination Type** is **Address**. Specify the Address object to use as the destination.
Destination Address Group	The **Destination Address Group** option is available if the **Destination Type** is **Address Group**. Specify the Address Group object to use as the destination.
Destination External Address Group	The **Destination External Address Group** option is available if the **Destination Type** is **External Resource**. Specify the external IP address list imported through the IP Address connector to use as the destination. For details, see IP Address Connector on page 106.
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.

Settings	Guidelines
	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.
Deny Log	The **Deny Log** option is available is the **Action** is **Deny**. Enable/disable logging for denied action.
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 835.
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	
	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 836.
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.

Intrusion Prevention

From the **Network Security > Intrusion Prevention** sub-menu, you can configure an Intrusion Prevention Profile that uses Intrusion Prevention System (IPS) technology to protect your network against cyberattacks by actively seeking and blocking external threats before they can reach potentially vulnerable network devices.

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.

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-set-up.

Predefined Profile	Comment
all_default	All predefined signatures with default setting.
all_default_pass	All predefined signatures with PASS action.
default	Prevents 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 **Network Security > 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.

To add IPS signatures to a Profile

1. Go to **Network Security > Intrusion Prevention**.
2. Select the IPS Profile to which you want to add the signature and click the edit icon.
3. Under IPS Signatures, click **Add Signature**.
4. Select one or more signatures from the list and click **Apply** to add them to the sensor.
   You can search through the available signatures using either the CVE Number and/or filter by Action, Application, Name, OS, Protocol, Severity, or Target.

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 **Network Security > Intrusion Prevention.**
2. Select the IPS Profile to which you want to add the signature and click the edit icon.
3. Under IPS Filters, select **Add Filter**.
   You can search through the available signatures using either the CVE Number and/or filters.
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**.
5. After selecting the IPS Filter, configure the Action from the drop-down.

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.

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
```

AntiVirus

The **Network Security > AntiVirus** sub-menu allows you to perform the following tasks to configure your anti-virus security:

- Creating an AV profile on page 842
- Setting AV quarantine policies on page 844
- Setting the AV service level on page 846

Overview

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.

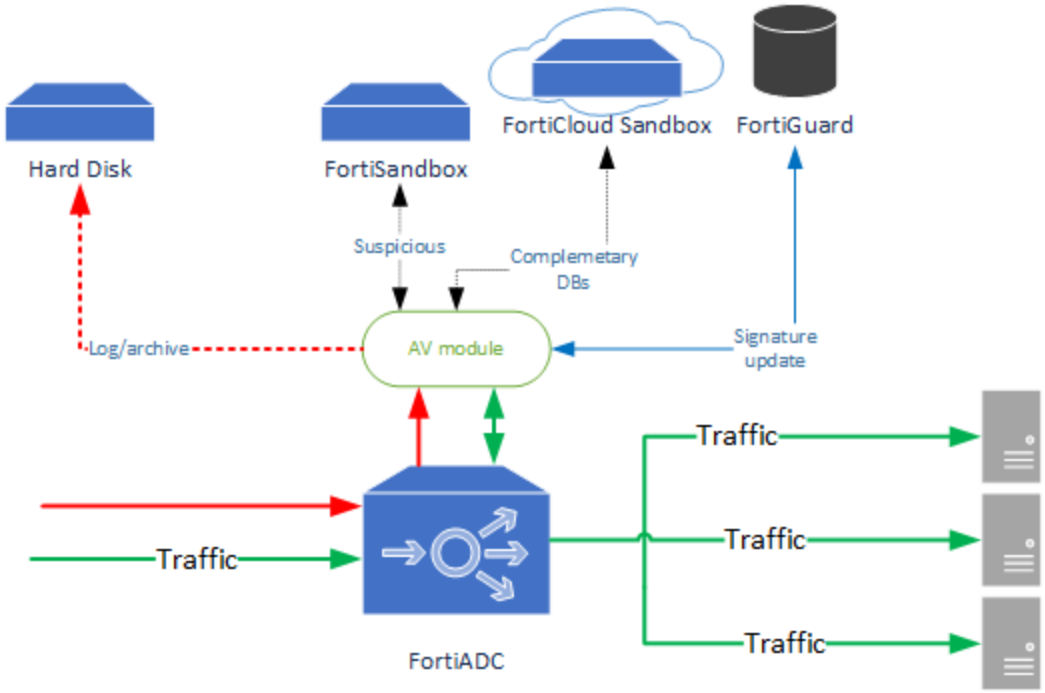
Using a suite of integrated security technologies, Anti-virus (AV) solutions provide protection against a variety of threats, including both known and unknown malicious codes (Malware) and Advanced Targeted Attacks (ATA).

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 842 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 842
- 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, HTTPS, or SMTP 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 > AntiVirus**.
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 843.
5. Click **Save** when done.

AV profile configuration

Settings	Description
Name	A unique name for the AV profile. An AV profile name can contain up to 63 alphanumeric characters.
Comments	A brief description of the profile. 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. The default is 2 MB. Valid values range from 1 to 2000 MB.
Uncomp Nest Limit	The maximum number of levels of nesting (compression) allowed for the system to decompress. The default is 2. Valid values range from 2 to 100.
Scan Bzip2	Scan archives using the bzip2 algorithm. This is disabled by default.
Streaming Content Bypass	Enable or disable bypass streaming content (rather than buffering it). This is enabled by default.
Oversize Limit	The maximum in-memory file size in KB to be scanned. The default is 1024 KB. Valid values range from 1 to 12000000 KB. **Note**: For AV files larger than 1000 KB, the device memory must be larger than 32 GB to support the scan.
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. The default option is Bypass.
Options	Select an option for the system to handle infected files: • AV Monitor — Block and log infected files. • Quarantine — Quarantine and log infected files. The default is AV Monitor.

Settings	Description
Emulator	Enable or disable the Win32 Emulator. This is 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. The default is Disable.
Analytics Max Upload	The maximum file size in KB allowed to upload to FortiSandbox. 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. This is disabled by default.
AV Virus Log	Enable or disable logging for anti-virus scanning. This is 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

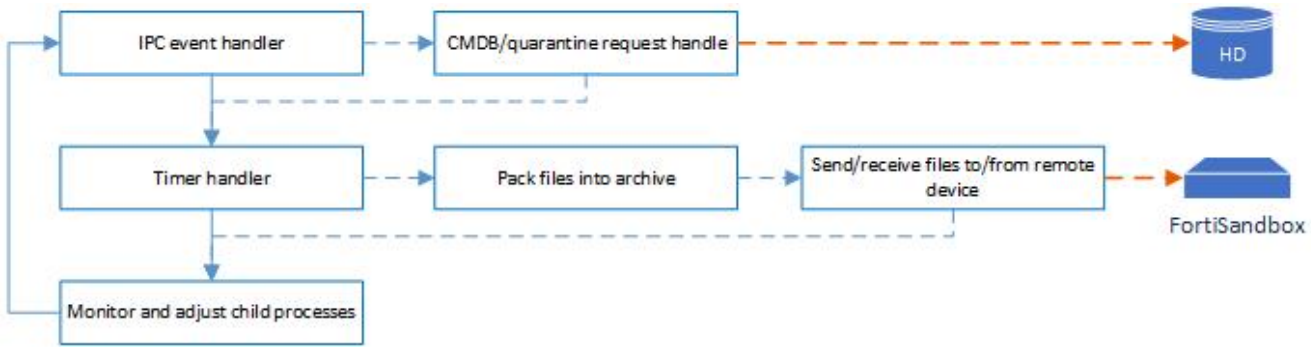
The “quarantined” daemon manages the infected or suspicious files. The quarantine destination can be either the local hard disk.

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 845.

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 > AntiVirus**.
2. Click the **Quarantine** tab.
3. Make the entries or selections as described in AV quarantine policy configuration on page 845.
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 form 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.

Settings	Description
	Note: Files larger than the set Max File Size will not be quarantined. In reality, this value is subject 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.
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
```

Setting the 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 > AntiVirus**.
2. Click the **Settings** tab.
3. Select a **Default DB** setting:

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.

4. Specify the **Ramdisk size** in GB. The default is 0 GB. Valid values range from 0 to 12 GB. Please ensure to enlarge the RAM disk size to scan large files.
   Note: When the changes to the RAM disk configuration is committed, you will be prompted to restart the device for the configuration to take effect.
5. Click **Save**.

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}
set ramdisk-size <integer>
end
```

IP Reputation

The **Network Security > IP Reputation** sub-menu allows you to perform the following tasks to configure your IP protection:

- Managing IP Reputation policy settings on page 848
- Configuring an IP reputation exception on page 849
- Configuring an IP reputation block list on page 849

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 848 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 848.
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: • Anonymous Proxy • Others

Settings	Guidelines
	• Block List
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.

Configuring an 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 849.
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	Is IP Range is selected in the Type field above, specify the starting address and ending address of the IP range.

Configuring an IP reputation block list

Upload the source IP's or CIDRs that you want the ADC to block in the IP reputation block list. When these source IP's try to access the VS, the connection will fail. You can create IP/Netmask or IP Range type block list, back up or restore files.

The content of IP reputation block 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. This operation will back up the current restored IP reputation block list, however, it does not back up user-configured entries.

You use the **Clean** utility to erase entries that were imported from the text file. This operation will erase the current restored IP reputation block list, however, it does not affect user-configured entries.

To create an IP Reputation block list:

1. Go to **Network Security > IP Reputation**.
2. Click the **IP Reputation Blocklist** tab.
3. Click **Create New** and configure a new block list as described in IP Reputation block list on page 850.
4. Click **Save**.

IP Reputation block 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.

Geo IP Protection

The **Network Security > Geo IP Protection** sub-menu allows you to perform the following tasks to configure your Geo IP protection:

- Using the Geo IP block list on page 850
- Using the Geo IP allowlist on page 852
- Special Geo codes on page 853

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 851 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 allowlist 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 allowlist 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 851.
4. Click **Save**.
5. Edit your new block list to add members as described in Geo IP block list configuration on page 851.
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.

Settings	Guidelines
	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.
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 allowlist

To configure a Geo IP allowlist:

1. Go to **Network Security > Geo IP Protection**.
2. Click the **Allowlist** tab to create a allowlist as described in Geo IP allowlist configuration on page 852.
3. Click **Save**.

Geo IP allowlist 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	

Settings	Guidelines
Type	Select and configure either of the following: 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 853 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)
O1	Other Country (Reserved for further use, and no IP address is assigned to this region)

Zero Trust Network Access (ZTNA)

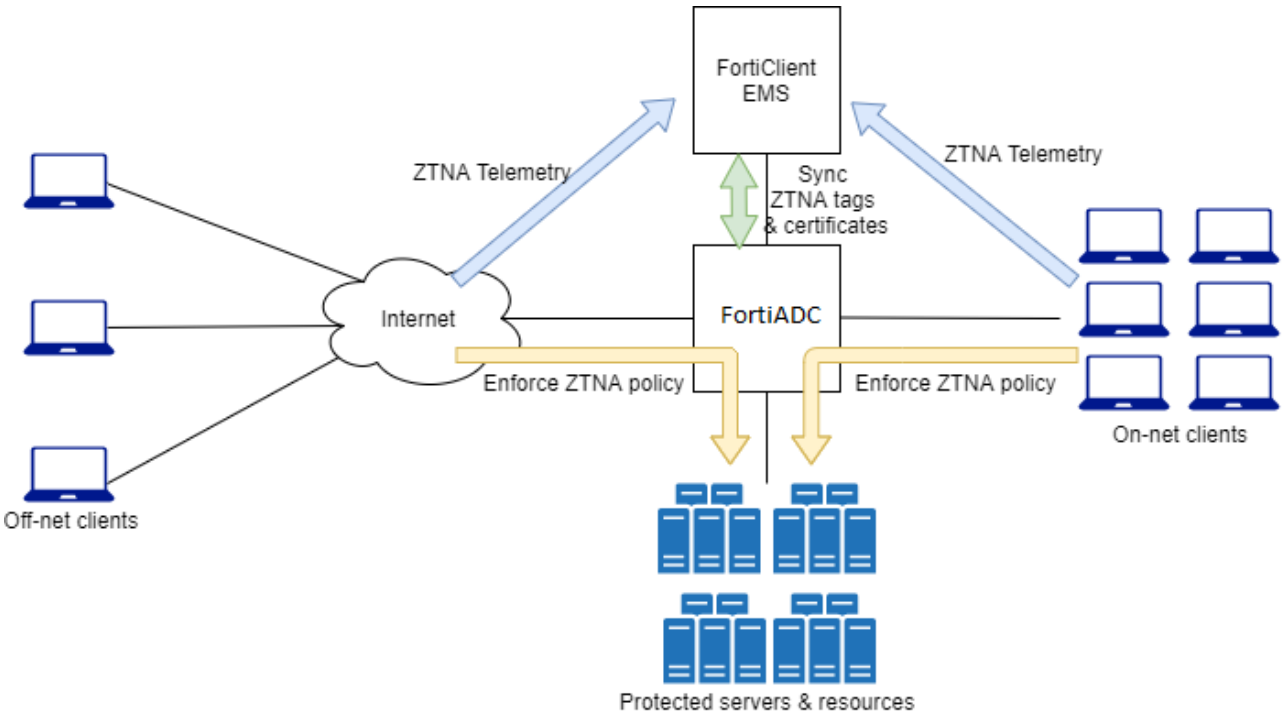
The **Network Security > ZTNA** sub-menu contains features and configurations relating to Zero Trust Network Access (ZTNA) in FortiADC:

- How device identity and trust context is established with FortiClient EMS on page 858
- Configuring FortiClient EMS Connector for ZTNA on page 859
- Verifying EMS CA certificate, ZTNA tag, and FortiClient endpoint synchronized from FortiClient EMS on page 863
- Configuring a ZTNA Profile on page 865
- ZTNA troubleshooting and debugging on page 866

ZTNA Overview

Protect your virtual server resources with the FortiADC Zero Trust Network Access (ZTNA) access control method that uses client device identification and Zero Trust tags to provide role-based application access. It provides administrators the flexibility to manage network access for On-net local users and Off-net remote users. Access to applications is granted only after verifying the device and user identity, and then performing context-based posture checks using Zero Trust tags.

ZTNA telemetry, tags, and policy enforcement



When On-net and Off-net FortiClient endpoints register to FortiClient EMS, device information, logged on user information, and security posture are all shared over ZTNA telemetry with the EMS server. Clients also make a certificate signing request to obtain a client certificate from the EMS that is acting as the ZTNA Certificate Authority (CA).

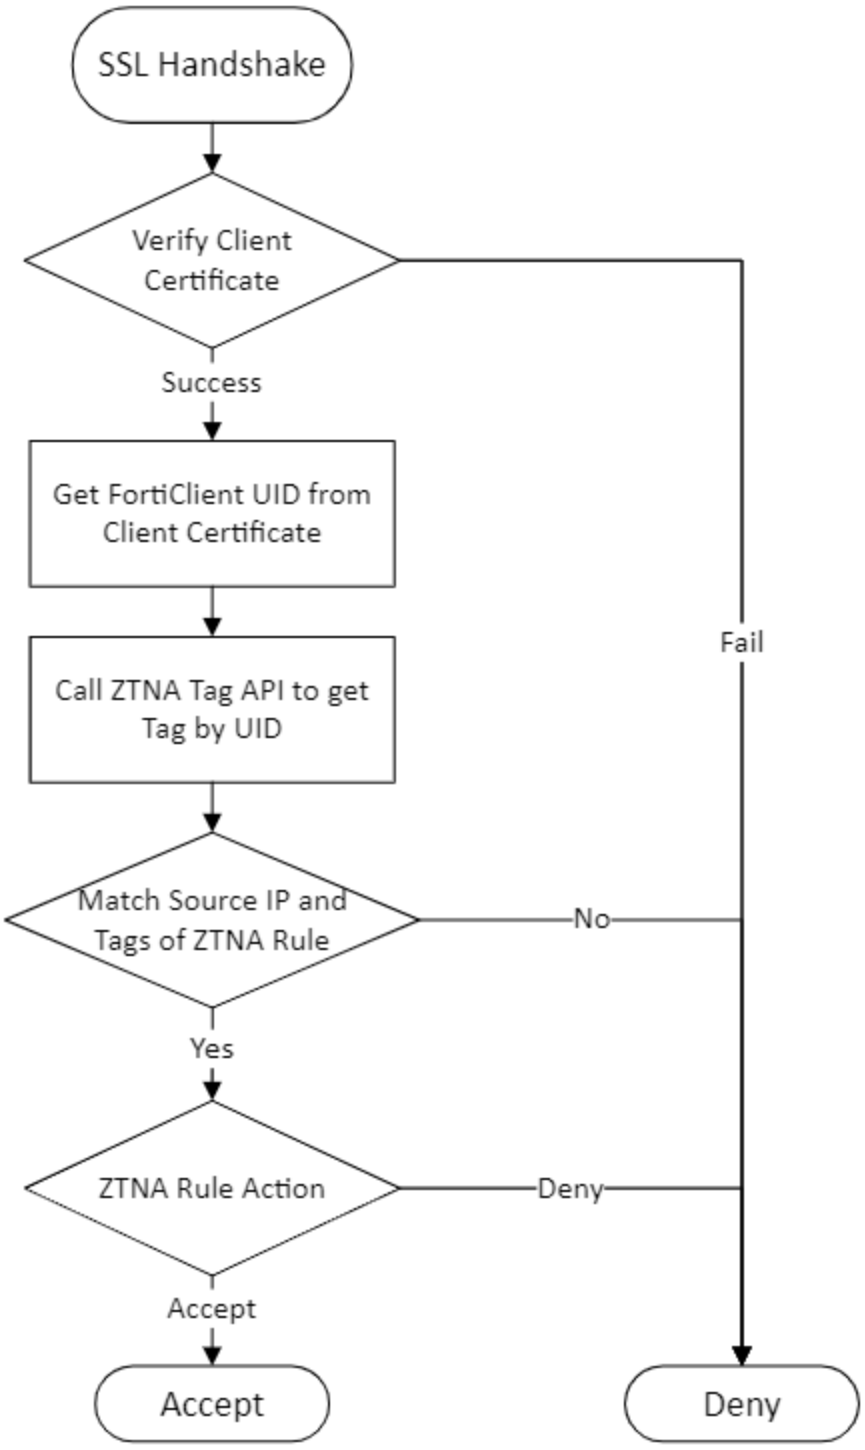
Based on the client information, EMS applies matching Zero Trust tagging rules to tag the clients. These tags, and the FortiClient endpoint information (including the device information, logged on user information, and security posture) are synchronized with the FortiADC in real-time. This allows the FortiADC to verify the client's identity using the client certificate, and grant access based on the ZTNA tags applied in the ZTNA security rule.

For more information, see How device identity and trust context is established with FortiClient EMS on page 858.

ZTNA in FortiADC server load balancing

The FortiADC ZTNA is a network security feature that allows users to securely access Layer 7 HTTPS and TCPS virtual server resources for server load balancing. Once the ZTNA security rule has been configured it can be referenced by Layer 7 HTTPS and TCPS virtual servers to implement role-based zero trust access by using the client certificate and ZTNA tags for identification and security posture check.

The chart below illustrates the FortiADC ZTNA logic flow.



Prerequisites

Before you begin to configure ZTNA on the FortiADC unit, you must have the following:

- FortiClient EMS running version 7.0.3 or later
- FortiClient running 7.0.1 or later
- FortiADC hardware, VM, or cloud platform that support FortiClient EMS.

  Supported hardware models:
  - FAD-120F
  - FAD-220F
  - FAD-300F
  - FAD-400F
  - FAD-1200F
  - FAD-2200F
  - FAD-4200F
  - FAD-5000F

  Supported cloud platforms with BYOL (PAYG FortiADC does not support FortiClient EMS):
  - AWS (Amazon Web Services)
  - Microsoft Azure
  - GCP (Google Cloud Platform)
  - OCI (Oracle Cloud Infrastructure)
  - Alibaba Cloud

  Supported VM environments:
  - VMware — ESXi 3.5, 4.x, 5.0, 5.1, 5.5, 6.0, 6.5, 6.7, 7.0
  - Microsoft Hyper-V — Windows Server 2012 R2, 2016 and 2019
  - KVM — Linux version 3.19.0 qemu-img v2.0.0, qemu-img v2.2
  - Citrix Xen — XenServer 6.5.0
  - Xen Project Hypervisor — 4.4.2, 4.5

  Note: The most recent certificate embedded license is required. If your license was issued prior to April 2021, please obtain a new certificate embedded license for your VM through Fortinet Customer Service & Support.
- Read-Write access permission for FortiADC Systems settings

Basic ZTNA configuration

To deploy FortiADC ZTNA, follow the basic workflow below:

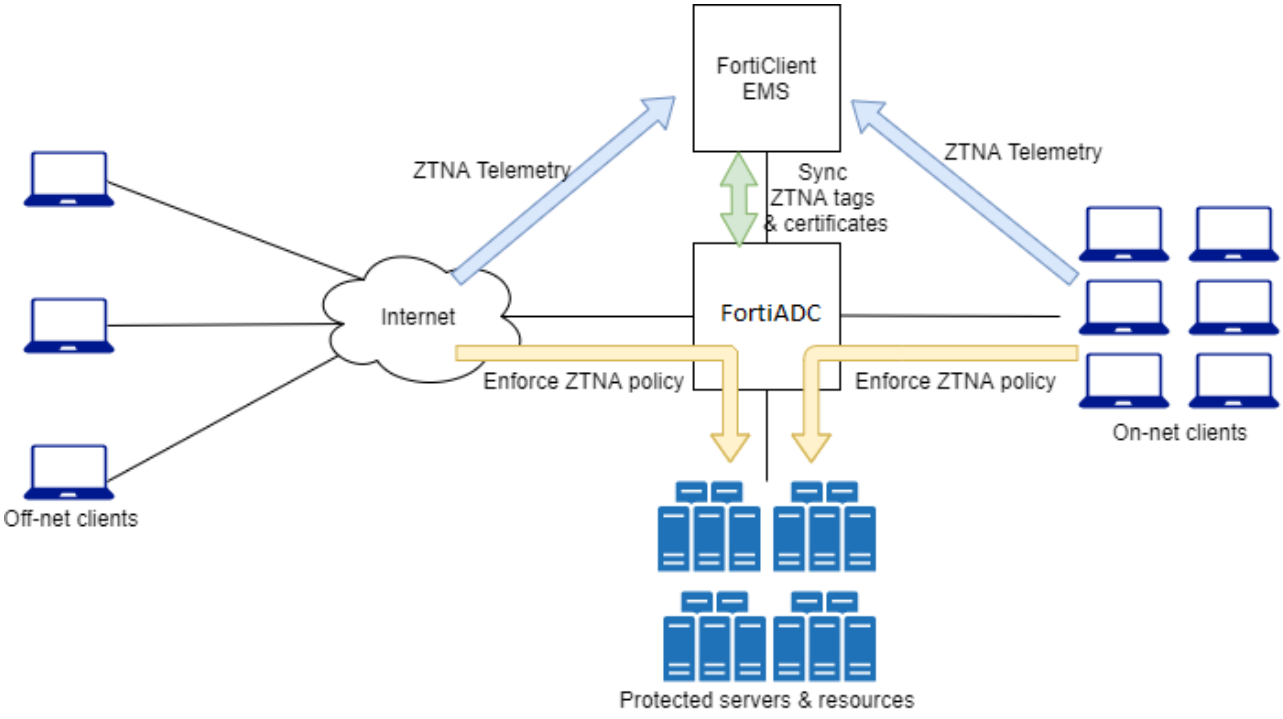
1. Configure a FortiClient EMS connector to register your FortiADC device as a Fabric Device in the FortiClient EMS. For details, see Configuring FortiClient EMS Connector for ZTNA on page 859.
2. Verify the information synchronized to FortiADC from FortiClient EMS. For details, see Verifying EMS CA certificate, ZTNA tag, and FortiClient endpoint synchronized from FortiClient EMS on page 863.
3. Configure a ZTNA profile to define the ZTNA rules. For details, see Configuring a ZTNA Profile on page 865
4. Apply the Security ZTNA profile to a Layer 7 HTTPS or TCPS virtual server to activate ZTNA for server load balancing. Ensure the corresponding Client SSL profile is enabled for client certificate verification. For details, see Configuring virtual servers on page 384 and Configuring Client SSL profiles on page 512.
5. Enable security logging for ZTNA. This is optional.

How device identity and trust context is established with FortiClient EMS

Integral to Zero Trust Network Access (ZTNA) is how device identity and trust context is established between FortiClient, FortiClient EMS, and the FortiADC.

Device roles

The following illustrates how the FortiClient, FortiClient EMS, and FortiADC each function and participate in the ZTNA workflow.



FortiClient

FortiClient endpoints provide the following information to FortiClient EMS when they register to the EMS:

- Device information (network details, operating system, model, and others)
- Logged on user information
- Security posture (On-net/Off-net, antivirus software, vulnerability status, and others)

FortiClient requests and obtains a client device certificate from the EMS ZTNA Certificate Authority (CA) when it registers to FortiClient EMS. The client uses this certificate to identify itself to the FortiADC.

FortiClient EMS

FortiClient EMS issues and signs the client certificate with the FortiClient UID, certificate serial number, and EMS serial number. EMS shares its EMS CA certificate with the FortiADC, so that the FortiADC can use it to verify the clients.

FortiClient EMS uses Zero Trust tagging rules to tag endpoints based on the information that it has on each endpoint. The tags are also shared with the FortiADC.

FortiADC

The FortiADC maintains a continuous connection to the EMS server to synchronize endpoint device information, including primarily:

- FortiClient UID
- Client certificate SN
- EMS SN
- Device credentials (user/domain)
- Network details (IP and MAC address and routing to the FortiADC)
- ZTNA tags

When a device's information changes, such as when a client moves from on-net to off-net, or their security posture changes, EMS is updated with the new device information and then updates the FortiADC. The FortiADC's virtual servers can use this information when processing ZTNA traffic. If an endpoint's security posture change causes it to no longer match the ZTNA rule criteria then it will be denied in a new session.

Configuring FortiClient EMS Connector for ZTNA

The FortiClient Endpoint Management Server (EMS) connector enables you to establish device identity through client certificates and device trust context between FortiClient, FortiClient EMS and the FortiADC as part of Zero Trust Network Access (ZTNA).

You can register your FortiADC device as a Fabric Device through the FortiClient EMS connector. When you create a FortiClient EMS connector, FortiADC sends a request to the FortiClient EMS server to obtain a EMS CA certificate to register your FortiADC device. From the FortiClient EMS, you can then authorize the FortiADC as a Fabric Device. Once authorized, the FortiClient EMS connector will display the status as **Connected**, indicating the device is registered. After the FortiADC connects to the FortiClient EMS, it automatically synchronizes ZTNA tags, the EMS CA certificate, and FortiClient endpoint information.

ZTNA tags are then generated from tagging rules configured on the FortiClient EMS. These tagging rules are based on various posture checks that can be applied on the endpoints.



In FortiClient EMS, do not use special characters such as `"`, `'`, and `\` in the ZTNA tag name. ZTNA tags that contain these special characters in their name may trigger unexpected behavior when referenced in the ZTNA Profile or in the security logs.

You can create a maximum of three FortiClient EMS connectors.

Requirements:

- FortiClient EMS running version 7.0.3 or later
- FortiClient running 7.0.1 or later

- FortiADC hardware, VM, or cloud platform that support FortiClient EMS.



FortiClient EMS is supported in most FortiADC platforms but not all of them. The following lists the hardware models, cloud platforms, and VM environments that support FortiClient EMS.

Hardware models:

- FAD-120F, FAD-220F, FAD-300F, FAD-400F, FAD-1200F, FAD-2200F, FAD-4200F, FAD-5000F

Cloud platforms with BYOL (PAYG FortiADC does not support FortiClient EMS):

- AWS (Amazon Web Services), Microsoft Azure, GCP (Google Cloud Platform), OCI (Oracle Cloud Infrastructure), Alibaba Cloud

VM environments:

- VMware, Microsoft Hyper-V, KVM, Citrix Xen, Xen Project Hypervisor
  Note: The most recent certificate embedded license is required. If your license was issued prior to April 2021, please obtain a new certificate embedded license for your VM through Fortinet Customer Service & Support.

- Read-Write access permission for FortiADC Systems settings

To create and configure a FortiClient EMS connector:

1. Go to **Security Fabric > Fabric Connectors**.
2. Click **Create New**.
3. Under **Core Network Security**, click **FortiClient EMS** to display the configuration editor.
4. Configure the following **FortiClient EMS** Settings:

Setting	Description
Name	Specify the FortiClient Enterprise Management Server (EMS) name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces.
IP/Domain name	Specify the server IPv4 address or the domain name of the FortiClient EMS FQDN. For example: 192.0.2.1
HTTPS Port	Specify the FortiClient EMS HTTPS access port number. Range: 1-65535, default: 443

5. Click **Save**.
   The **Verify EMS server certificate** dialog displays the following message:
   In order for the FortiClient EMS and FortiADC to communicate, the following certificate provided by the FortiClient EMS must be reviewed for correctness, and accepted if deemed valid.
   Do you wish to Accept the certificate as detailed below?
6. After you have verified the EMS server certificate information displayed, click **OK** to accept the EMS server certificate.
   The **Verify completed** dialog displays the following message:
   This FortiADC is not authorized on FortiClient EMS yet. Please let FortiClient EMS to authorize it.
   Note: This message will only appear if the FortiADC device has not yet been authorized as a Fabric Device through FortiClient EMS.
7. Click **OK**.

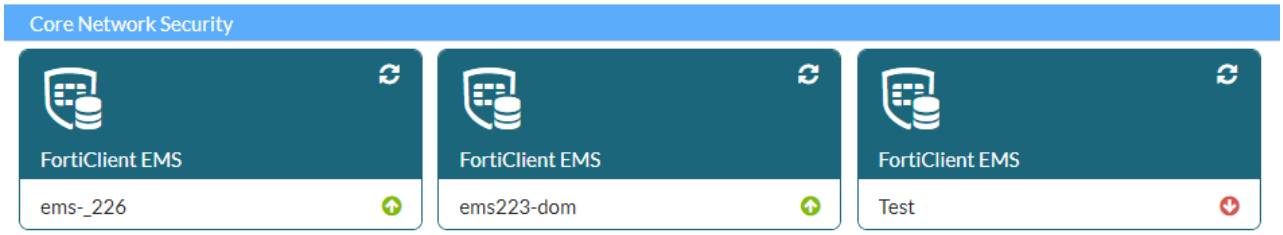
The newly created FortiClient EMS connector is added to the **Security Fabric > Fabric Connectors** page, under the **Core Network Security** section. The FortiClient EMS connector will not be connected until the FortiADC has been authorized as a Fabric Device in FortiClient EMS.

To authorize the FortiADC as a Fabric Device in FortiClient EMS:

1. Login to FortiClient EMS.
2. From the FortiClient EMS landing page, the **Fabric Device Authorization Requests** pop-up displays the Serial Number and IP information of the FortiADC device. Click **Authorize**.
3. Alternatively, you can go to **Administration > Fabric Devices** and select the Fabric device you want to authorize.

To check and troubleshoot the FortiClient EMS connector connection:

1. Go to **Security Fabric > Fabric Connectors**.
2. Under the **Core Network Security** section, locate the FortiClient EMS connector configurations.

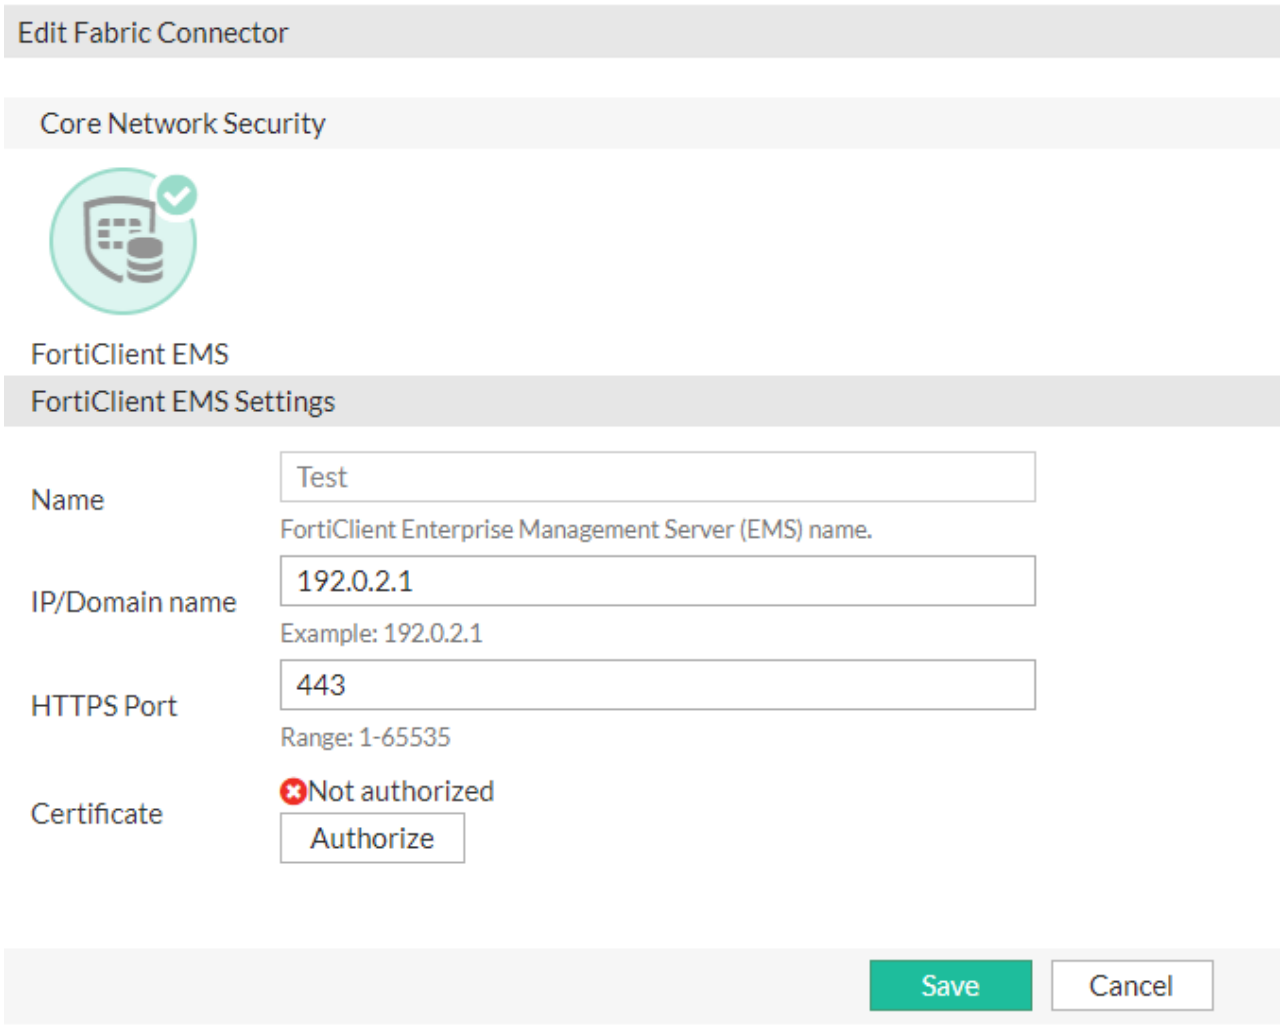


3. The and icons indicate whether FortiClient EMS has successfully authorized the FortiADC Fabric Device for the corresponding FortiClient EMS connector. Hover over the FortiClient EMS connector to see the status details. The table below lists the possible connection statuses for the FortiClient EMS connector.

Icon	EMS Status	Description
	Connected	The FortiADC has been successfully authorized as a Fabric Device through FortiClient EMS.
	Cert unauthorized	FortiADC does not verify the EMS server's CA certificate. You can edit the FortiClient EMS connector configuration and restart the verification to accept the EMS CA certificate.
	Auth failed	The EMS server does not authorize the FortiADC, indicating the request is either denied or pending authorization. If pending authorization, the status will change to **Connected** once authorization is successful on the EMS server.
	Not reachable	The EMS server was not reachable. Ensure the EMS server IP and system router is properly configured.
	EMS server connection failed	The EMS server connection failed with unknown issue. For example, an incorrect EMS server port may cause this issue.
	No compatible	The EMS server connection failed because the server is not compatible with FortiADC.
	Not sent	The EMS domain name cannot resolve. Ensure proper configuration for the DNS server setting, domain name, and system router.

If the status is not Connected, edit the FortiClient EMS connector accordingly to troubleshoot the connection issue.

4. Locate the newly created FortiClient EMS connector, click the FortiClient EMS connector configuration then click **Edit**, or double click the configuration object to display the configuration editor.



5. Edit the configuration to troubleshoot the connection issue then click **Authorize** to restart the verification to accept the EMS CA certificate.
A request is resent to the FortiClient EMS to authorize the FortiADC as a Fabric Device in FortiClient EMS. The FortiClient EMS connector will not be connected until the FortiADC has been authorized as a Fabric Device in FortiClient EMS.

FortiClient EMS in virtual domain configurations

Virtual domains (VDOMs) are full FortiADC instances configured within a FortiADC device. Once the FortiADC device is registered to the FortiClient EMS, the configuration settings are applied globally. Using the same ZTNA tag, each VDOM can then configure ZTNA security rules that apply individually at the VDOM level.

FortiClient EMS for High Availability configurations

In a High Availability group, both the FortiADC units must be registered to the FortiClient EMS as individual Fabric devices. However, you only need to configure the FortiClient EMS connector on one of the FortiADC units. Once the FortiClient EMS connector configuration has been completed in one of the FortiADC units in the HA group, the configuration will be synchronized to the second FortiADC unit. The Fabric Device authorization request for both FortiADC units are sent to FortiClient EMS to complete the device registration.

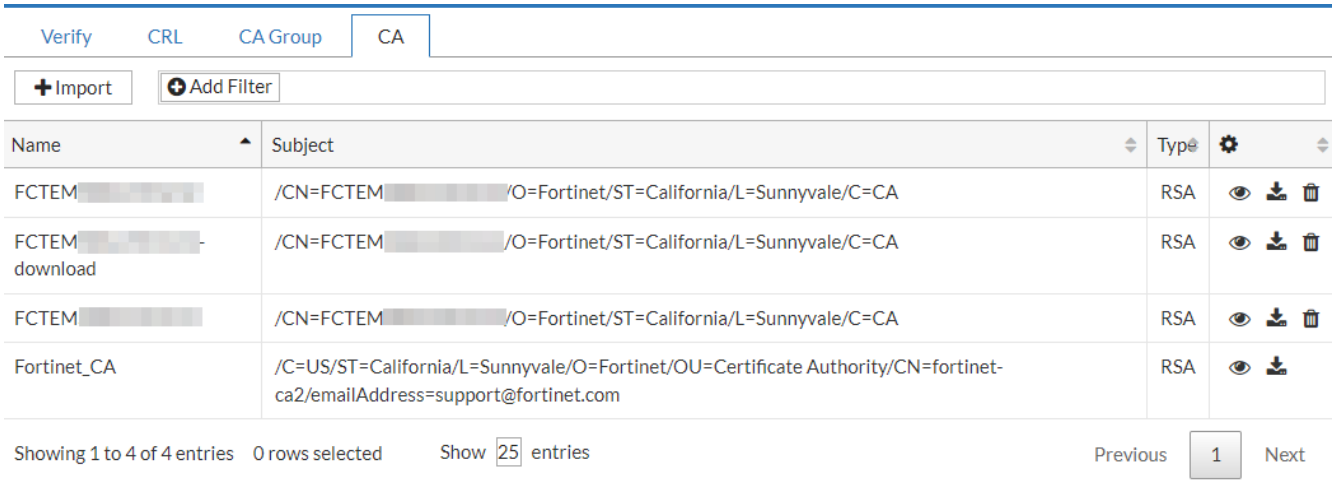
Verifying EMS CA certificate, ZTNA tag, and FortiClient endpoint synchronized from FortiClient EMS

After the FortiADC device connects to the FortiClient EMS, it automatically synchronizes ZTNA tags, the EMS CA certificate and the FortiClient endpoint information from the FortiClient EMS. Verify all the information have successfully synchronized from FortiClient EMS to FortiADC in the following:

- **Systems > Verify > CA** tab to view the EMS CA certificate
- **Network Security > ZTNA > ZTNA Tags** tab to view the ZTNA tags
- **FortiView > ZTNA** to view the FortiClient endpoint information and status

Systems > Verify > CA tab

The EMS CA certificate is synchronized to **Systems > Verify > CA** tab. When a ZTNA profile is referenced in a VS configuration, the corresponding Client SSL profile must enable verification of the client device certificate.



Click the (View icon) to see the EMS CA certificate details.

Network Security > ZTNA > ZTNA Tags tab

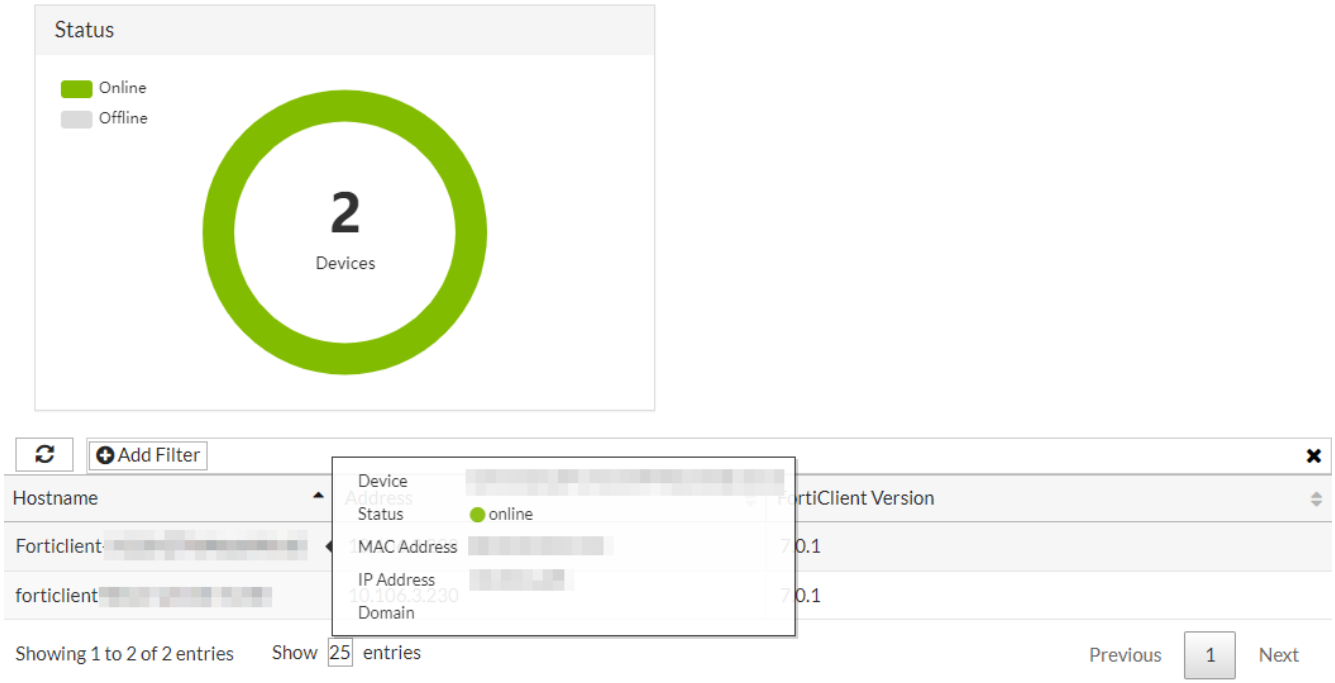
ZTNA tags are synchronized to the **Network Security > ZTNA > ZTNA Tags** tab. After the FortiClient EMS connector has successfully connected, check the **ZTNA Tags** page to ensure the corresponding ZTNA tag has been synchronized.



The **Ref.** column indicates the number of ZTNA Profile rules that have referenced the ZTNA tag.

FortiView > ZTNA

The FortiClient endpoint information and status are synchronized to the **FortiView > ZTNA**. From here, you can monitor the real-time status of the endpoints registered to FortiClient EMS.



You can hover over the **Hostname** column to view the device details synchronized from the FortiClient EMS.

Configuring a ZTNA Profile

The ZTNA Profile is the ZTNA policy used to enforce access control to Layer 7 HTTPS and TCPS virtual servers. ZTNA profiles consist of one or more ZTNA rule that determine the Source IP and ZTNA tags that are allowed access, and the resulting action to take.

After you have created a ZTNA profile, you can reference the ZTNA profile in Layer 7 HTTPS and TCPS virtual server Security configurations.

Before you begin:

- You must have registered the FortiADC device through the FortiClient EMS connector. For more information, see Zero Trust Network Access (ZTNA) on page 854 and Configuring FortiClient EMS Connector for ZTNA on page 859.
- You must have Read-Write permission for System settings.

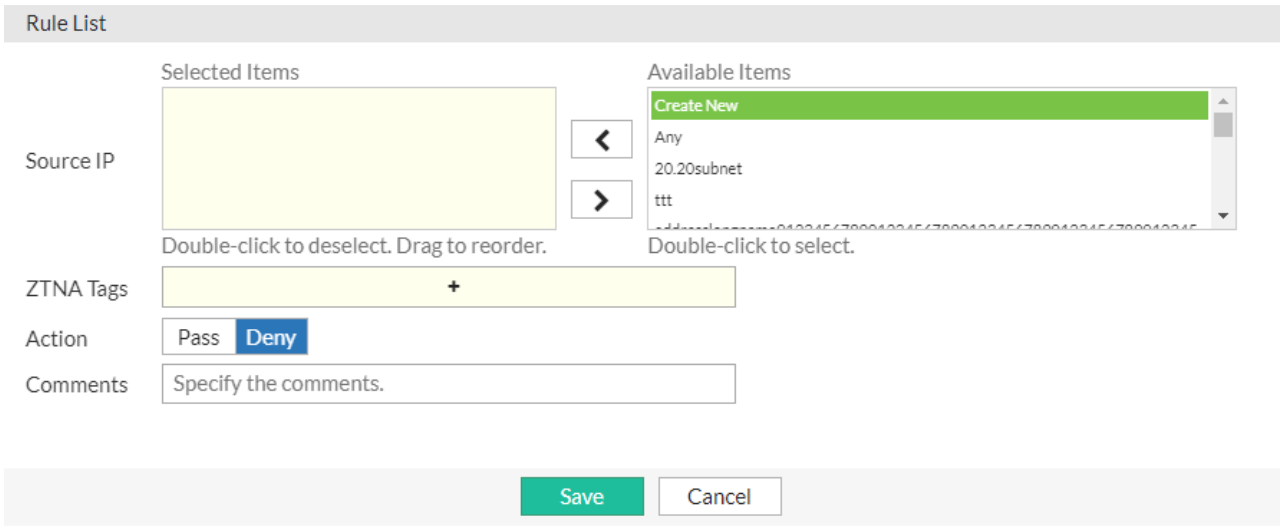
To create and configure a ZTNA Profile:

1. Go to **Network Security > ZTNA**.
2. Click the **ZTNA Profile** tab.
3. Click **Create New** to display the configuration editor.
4. Configure the following:

Parameter	Description
Name	Specify the ZTNA Profile name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces.
Log Status	Enable/disable logging.
Comments	Optionally, enter comments about the ZTNA Profile.

5. Click **Save**.
6. Under **Rule List**, click **Create New** to display the configuration editor.
7. Configure the following:

Parameter	Description
Source IP	Select the source IPs.
ZTNA Tags	Select the ZTNA tags.
Action	Select either of the following actions: • Pass • Deny **Deny** is the default action.
Comments	Optionally, enter comments about the ZTNA rule.



8. Click **Save**.

Apply the Security ZTNA profile to a Layer 7 HTTPS or TCPS virtual server to activate ZTNA for server load balancing. Ensure the corresponding Client SSL profile is enabled for client certificate verification. For details, see Configuring virtual servers on page 384 and Configuring Client SSL profiles on page 512.

ZTNA troubleshooting and debugging

The following CLI commands can be used to troubleshoot ZTNA issues:

Command	Description
`# execute fctems test-connectivity <EMS>`	Verify the FortiADC to FortiClient EMS connectivity. This provides the connection status of your FortiClient EMS connector configuration. If the connection is not successful, further detail is provided for the status condition.
`# execute fctems is-verified <EMS>`	Check if the configured EMS server has a verified certificate.
`# diagnose debug module fcnacd`	View information about your FortiClient NAC daemon (fcnacd), which handles FortiADC to FortiClient EMS connectivity.
`# diagnose endpoint-control client list`	List the FortiClient endpoints synchronized to FortiADC from FortiClient EMS.
`# diagnose endpoint-control tag list`	List the ZTNA tags synchronized to FortiADC from FortiClient EMS.
`# diagnose debug module httproxy ztna`	View information about your Layer 7 HTTPS virtual server that has referenced a ZTNA Profile.

For details of the above commands, see the FortiADC CLI Reference document.

DoS Protection

The **DoS Protection** menu contain features and configurations that allow you to protect your applications against DoS (denial-of-service) attacks.

This section is organized into the following sub-menu topics:

- DoS Protection Profile on page 867
- Application on page 869
- Networking on page 879

DoS Protection Profile

From the **DoS Protection > DoS Protection Profile** sub-menu, you can configure a DoS Protection profile that references the DoS policies that are to be enforced.

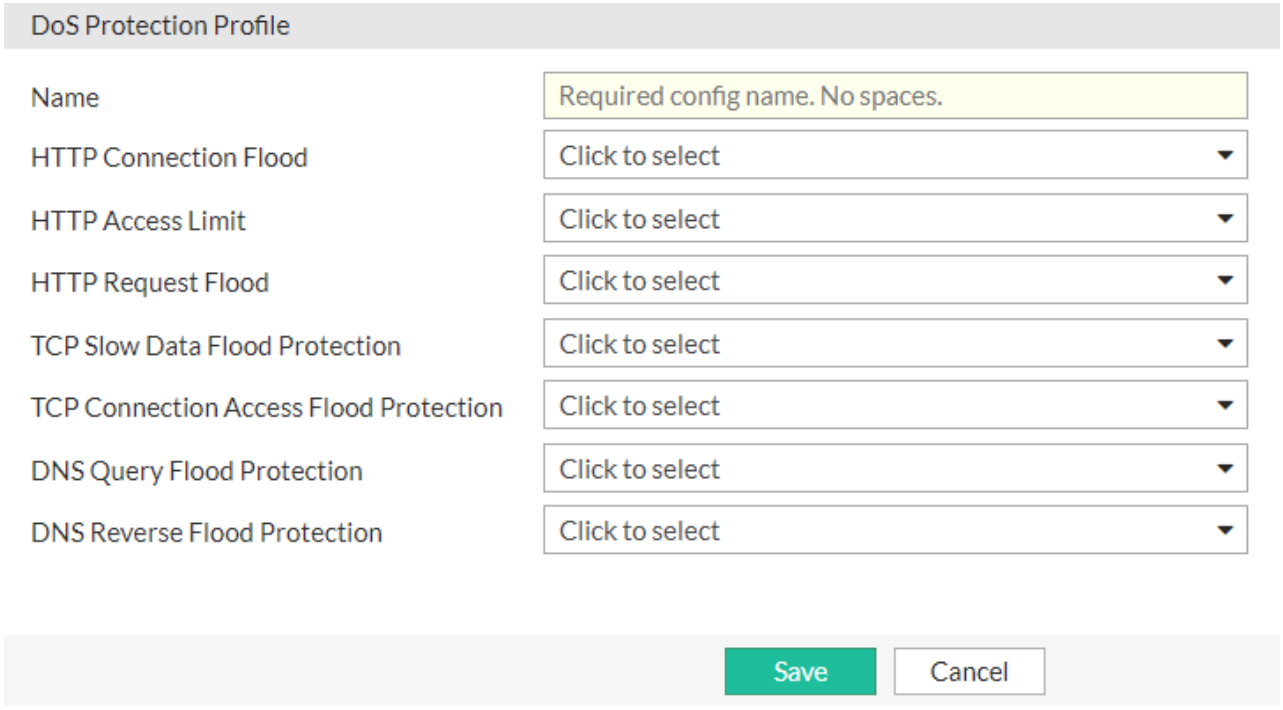
Before you begin:

- 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 FortiADC 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.
TCP Connection Access Flood Protection	A TCP connection flood attempts to prevent legitimate requests from being established by flooding the server with requests for new connections. By setting a threshold limit for TCP requests, FortiADC can detect and take action to protect against a TCP connection flood.
DNS Query Flood Protection	The DNS Query Flood Protection policy can limit the number of DNS request per second to mitigate against DNS query flood attacks that aim to overwhelm DNS servers with high volumes of illegitimate DNS queries.
DNS Reverse Flood Protection	The DNS Reverse Flood Protection policy can limit the number of ANY type DNS requests per second to mitigate against DNS reverse flood attacks that aim to overwhelm network resources with high volumes of DNS responses.

4. **Save** the configuration.

Application

The **DoS Protection > DoS Protection Profile** sub-menu allows you to perform the following tasks to configure the application DoS protection:

- Configuring an HTTP Access Limit policy on page 870
- Configuring an HTTP Connection Flood policy on page 871
- Configuring an HTTP Request Flood policy on page 873

- Configuring a DNS Query Flood Protection policy on page 874
- Configuring a DNS Reverse Flood Protection policy on page 877

Configuring an HTTP Access Limit policy

HTTP Access Limit policies can control the rate of HTTP requests from a source IP. The "Limit Request by IP" feature restricts the number of HTTP requests per second per source IP address. This helps prevent HTTP request floods involving multiple URLs and can identify source IP addresses shared by multiple clients. Even when clients do not support cookies, the system can intelligently apply separate request rate limits for those IPs.

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 an HTTP Access Limit policy:

1. Go to **DoS Protection > Application**.
   The configuration page displays the **HTTP Access Limit** tab.
2. Click **Create New** to display the configuration editor.

3. Configure the following settings:

Setting	Description
Name	Specify the configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces.
Status	Enable or disable this HTTP Access Limit policy. This is disabled by default.
HTTP Request Limit	Specify the maximum number of HTTP requests allowed per second from a specific IP address. A value of 0 disables the limit, allowing unlimited HTTP requests. The default value is set to 0, with an acceptable range of 0-65535. When FortiADC receives multiple HTTP requests from the same source IP within one second, it checks whether the request count exceeds the specified limit. If the limit is reached, FortiADC takes the configured action. The one-second timer starts when the first request is received, and the count block is automatically cleared after the one-second period elapses.
Action	Specify the action to take when the HTTP Connection Flood policy is triggered. • 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. • Captcha—The client must successfully complete the CAPTCHA request, which is triggered when a DDoS attack is detected. The CAPTCHA is applied based solely on the client's IP address.
Period Block	The **Period Block** option is available if **Action** is **Period Block**. Specify the duration (in seconds) for blocking connections. The default value is 60 seconds, with an acceptable range of 1 to 3600 seconds. During this period, HTTP requests containing the blocked cookie will result in the associated TCP connection being terminated. This blocking action persists even after the FortiADC is rebooted.
Log	Enable or Disable logging. When enabled, the Action will be logged.
Severity	Select the severity level. • 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 HTTP Connection Flood policy

The HTTP Connection Flood policy enforces a limit on the number of TCP connections from a client, identified by a session cookie. This policy operates at the application layer, leveraging the cookie to track and restrict the connection rate for each client.

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 an HTTP Connection Flood policy:

1. Go to **DoS Protection > Application**.
2. Click the **HTTP Connection Flood** tab.
3. Click **Create New** to display the configuration editor.
4. Configure the following settings:

Setting	Description
Name	Specify the configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces.
Status	Enable or disable this HTTP Connection Flood policy. This is disabled by default.
HTTP Connection Number Limit	Define the maximum number of TCP connections allowed per HTTP session. This limit applies to connections sharing the same session cookie. The default value is set to 1, with an acceptable range of 1-1024.
Action	Specify the action to take when the HTTP Connection Flood policy is triggered. • 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. • Captcha—The client must successfully complete the CAPTCHA request, which is triggered when a DDoS attack is detected. The CAPTCHA is applied based solely on the client's IP address.
Period Block	The **Period Block** option is available if **Action** is **Period Block**. Specify the duration (in seconds) for blocking connections. The default value is 60 seconds, with an acceptable range of 1 to 3600 seconds. During this period, HTTP requests containing the blocked cookie will result in the associated TCP connection being terminated. This blocking action persists even after the FortiADC is rebooted.
Log	Enable or Disable logging. When enabled, the Action will be logged.
Severity	Select the severity level. • High—Log as high severity events. • Medium—Log as a medium severity events. • Low—Log as low severity events. The default value is High.

5. **Save** the configuration.

Configuring an HTTP Request Flood policy

The HTTP Request Flood policy can limit the rate of HTTP requests from a client based on a session cookie, restricting the number of requests per second per session. This feature operates at the application layer, utilizing session cookies, rather than relying solely on TCP/IP connections at the network layer. This allows the system to distinguish between multiple clients that may share the same source IP address, such as in cases where network address translation (NAT) is used.

For this feature to function, the client must support cookies.

Primarily designed to protect against HTTP request-based DDoS attacks, the policy works by adding a Set-Cookie header to the HTTP response before forwarding the backend real-server's response to the client. As a result, subsequent client requests will include the relevant cookie (with most browsers supporting cookies). Using these cookies as keys, FortiADC can assess whether the request rate exceeds the configured limits, helping to identify and mitigate potential DDoS attacks.

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 an HTTP Request Flood policy:

1. Go to **DoS Protection > Application**.
2. Click the **HTTP Request Flood** tab.
3. Click **Create New** to display the configuration editor.

4. Configure the following settings:

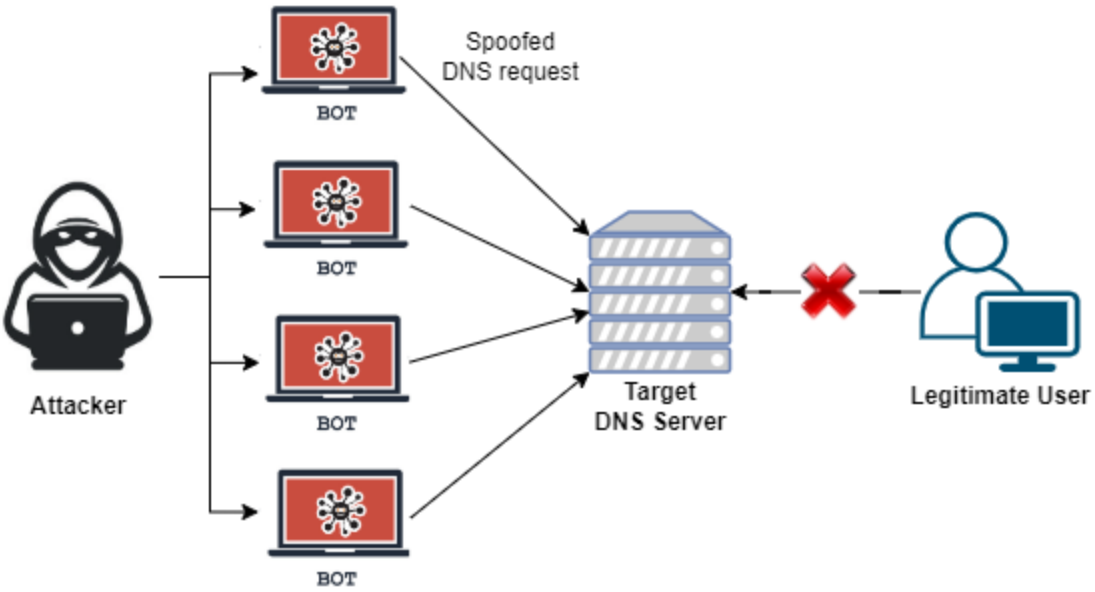
Setting	Description
Name	Specify the configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces.
Status	Enable or disable this HTTP Request Flood policy. This is disabled by default.
HTTP Request Limit	Specify the maximum number of HTTP requests allowed per second for a given session cookie. A value of 0 indicates no limit on HTTP requests. The default value is set to 0, with an acceptable range of 0-65535. When FortiADC receives an HTTP request, it first checks for a match between the URL and host. If a match is found, FortiADC inserts a cookie into the response header. For subsequent requests carrying the cookie inserted by FortiADC, the system tracks all TCP connections associated with the same cookie. If the number of connections reaches the specified limit, FortiADC enforces the configured action.
Action	Specify the action to take when the HTTP Connection Flood policy is triggered. • 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. • Captcha—The client must successfully complete the CAPTCHA request, which is triggered when a DDoS attack is detected. The CAPTCHA is applied based solely on the client's IP address.
Period Block	The **Period Block** option is available if **Action** is **Period Block**. Specify the duration (in seconds) for blocking connections. The default value is 60 seconds, with an acceptable range of 1 to 3600 seconds. During this period, HTTP requests containing the blocked cookie will result in the associated TCP connection being terminated. This blocking action persists even after the FortiADC is rebooted.
Log	Enable or Disable logging. When enabled, the Action will be logged.
Severity	Select the severity level. • High—Log as high severity events. • Medium—Log as a medium severity events. • Low—Log as low severity events. The default value is High.

5. **Save** the configuration.

Configuring a DNS Query Flood Protection policy

The DNS Query Flood Protection policy can limit the number of DNS request per second to mitigate against DNS query flood attacks.

A DNS query flood is a type of DDoS attack that targets the DNS server using DNS query and response traffic. A DNS query flood attack aims to overwhelm the target's resources with an excessive volume of fake requests, causing service unavailability and disrupting legitimate user access to the online service.

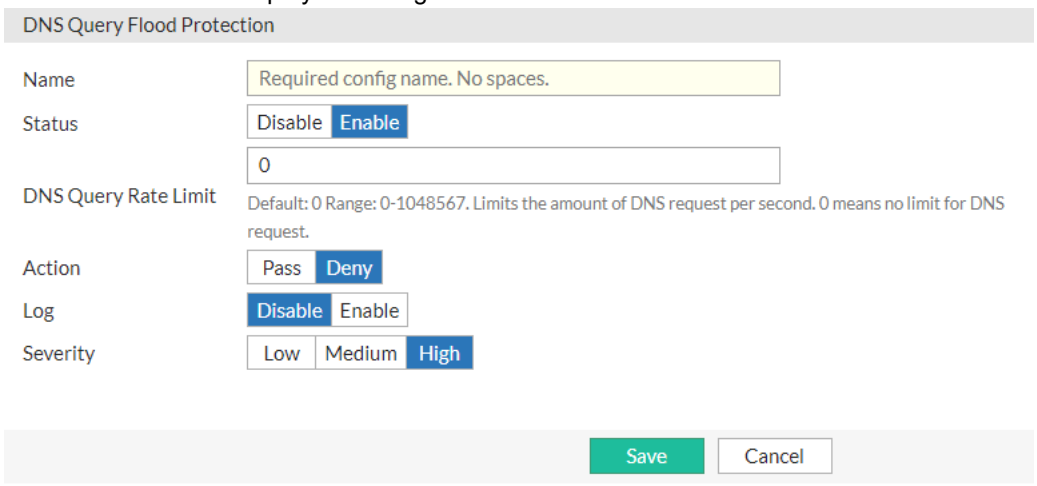


Using the DNS Query Flood Protection policy, you can set a DNS Query Rate Limit to restrict the number of queries that can be made per second. Once the query rate exceeds the limit, it will trigger a corresponding action (Pass or Deny).

After you have configured a DNS Query Flood Protection policy, you can apply it in a DoS Protection Profile.

To configure a DNS Query Flood Protection policy:

1. Go to **DoS Protection > Application**.
2. Click the **DNS Query Flood Protection** tab.
3. Click **Create New** to display the configuration editor.



4. Configure the following DNS Query Flood Protection settings:

Setting	Description
Name	Specify a name for the DNS Query Flood Protection policy. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No spaces. The configuration name cannot be edited once it has been saved.
Status	Enable/disable the status of this DNS Query Flood Protection policy.
DNS Query Rate Limit	Specify the allowable number of DNS requests per second. The range is 0 to 1048567. The default is 0, which means that no limit is placed on the number of DNS queries that can be made per second. **Note**: Multiple "rate limit" type of operations may be executed through various configurations, however, they cannot be executed all at once. Priority is given to certain rate limit operations. The following lists the execution sequence. **1.** Transaction Rate Limit (from the virtual server configuration). **2.** DNS Query Rate Limit (from DNS Query Flood Protection policy). **3.** ANY Query Rate Limit (from DNS Reverse Flood Protection policy).
Action	Select the corresponding action to take when the DNS Query Rate Limit is exceeded: • Pass — Allow the traffic. • Deny — Drop the traffic, send a 400 Bad request to the client. **Deny** is the default option.
Log	Enable/disable logging for the Action. This is disabled by default.
Severity	Select the event severity to log when the DNS Query Flood Protection policy is triggered: • High — Log as high severity events. • Medium — Log as a medium severity events. • Low — Log as low severity events. The default is **High**.

5. Click **Save**.
After the new DNS Query Flood Protection policy has been saved, it will appear in the DNS Query Flood Protection page. You can now apply this DNS Query Flood Protection policy to a DoS Protection Profile configuration.



Statistical data of DNS Query Flood attacks are recorded in Security logs in detail (from FortiView and Log & Report), and in the Dashboard Security widget as event counts.

Typically, Security logs provide the attack count as well, however, due to the speed at which DNS Query Flood attacks can occur (within milliseconds), the Security logs cannot accurately count each attack as the logs are counted every 1 second. This means that when a DNS Query Flood attack occurs, Security logs can only capture the attack details but not the accurate count of the number of packets that has exceeded the DNS Query Rate Limit.

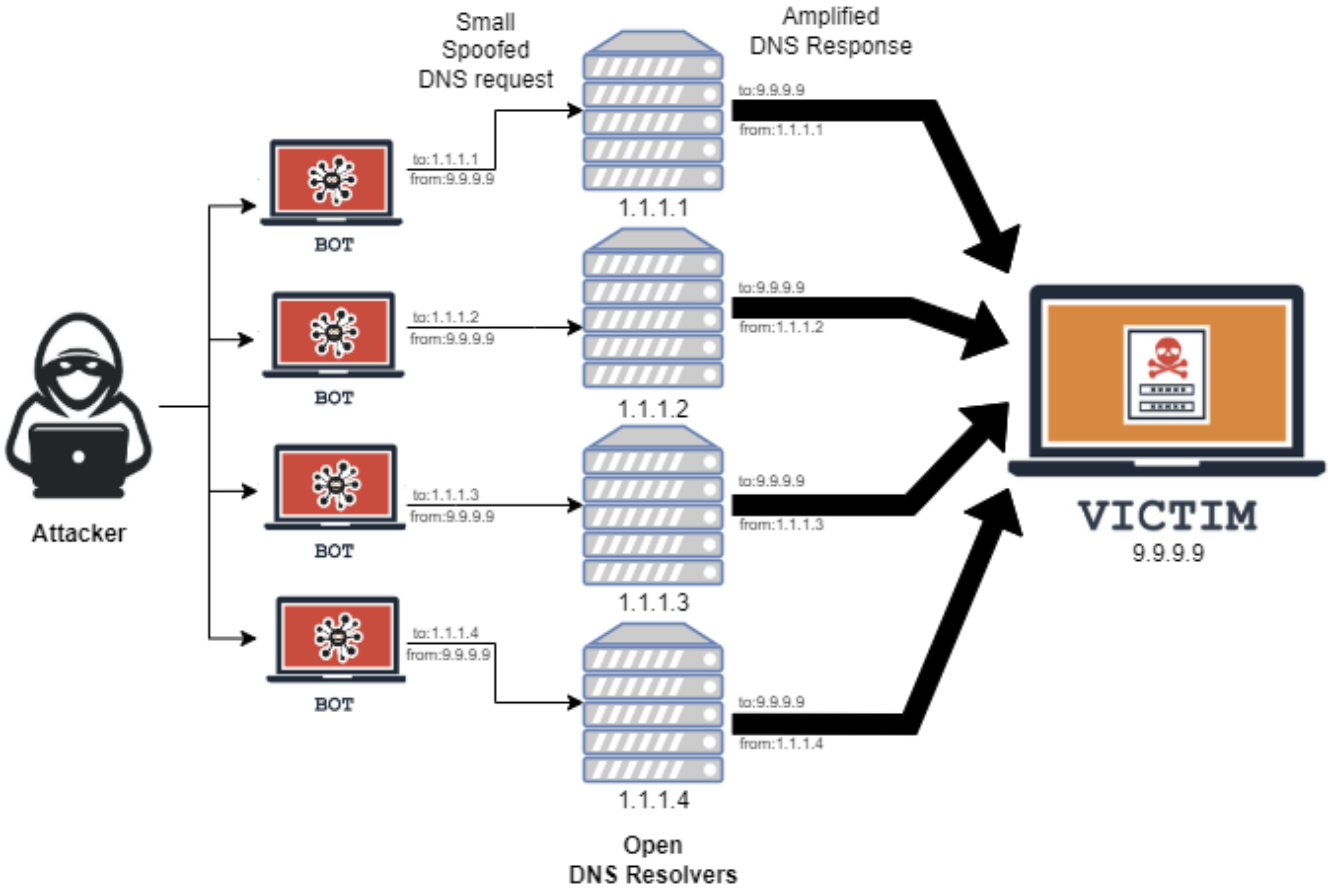
To view the correct count of each DNS Query Rate Limit excess event, you can reference the Security widget from the Dashboard that is a dedicated security event counter.

Configuring a DNS Reverse Flood Protection policy

The DNS Reverse Flood Protection policy can limit the number of ANY type DNS requests per second to mitigate against DNS reverse flood attacks.

A DNS reverse flood is a type of DDoS attack that uses high volumes of DNS responses to overwhelm network resources. When a small request has a disproportionately larger response, the outbound network pipe can easily become congested and disrupt traffic responses for legitimate requests and/or other applications. And since the transport protocol is UDP, these types of requests can be easily spoofed.

One popular form of DNS reverse flood attack is DNS amplification attack. The Attacker sends out a DNS query with a forged IP address (the Victim's) to an open DNS resolver, prompting it to reply back to that address with a DNS response. With numerous fake queries being sent out, and with several DNS resolvers replying back simultaneously, the Victim's network can easily be overwhelmed by the high volume of DNS responses.

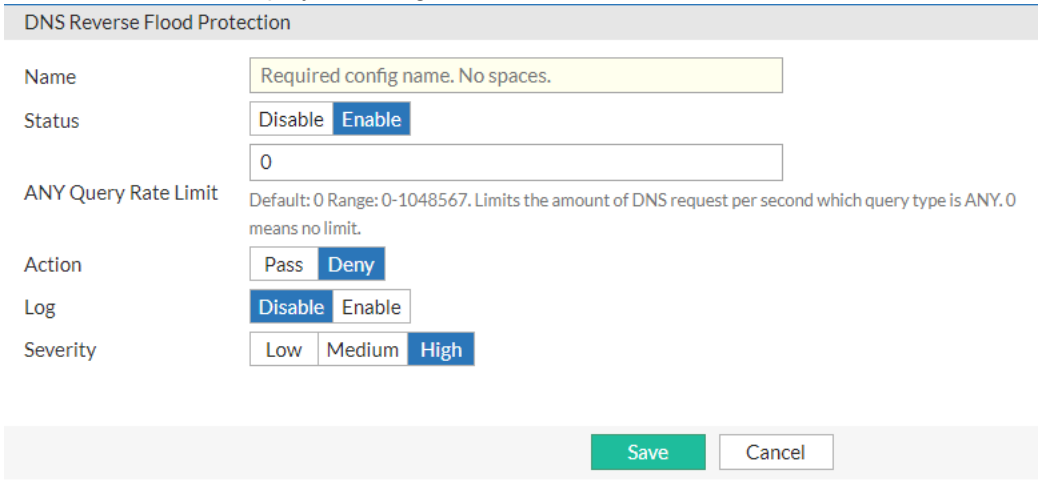


Using the DNS Reverse Flood Protection policy, you can set a ANY Query Rate Limit to restrict the number of all types (ANY) of DNS queries that can be made per second. Once the query rate exceeds the limit, it will trigger a corresponding action (Pass or Deny).

After you have configured a DNS Reverse Flood Protection policy, you can apply it in a DoS Protection Profile.

To configure a DNS Reverse Flood Protection policy:

1. Go to **DoS Protection > Application**.
2. Click the **DNS Reverse Flood Protection** tab.
3. Click **Create New** to display the configuration editor.



4. Configure the following DNS Reverse Flood Protection settings:

Setting	Description
Name	Specify a name for the DNS Reverse Flood Protection policy. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No spaces. The configuration name cannot be edited once it has been saved.
Status	Enable/disable the status of this DNS Reverse Flood Protection policy.
ANY Query Rate Limit	Specify the allowable number of DNS requests per second, with query type ANY. The range is 0 to 1048567. The default is 0, which means that no limit is placed on the number of DNS queries that can be made per second. **Note**: Multiple "rate limit" type of operations may be executed through various configurations, however, they cannot be executed all at once. Priority is given to certain rate limit operations. The following lists the execution sequence. **1.** Transaction Rate Limit (from the virtual server configuration). **2.** DNS Query Rate Limit (from DNS Query Flood Protection policy). **3.** ANY Query Rate Limit (from DNS Reverse Flood Protection policy).
Action	Select the corresponding action to take when the ANY Query Rate Limit is exceeded: • Pass — Allow the traffic. • Deny — Drop the traffic, send a 400 Bad request to the client. **Deny** is the default option.
Log	Enable/disable logging for the Action. This is disabled by default.
Severity	Select the event severity to log when the DNS Reverse Flood Protection policy is triggered:

Setting	Description
	• High — Log as high severity events. • Medium — Log as a medium severity events. • Low — Log as low severity events. The default is High.

5. Click **Save**.
After the new DNS Reverse Flood Protection policy has been saved, it will appear in the DNS Reverse Flood Protection page. You can now apply this DNS Reverse Flood Protection policy to a DoS Protection Profile configuration.



Statistical data of DNS Reverse Flood attacks are recorded in Security logs in detail (from FortiView and Log & Report), and in the Dashboard Security widget as event counts.

Typically, Security logs provide the attack count as well, however, due to the speed at which DNS Reverse Flood attacks can occur (within milliseconds), the Security logs cannot accurately count each attack as the logs are counted every 1 second. This means that when a DNS Reverse Flood attack occurs, Security logs can only capture the attack details but not the accurate count of the number of packets that has exceeded the ANY Query Rate Limit.

To view the correct count of each ANY Query Rate Limit excess event, you can reference the Security widget from the Dashboard that is a dedicated security event counter.

Networking

The **DoS Protection > Networking** sub-menu allows you to perform the following tasks to configure the network DoS protection:

- Configuring an IP fragmentation policy on page 879
- Configuring a TCP SYN flood protection policy on page 880
- Configuring a TCP slow data flood protection policy on page 881
- Configuring a TCP connection access flood protection policy on page 882

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**.
   The configuration page displays the **IP Fragmentation Protection** tab.
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 attacks. 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:

1. Go to **DoS Protection > Networking**.
2. Click the **TCP SYN Flood Protection** tab.
3. Click **Edit** to display the configuration editor.
4. 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.

5. **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 TCP Slow Data Flood Protection policy:

1. Go to **DoS Protection > Networking**.
2. Click the **TCP Slow Data Flood Protection** tab.
3. Click **Create New** to display the configuration editor.
4. 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

5. **Save** the configuration.

Configuring a TCP connection access flood protection policy

A TCP connection flood attempts to prevent legitimate requests from being established by flooding the server with requests for new connections. By setting a threshold limit for TCP requests, FortiADC can detect and take action to protect against a TCP connection flood.

Before you begin:

- You must have Read-Write permission for Security settings.

To configure a TCP Connection Access Flood Protection policy:

1. Go to **DoS Protection > Networking**.
2. Click the **TCP Connection Access Flood Protection** tab.
3. Click **Create New** to display the configuration editor.
4. Configure the following settings:

Setting	Description
Name	Enter a unique name for the TCP Connection Access Flood Protection configuration. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces.
Status	If enabled, this policy will be activated, otherwise it is inactive. This is enabled by default.
Limit	Limit the number of TCP connections per source IP address. Default: 0 Range: 0-65535. 0 means no TCP request limit.
Action	Specify the action to take when the TCP connection number exceeds the limit. • Pass — allow the new connection from this IP address. • Deny — deny the new connection from this IP address • Period Block — deny the new connection from this IP address for a period of time. The default action is **Deny**.
Period Block	The **Period Block** option is available if the **Action** is **Period Block**. Specify the number of seconds to block new TCP connections from being established. During this period, new TCP connection requests will be aborted. Default: 60 Range: 1-3600 seconds.
Log	Enable or disable log. This is disabled by default.
Severity	Specify the log severity level. • Low — Log as low severity events. • Medium — Log as a medium severity events. • High — Log as high severity events. The default severity level is **High**.

5. Click **Save**.

User Authentication

The **User Authentication** menu contain features and configurations relating to user authentication.

This section is organized into the following sub-menu topics:

- Authentication Policy on page 883
- User Group on page 885
- Local User on page 892
- Remote User on page 893
- Authentication Relay on page 908
- SAML on page 913
- AD FS Proxy on page 917
- OAuth Proxy on page 920

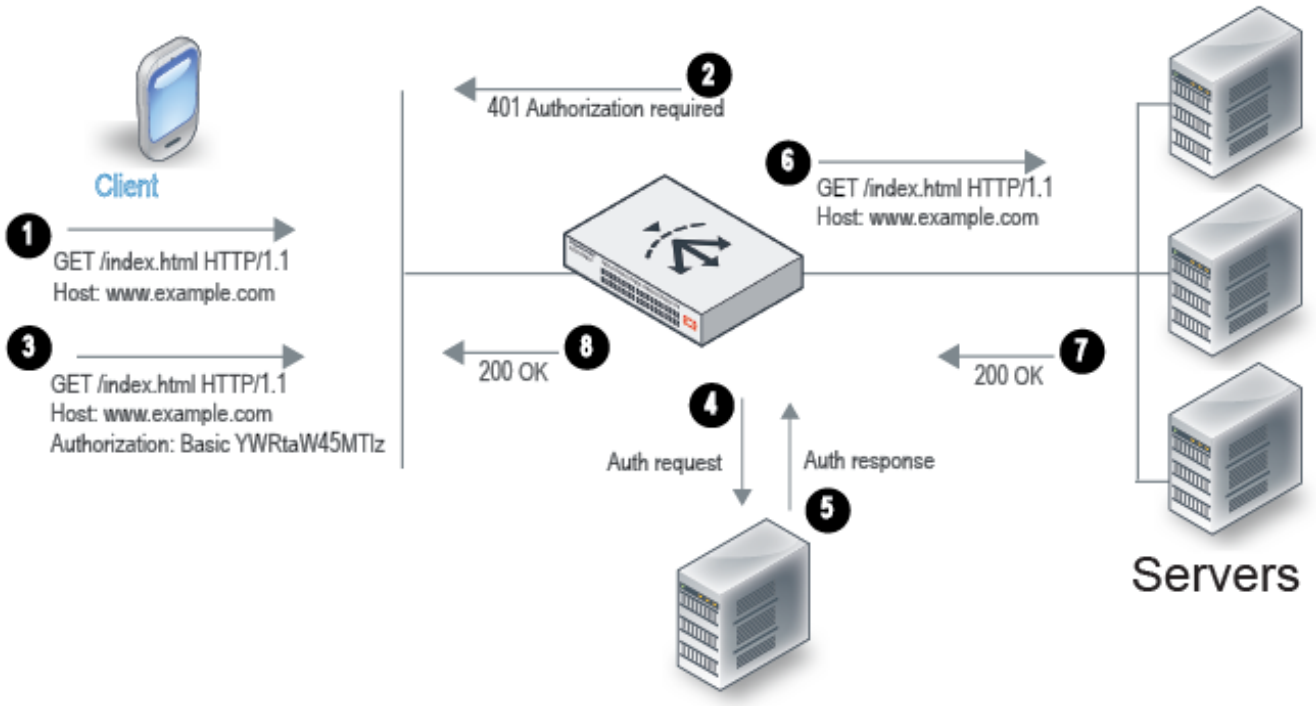
Authentication Policy

From the **User Authentication > Authentication Policy** sub-menu, you can configure authentication policies to 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 883 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 885.
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, `www.example.com` matches `www.example.com` but not `www.example.com.hk`.
Type	Select either of the following: • Standard • SAML • OAuth
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. Available only if **Standard** is selected as the **Type**.
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, `/abc` matches `http://www.example.com/abcd` and `http://www.example.com/abc/11.html` but not `http://www.example.com/1abcd`.
User Group	Select the user group that is authorized to access the protected resource. Available only if **Standard** is selected as the **Type**.
SAML SSO ID	Select the SAML SSO ID that is authorized to access the protected resource. Available only if **SAML** is selected as the **Type**.
OAuth Policy	Select the OAuth policy that is authorized to access the protected resource. Available only if **OAuth** is selected as the **Type**.

User Group

The **User Authentication > User Group** sub-menu allows you to perform the following tasks to configure user groups for user authentication:

- Configuring user groups on page 886
- Configuring customized authentication forms on page 888

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, RADIUS, NTLM, and TACACS+ 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:

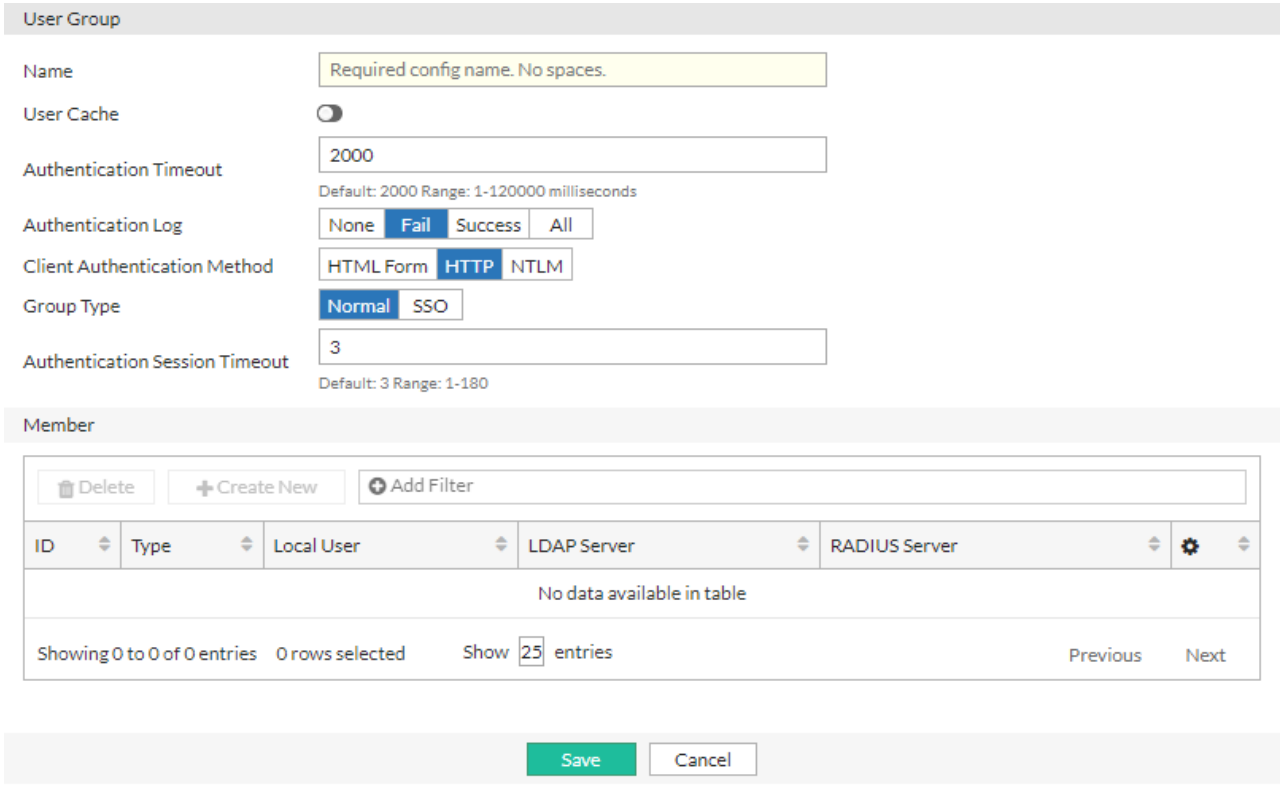
- You must have created configuration objects for any LDAP, RADIUS, NTLM, and/or TACACS+ servers you want to use, and you must have created user accounts for local users.
- You must have read-write permission for System and User settings.

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**.
   The configuration page displays the **User Group** tab.

2. Click **Create New** to display the configuration editor.



3. Configure the following User Group settings:

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, TACACS+) once they are authorized.
User Cache Timeout	The **User Cache Timeout** option is available if **User Cache** is enabled. 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 1-120,000 milliseconds.
Authentication Log	Specify one of the following logging options for authentication events: • None — No logging. • Fail — Log failed attempts. • Success — Log successful attempts. • All — Log all (both failed and successful attempts).
Client Authentication Method	• HTML Form • HTTP • NTLM (only if you want to use NTLM server as a authentication server)

Settings	Guidelines
Use Default Form	The **Use Default Form** option is available if **Client Authentication Method** is **HTML Form**. Enabled by default to use the default authentication form. Disable to use a customized authentication form.
Customized Authentication Form	The **Customized Authentication Form** option is available if **Client Authentication Method** is **HTML Form** and **Use Default Form** is **disabled**. Select a Customized Authentication Form object or create new.
Group Type	• Normal — Default. No action is needed. • SSO — Select to enable Single Sign-On (SSO).
Authentication Relay	The **Authentication Relay** option is available if **Group Type** is **SSO**. 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 Cross Domain Support	The **SSO Cross Domain Support** option is available if **Group Type** is **SSO**. Disabled by default. When enabled, you must specify the SSO domain. **Note**: Authentication policies cannot be applied to multiple virtual servers. Due to security reasons, such as protection against XSS attacks, there is no shared mechanism between virtual servers to decrypt cookies. As a result, you cannot log into a second virtual server while already logged into the first virtual server as the virtual servers are independent from each other. SSO Cross Domain Support allows you to have multiple domain names on the same virtual server (the virtual host), where you can specify a first-level domain name to enable the second-level domain names on the virtual server to decrypt cookies at the same time.
SSO Domain	The **SSO Domain** option is available if **Group Type** is **SSO** and **SSO Cross Domain Support** is **enabled**. Specify the SSO domain.
Log Off URL	The **Log-off URL** option is available if **Group Type** is **SSO**. Specify the log-off URL.

4. Click **Save**.
   Once the **User Group** configuration is saved, the **Member** section becomes available for configuration.
5. Under the **Member** section, click **Create New** to display the configuration editor.
6. Configure the following Member settings and save the configuration:
   a. Select the **Type**: Local, LDAP, RADIUS, NTLM, or TACACS+.
   b. Select the corresponding configuration based on the selected Type.
7. Click **Save** again to save the Member added to the User Group configuration.

Configuring customized authentication forms

FortiADC allows you to customize your login page with your company brand images and modify the layout/text.

Similar to the GUI of the error page, we allow the user to upload a zip/tar/tar.gz file. For details, see Importing a customized authentication form on page 889.

After you have configured your customized authentication form, you can apply it in the Authentication User Group.

To configure a Customized Authentication Form:

1. Go to **User Authentication > Customized Authentication Form**.
2. Click **Create New** to display the configuration editor.

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.
File	File package for customize Authentication form page. Click 'Choose File' to upload.
Username Field Name	The username field name in customize form
Password Field Name	The password field name in customize form
Virtual Path	Virtual path of customized authentication form function. This path is running on VS, so it will conflict with other configure like error page's vpath and Captcha.

3. Click **Save** to commit the changes.

Importing a customized authentication form

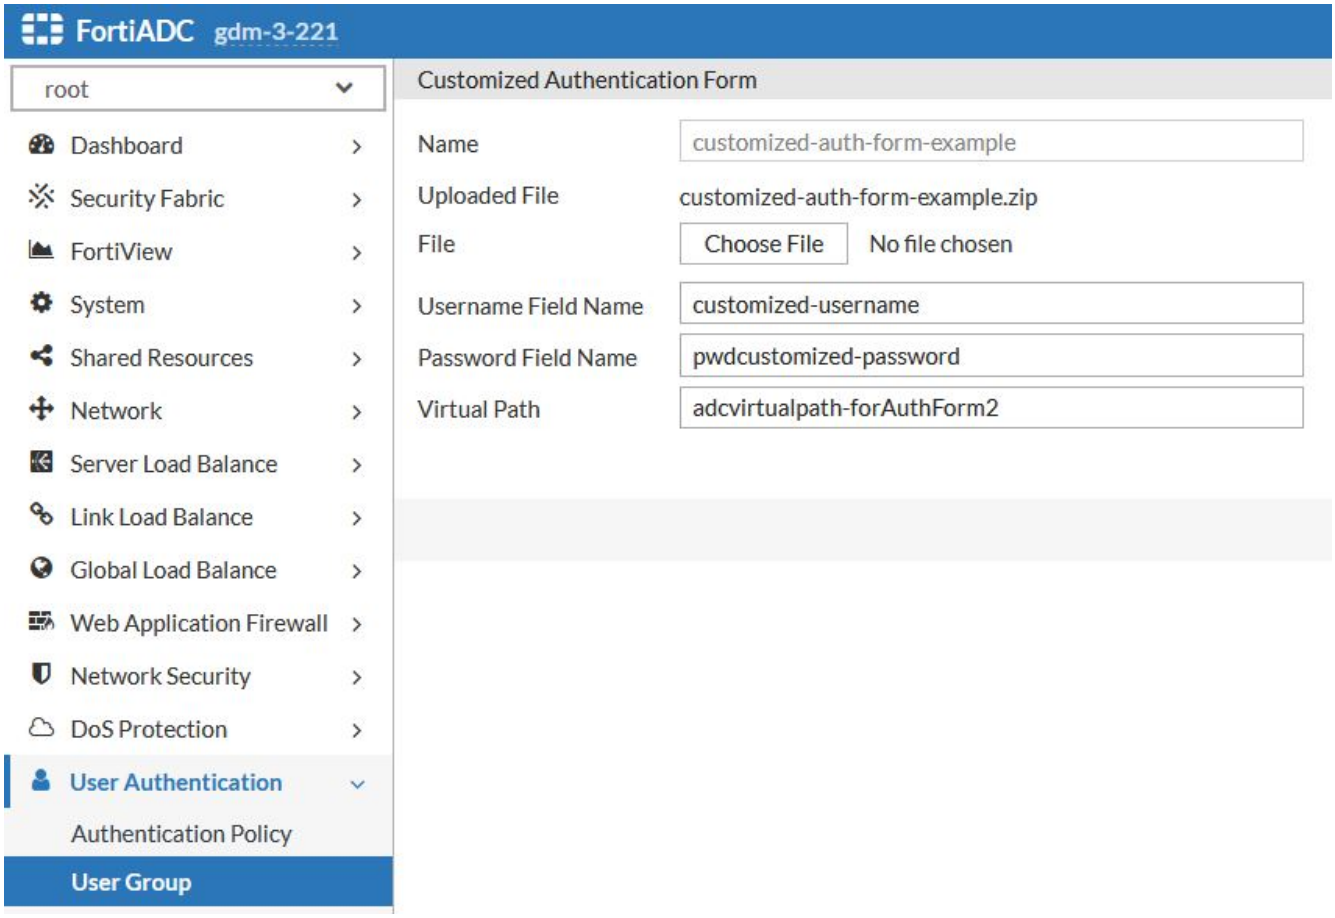
You can import a customized authentication page package in zip, tar, or tar.gz format. This package must include two files: `login.html` and `login_failed.html`. Both files must share the same username and password field IDs. The `login_failed.html` page is displayed when the user enters an incorrect username or password. Ensure that the form in these pages includes the attribute `onsubmit="return Fsb(event)"` and the placeholder tag `%%auth_script%%` to enable proper authentication functionality.

The following is an example of a login.html file:

```
<html>
   <head>
      <style type="text/css">html, body, div, h1, p, form, section,
           {fontsize:100%;font:inherit;vertical-align:baseline;} section {display:block;}
           body {lineheight:1;font:13px/20px 'Lucida Grande',Tahoma,Verdana,sansserif;
           color:#404040;background:#fff;} .cter {margin:80px auto;width:640px;} .lgin
           {position:relative;margin:0 auto;padding:20px 20px
           20px;width:270px;background:#f2f2f2;border-radius:3px;} .lgin h1 {margin:-20px -
           20px 21px;line-height:40px;font-size:15px;font-
           weight:bold;color:white;textalign:center;background:#555555;border-bottom:1px
           solid #cfcfcf;border-radius:3px 3px 0 0;} .lgin p {margin:20px 0 0;} .lgin
           p.submit {text-align:center;} input[type=text], input[type=password]
           {margin:5px;padding:0
           10px;width:160px;height:25px;color:#404040;background:white;border:1px
           solid;bordercolor:#c4c4c4 #d1d1d1 #d4d4d4;outline:3px solid #eff4f7;} input
           [type=text]:focus, input [type=password]:focus {border-color:#7dc9e2;outline-
           color:#dceefc;outline-offset:0;} input[type=submit] {padding:0
           18px;height:29px;font-
           size:12px;fontweight:bold;color:#527881;background:#cde5ef;border:1px
           solid;border-color:#b4ccce #b3c0c8 #9eb9c2;border-radius:8px;outline:0;} input
```

```
            [type=submit]:active {background:#dfe7f2;border-color:#9eb9c2 #b3c0c8
            #b4ccce;}</style>
   </head>
%%auth_script%%
   <body >
      <section class="cter">
         <div class="lgin">
            <h1>Web Authentication</h1>
            <form method="post" action="/" onsubmit="return Fsb(event)">
               <p>UserName: <input type="text" id="customize-username" required></p>
               <p>Password : <input type="password" id="pwdcustomized-password" required></p>
               <p class="submit"><input type="submit" value="Login"></p>
            </form>
         </div>
      </section>
   </body>
</html>
```

Using the information from the example login.html file, the following is the corresponding Customized Authentication Form configuration in the GUI.



The following is an example of a login_failed.html file:

```
<html>
    <head>
```

```
        <style type="text/css">html, body, div, h1, p, form, section,{font-
size:100%;font:inherit;vertical-align:baseline;} section {display:block;} body {line-
height:1;font:13px/20px 'Lucida Grande',Tahoma,Verdana,sans-
serif;color:#404040;background:#fff;} .cter {margin:80px auto;width:640px;} .lgin
{position:relative;margin:0 auto;padding:20px 20px
20px;width:270px;background:#f2f2f2;border-radius:3px;} .lgin h1 {margin:-20px -20px
21px;line-height:40px;font-size:15px;font-weight:bold;color:white;text-
align:center;background:#555555;border-bottom:1px solid #cfcfcf;border-radius:3px 3px 0 0;}
.lgin p {margin:20px 0 0;} .lgin p.submit {text-align:center;} input[type=text], input
[type=password] {margin:5px;padding:0
10px;width:160px;height:25px;color:#404040;background:white;border:1px solid;border-
color:#c4c4c4 #d1d1d1 #d4d4d4;outline:3px solid #eff4f7;} input[type=text]:focus, input
[type=password]:focus {border-color:#7dc9e2;outline-color:#dceefc;outline-offset:0;} input
[type=submit] {padding:0 18px;height:29px;font-size:12px;font-
weight:bold;color:#527881;background:#cde5ef;border:1px solid;border-color:#b4ccce #b3c0c8
#9eb9c2;border-radius:8px;outline:0;} input[type=submit]:active {background:#dfe7f2;border-
color:#9eb9c2 #b3c0c8 #b4ccce;}</style>
    </head>
%%auth_script%%
    <body >
        <img align="middle" src="goldbrig.jpg">
        <section class="cter">
            <div class="lgin">
                <h1>Web Authentication-Mick Fail login page</h1>
                <form method="post" action="/" onsubmit="return Fsb(event)">
                    <p>UserName-26653: <input type="text" id="customize-username"
required></p>
                    <p>Password-26653 : <input type="password" id="pwdcustomized-password"
required></p>
                    <p class="submit"><input type="submit" value="Login"></p>
                </form>
            </div>
        </section>
    </body>
</html>
```

Using the customized authentication form for 2FA token

To use a token for 2FA, you can upload a zip/tar/tar.gz file. The file must include a `2fatoken.html` file that contains the tag `%%auth_script%%` that you can use in your web page and use `onsubmit="return Fsb(event)"` in your form.

The following is an example of 2fatoken.html file for 2FA token:

```
<html>
   <head>
      <style type="text/css">html, body, div, h1, p, form, section,{font-
size:100%;font:inherit;vertical-align:baseline;} section {display:block;} body {line-
height:1;font:13px/20px 'Lucida Grande',Tahoma,Verdana,sans-
serif;color:#404040;background:#fff;} .cter {margin:80px auto;width:640px;} .lgin
{position:relative;margin:0 auto;padding:20px 20px
20px;width:270px;background:#f2f2f2;border-radius:3px;} .lgin h1 {margin:-20px -20px
21px;line-height:40px;font-size:15px;font-weight:bold;color:white;text-
align:center;background:#555555;border-bottom:1px solid #cfcfcf;border-radius:3px 3px 0 0;}
.lgin p {margin:20px 0 0;} .lgin p.submit {text-align:center;} input[type=text], input
[type=password] {margin:5px;padding:0
10px;width:160px;height:25px;color:#404040;background:white;border:1px solid;border-
```

```
color:#c4c4c4 #d1d1d1 #d4d4d4;outline:3px solid #eff4f7;} input[type=text]:focus, input
[type=password]:focus {border-color:#7dc9e2;outline-color:#dceefc;outline-offset:0;} input
[type=submit] {padding:0 18px;height:29px;font-size:12px;font-
weight:bold;color:#527881;background:#cde5ef;border:1px solid;border-color:#b4ccce #b3c0c8
#9eb9c2;border-radius:8px;outline:0;} input[type=submit]:active {background:#dfe7f2;border-
color:#9eb9c2 #b3c0c8 #b4ccce;}</style>
   </head>
   %%auth_script%%
   <body >
      <section class="cter">
      <img align="middle" src="goldbrig.jpg">
         <div class="lgin">
            <h1>2FA-Mick-page</h1>
            <form method="post" action="/" onsubmit="return Fsb(event)">
               <p>Token: <input type="text" id="token" required></p>
               <p class="submit"><input type="submit" value="Login"></p>
            </form>
         </div>
      </section>
   </body>
</html>
```

Local User

From the **User Authentication > Local User** sub-menu, you can configure a local authentication server to authenticate destination server user logins. FortiADC uses FortiToken Cloud as the remote authentication server which provides the security token needed for two-factor authentication on FortiADC.

To assign a FortiToken Cloud to a local server, the device must be registered on the same account as the FortiToken Cloud contracts; see Fortinet Customer Service & Support.

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.

The FortiToken Cloud User is only supported if the Client Authentication Method in the User group configuration is HTML Form.

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 893.
4. Save the configuration.

Local authentication server configuration

Settings	Guidelines
Name	Name of the user account, such as `user1` or `user1@example.com`. Do not use spaces or special characters except the 'at' symbol (`@`) or dot (`.`). 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.
Two-factor Authentication	• None — Default. Use the local authentication • FortiToken Cloud — Enable to 2FA authentication using FortiToken Cloud service. **Note**: FortiADC does not support FortiToken Cloud functionality in HA condition.
Email Address	The email is the email address that will receive the OTP. We will send the registration information including the QR code to help the user to register on the FortiToken app.
County Dial Code	The phone of the country code.
Phone Number	Use this phone number to send the OTP in an SMS text message to the mobile device
FortiToken Mobile Push	Enable two-factor push notifications to your mobile app for fast and secure access.

Remote User

The **User Authentication > Remote User** sub-menu allows you to perform the following tasks to configure remote servers for user authentication:

- Using an LDAP authentication server on page 893
- Using a RADIUS authentication server on page 903
- Using an NTLM authentication server on page 905
- Using a TACACS+ authentication server on page 906

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.

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 or FQDN and the 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. Click 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 895.
5. Click **Test Connectivity** to validate the configuration.
6. 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 or FQDN of the LDAP server. **Note**: When enabling LDAPS or StartTLS, please use the common name (CN) of the LDAP server certificate in this field if you want to identify the Server with the **CA profile** (see CA Profile on page 896).
Port	Port number for the server. Port 389 is typically used for non-secured connections or for StartTLS-secured connections, and port 636 is typically used for SSL-secured (LDAPS) connections.
Common Name Identifier	Enter the identifier for the common name (CN) attribute (also called the CNID) whose value is the user name. Identifiers vary based on the schema of your LDAP directory. This is often `cn` or `uid`. For Windows Active Directory, it is often the attribute `sAMAccountName`. For example, in a default OpenLDAP directory, if a user object is `uid=fortiadc,cn=users,dc=fortinet,dc=com` then the CNID is `uid`.

Distinguished Name	Specifies the Base DN from which the LDAP query starts. This DN is the full path in the directory to the user account objects.
	For example:
	`ou=People,dc=example,dc=com`
	or
	`cn=users,dc=example,dc=com`
	You can use the **Fetch DN** function to get the entire Directory Information Tree, and select the DN of the LDAP query starting entry.
	Note: When using Windows Active Directory as the LDAP server, you may need to use regular bind for FortiADC to get access permission for LDAP entries when using the **Fetch DN** function.

Bind Type	• Simple—bind without user search. It can be used only if all the users belong to the same “branch”. • 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**. Enter the bind DN of an LDAP user account with permissions to query the Distinguished Name (see Distinguished Name on page 895). The maximum length is 256 characters. For example: `cn=fortiadc,cn=users,dc=fortinet,dc=com`

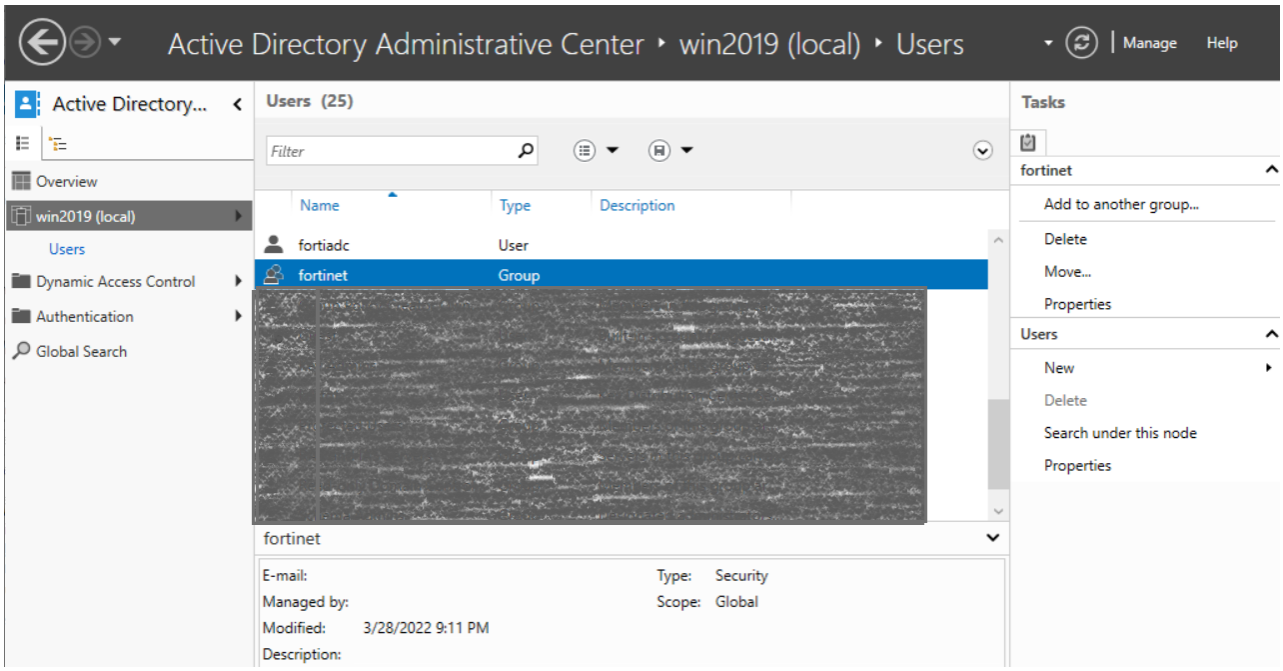
Settings	Guidelines
	For Windows Active Directory, the UPN (User Principle Name) is often used instead of a bind DN (for example, `user@domain.com`). This field can be optional if your LDAP server does not require the FortiADC appliance to authenticate when performing queries.
Password	Enter the password of the User DN.
Secure Connection	• Disable • LDAPS • STARTTLS
CA Profile	Available only when **Secure Connection** is set to **LDAPS** or **STARTTLS**, regardless of the Bind type being selected. Select a CA profile to identify the server certificate or you can leave the field blank. For details on how to import the CA profile, see Importing CAs on page 296.
Group Authentication	Available only when **Bind Type** is **Regular**. Enable to filter the query results, only allowing users to authenticate if they are members of the LDAP group that you define in the **Group DN** field. Users that are not members of that group are not allowed to authenticate.
Group Type	Available only when **Bind Type** is **Regular**. Indicate the schema of your LDAP directory as one of the following: • OpenLDAP — The directory uses a schema where each user object's group membership is recorded in an attribute named `gidNumber` or a virtual attribute named `memberOf`. • Windows-AD — The directory uses a schema where each user object's group membership is recorded in an attribute named `memberOf`. • FortiAuthenticator — Group membership attributes may have different names depending on the LDAP directory schema. The FortiADC appliance will use the group membership attribute that matches your directory's schema when querying the group DN.
Group DN	Available only when **Bind Type** is **Regular**. Enter the value of the group membership attribute that query results must have in order to be able to authenticate. For example: `ou=Groups,dc=example,dc=com` or a group ID (GID) such as `100` The value may vary based on your directory's schema. For details, see Setting the LDAP group on the LDAP server on page 896.

Setting the LDAP group on the LDAP server

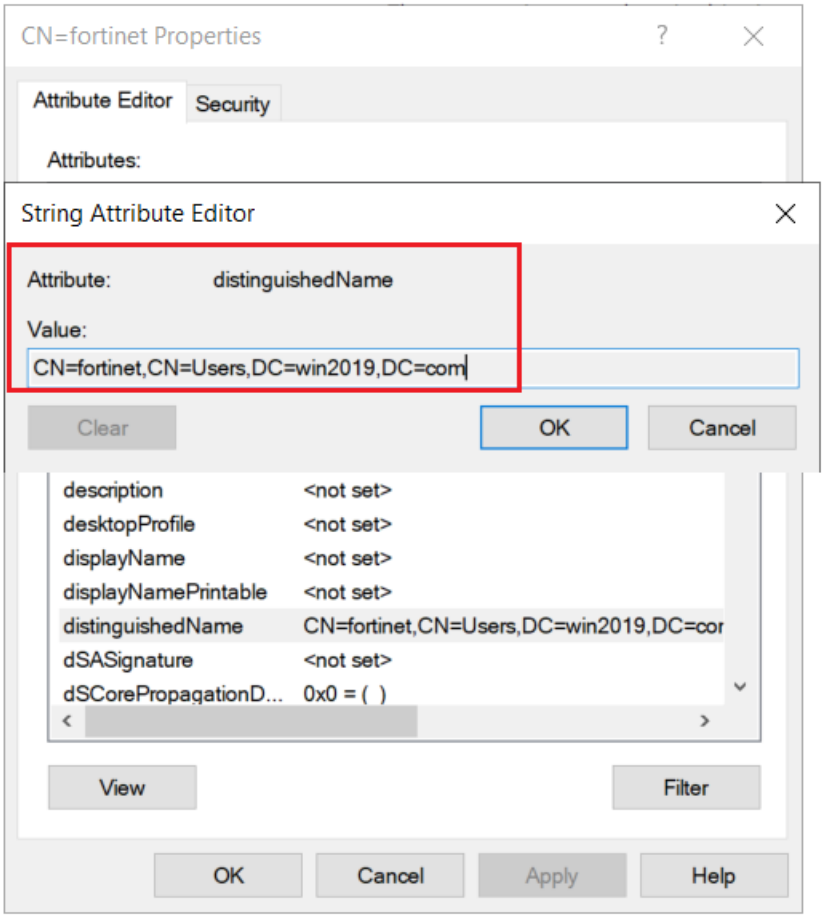
Using Windows Active Directory

Set the LDAP group on the LDAP server for when the **Group Type** is **Windows-AD**.

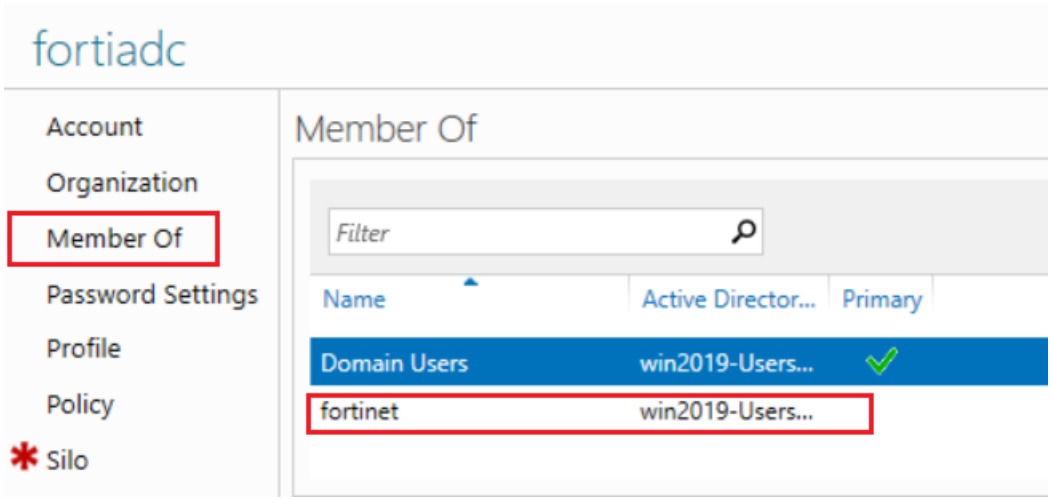
1. Open the Windows Active Directory Administrator Center and create a new group.
   For example, add a new group named "fortinet".



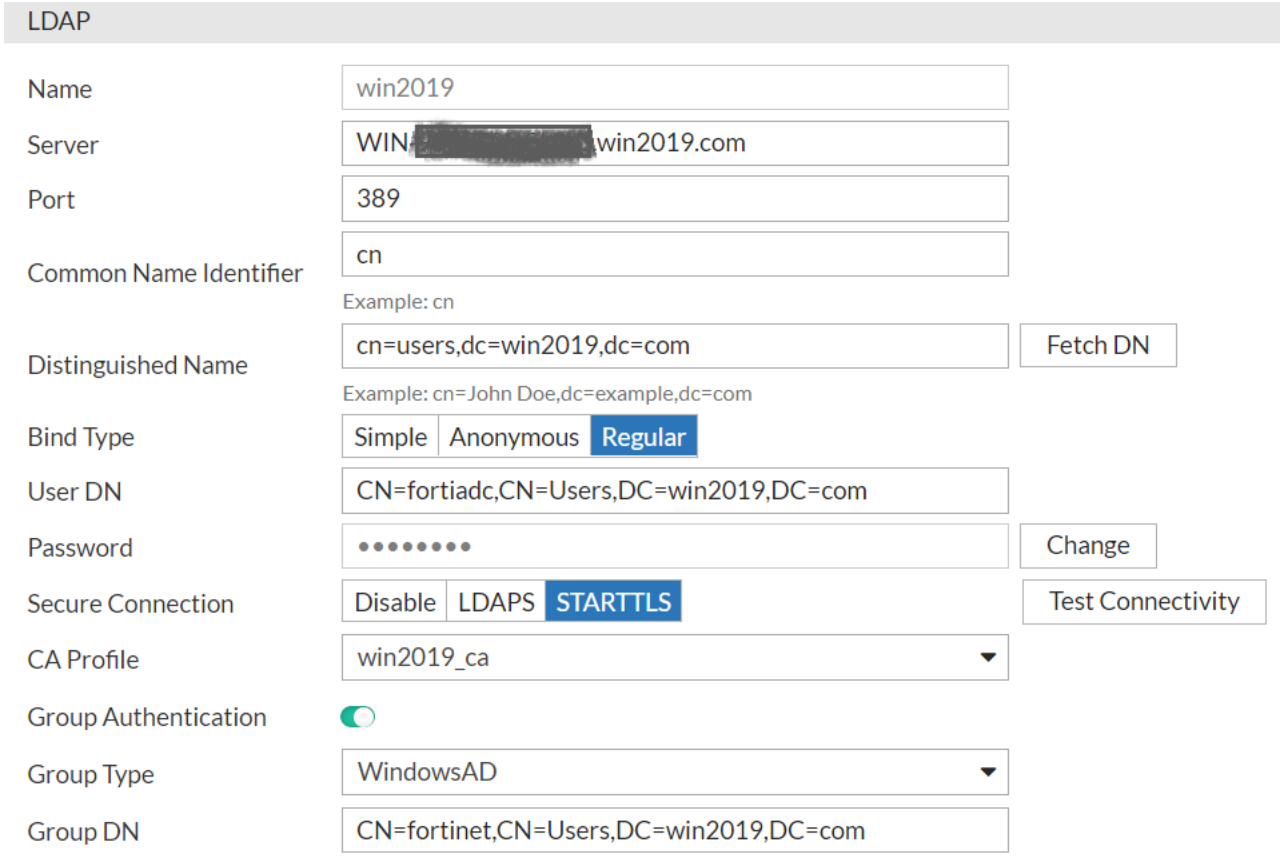
2. Find the Group Distinguished Name. This value is used for the **Group DN** field of the FortiADC LDAP server configuration.



3. Add the user to the group. Ensure you **do not** mark this group as the primary group for login user.
For example, add the User "fortiadc" to the Group "fortinet".



4. In FortiADC, configure the LDAP Group DN settings using the Distinguished Name value recorded from step 2.



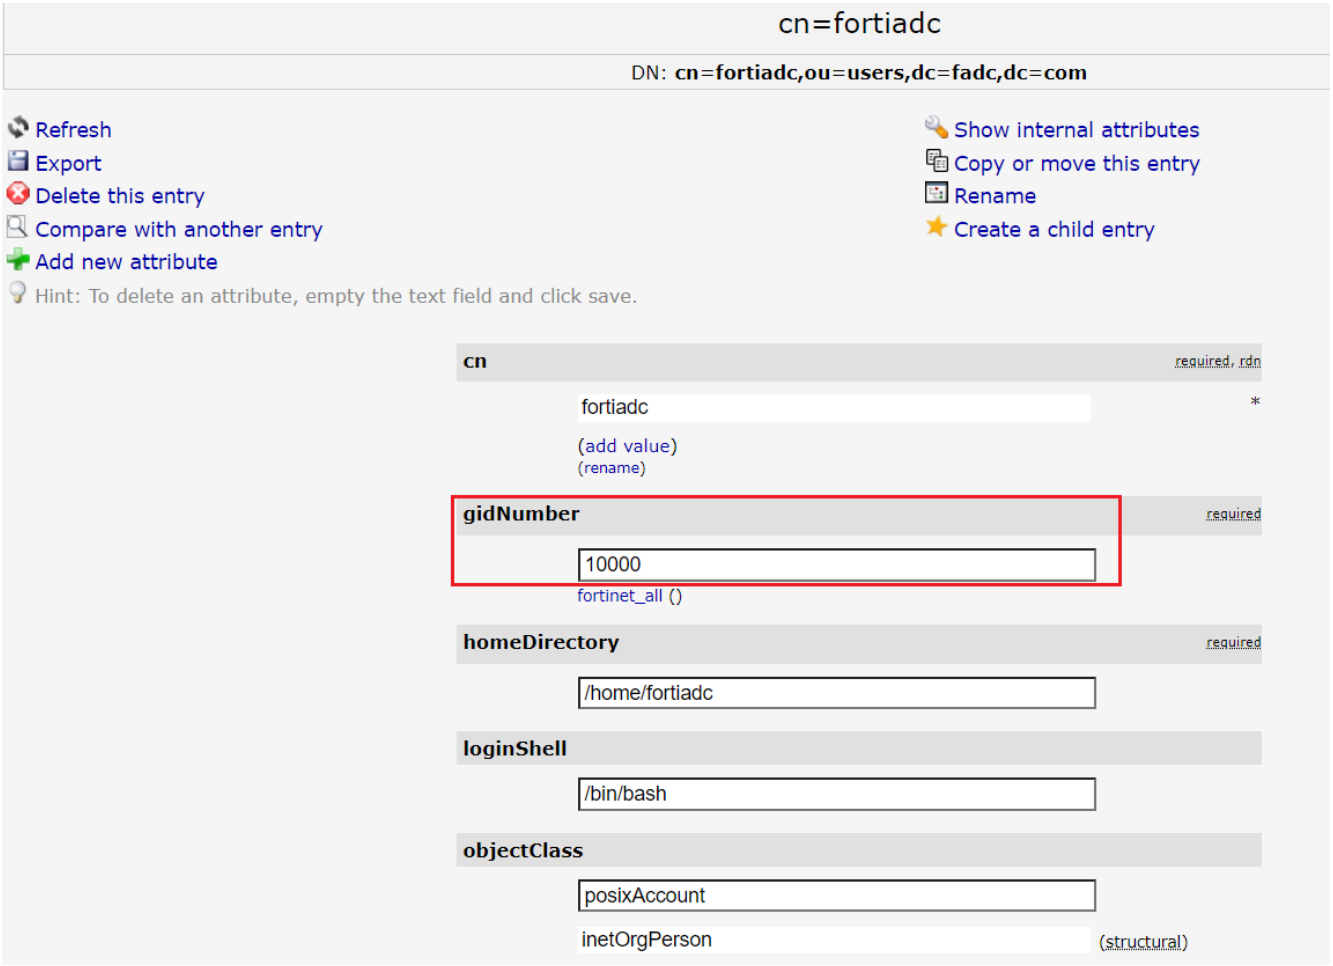
FortiADC will check if `{Common Name Identifier}={login admin name}` is the member of the group you specified.
For example, when logging into FortiADC with the admin name "fortiadc", FortiADC will check if `cn=fortiadc` is the member of the group `cn=fortinet,cn=users,dc=win2019,dc=com` with the search base dn `cn=users, dc=win2019, dc=com`. If the entry exists, FortiADC gets the DN of the entry `cn=fortiadc, cn=users, dc=vm, dc=fadc` and binds this entry and its password to Windows AD.

Using OpenLDAP

Set the LDAP group on the LDAP server for when the **Group Type** is **OpenLDAP**. There are two methods to adding a user in a group.

Method 1:

Create a user with the attribute `gidNumber` which points to the group.
Specify the GID number in the **Group DN** field in the FortiADC LDAP server configuration. For example: `10000`.



Method 2:

Create a group and add the user as the member. Prior to doing this, the `memberof` overlay must be enabled.

The following is an example of `OpenLDAP(sladp)` with MDB database installed on Ubuntu. You can reference the steps below using parameters applicable to your environment.

1. Enable the `memberof` module with the following command:

```
ldapmodify -Y EXTERNAL -H ldapi:/// <<EOL


dn: cn=module{0},cn=config
```

```
add: olcModuleLoad
olcModuleLoad: memberof
EOL
```

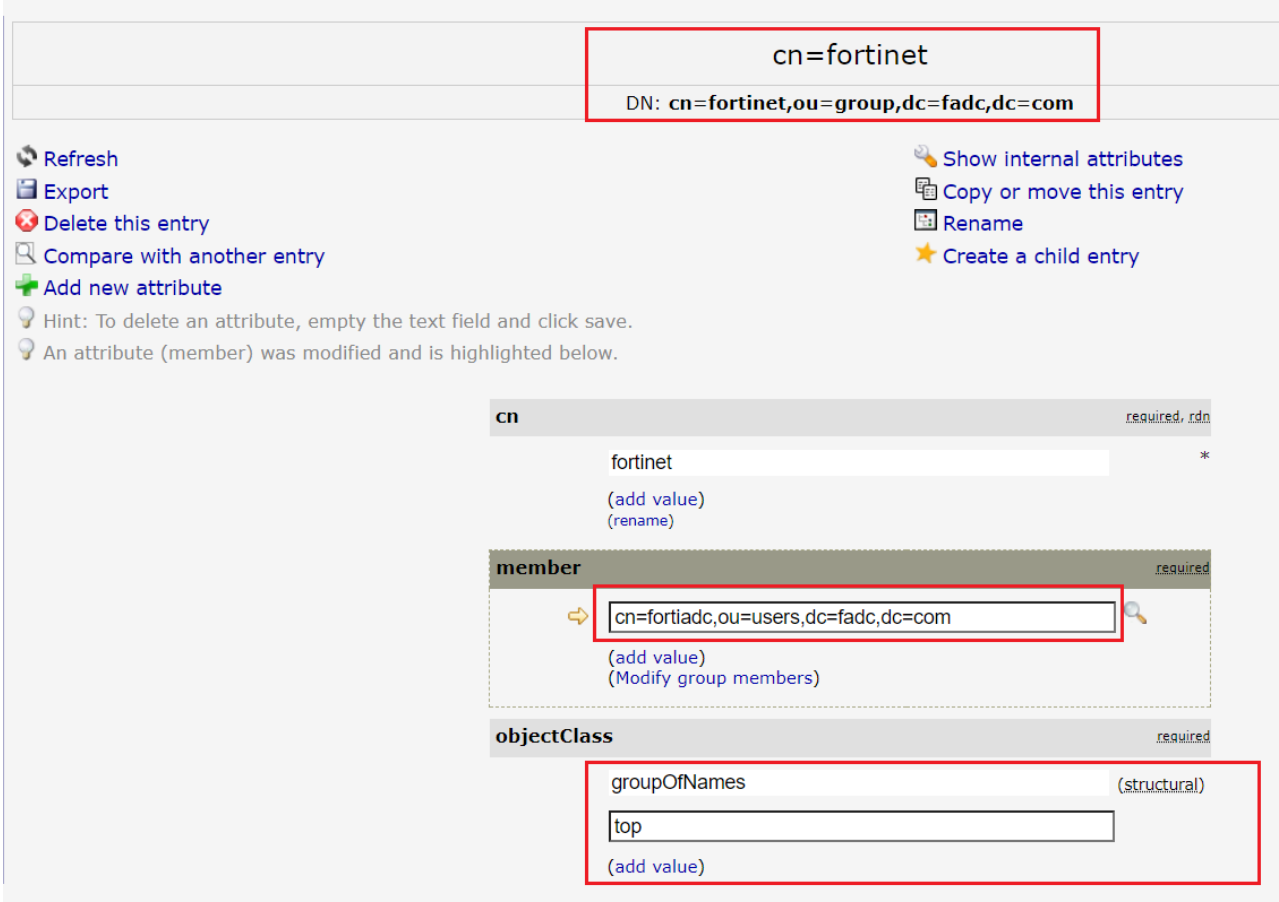
2. Configure OpenLDAP to use the `memberof` module with the following command:

```
ldapadd -Y EXTERNAL -H ldapi:/// <<EOL
dn: olcOverlay={0}memberof,olcDatabase={1}mdb,cn=config
objectClass: olcConfig
objectClass: olcOverlayConfig
olcOverlay: memberof
EOL
```

3. Restart `sladp` and check if `memberof` module is loaded.

```
service slapd restart
slapcat -n 0 | grep olcModuleLoad
olcModuleLoad: {0}back_mdb
olcModuleLoad: {0}memberof
```

4. Create the group with `objectClass` "`groupOfNames`" and add the user member in it.
Specify the DN of the group in the Group DN field in the FortiADC LDAP server configuration.
For example: `cn=fortinet,ou=group,dc=fadc,dc=com`



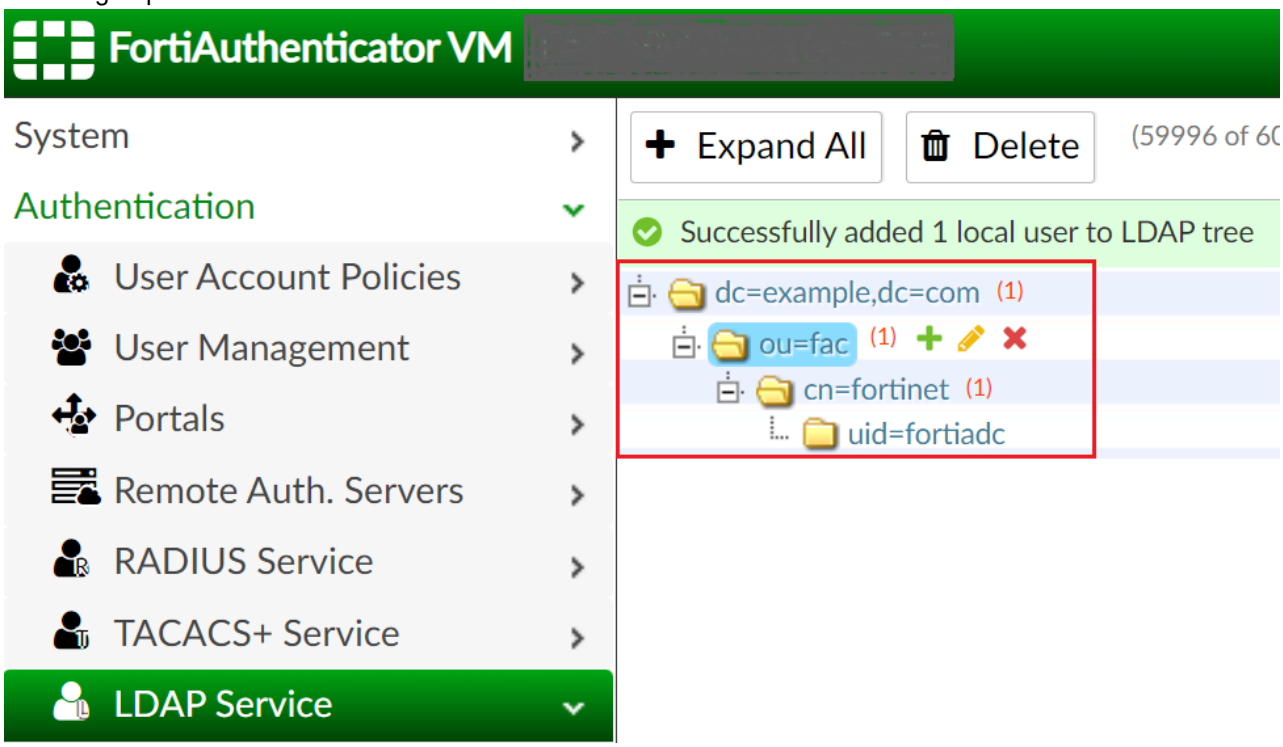
Using FortiAuthenticator

Set the LDAP group on the LDAP server for when the **Group Type** is **FortiAuthenticator**.

1. In FortiAuthenticator, create a group and add a user to that group.
   In the example below, the group "fortinet" is created and the user "fortiadc" is a member of this group.



2. Add the group and the user to the LDAP tree.



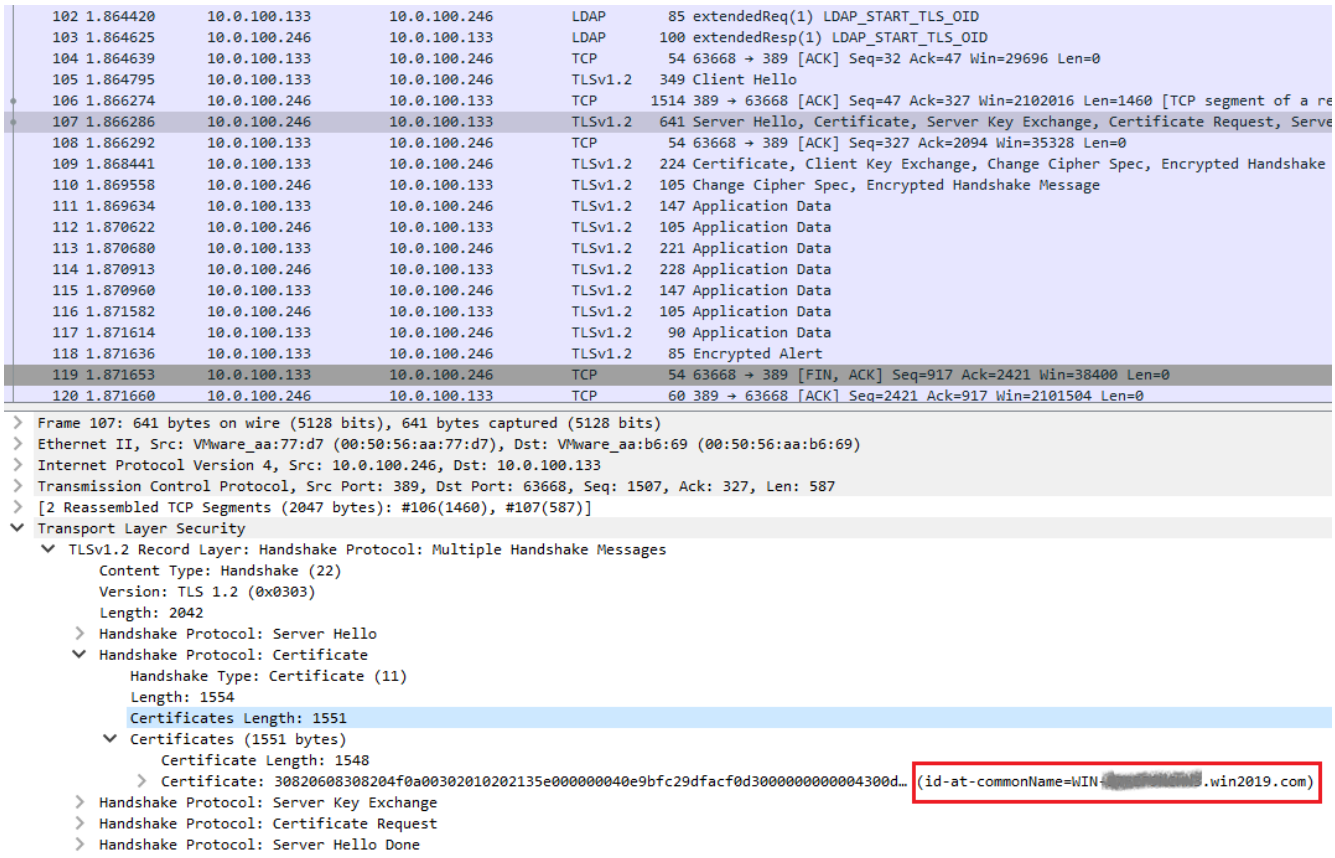
3. Specify the Group DN in the FortiADC LDAP server configuration.
   In the example, the value is `cn=fortinet,ou=fac,dc=example,dc=com`.

FAQs when using an LDAP authentication server

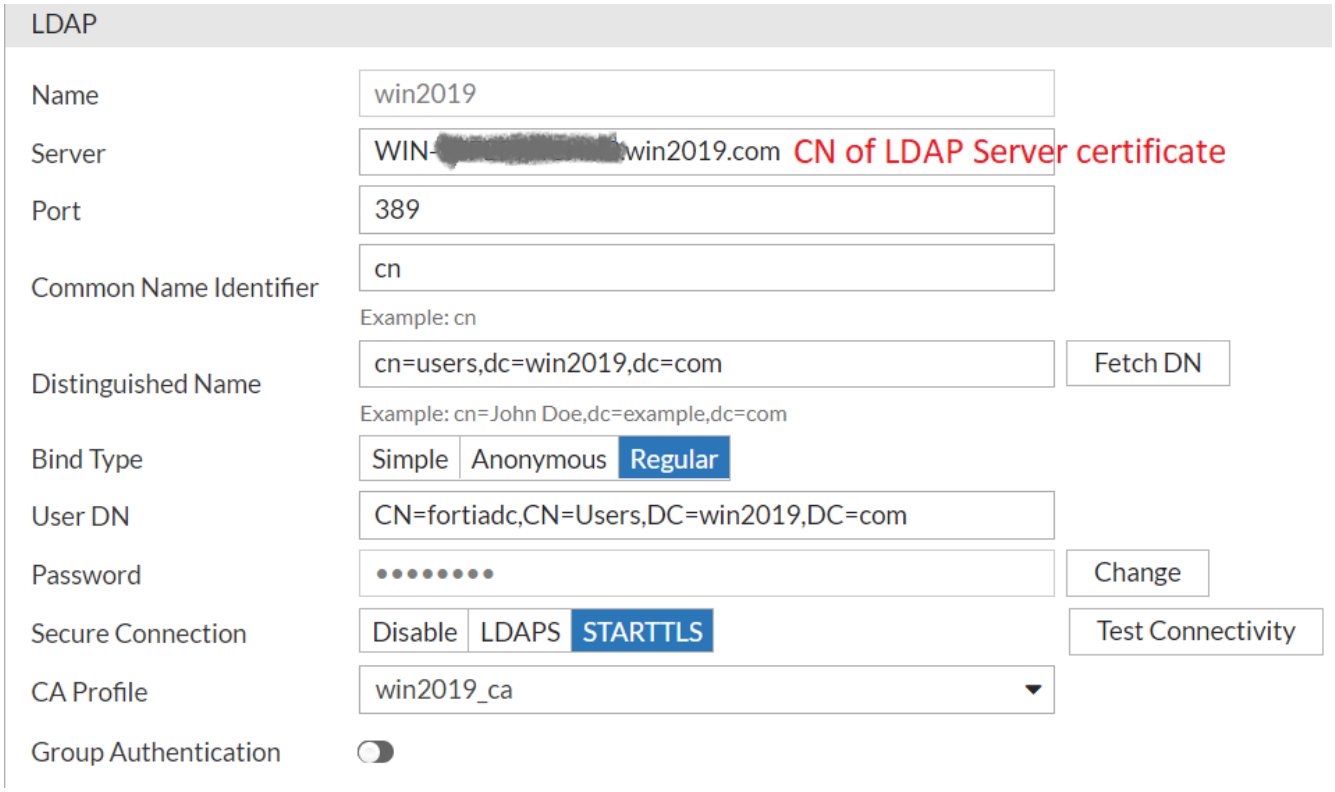
Why does LDAPS or StartTLS not work with Windows AD when a CA profile is selected?

When a CA profile is selected, the CN of the LDAP server certificate must be the same value as the Server field in the FortiADC LDAP configuration. Below is an example configuration for users using Windows AD with StartTLS.

The CN of the Windows AD certificate:



LDAP Configuration on FortiADC:



How do I debug "Test Connectivity" or "Fetch DN" fails when using Windows AD as the LDAP server?

You can install the LDAP Admin tool on the Windows server to verify whether the configuration on Windows AD is correct.

Using 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 903.
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 FQDN 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.

Settings	Guidelines
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 904.
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 an NTLM authentication server

You can use a NTLM authentication server to authenticate user login to destination server.

Before you begin:

- You must know the IP address, port, used to access the NTLM server.
- You must have Read-Write permission for User settings.

Basic steps:

1. Configure a connection to an NTLM server that can authenticate user login.
2. Select the NTLM server configuration when you add users or user groups.

To create a NTLM server configuration:

1. Go to **User Authentication > Remote Server**.
2. Select the **NTLM Server** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described below.
5. Save the configuration.

Settings	Guidelines
Name	Specify a unique name for the NTLM server configuration. Valid characters are A-Z, a-z, 0- 9, _, and -. No spaces allowed. After you save the configuration, you cannot edit the name.
Server	IP address of the remote NTLM server.
Port	The listening port of the NTLM server. The commonly used port for an NTLM server is 445.

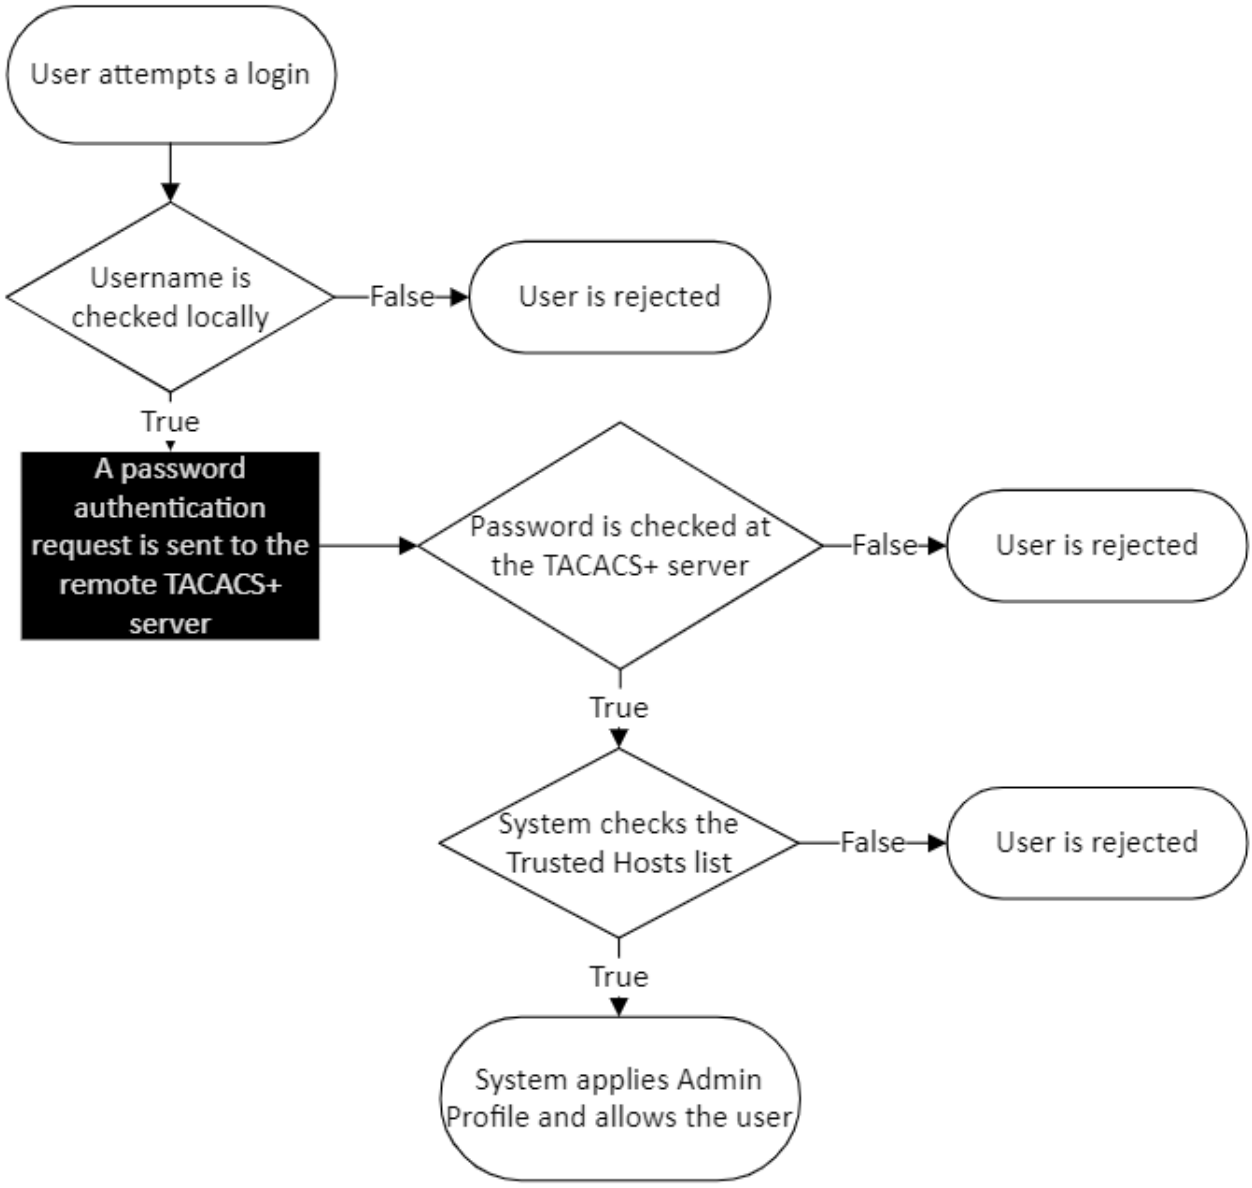
After configuring an NTLM server, configure a user group and add a member of NTLM type. This makes it possible for related authentication policy and virtual server to work under NTLM authentication.

Note: For user groups with “Client Authentication Method” set to “NTLM”, only allow use of an NTLM server as member; for “Client Authentication Method” set to “HTML form” and “HTTP”, use “NTLM server” is also allowed. Only NTLM version 1 is supported.

Using a TACACS+ authentication server

Terminal Access Controller Access-Control System Plus (TACACS+) is a remote authentication protocol that provides access control for routers, network access servers, and other network computing devices through one or more centralized servers. TACACS+ allows FortiADC to accept a user name and password and send a query to a TACACS authentication server. The server host determines whether to accept or deny the request and sends a response back that allows or denies access to the FortiADC user. The default TCP port for a TACACS+ server is 49.

Once TACACS+ is enabled, a series of checks is performed locally and at the TACACS+ server level. The diagram below illustrates the TACACS+ authentication flow.



To use a TACACS+ server to authenticate administrators, the server must be configured before configuring the administrator accounts that will use it.

Basic steps:

1. Configure a connection to a TACACS+ server that can authenticate administrator or user logins.
2. Select the TACACS+ 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 TACACS+ server.
- You must have Read-Write permission for System settings.

To configure a TACACS+ server:

1. Go to **User Authentication > Remote Server**.
2. Click the **TACACS+ Server** tab.
3. Click **Create New** to display the configuration editor.

4. Configure the following settings:

Setting	Description
Name	Specify a unique name for the TACACS+ 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.
Authentication Protocol	Select the authentication protocol used for the TACACS+ server: • Auto — FortiADC tries all authentication protocols in order: MS-CHAP → CHAP → PAP → ASCII. • MS-CHAP — Microsoft version of CHAP (Challenge Handshake Authentication Protocol). • CHAP — Challenge Handshake Authentication Protocol (defined in RFC 1994). • PAP — Password Authentication Protocol. • ASCII — American Standard Code for Information Interchange. **Auto** is the default option.
Timeout	Specify the amount of time that FortiADC must wait for responses from the remote TACACS+ server before it times out the connection. Valid values are from 5 to 60 seconds. The default is 5 seconds.
Shared Secret	Shared secret string used when connecting to the TACACS+ server.
Server	Enter the IP address or FQDN of the TACACS+ server.

Test Connectivity	Tests the connectivity of the TACACS+ server.
	To test connectivity, you must have configured a user on the TACACS+ server with the credentials `test/123456` as follows:
	`user = test {`
	`  login = cleartext "123456"`
	`  chap = cleartext "123456"`
	`  member = noaccess`
	`}`

5. Click **Save**.

Authentication Relay

From the **User Authentication > Authentication Relay** sub-menu, you can configure a Kerberos authentication relay policy.

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 pragraph.
- User Group. See Configuring user groups on page 886.
- Authentication Policy. See Authentication Policy on page 883
- Virtual Server. See Configuring virtual servers on page 384

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 910.
4. Click Save when done.

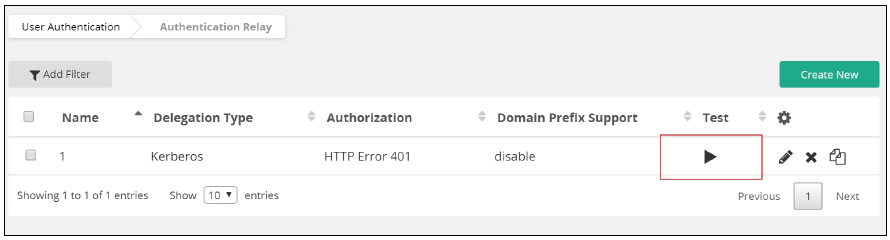
Kerberos authentication configuration

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.

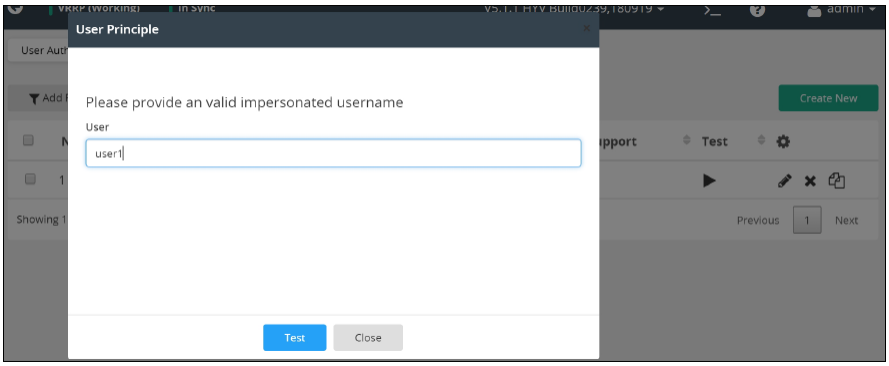
Settings	Guidelines
Delegator Account	Specify the delegator account. Required.
Delegator Password	Specify the delegator password. Required.
Authorization	• HTTP Error 404 • Always
Delegated SPN	Specify the delegated SPN. Required.
Add Default Domain	Disabled by default. When selected, specify the default domain below.
Default Domain	Enter the default domain.

Kerberos Connectivity Test

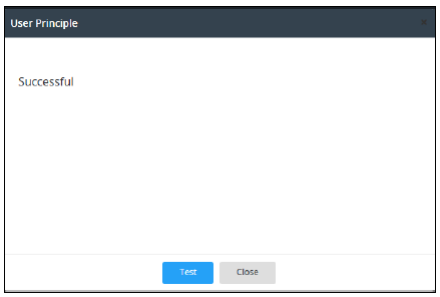
After creating a Kerberos Authentication relay, the **Test** function will appear:



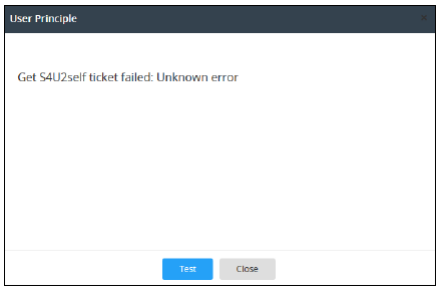
Click on the **Test** symbol, and it will open the **User Principle** dialogue. Set a user name.



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.



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 912 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 912 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 913.
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

The **User Authentication > SAML** sub-menu allows you to perform the following tasks to configure SAML (Security Assertion Markup Language) settings for user authentication:

- Configure an SAML service provider on page 914
- Import IDP Metadata on page 917

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 an SAML service provider

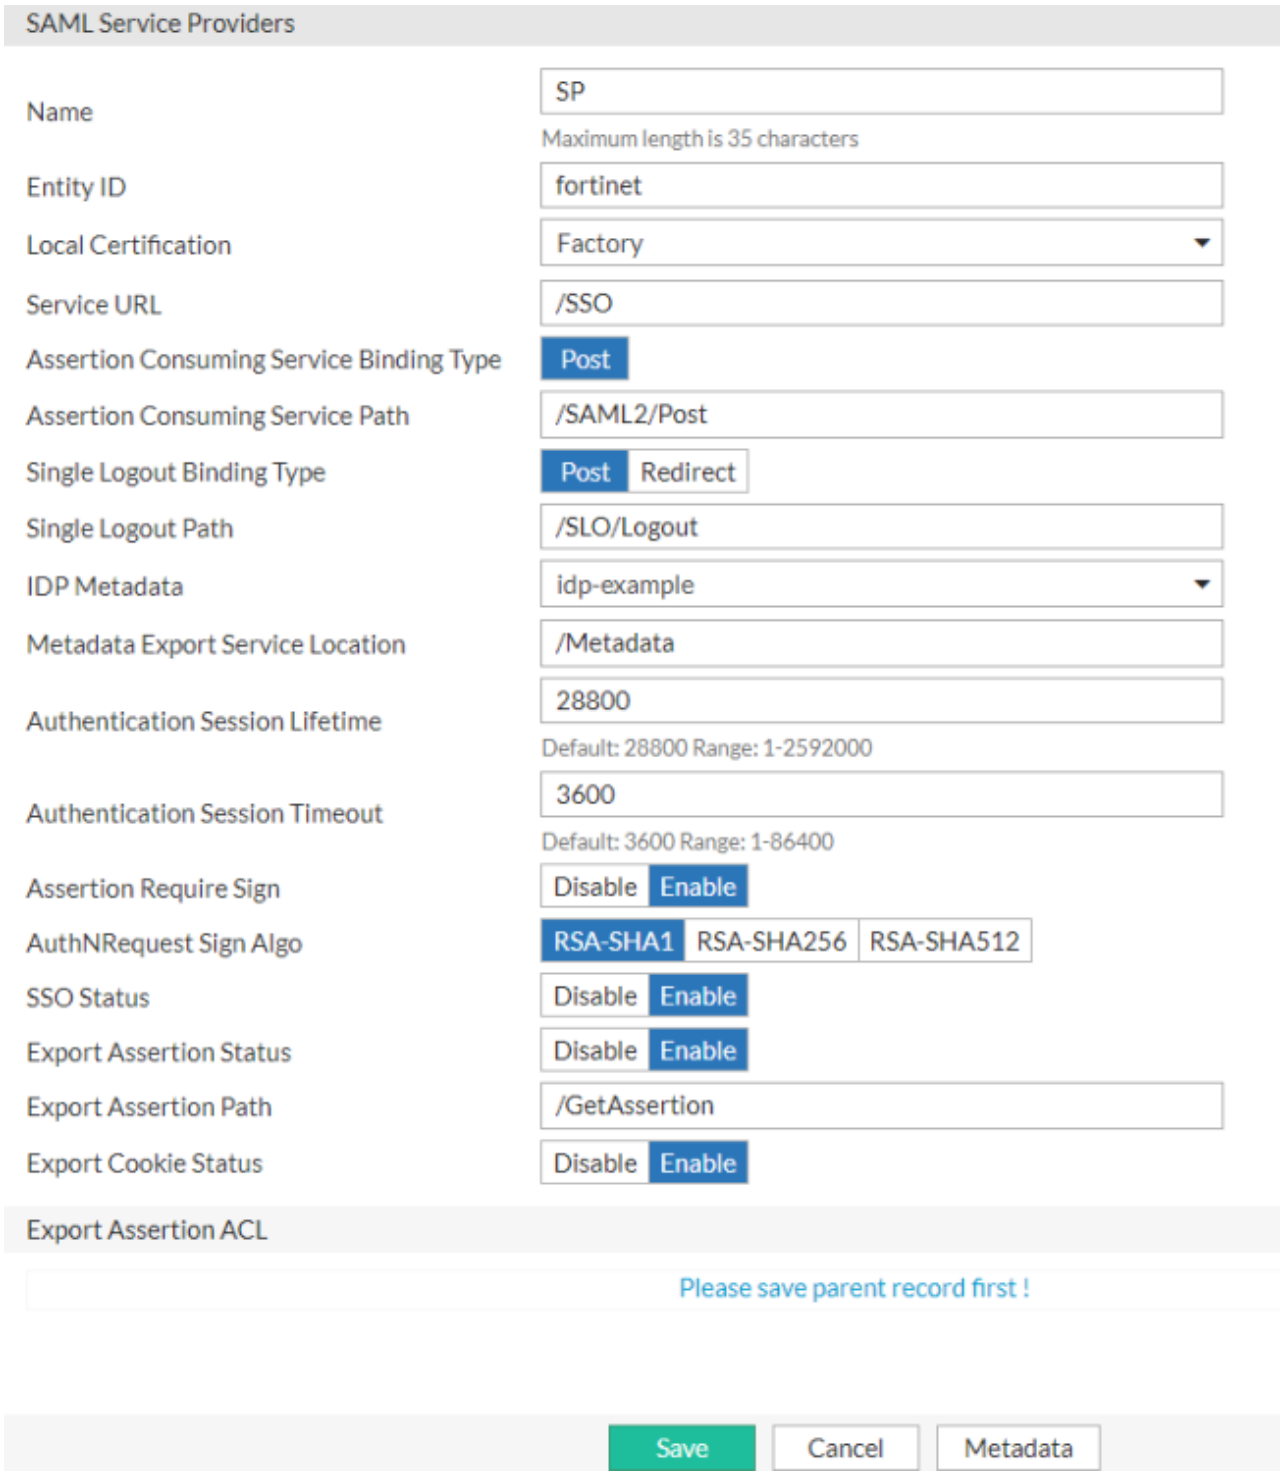
You must configure your SPs in order to use SAML authentication. To configure an SP, you must have the required IDP metadata file imported into FortiADC ahead of time. See Import IDP Metadata on page 917 for more information.

Once you have imported the needed IDP metadata file into FortiADC, you can use the following steps to configure a SAML service provider:

1. Go to **User Authentication > SAML**.
   The configuration page displays the **SAML Service Providers** tab.
2. Click **Create New** to open the SAML Service Providers configuration editor.
3. Configure the following settings.

Parameter	Description
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 a Local Certification from the drop-down. The default is Factory.
Service URL	Specify the SAML service URL. The default value is **/SSO**.
Assertion Consuming Service Binding Type	Specify the Assertion Consuming Service Binding Type. The default value is **Post**.
Assertion Consuming Service Path	Specify the Assertion Consuming Service Path. The default value is **/SAML2/Post**.
Single Logout Binding Type	Select either of the following Single Logout Binding Type: • Post • Redirect The default value is **Post**.
Single Logout Path	Specify the Single Logout Path. The default value is **/SLO/Logout**.

Parameter	Description
IDP Metadata	Select an IDP metadata file from the drop-down. Note: You must have the IDP metadata file imported into FortiADC ahead of time.
Metadata Export Service Location	Specify the Metadata Export Service Location. The default value is **/Metadata**.
Authentication Session Lifetime	Specify the Authentication Session Lifetime in seconds. (Range: 1-2592000, Default: 28800)
Authentication Session Timeout	Specify the Authentication Session Timeout in seconds. (Range: 1-86400, Default: 3600)
Assertion Require Sign	Enable/disable the AuthNRequest algorithm to allow FortiADC to sign the SAML authentication request.
AuthNRequest Sign Algo	Select either of the following AuthNRequest algorithm: • RSA-SHA1 • RSA-SHA256 • RSA-SHA512 The default value is **RSA-SHA1**.
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	Specify the Export Assertion Path. The default value is **/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.



4. Click **Save** when done.
5. Optional: Click **Metadata** to export the SP Metadata.
   a. Specify the SP Root URL.
   b. Click **Export**.



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 of 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. Go to **User Authentication > SAML**.
2. Click the **IDP Metadata** tab.
3. Click **Import** to open the IDP Metadata configuration editor.
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 SAMLservice providers and import the new SP metadata XML file.

AD FS Proxy

The **User Authentication > AD FS Proxy** sub-menu allows you to perform the following tasks to configure Microsoft AD FS (Active Directory Federation Services) settings for user authentication:

- Adding an AD FS Publish on page 918
- Adding an AD FS Proxy on page 919

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 Websites 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.

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. After attaching an AD FS Publish, go to **Server Load Balance > Scripting**.
2. Find the script whose name format is "ADFS_virtual server name_AD FS Publish name." Then clone it.
3. Detach the AD FS Published Service for the virtual server.
4. 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.
5. Attach the content routing created in step 5 to virtual server.
6. Add the cloned script in step 3 into virtual server.
7. Save the configuration.

Adding an AD FS Publish

1. Go to User **Authentication > AD FS Proxy**.
   The configuration page displays the **Publish** tab.
2. Click Create New to display the AD FS Publish configuration editor.
3. Make the desired entries or selections, as described in the table below.
4. 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.

Parameter	Description
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/

Adding an AD FS Proxy

1. Go to **User Authentication > AD FS Proxy**.
2. Click the **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 552.

Parameter	Description
	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.

OAuth Proxy

From the **User Authentication > OAuth Proxy** sub-menu, you can configure OAuth 2.0 authentication.

OAuth (Open Authorization) is an authorization framework that can provide client applications with a secure delegated access to server resources on behalf of a resource owner. OAuth works over HTTPS and authorizes third-party clients such as devices, APIs, servers, and applications with access tokens rather than credentials with the approval of the resource owner. The third party then uses the access tokens to access the protected resources hosted by the resource server. This enables applications to obtain limited access to HTTP services on behalf of a user by delegating the user authentication to the service that hosts the user account, and authorizing the third-party application to access the user account.

Through FortiADC's OAuth 2.0 feature, you can:

- Mandate the authentication to a third party that you trust.
- Enable your web application (RealServer) to access resources that belong to the user and do not belong to the web application (RealServer).
- Implement an alternative to Single SignOn.



FortiADC will only be supporting OAuth 2.0 which is the most widely used form of OAuth. There will be no backwards compatibility between OAuth 1.0 and OAuth 2.0 as their specifications are so different that they cannot be used together.

Deploying OAuth 2.0 authentication

To deploy OAuth 2.0 in FortiADC, you need to first set up the OAuth policy to establish the authorization flow between FortiADC, the token server, and the authorization server. The OAuth policy serves to obtain the authorization code and access token. After the OAuth policy is set up, it is then applied in the authentication policy, in which you will apply to the virtual server to complete the OAuth deployment.

To configure the OAuth policy:

1. Go to **User Authentication > OAuth Proxy**.
2. Click **Create New** to display the configuration editor.

3. Configure the following settings:

Settings	Description
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.
Client ID	The client ID for your application.
Client Secret	The secret used to apply for the access token.
Authentication URL	The URL of the authorization server.
Token URL	The URL of the token server.
Redirect URL	The URL of the redirected server.
Logout URL	The URL will trigger a logout event in which the user will be logged out and FortiADC will delete the cookie. For the next access, the OAuth 2.0 process will need to be conducted again. The value is parsed as a match string prefix. For example, `/abc` matches `http://www.example.com/abcd` and `http://www.example.com/abc/11.html` but not `http://www.example.com/1abcd`.
HTTP Method	The HTTP method used for the OAuth transaction. Select from the following: • POST • GET
Relay Mode	Enable/disable relay mode allows FortiADC to add an Authorization Header to the HTTP request after verifying the token.
Include Granted Scopes	Select from the following: • True • False • None This enables applications to use incremental authorization to request access to additional scopes in context. If you set this parameter's value to **True** and the authorization request is granted, then the new access token will also cover any scopes to which the user previously granted the application access.
Prompt	Select from the following: • Disable — Disable prompts. • None — Do not display any authentication or consent screens. • Consent — Prompt the user for consent. • Select Account — Prompt the user to select an account.
Token Timeout	The amount of time in seconds the token will be valid. (Range: 120-86,400, default = 3600). The client will not be allowed to access the scope after this time has elapsed.

4. Click **Save**.

To configure the authentication policy:

1. Go to **User Authentication > Authentication Policy**.
2. Click **+Create New**.
3. Name the new authentication policy.
4. Configure the following to set the **Member** as the OAuth for the new OAuth policy.
   a. Set the **Type** as **OAuth**.
   b. Specify the pathname.
   c. Select the applicable **OAuth Policy**.
5. Click **Save**.

To apply the authentication policy to the virtual server:

1. Go to **Server Load Balance > Virtual Server**.
2. Click **+Create New** and select **Advanced Mode**. Or double-click an existing virtual server configuration from the list.
3. In the **General** tab, under the **Resources** section, select the applicable **Auth Policy**.
4. Click **Save**.

OAuth 2.0 scopes

Scopes enable your application to only request access to the resources that it needs while also enabling users to control the amount of access that they grant to your application. As part of the OAuth policy configuration, add the scopes to the **Scope List** to identify the resources your application could access on the user's behalf. These will be shown to the user to obtain their consent when they access the resource server. However, there is an inverse relationship between the number of scopes requested and the likelihood of obtaining user consent; the user must consent to all or none of the requests within the scope.

Note: OAuth 2.0 scopes are restricted to each application. Please refer to the applicable guidelines for your application to ensure the scopes are valid.

Log & Report

The **Log & Report** menu contain features and configurations that allow you to view various logs and reports.

The menu is further divided into sub-menus that allow you to configure log and report settings:

- Log Setting on page 937
- Report Setting on page 946

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 > Traffic Log** to display the traffic log.
2. Click **Add Filter** 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 926 to GLB log on page 931 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 security log

The Security Log displays logs related to the following FortiADC security features:

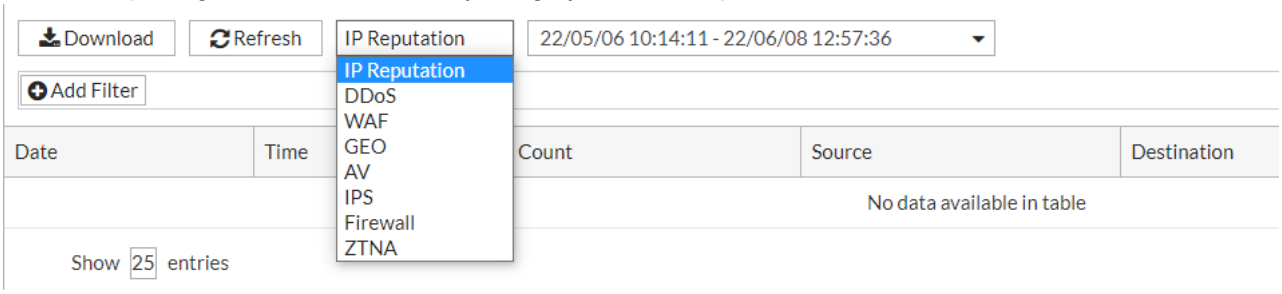
- IP Reputation — Traffic logged by the IP Reputation feature.
- DDoS — Traffic logged by the DoS Protection feature.
- WAF — Traffic logged by the Web Application Firewall feature.
- GEO — Traffic logged by the Geo IP block list feature.
- AV — Traffic logged by the Anti Virus module.
- IPS — Traffic logged by the IPS feature.
- Firewall — Traffic logged by the Firewall module.
- ZTNA - Traffic logged by the ZTNA feature.

Before you begin:

- You must have Read-Write permission for Log & Report settings.
- Have enabled to write security logs on the FortiADC log disk in Log & Report > Log Setting > Local Log.
- Have enabled or disabled related security logs in Log & Report > Log Setting > Local Log.

To view and filter the log:

1. Go to **Log & Report > Security Log**.
2. From the top navigation, select the security category from the drop-down menu.



The log page displays with the log columns and data specific to the security category.

The following lists the log columns in the order in which they appear in each security log. Use the below links to navigate to the security log of your choosing:

- IP Reputation log on page 933
- DDoS log on page 934
- WAF log on page 934
- GEO log on page 935
- AV log on page 935
- IPS log on page 936
- Firewall log on page 936
- ZTNA log on page 936

For additional detail on each log, click the (Detail icon) for any log. For further description of each log message, see the FortiADC Log Reference.

IP Reputation log

Column	Description
Date	Log date.
Time	Log time.
Count	Rule match count.
Source	Source IP address.
Destination	Destination IP address.
Action	Action type that was taken as a result.

Column	Description
(Detail icon)	Click the (Detail icon) for the log details. For further description of each log message, see the FortiADC Log Reference.

DDoS log

Column	Description
Date	Log date.
Time	Log time.
Count	Rule match count.
Source	Source IP address.
Destination	Destination IP address.
Action	Action type that was taken as a result.
(Detail icon)	Click the (Detail icon) for the log details. For further description of each log message, see the FortiADC Log Reference.

WAF log

Column	Description
Date	Log date.
Time	Log time.
WAF Subcategory	Web Application Firewall subcategory.
Severity	Security level.
Source	Source IP address.
Destination	Destination IP address.
Action	Action type that was taken as a result.
(Detail icon)	Click the (Detail icon) for the log details. For further description of each log message, see the FortiADC Log Reference. The following actions may be performed directly from the WAF log details: • Add Exception — You can add WAF Exceptions directly from the WAF log. This option appears only for WAF subcategories that support WAF Exceptions. For details, see Configuring WAF Exception objects on page 673. • Disable Signature — You can disable WAF signature profiles directly from the WAF log. This option appears only for Attacks Signature WAF subcategories. Disable Signature can only be successful if the WAF

Column	Description
	signature profile exists, otherwise the disable will fail with the error message "Entry not found". • View Signature — You can view the WAF signature status and information directly from the WAF log. This option appears only for Attacks Signature WAF subcategories.

GEO log

Column	Description
Date	Log date.
Time	Log time.
Count	Rule match count.
Severity	Security level.
Source	Source IP address.
Destination	Destination IP address.
Action	Action type that was taken as a result.
 (Detail icon)	Click the  (Detail icon) for the log details. For further description of each log message, see the FortiADC Log Reference.

AV log

Column	Description
Date	Log date.
Time	Log time.
Source	Source IP address.
Destination	Destination IP address.
Service	Service type.
Severity	Security level.
Virus Category	Virus category.
Action	Action type that was taken as a result.
 (Detail icon)	Click the  (Detail icon) for the log details. For further description of each log message, see the FortiADC Log Reference.

IPS log

Column	Description
Date	Log date.
Time	Log time.
Source	Source IP address.
Destination	Destination IP address.
Service	Service type.
Severity	Security level.
Rule Name	Security rule name
Action	Action type that was taken as a result.
 (Detail icon)	Click the  (Detail icon) for the log details. For further description of each log message, see the FortiADC Log Reference.

Firewall log

Column	Description
Date	Log date.
Time	Log time.
Log Level	Log level.
Policy	Firewall policy.
Message	Security rule name, category, subcategory, and description of the attack.
 (Detail icon)	Click the  (Detail icon) for the log details. For further description of each log message, see the FortiADC Log Reference.

ZTNA log

Column	Description
Date	Log date.
Time	Log time.
Severity	Security level.
Source	Source IP address.
Destination	Destination IP address.
Action	Action type that was taken as a result.

Column	Description
(Detail icon)	Click the (Detail icon) for the log details. For further description of each log message, see the FortiADC Log Reference.

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.

Log Setting

The **Log & Report > Log Setting** sub-menu allows you to perform the following tasks to configure settings for various logs:

- Configuring local log settings on page 937
- Configuring syslog settings on page 939
- Configuring OFTP settings for FortiAnalyzer logs on page 943
- Configuring fast stats log settings on page 945

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 938.
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: • 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.

Settings	Guidelines
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. • IPS—IPS logs • FW—Firewall logs • ZTNA—ZTNA 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.

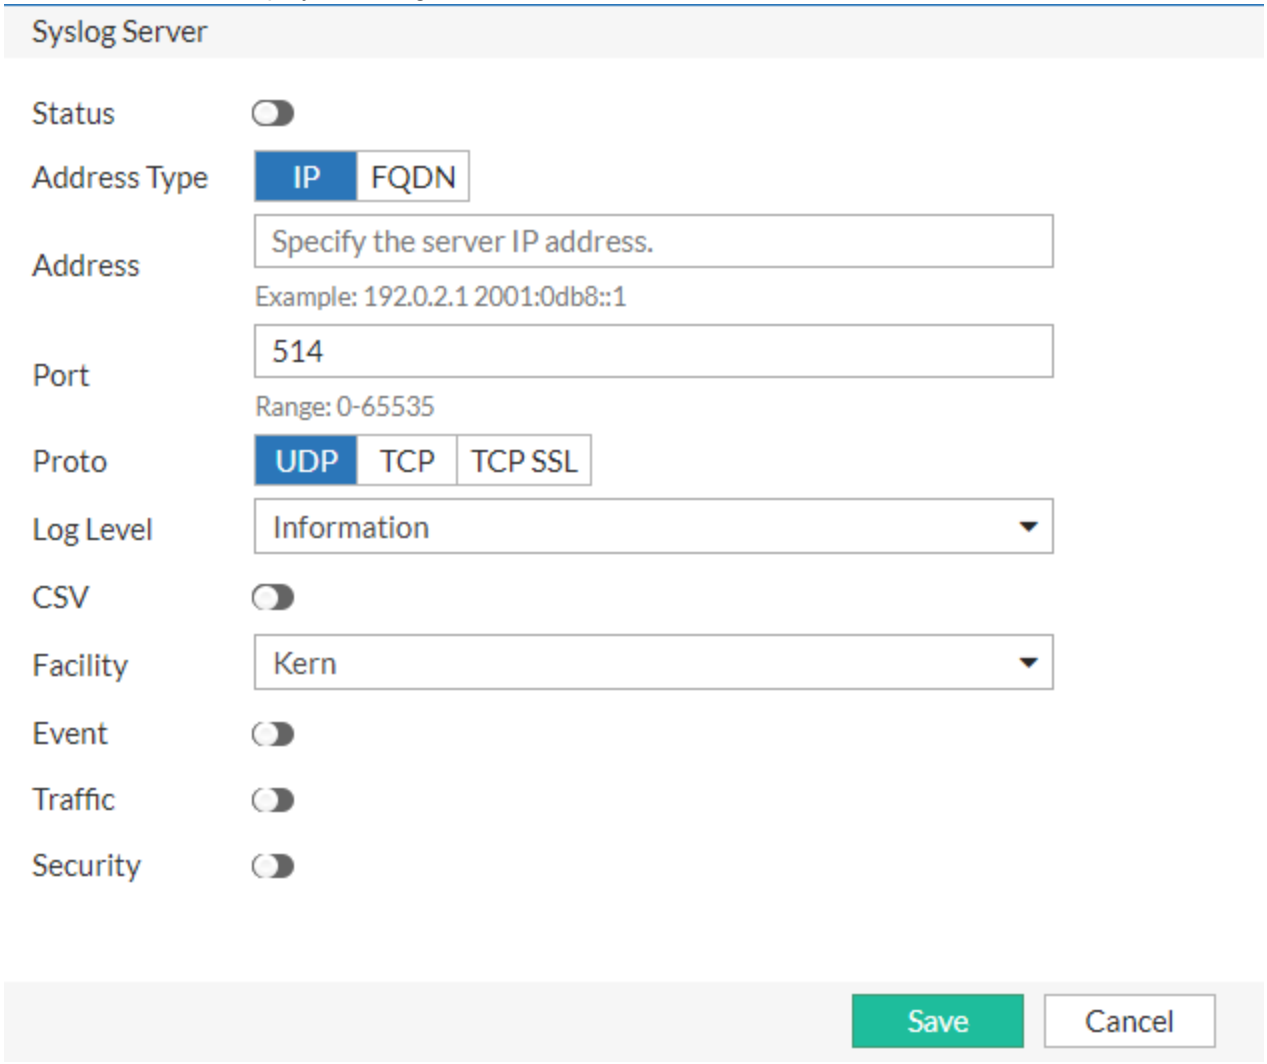
Before you begin:

- 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. If the VDOM is enabled, enable/disable **Override** to determine which server list to use.
   Enable **Override** to allow the syslog to use the VDOM FortiAnalyzer server list. Otherwise, disable **Override** to use the Global syslog server list.

4. Click **Create New** to display the configuration editor.



5. Configure the following settings:

Setting	Description
Status	Enable/disable the configuration.
Address Type	Select the Address Type of the syslog server: • IP • FQDN
Address	The **Address** option is available if the **Address Type** is **IP**. Specify the IP address of the syslog server. **Note**: In IPv6, certain reserved or designated addresses cannot function as globally unique addresses. Users can configure these invalid IPs in remote settings without errors, but this leads to failed IPv6 communication. Below are examples of unusable IPv6 address types.

Setting	Description
	• **Unspecified Address (`::/128`)** **Usage**: Represents an unspecified address (analogous to `0.0.0.0` in IPv4). Typically used during the initialization phase before an IP address is assigned. Not valid for routing or general use. • **Loopback Address (`::1/128`)** **Usage**: Used for intra-device communication (equivalent to `127.0.0.1` in IPv4). Packets addressed to this cannot leave the local device. • **Link-Local Addresses (`fe80::/10`)** **Range**: `fe80::` to `febf:ffff:ffff:ffff:ffff:ffff:ffff:ffff` **Usage**: Used for communication within a single local link. Not routable beyond the local subnet and typically used for auto-configuration and neighbor discovery. • **Multicast Addresses (`ff00::/8`)** **Range**: ff00:: to ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff **Usage**: Facilitates one-to-many communication. Cannot be used as a unicast address. Replaces IPv4 broadcast functionality. • **Unique Local Addresses (fc00::/7)** **Range**: `fc00::` to `fdff:ffff:ffff:ffff:ffff:ffff:ffff:ffff` **Usage**: Analogous to private IPv4 addresses (`10.0.0.0/8`, `192.168.0.0/16`). Used within local networks and not routable globally. • **Documentation Prefix (`2001:db8::/32`)** **Range**: `2001:db8::` to `2001:db8:ffff:ffff:ffff:ffff:ffff:ffff` **Usage**: Reserved for documentation and example purposes. Not valid for real-world network deployment. • **Reserved Address Space (Various Prefixes)** **Usage**: Reserved for future use or specific protocols. These addresses should not be used unless explicitly defined. For example, `2002::/16` is used for 6to4 tunneling. • **Global Unicast Addresses (`2000::/3`)** **Range**: `2000::` to `3fff:ffff:ffff:ffff:ffff:ffff:ffff:ffff` **Usage**: Globally routable, but some portions may be reserved or not allocated. Ensure addresses used within this range are valid.
FQDN	The **FQDN** option is available if the **Address Type** is **FQDN**. Specify the FQDN of the syslog server.
Port	Specify the port that FortiADC uses to communicate with the log server. This is the listening port number of the syslog server. Usually this is UDP port 514.
Proto	Select the protocol used for log transfer from the following: • UDP • TCP • TCP SSL
TCP Framing	If **Proto** is **TCP** or **TCP SSL**, the **TCP Framing** options appear.

Setting	Description
	Select one of the following options: • Traditional • Octet Counted
Log Level	Select the lowest severity to log from the following options: • 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. The exported logs will include the selected severity level and above. 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 severity level Alert and Emergency.
CSV	Enable to export the logs as a CSV file.
Facility	Select the source facility of the logs. We only support the local use facilities which are not reserved and are available for general use.
Event	Enable/disable logging for events.
Event Category	If **Event** is enabled, the **Event Category** options appear. Select one or more of the following event categories to include in the event logs export: • 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	Enable/disable logging for traffic processed by the load-balancing modules.
Traffic Category	If **Traffic** is enabled, the **Traffic Category** options appear. Select one or more of the following traffic categories to include in the traffic logs export: • SLB — Server Load Balancing traffic logs related to sessions and throughput.

Setting	Description
	• GLB — Global Load Balancing traffic logs related to DNS requests. • LLB — Link Load Balancing traffic logs related to session and throughput.
Security	Enable/disable logging for traffic processed by the security modules.
Security Category	If **Security** is enabled, the **Security Category** options appear. Select one or more of the following security categories to include in the security logs export: • DDoS — DoS protection logs. • IP Reputation — IP Reputation logs. • WAF — WAF logs. • GEO — Geo IP blocking logs. • AV — AV logs. • IPS — IPS logs. • FW — Firewall logs. • ZTNA — ZTNA logs.

6. Click **Save**.

Configuring OFTP settings for FortiAnalyzer logs

The Optimized Fabric Transfer Protocol (OFTP) is used when information is synchronized between FortiAnalyzer and FortiADC, as well as for other Fortinet products. Remote logging and archiving can be configured on the FortiADC to send logs to a FortiAnalyzer unit.

OFTP listens on port TCP/514.

You can configure the OFTP settings from **Log & Report > Log Setting**, or directly in your FortiAnalyzer connector configuration (for details, see FortiAnalyzer Connector on page 78).

Requirements:

- Read-Write permission for Log & Report settings.
- The FortiAnalyzer service is required to be exposed on External IP.



FortiADC supports integration with FortiAnalyzer versions 7.0.2 or later. As earlier versions of FortiAnalyzer is not optimally compatible with FortiADC, unexpected behavior may occur.

To configure OFTP settings for FortiAnalyzer logs:

1. Go to **Log & Report > Log Setting**.
2. Click the **FortiAnalyzer** tab.
3. If the VDOM is enabled, enable/disable **Override** to determine which server list to use.
   Enable **Override** to use the VDOM FortiAnalyzer server list. Otherwise, disable **Override** to use the Global FortiAnalyzer server list.

4. Click **Create New** to display the configuration editor.
5. Configure the following settings:

Setting	Description
Status	Enable/disable the Fabric Connector object.
Address	Specify the IP address of the FortiAnalyzer Log server.
Log Level	Select the lowest severity to log from the following options: • 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. The exported logs will include the selected severity level and above. For example, if you select **Error**, the system collects logs with severity level Error, Critical, Alert, and Emergency. If you select **Alert**, the system collects logs with severity level Alert and Emergency.
Event	Enable/disable logging for events.
Event Category	If **Event** is enabled, the **Event Category** options appear. Select one or more of the following event categories to include in the event logs export: • 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	Enable/disable logging for traffic processed by the load-balancing modules.
Traffic Category	If **Traffic** is enabled, the **Traffic Category** options appear. Select one or more of the following traffic categories to include in the traffic logs export: • 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.

Setting	Description
Security	Enable/disable logging for traffic processed by the security modules.
Security Category	If **Security** is enabled, the **Security Category** options appear. Select one or more of the following security categories to include in the security logs export: • DDoS — DoS protection logs. • IP Reputation — IP Reputation logs. • WAF — WAF logs. • GEO — Geo IP blocking logs. • AV — AV logs. • IPS — IPS logs. • FW — Firewall logs.

6. Optionally, click **Test Connectivity** after entering the **Address** to check the FortiAnalyzer OFTP connectivity. The **Connection Status** appears showing the OFTP connection status.
There are three possible OFTP connection statuses:

Icon	OFTP Status	Description
	Connected	The FortiADC has successfully connected to FortiAnalyzer and is authorized by FortiAnalyzer as a Fabric Device. FortiADC can now send log data to FortiAnalyzer.
	Disconnected	The FortiADC cannot connect to FortiAnalyzer. Ensure there are no network connectivity issues.
	Need authorization	The FortiADC has successfully connected to FortiAnalyzer but is not authorized by FortiAnalyzer as a Fabric Device. This status may indicate the authorization is either denied or pending. If pending authorization, the status will change to **Connected** once authorization is successful on the FortiAnalyzer server.

If the status is not Connected, edit the FortiAnalyzer connector accordingly to troubleshoot the connection issue.

7. Click **Save**.

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:

1. Go to **Log & Report > Log Setting**.
2. Click the **Fast Stats** tab.

3. Complete the configuration as described in Fast stats log configuration on page 946.
4. 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.
Attack	Enable/disable fast statistics for attack logs. Disabled by default.
Attack Category	Enable/disable fast statistics for attack categories.
Security	Select to enable logging for traffic processed by the security modules. Disabled by default.

Report Setting

The **Log & Report > Report Setting** sub-menu allows you to perform the following tasks to configure settings for various reports:

- Configuring reports on page 946
- Configuring report emails on page 947
- Configuring report queries on page 948
- Configuring fast reports on page 950

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 Setting**.
   The configuration page displays the **Report** tab.
2. Click **Create New** to display the configuration editor.
3. Complete the configuration as described in Report configuration on page 947.
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 emails

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. Go to **Log & Report > Report Setting**.
2. Click the **Report Email** tab.
3. Click the **Create New** to display the configuration editor.
4. Complete the configuration as described in Report mail configuration on page 948.
5. 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 report queries

The predefined list of queries covers the most common administrator and stakeholder interests. It includes some of 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 Setting**.
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 949.
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	

Settings	Guidelines
DNS Sort Type	Only count is applicable.
DNS Subtype	• Top_Policy • 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.

There are two ways to configure a fast report, through **Log & Report** or **FortiView**.

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 fast reports through Log & Report:

1. Go to **Log & Report > Report Setting**.
2. Click the **Fast Report** tab.
3. Click **Create New** to display the configuration editor.
4. Complete the configuration as described in Fast report configuration on page 951.
5. Save the configuration.

To configure fast reports through FortiView:

1. Go to **FortiView > Data Analytics**.
2. Click the **Add Widget** button in the far right. The **Fast Report** dialogue will display. It will be the same as in the Log & Report route.
3. Complete the configuration as described in Fast report configuration on page 951.
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 • Attack
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
Attack 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 • Top platform name by dest for AV • Top platform name by destcountry for AV • Top platform name by src for AV • Top platform name by srccountry for AV • Top platform name by vs for AV • Top reference by dest for AV • Top reference by destcountry for AV • Top reference by src for AV • Top reference by srccountry for AV • Top reference by vs for AV • Top src for IPS • Top srccountry for IPS
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

AI Threat Analytics

Through the FortiADC integration with Fortinet AI Threat Analytics, you can forward FortiADC security logs to FortiWeb Cloud where the AI-based Threat Analytics engine identifies unknown attack patterns by parsing through all FortiADC security logs and then aggregating similar or related security logs into into meaningful security incidents. This allows you to use these identified attack patterns to protect your application against the identified threats.



AI Threat Analytics is not supported in AWS, AWS On-Demand, GCP On-Demand, and Azure On-Demand platforms.

Prerequisites for using AI Threat Analytics for FortiADC security logs:

- You must have a valid AI Threat Analytics service license.
- You must have the AI Threat Analytics service enabled in FortiADC.

14-day Evaluation license

A 14-day Evaluation license is offered to customers who would want to evaluate the AI Threat Analytics service. This 14-day Evaluation license can only be used once. To activate the 14-day Evaluation license, enable the **Threat Analytics** connector from **Security Fabric > Fabric Connectors**. During this 14-day trial period, you can disable and re-enable AI Threat Analytics anytime. The 14-day trial period starts from the first time Threat Analytics is enabled.

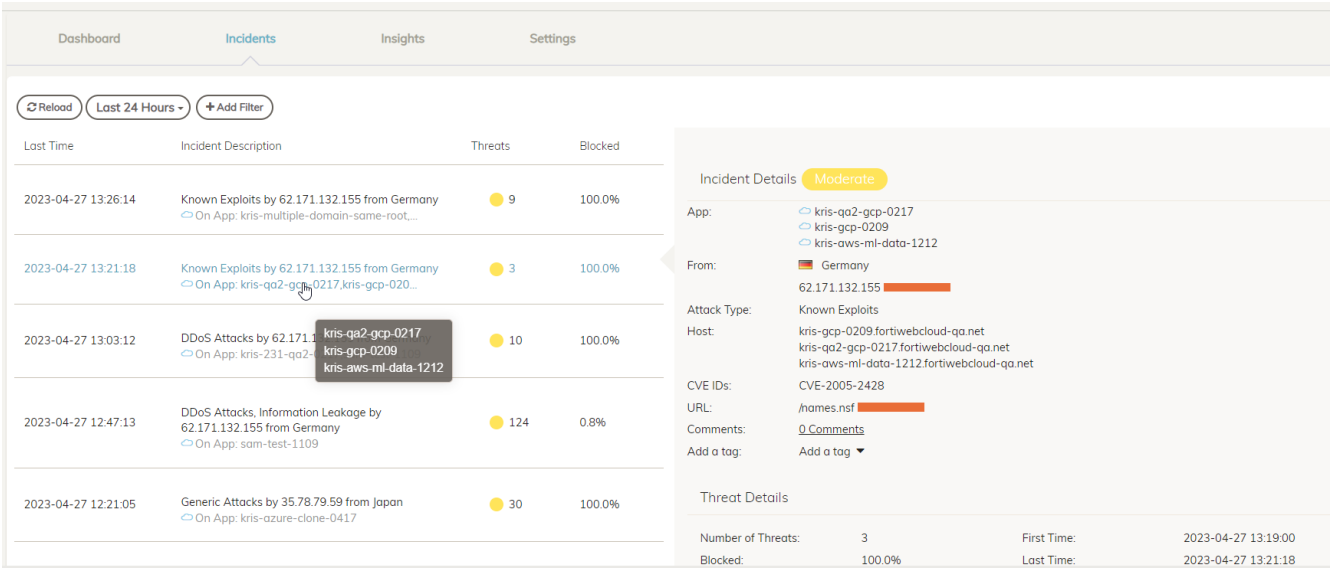
When you are ready to purchase the full license with the Threat Analytics service, contact the Fortinet Sales team.

Overview

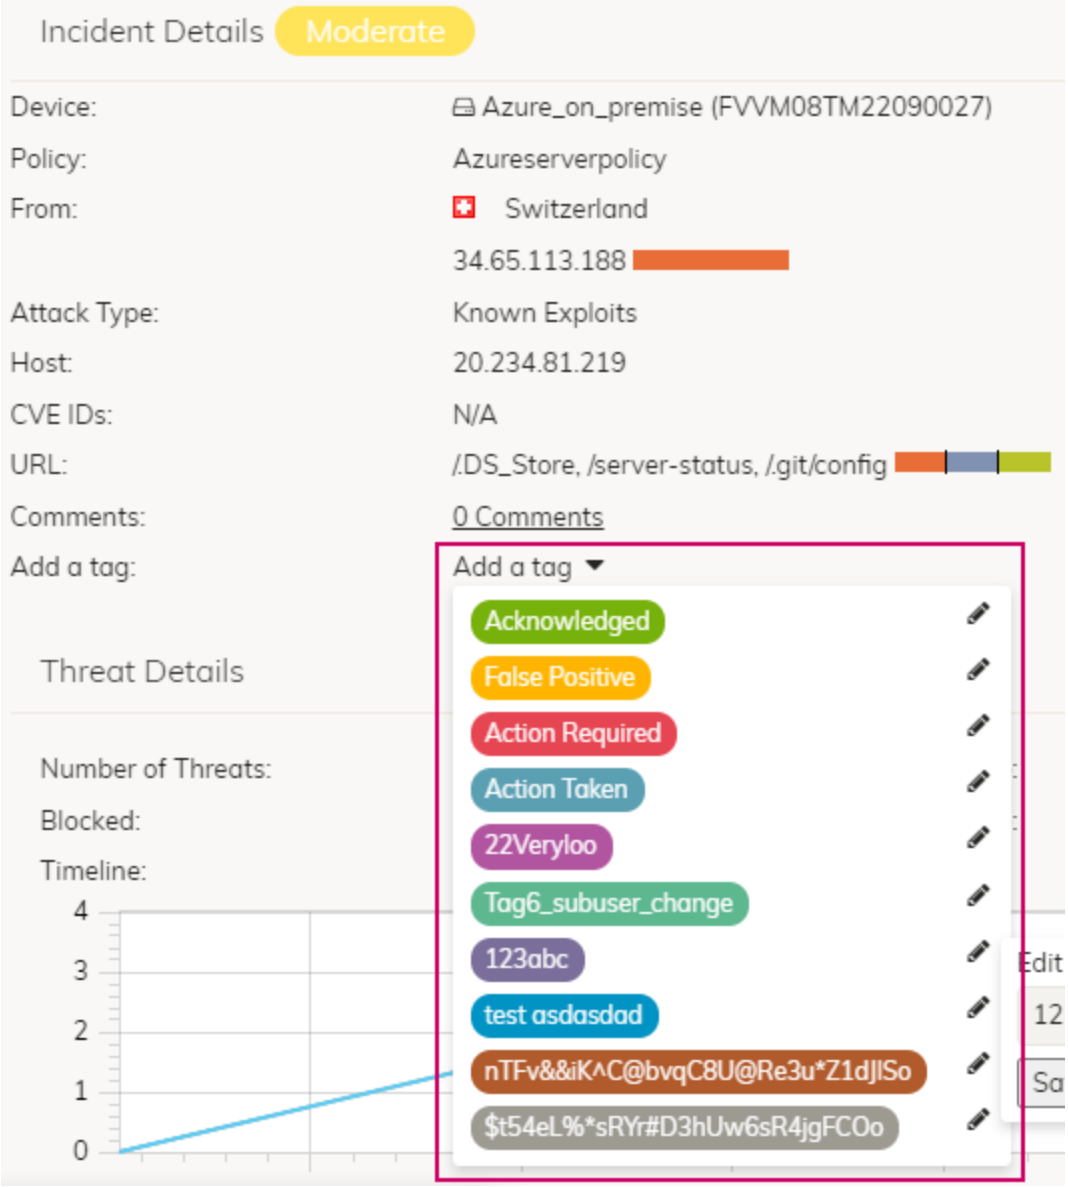
Fortinet AI Threat Analytics leverages machine learning algorithms to identify attack patterns across your entire application attack surface and aggregate them into comprehensible security incidents. The solution separates significant threats from informational alerts and false positives by identifying patterns and assigning a severity to help your security team focus on the threats that matter.

AI Threat Analytics parses through all FortiADC attack logs and aggregates the attack events, grouping them into incidents by common characteristics. These groupings can allow you to identify attack patterns such as which attack types occur the most frequently, or which source IP is the most malicious.

You can drill down to view incident details by clicking the incident number. Incident Details include information such as the attack type, the target application, and source IPs.



You can use predefined tags for AI Threat Analytics incidents. You can edit the tag name according to your needs to help in labeling the incidents for future usage, such as sorting, filtering, and acknowledging incidents.



To enable AI Threat Analytics:

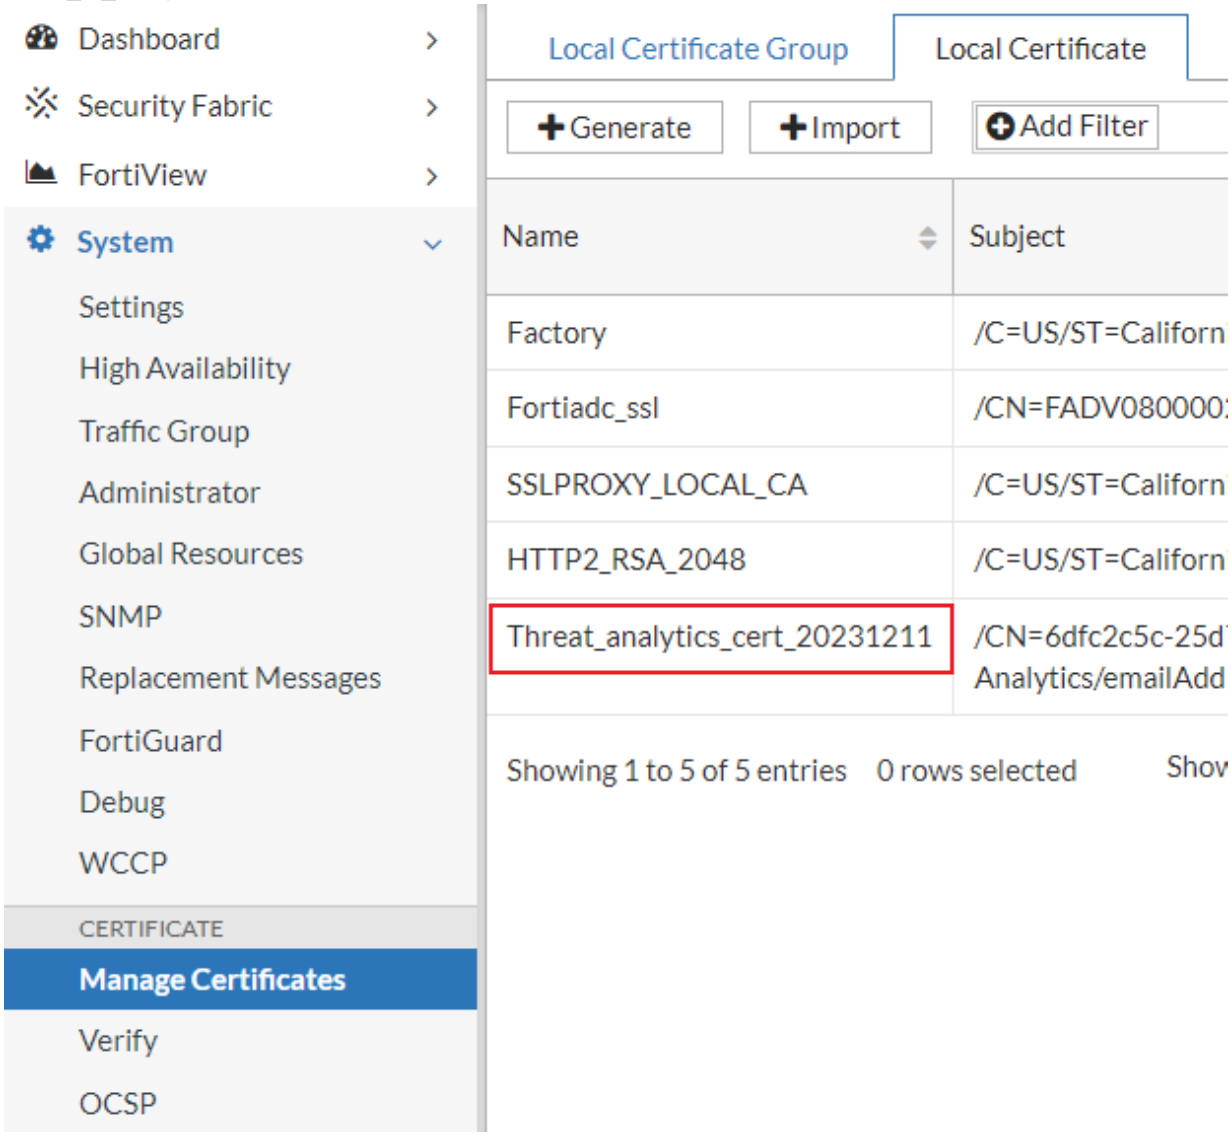
1. Contact the Fortinet Sales team to purchase a license with the Fortinet AI Threat Analytics service, then register the license on the Fortinet Support site: https://support.fortinet.com/.
2. Log in to FortiADC.
3. Go to **System > FortiGuard**, log in to your Fortinet Support Contract.
   You must be logged into your Fortinet Support Contract to connect FortiADC with the AI Threat Analytics service as FortiWeb Cloud requires your Email ID to connect.
4. In the **Dashboard > Status License** widget, check the status of AI Threat Analytics. The status should be displayed as Valid.
5. Go to **Security Fabric > Fabric Connectors**. Under **Other Fortinet Products** section, locate the **Threat Analytics** connector.

6. **Enable** the Threat Analytics connector.

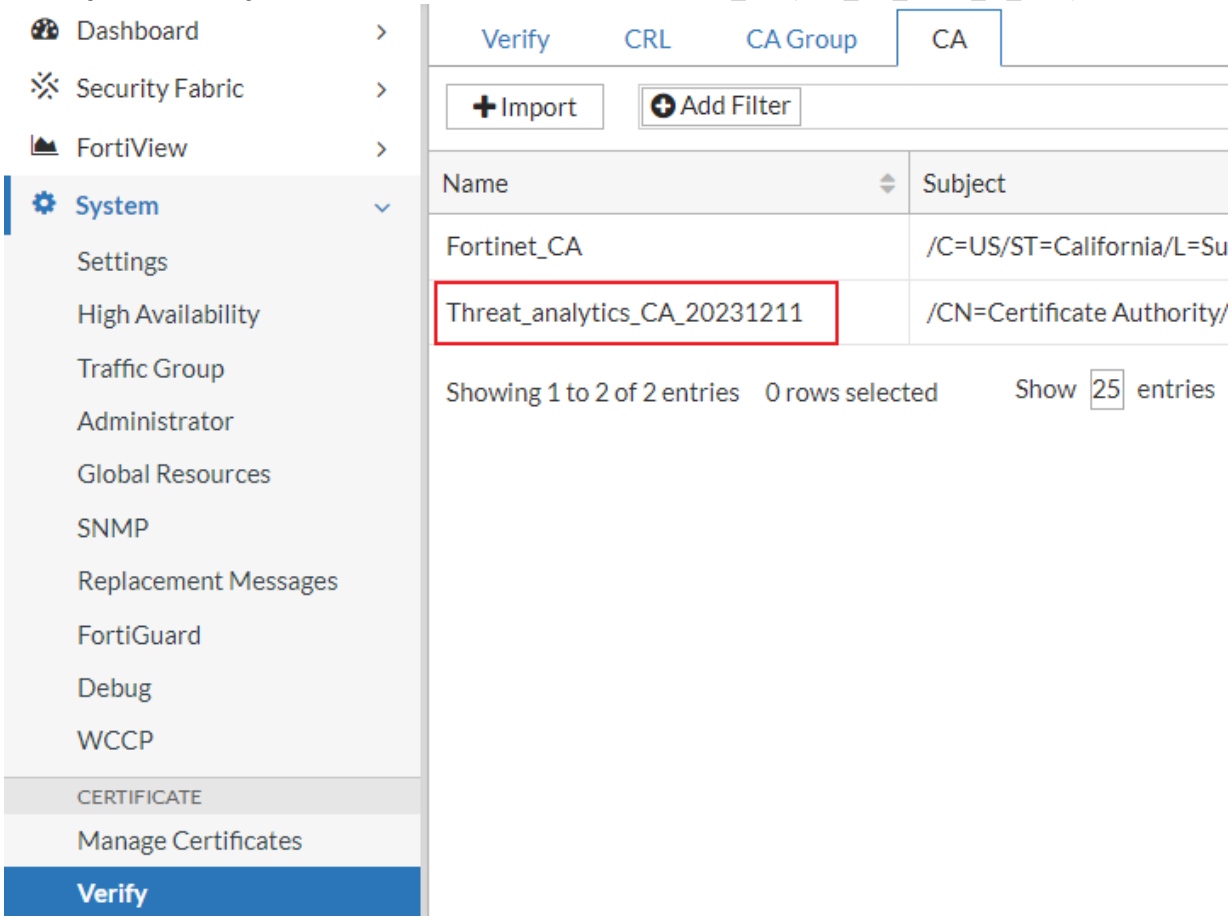


Note: When enabling the Threat Analytics connector for the first time, it may take 5-10 minutes to connect.

7. Once the Threat Analytics connector successfully connects FortiADC to the Fortinet AI Threat Analytics service, a new local certificate and CA will be created. Check the certificates and CA to ensure they are present.
   a. Go to **System > Manage Certificates** to locate the new local certificate with the name *Threat_analytics_cert_<date_of_today>*.



b. Go to **System > Verify** to locate the new CA with the name *Threat_analytics_CA_<date_of_today*.



c. A new syslog global_remote server will be created with the FQDN address type and with the comment "fweb_cloud".
The remote syslog server capacity is three. Prior to enabling AI Threat Analytics, ensure there is at least one syslog server entry available for the new remote syslog server created when AI Threat Analytics is enabled.

8. Wait to allow FortiADC to generate attack logs and forward them to FortiWeb Cloud.
9. Log in to FortiWeb Cloud with the account you used when registering your license on the Fortinet Support site.



Do not delete or modify the syslog remote and certificate/CA entry. AI Threat Analytics cannot be functional without these configurations.

AI Threat Analytics in VDOM

When AI Threat Analytics is enabled in VDOMs, Override in the Syslog Server configuration will be disabled in order to use the global syslog server. If you have previously enabled Override in the Syslog Server configuration, then the default global syslog server list would be removed and you may use a new syslog server list specifically defined in the VDOM. By default, the new syslog remote server would also be created in all the VDOMs with AI Threat Analytics enabled, which disables Override in order to use the global syslog server.

When AI Threat Analytics is enabled, it will always use the global or root DNS, and not the VDOM's DNS.

AI Threat Analytics in HA

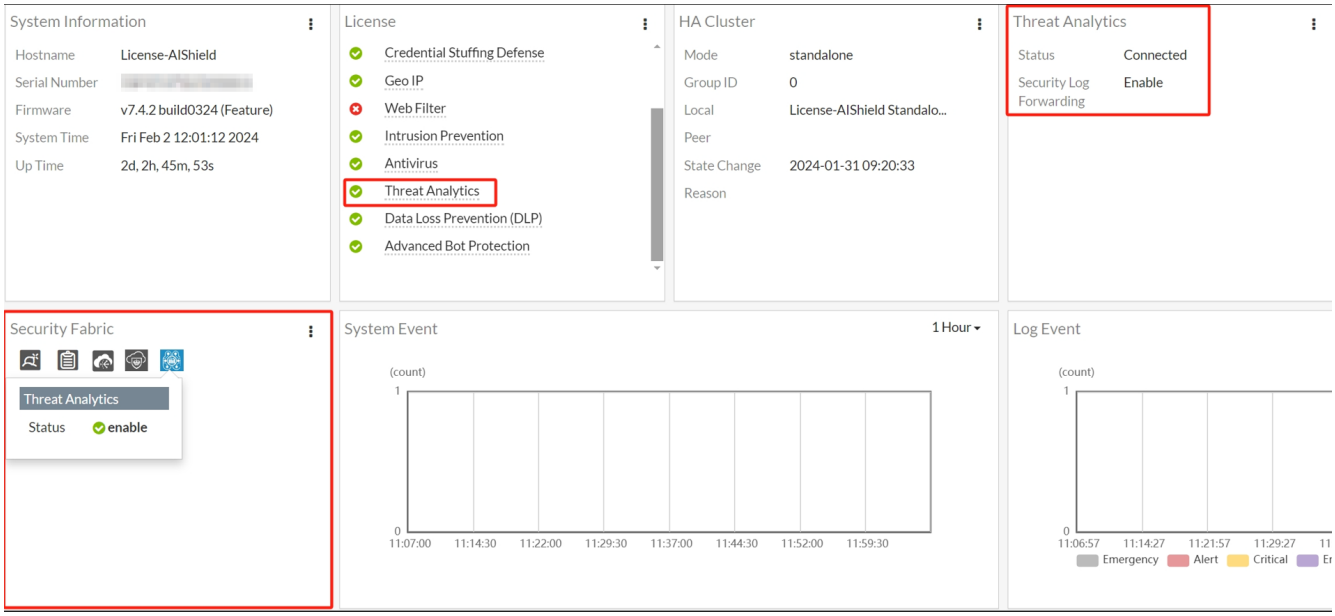
In HA mode, only the primary node is connected to the FortiWeb Cloud server for the Fortinet AI Threat Analytics service. The certification and syslog configurations from the primary unit are then synchronized to the secondary unit. This workflow is designed to prevent HA synchronization issues that can arise with having both the primary and secondary nodes connect to the FortiWeb Cloud at the same time. Once the primary node is connected to the AI Threat Analytics service, the secondary node will synchronize the connection status and show as "Connected".

AI Threat Analytics troubleshooting and debugging

You can use the following tools to diagnose and troubleshoot Threat Analytics issues in FortiADC.

Dashboard

From the FortiADC main dashboard, statuses relating to Threat Analytics is displayed in several widgets.



License

From the **License** widget, you can check the status of your Threat Analytics service license. You must have a Threat Analytics service license in order to use this integrated service for FortiADC logs. You will not be able to connect to the FortiWeb Cloud server without a valid license.

Security Fabric

From the **Security Fabric** widget, you can check the status of the Threat Analytics connector.

Threat Analytics

The Threat Analytics dashboard widget displays the connection status to the FortiWeb Cloud server and the status of the FortiADC attack log forwarding.

Threat Analytics connector

When you enable the Threat Analytics connector, the Fortinet AI Threat Analytics service license status will display.



The and icons indicate whether the Threat Analytics connector has successfully connected to the FortiWeb Cloud server. If the connection is down , FortiADC will first perform an inspection of the Fortinet AI Threat Analytics license status to determine whether the connection issue is caused by an invalid license. If a valid Fortinet AI Threat Analytics license exists, then further troubleshooting may be required to determine the root cause of the Threat Analytics connector issue.

Icon	Threat Analytics connector status	Guidelines
	Connected	The FortiADC is successfully connected to the FortiWeb Cloud server.
	Valid License	A valid license for Fortinet AI Threat Analytics service is present. The Threat Analytics connector is unable to connect to the FortiWeb Cloud server due to issues other than the license status. Further troubleshooting is recommended, such as checking your network settings.
	14-day Trial Started	The 14-day evaluation license for Fortinet AI Threat Analytics service has been activated, but has not yet expired. The Threat Analytics connector is unable to connect to the FortiWeb Cloud server due to issues other than the license status. Further troubleshooting is recommended, such as checking your network settings.
	14-day Trial Expired	The 14-day evaluation license for Fortinet AI Threat Analytics service has expired. The Threat Analytics connector is unable to connect to the FortiWeb Cloud server due to invalid license. Please contact the Fortinet Sales team to purchase the Fortinet AI Threat Analytics service license to continue using AI Threat Analytics.
	14-day Trial Not Started	The 14-day evaluation license for Fortinet AI Threat Analytics service did not activate because FortiWeb Cloud was unable to identify the connecting FortiADC through the Fortinet Support Contract Email ID.

Icon	Threat Analytics connector status	Guidelines
		FortiWeb Cloud requires the Email ID registered to the Fortinet Support Contract to identify and connect your FortiADC to the AI Threat Analytics service. When you attempt to enable the Threat Analytics connector before logging into your Fortinet Support Contract from FortiADC, the Threat Analytics connector will fail to connect. Please ensure to log into your Fortinet Support Contract from the **System > FortiGuard** page.
	No license.	There is no basic FortiADC license. If the FortiADC is on a trial license, you also cannot activate the 14-day Evaluation License. Please contact the Fortinet Sales team to purchase a FortiADC license.

CLI commands to view debug logs relating to AI Threat Analytics

Command	Guidelines
`diagnose debug module wassd`	To view the debug information of the wassd daemon. The wassd daemon forms the connection between FortiADC and FortiWeb Cloud and performs several integral functions when AI Threat Analytics is enabled. This includes the following: • Establishing a web socket connection with the FortiWeb Cloud using a token. The wassd identifies whether a CA exists before registering to the FortiWeb Cloud. If a CA does exist, then the wassd will send the issue date of the CA certificate to the FortiWeb Cloud. • Updating FortiWeb Cloud with FortiADC configuration changes, such as HA status changes, member updates, or mode modification. • Updating device certificates received from the FortiWeb Cloud. If wassd registered to the FortiWeb Cloud without the issue date of the CA or that the certificate has expired, then FortiWeb Cloud will send new certificates (including the certificate, key, and CA) to wassd. The wassd will update to the local certificate and CA table, and register to FortiWeb Cloud again with the latest CA issue date. • Starting the forwarding of FortiADC attack logs to FortiWeb Cloud. If wassd has successfully registered to FortiWeb Cloud, then it will start the action with the log server and port from the FortiWeb Cloud. **Note**: The wassd daemon is create for AI Threat Analytics and executes the `wassd_ws` Python script when AI Threat Analytics is enabled. The backend log for the Python script is stored in `/var/log/wassd.log`.
`diagnose debug module miglogd syslog`	To view the debug information for the miglogd syslog.

Command	Guidelines
	AI Threat Analytics functionality requires FortiADC to send the attack logs to the FortiWeb Cloud via syslog (system logging protocol) that uses TCP SSL. In the case where communication issues arise between FortiADC and FortiWeb Cloud, you can use the `diagnose debug module miglogd syslog` command to print out the key information about these communications to find the cause.
`diagnose system threat-analytics info`	To view the system information for AI Threat Analytics.

SSL Advanced Services

This section includes the following topics:

- SSL offloading on page 962
- SSL decryption by forward proxy on page 964
- SSL profile configurations on page 967
- Certificate guidelines on page 970
- SSL/TLS versions and cipher suites on page 970
- Exceptions list on page 975
- SSL traffic mirroring on page 976
- HSM Integration on page 979

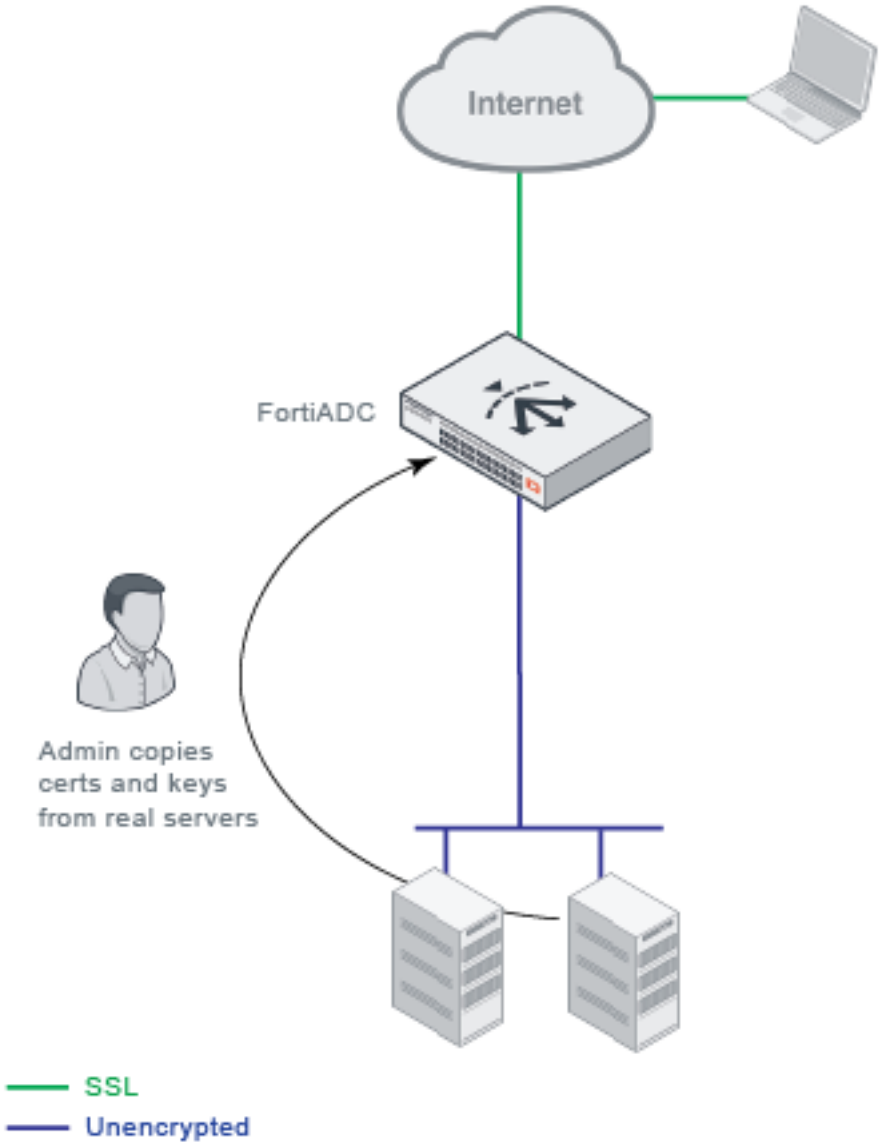
SSL offloading

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 963. 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.

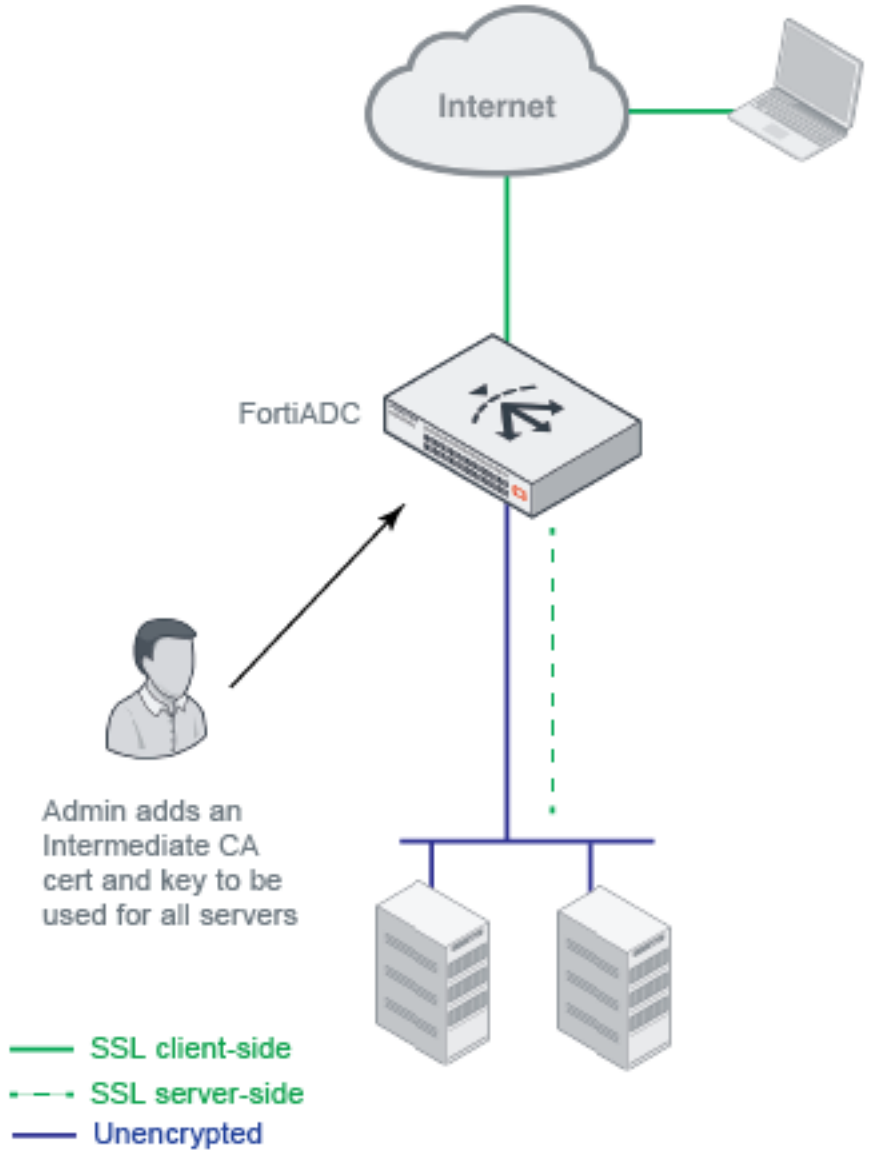
Basic steps:

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 965 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 965 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

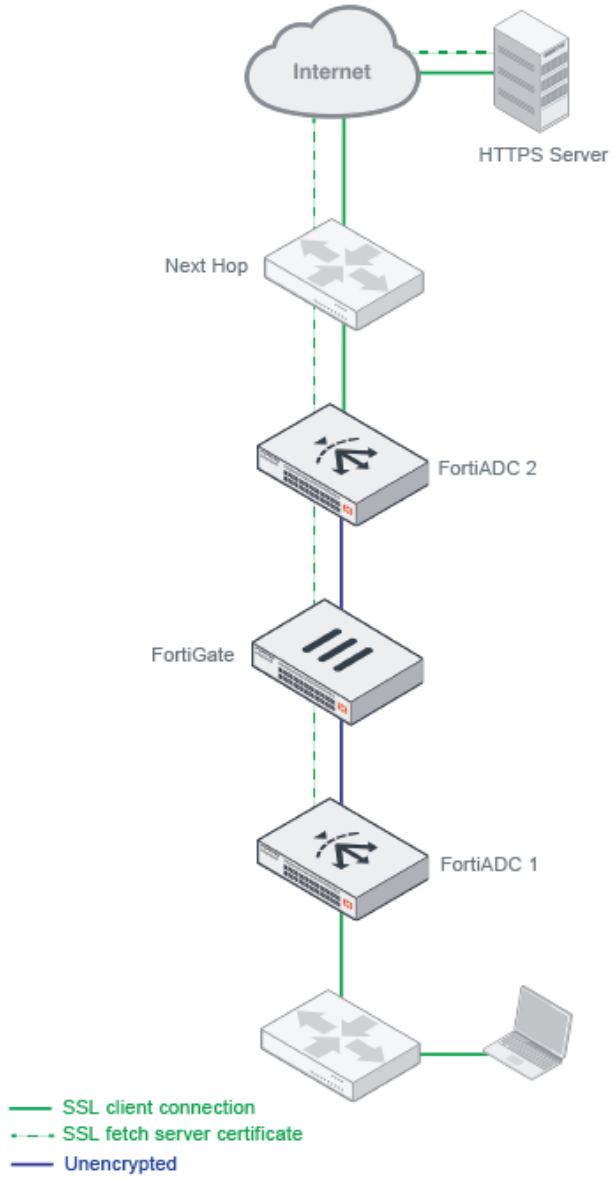
You can use FortiADC in a Layer 2 sandwich toplogy to offload SSL decryption tasks from FortiGate.

Layer 2 SSL decryption by forward proxy on page 966 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

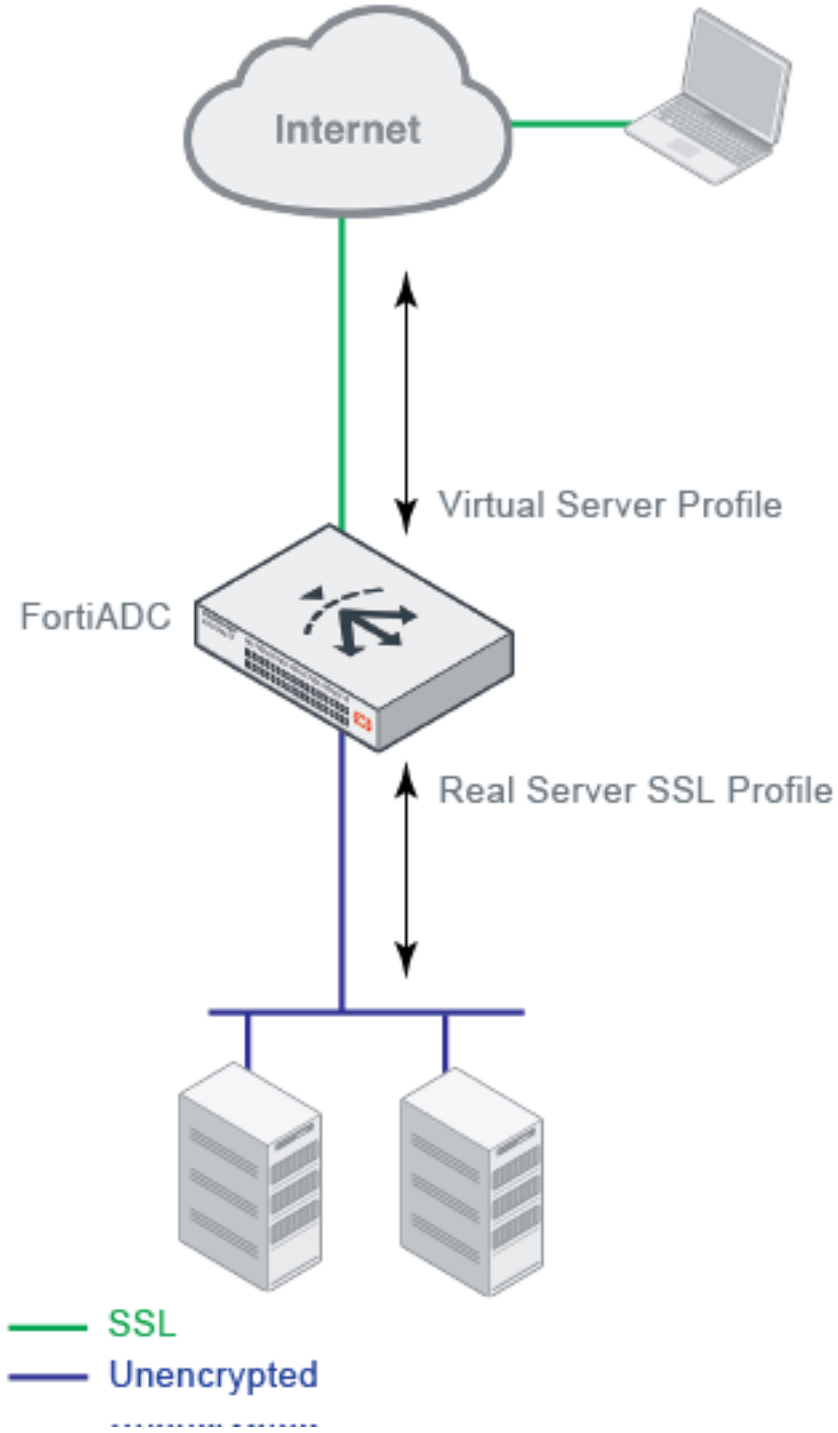


SSL profile configurations

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 967 illustrates the basic idea of client-side and server-side profiles.

SSL profiles



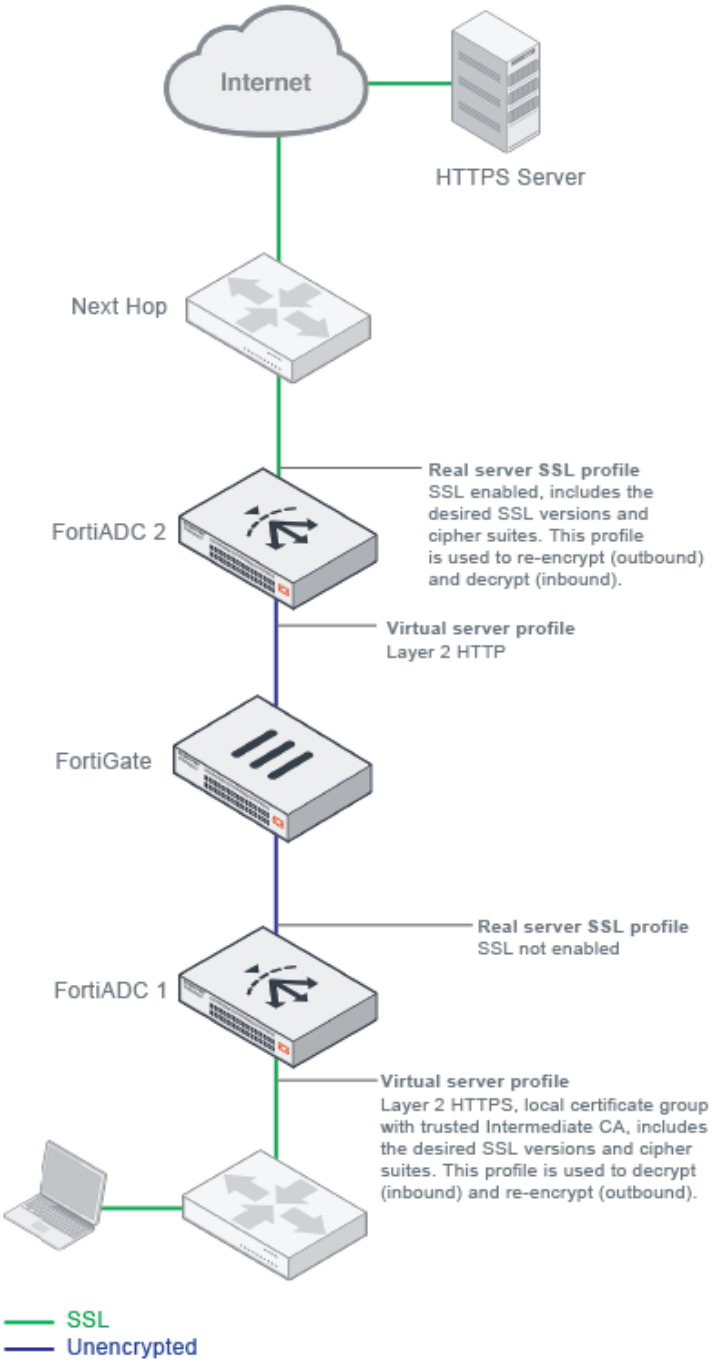
The call-outs in Layer 2 sandwich profiles on page 969 have guidance for the two types of profiles used in a Layer 2 sandwich deployment.

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 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, Camellia ciphers, and eNull ciphers.

Cipher suites with RSA key exchange on page 971 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

*These ciphers are fully supported by hardware SSL (in 400F, 420F, 1200F, 2200F, 4200F and 5000F).

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 972 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
*These ciphers are fully supported by hardware SSL (in 400F, 420F, 1200F, 2200F, 4200F and 5000F).						

Cipher suites with EDCHE key exchange on page 972 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 EDCHE 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

Abbreviation	Cipher Suite	Protocol	Kx	Au	Enc	MAC
**ECDHE-ECDSA-AES128-SHA256*	TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256	TLSv1.2	ECDH	ECDSA	AES128	SHA256
**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

*These ciphers are fully supported by hardware SSL (in 400F, 420F, 1200F, 2200F, 4200F and 5000F).

The Camellia is a symmetric key block cipher with a block size of 128 bits and key sizes of 128, 192 and 256 bits. The Camellia cipher has been approved for use by the ISO/IEC, the European Union's NESSIE project and the Japanese CRYPTREC project. It has security levels and processing abilities comparable to the Advanced Encryption Standard.

Camellia Cipher suites on page 974 lists supported Camellia ciphers with ECDHE and DHE key exchange. The supported suites include both the ECDSA and RSA key authentication algorithms.

Camellia Cipher suites

Abbreviation	Cipher Suite	Protocol	Kx	Au	Enc	MAC
ECDHE-ECDSA-CAMELLIA256-SHA384	TLS_ECDHE_ECDSA_WITH_CAMELLIA_256_CBC_SHA384	TLS 1.2	ECDH	ECDSA	CAMELLIA256	SHA384
ECDHE-RSA-CAMELLIA256-SHA384	TLS_ECDHE_RSA_WITH_CAMELLIA_256_CBC_SHA384	TLS 1.2	ECDH	RSA	CAMELLIA256	SHA384
DHE-RSA-CAMELLIA256-SHA256	TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA256	TLS 1.2	DH	RSA	CAMELLIA256	SHA256
ECDHE-ECDSA-CAMELLIA128-SHA256	TLS_ECDHE_ECDSA_WITH_CAMELLIA_128_CBC_SHA256	TLS 1.2	ECDH	ECDSA	CAMELLIA128	SHA256
ECDHE-RSA-CAMELLIA128-SHA256	TLS_ECDHE_RSA_WITH_CAMELLIA_128_CBC_SHA256	TLS 1.2	ECDH	RSA	CAMELLIA128	SHA256
DHE-RSA-CAMELLIA128-SHA256	TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA256	TLS 1.2	DH	RSA	CAMELLIA128	SHA256
DHE-RSA-CAMELLIA256-SHA	TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA	TLS 1.2, 1.1, 1.0	DH	RSA	CAMELLIA256	SHA

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/manprimary/apps/ciphers.html

TLS 1.3 Cipher Suites

TLS 1.3 introduces major enhancements to cryptographic security and protocol efficiency. It removes outdated and insecure algorithms, providing a streamlined and robust cryptographic baseline. Forward secrecy is enforced by default through the exclusive use of ephemeral key exchanges in all cipher suites, ensuring that the compromise of a server's private key does not jeopardize the confidentiality of past sessions. Additionally, TLS 1.3 optimizes the handshake process by reducing round-trip exchanges, resulting in lower latency and improved connection performance without sacrificing security.

Key Exchange and Cipher Suites in TLSv1.3

Abbreviation	Cipher Suite	Protocol	Kx	Au	Enc	MAC
TLS_AES_256_GCM_SHA384	TLS_AES_128_CCM_8_SHA256	TLS 1.3	any	any	AESGCM(256)	AEAD
TLS_AES_128_GCM_SHA256	TLS_AES_128_GCM_SHA256	TLS 1.3	any	any	AESGCM(128)	AEAD
TLS_CHACHA20_POLY1305_SHA256	TLS_CHACHA20_POLY1305_SHA256	TLS 1.3	any	any	CHACHA20/POLY1305 (256)	AEAD
TLS_AES_128_CCM_SHA256	TLS_AES_128_CCM_SHA256	TLS 1.3	any	any	AESCCM(128)	AEAD
TLS_AES_128_CCM_8_SHA256	TLS_AES_128_CCM_8_SHA256	TLS 1.3	any	any	AESCCM8(128)	AEAD

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

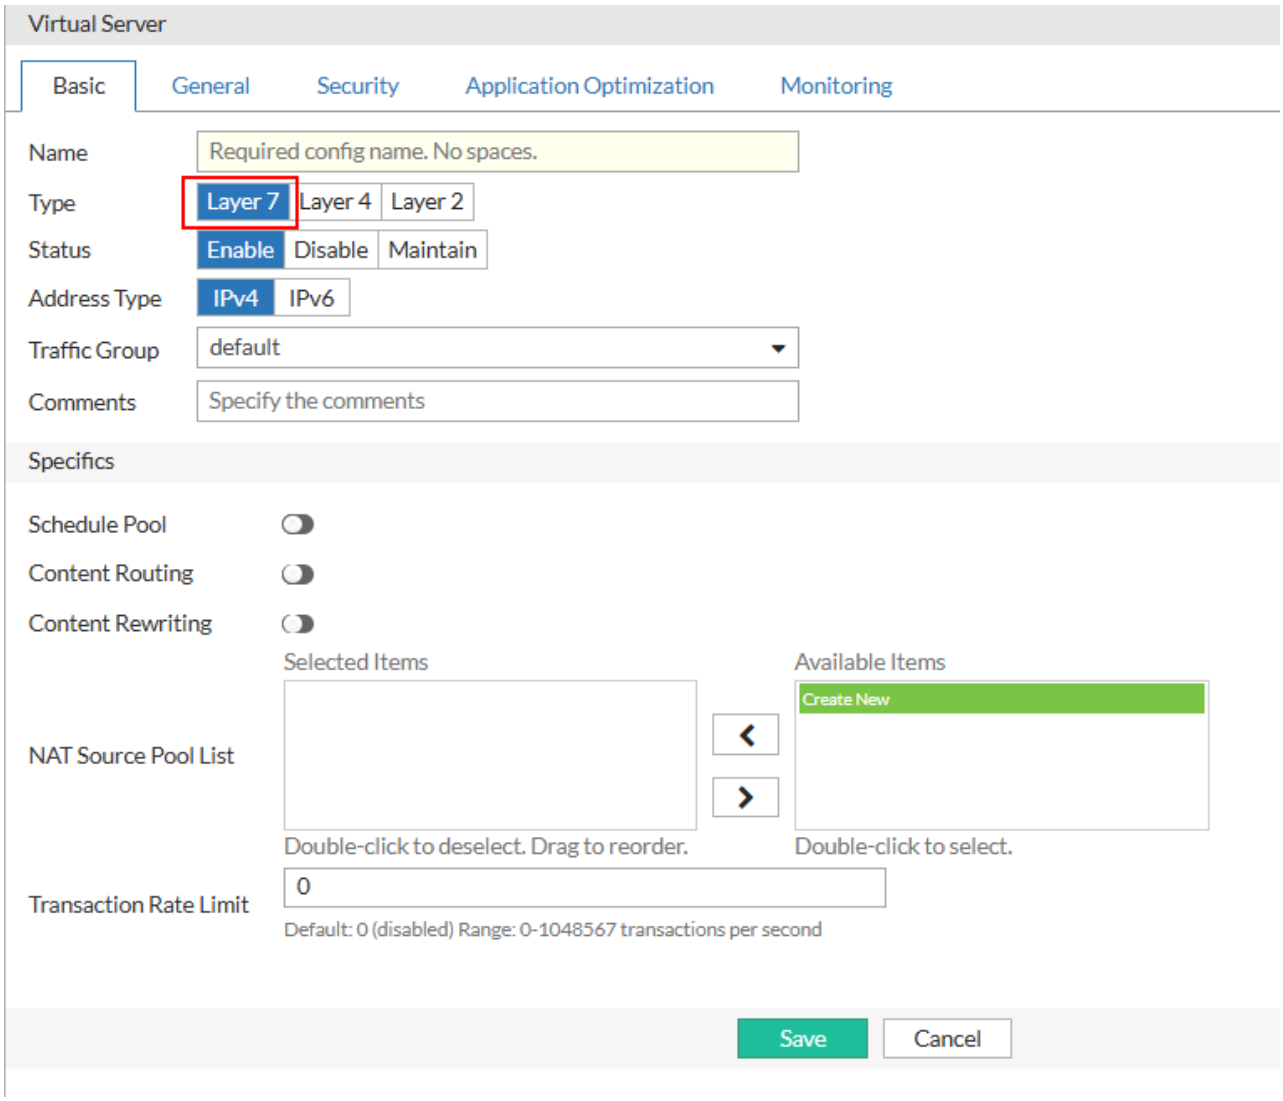
FortiADC supports mirroring packets (HTTPS/TCPS) to specified network interfaces. When the feature is enabled, SSL traffic will be mirrored to the specified ports by the virtual server after it has been decrypted. See the following figures.

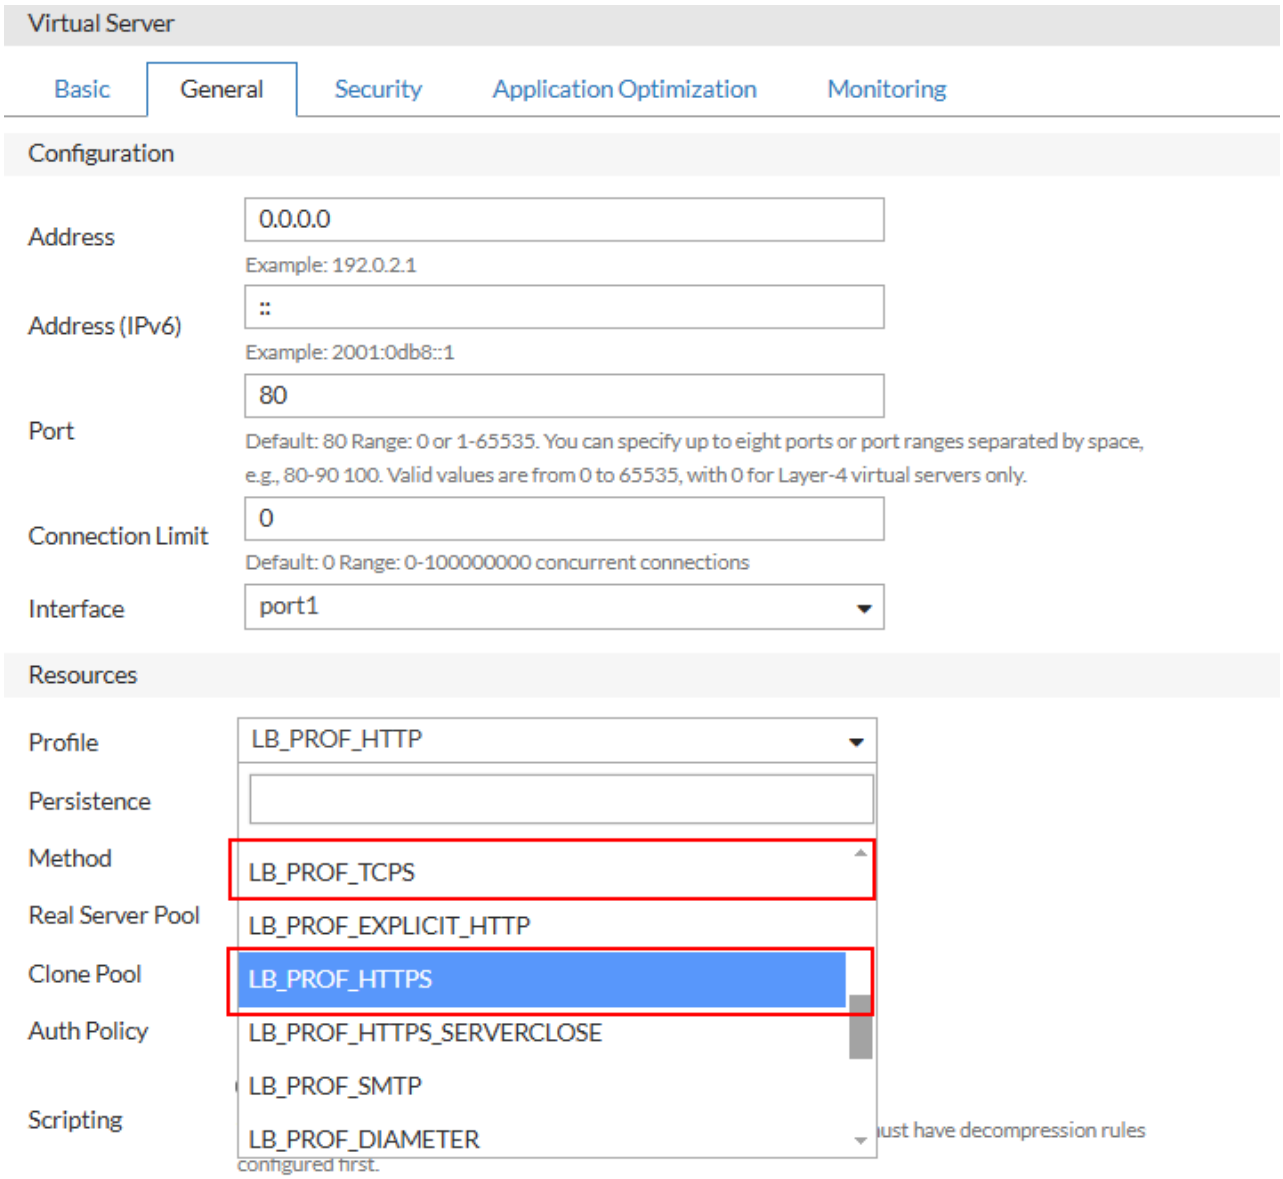
The feature supports both IPv4 and IPv6. FortiADC can send traffic to up to four outgoing interfaces, including aggregated and VLAN interfaces. Mirrored traffic is transmitted as a single packet stream, using the original client-side source and destination IP address and port numbers. The source and 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.

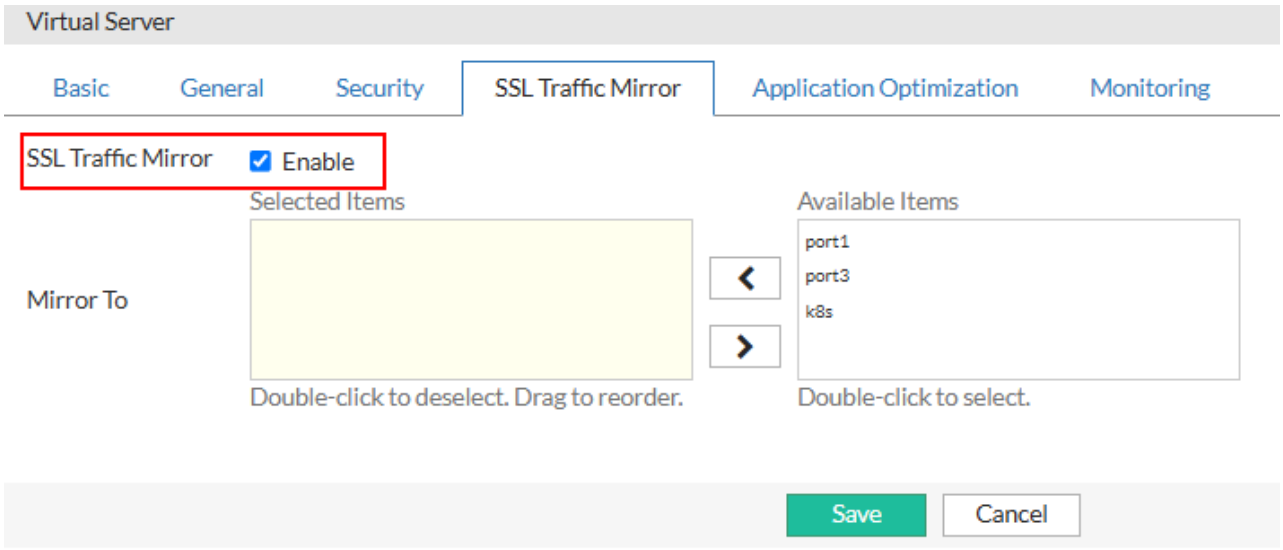
To configure SSL traffic mirroring

1. Navigate to **Server Load Balance > Virtual Server**.
   The configuration page displays the **Virtual Server** tab.
2. Click **Create New > Advanced Mode** to open the Advanced Mode configuration editor.
3. Configure the following key parameters:

Tab	Parameter	Guideline
Basic	Type	Select Layer 7.
General	Profile	Select an HTTPS or TCPS profile. **Note**: The SSL Traffic Mirror tab becomes configurable once an HTTPS or TCPS profile is selected.
SSL Traffic Mirror	SSL Traffic Mirror	Enable this option.


Virtual Server
Basic
General
Security
Application Optimization
Monitoring
Name
Required config name. No spaces.
Type
Layer 7
Layer 4
Layer 2
Status
Enable
Disable
Maintain
Address Type
IPv4
IPv6
Traffic Group
default
Comments
Specify the comments
Specifics
Schedule Pool
Content Routing
Content Rewriting
Selected Items
Available Items
Create New
NAT Source Pool List
Double-click to deselect. Drag to reorder.
Double-click to select.
Transaction Rate Limit
0
Default: 0 (disabled) Range: 0-1048567 transactions per second
Save
Cancel


Virtual Server
Basic
General
Security
Application Optimization
Monitoring
Configuration
Address
0.0.0.0
Example: 192.0.2.1
Address (IPv6)
::
Example: 2001:0db8::1
Port
80
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.
Connection Limit
0
Default: 0 Range: 0-100000000 concurrent connections
Interface
port1
Resources
Profile
LB_PROF_HTTP
Persistence
Method
Real Server Pool
Clone Pool
Auth Policy
Scripting
LB_PROF_TCPS
LB_PROF_EXPLICIT_HTTP
LB_PROF_HTTPS
LB_PROF_HTTPS_SERVERCLOSE
LB_PROF_SMTP
LB_PROF_DIAMETER
ust have decompression rules
configured first.



4. Click **Save** to commit the changes.

To enable this feature in a policy via the CLI, execute the following command:

```
config load-balance virtual-server
  edit vs-name
    set ssl-mirror enable
    set ssl-mirror-intf port1 port2
  next
end
```

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 981. Then move on to Generating a certificate-signing request on FortiADC on page 982.

HSM Configuration Parameters

Parameter	Description
Client Certificate	
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 **Choose File** to browse for the server certificate file that you retrieved earlier.
Server FIPS Support	Enable/disable server FIPS support.
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., `/example.conf`
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 982. Then move on to Downloading and uploading the certificate request (.csr) file on page 984.

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.

Parameter	Description
	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.
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 `config system global` command. For more information, see Integrating FortiADC with SafeNet Network HSM on page 979.

Parameter	Description
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. *This is the only enrollment method if HSM is selected.* • 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. *This option is not available if HSM is selected.*

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 985, 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 [[[Undefined variable FortiADCVariables.FortiProduct]]] . See HSM Integration on page 979. 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.

Best Practices and Fine Tuning

This section is a collection of best practice tips and fine-tuning guidelines. It includes the following topics:

- Regular backups on page 986
- Security on page 987
- Performance tips on page 989
- High availability on page 989

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.

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 to restore the entire device to the original out-of-the-box configuration:

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 perform a factory reset to restore factory defaults but retain the interface and VDOM configuration:

From the CLI console, enter the following commands:

```
execute factoryreset2
```

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.

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
```

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 to traffic 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.

FortiADC Use Cases

The following section provides detailed implementation scenarios that demonstrate FortiADC's versatile capabilities in managing traffic and optimizing application delivery. Each use-case highlights practical solutions for real-world challenges, showcasing how FortiADC enhances performance, security, and scalability across various environments. Whether you're optimizing load balancing, ensuring high availability, or protecting critical applications, these scenarios serve as a guide for effectively leveraging FortiADC to meet your infrastructure's demands.

DNS Load Balancing and DDoS Protection on page 992

Global Server Load Balancing in Hybrid Environments on page 999

Selecting the Optimal Mode: VDOM vs. ADOM on page 1019

SAP Application Integration on page 1037

Multi-Cloud Application Load Balancing on page 1052

Leveraging DevOps Tools (Ansible/Terraform) to Manage FortiADC on page 1066

DNS Load Balancing and DDoS Protection

As a fundamental component of internet architecture, the Domain Name System (DNS) plays a pivotal role in directing user traffic to web services and applications. For GlobalNet, a leading online service provider managing a global DNS infrastructure, ensuring the reliability and performance of DNS services is critical. However, the increasing prevalence of DDoS (Distributed Denial of Service) attacks poses significant risks to DNS availability and responsiveness. This section delves into the deployment of FortiADC to optimize DNS traffic management through intelligent load balancing and distribution across multiple DNS servers. By integrating advanced features such as real-time traffic analysis, health checks, and DDoS mitigation strategies, FortiADC enhances the resilience of DNS services, ensuring high availability and robust performance against malicious disruptions. This technical overview provides insights into configuring FortiADC to safeguard DNS infrastructure while maintaining seamless access for users across diverse geographical regions.

User Scenario: DDoS Mitigation and DNS Reliability for Global Online Services

GlobalNet, a major online service provider, operates a distributed DNS infrastructure supporting essential services like email, web hosting, and online applications. To ensure high performance and service availability, safeguarding its DNS infrastructure from DDoS attacks is critical. FortiADC is deployed as the primary traffic management and DDoS mitigation solution, positioned in front of the DNS servers to manage incoming DNS queries. This deployment enables efficient traffic distribution using load balancing algorithms and DNS caching to optimize query handling and improve overall performance.

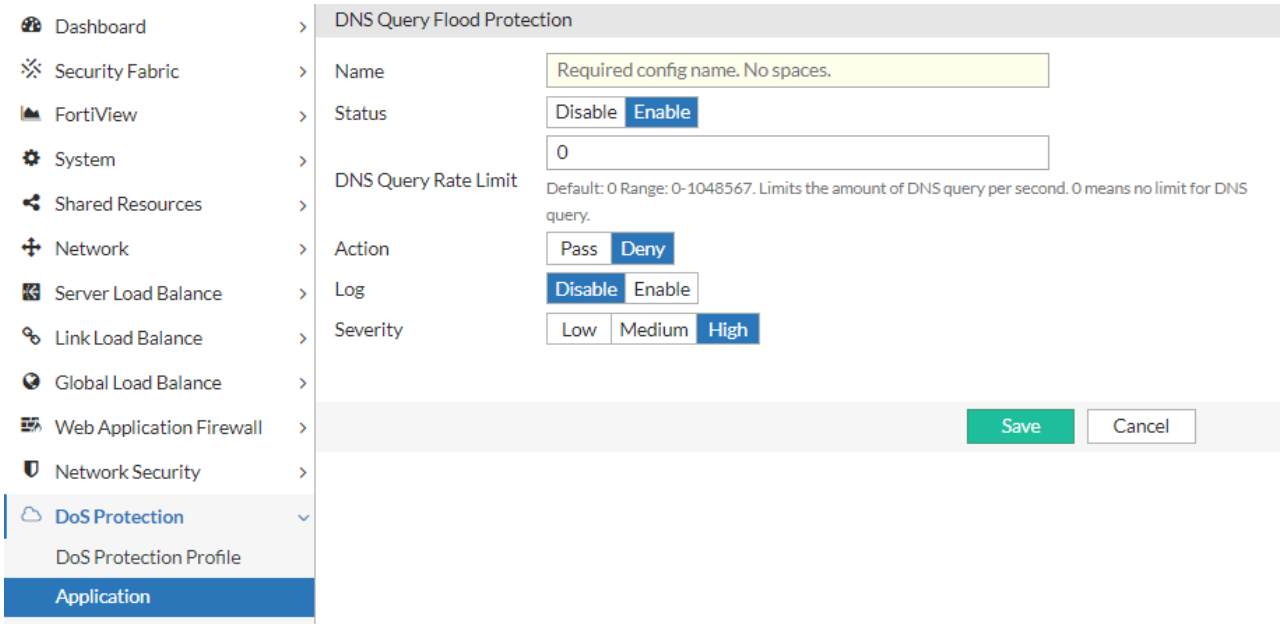
To protect against DDoS threats, FortiADC is configured with DNS Query Flood Protection and DNS Reverse Flood Protection policies, which help mitigate volumetric attacks targeting DNS services. Additionally, traffic filtering mechanisms are employed to detect and block malformed queries, ensuring that only legitimate requests are processed. This configuration enhances GlobalNet's DNS resilience, ensuring high availability and consistent service performance in a dynamic and evolving threat landscape.

Step 1: DNS Query Flood and Reverse Flood Protection

This step outlines the deployment of DNS Query Flood Protection and DNS Reverse Flood Protection mechanisms within FortiADC. By imposing rate limits on DNS queries per second, FortiADC effectively mitigates the risk of service disruptions induced by excessive query volumes. These policies will be configured in response to real-time threat intelligence, ensuring sustained availability and optimal performance of DNS services in the face of evolving attack vectors.

To configure the DNS Query Flood Protection policy:

1. Navigate to **DoS Protection > Application**, and click the **DNS Query Flood Protection** tab
2. Click **Create New** to display the configuration editor. Configure the required settings.

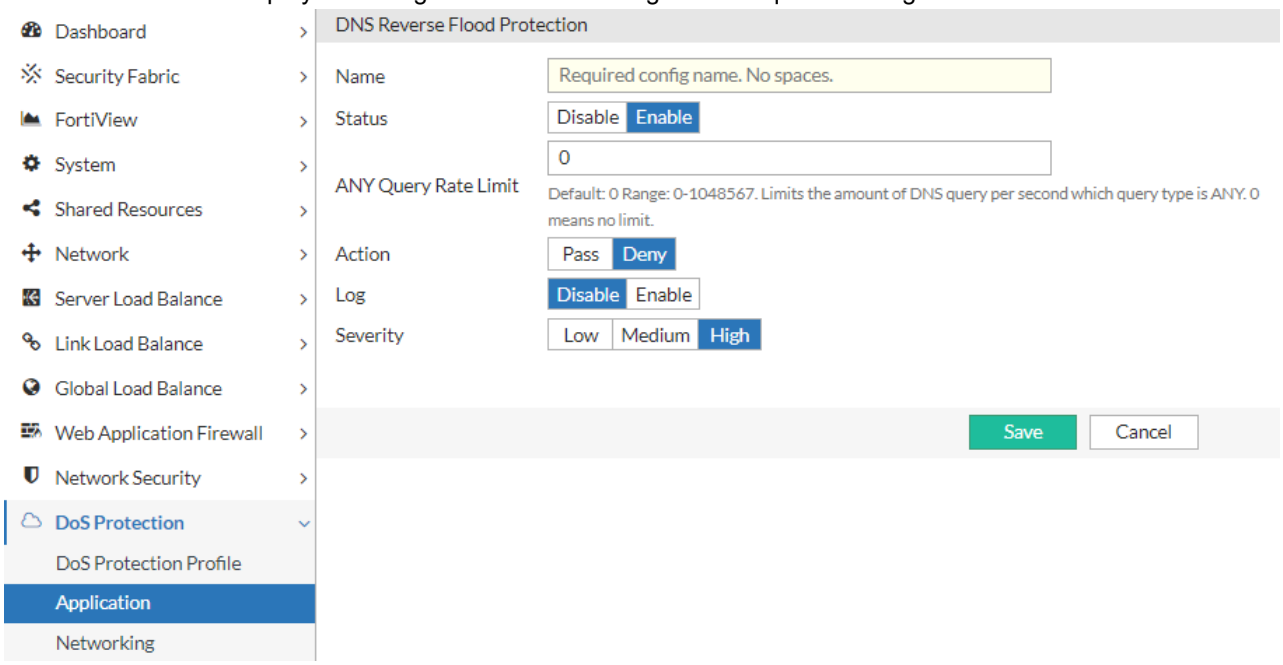


For configuration details, see Configuring a DNS Query Flood Protection policy on page 874.

3. Save the configuration.

To configure the DNS Reverse Flood Protection policy:

1. Navigate to **DoS Protection > Application**, and click the **DNS Reverse Flood Protection** tab
2. Click **Create New** to display the configuration editor. Configure the required settings.

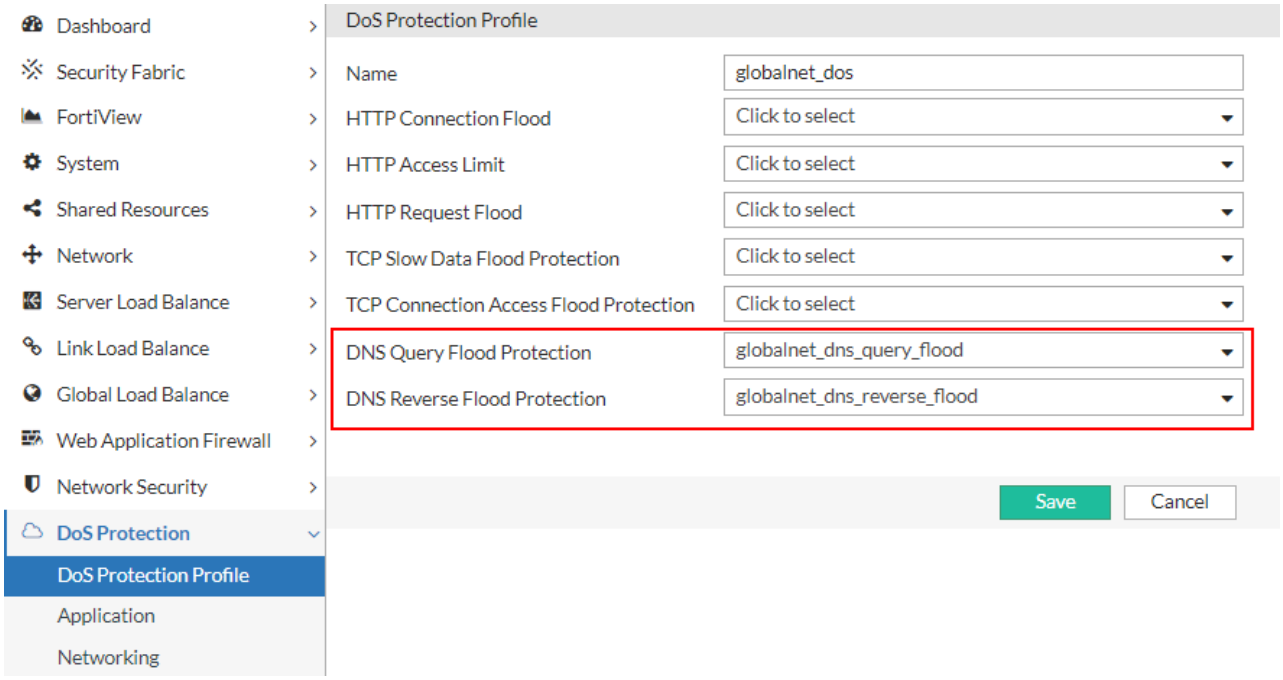


For configuration details, see Configuring a DNS Reverse Flood Protection policy on page 877.

3. Save the configuration.

Create a DoS Protection Profile:

1. Navigate to **DoS Protection > DoS Protection Profile** and click **Create New** to display the configuration editor.
2. Select the DNS Query Flood Protection and DNS Reverse Flood Protection policies previously created.



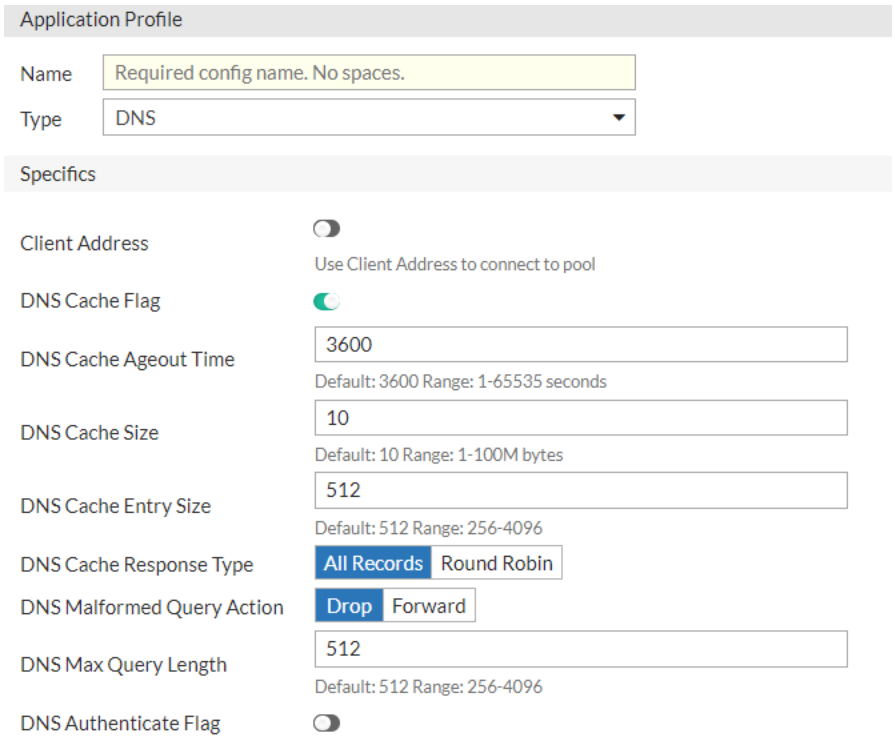
3. Save the configuration.

Step 2: DNS Traffic Management

Implementing robust DDoS protection involves proactive traffic analysis and anomaly detection, which FortiADC supports through its Server Load Balancing (SLB) and DNS Application Profiles. By establishing baseline DNS traffic profiles using historical data, FortiADC can apply precise rate limiting to mitigate volumetric attacks by controlling query rates per IP address. Advanced DNS caching features offload backend servers by serving repetitive queries from the cache, with tunable parameters like TTL ("DNS Cache Ageout Time") and cache size to optimize performance. Additionally, detailed traffic filtering within the DNS profile allows for the detection and mitigation of malformed or anomalous DNS queries through configurable actions such as dropping, redirecting, or logging based on query attributes like type, size, and header integrity. These combined measures enhance the resilience of GlobalNet's DNS infrastructure against DDoS attacks and service disruptions.

To configure a DNS Application Profile:

1. Navigate to **Server Load Balance > Application Resources**.
   The configuration page displays the **Application Profile** tab.
2. Click **Create New** to display the configuration editor.
3. Enter a unique name for the custom profile, and select the Type as DNS. Configure the required settings to handle DNS traffic:



Parameter	Description
Client Address	Enable/disable to use the original client IP address as the source address when connecting to the real server.
DNS Cache Flag	Enable/disable the cache for the DNS virtual server.
DNS Cache Ageout Time	Specify the cache age-out time (in seconds). The default is 3,600. The valid range is 0 to 65,535.
DNS Cache Size	Specify the maximum cache size (in Megabytes). The default is 10. The valid range is 1 to 100.
DNS Cache Entry Size	Specify the maximum cache entry size. The default is 512. The valid range is 256 to 4,096.
DNS Cache Response Type	Select either of the following cache response types: • All Records • Round Robin
DNS Malform Query Action	Select either of the following reactions for the malformed requests: • Drop • Forward
DNS Max Query Length	Specify the maximum query length. The default is 512. The valid range is 256 to 4,096.
DNS Authentication Flag	Enable/disable to authenticate client by redirecting UDP query to TCP.

4. Save the configuration.

Alternatively, you can also use the predefined **LB_PROF_DNS** profile that has the following settings:

Parameter	Default value
Client Address	Off
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
DNS Max Query Length	512
DNS Authentication Flag	Disabled

Step 3: Applying DNS DDoS Protection to SLB

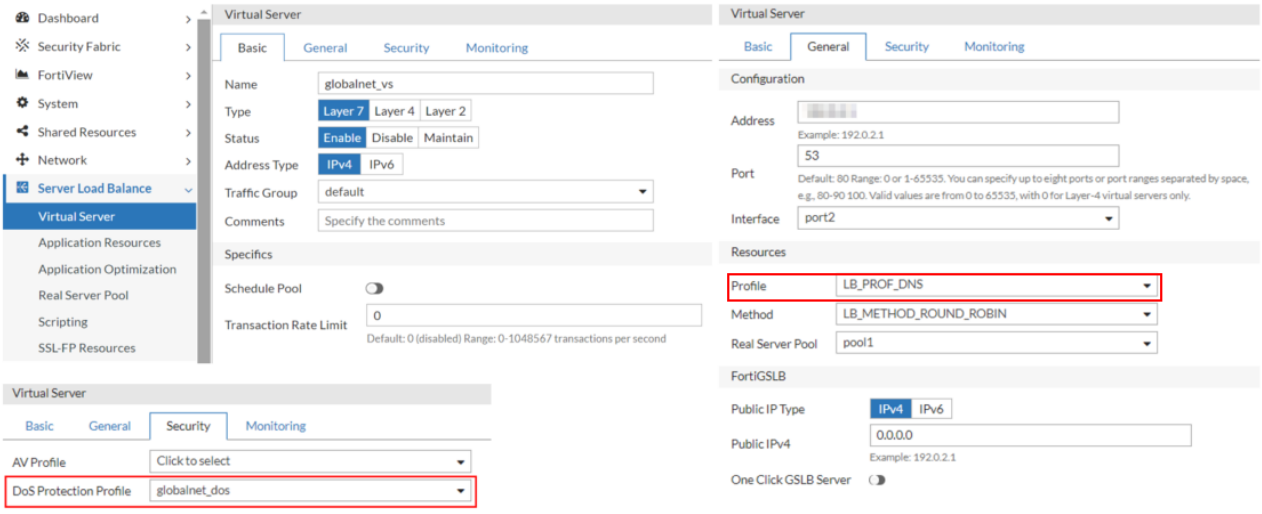
After configuring the DoS Protection Profile and the DNS Application Profile, the next step is to associate them with your Server Load Balancing (SLB) configuration. This integration ensures that all incoming DNS traffic directed to backend servers is monitored and managed according to the predefined traffic management rules and DoS protection policies. By binding the DNS Profile and DoS Protection Profile to the SLB configuration, FortiADC can effectively detect and mitigate potential DDoS attacks while maintaining consistent traffic distribution and high availability across the server pool.

To apply the DNS Profile and DoS Protection Profile to a virtual server:

1. Navigate to **Server Load Balance > Virtual Server**.
   The configuration page displays the **Virtual Server** tab.
2. Click **Create New** and select the **Advanced Mode** to display the configuration editor. Configure the key settings:

Setting	Guideline
Name	Enter a unique name for the Virtual Server.
Type	Select **Layer 7**.
Address	Enter the IP address for the Virtual Server.
Port	Enter **53** for the DNS common port
Interface	Select the interface to which you bind the VS.
Profile	Select a custom user-defined DNS Application Profile or the predefined **LB_PROF_DNS** Application Profile. For details, see Configuring Application profiles on page 451.

Setting	Guideline
Method	Select the **LB_METHOD_ROUND_ROBIN** predefined profile for the load-balancing method. The Round Robin method selects the next server in the series: server 1, then server 2, then server 3, and so on. For details, see Configuring load-balancing (LB) methods on page 526.
Real Server Pool	Select a real server pool configuration.
DoS Protection Profile	Select the DNS Protection Profile previously created.

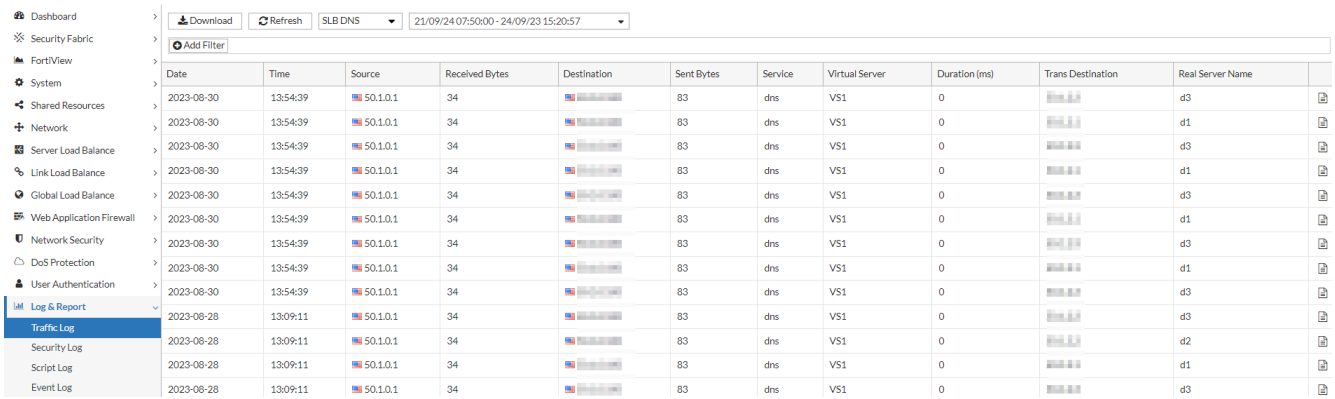


3. Save the configuration.

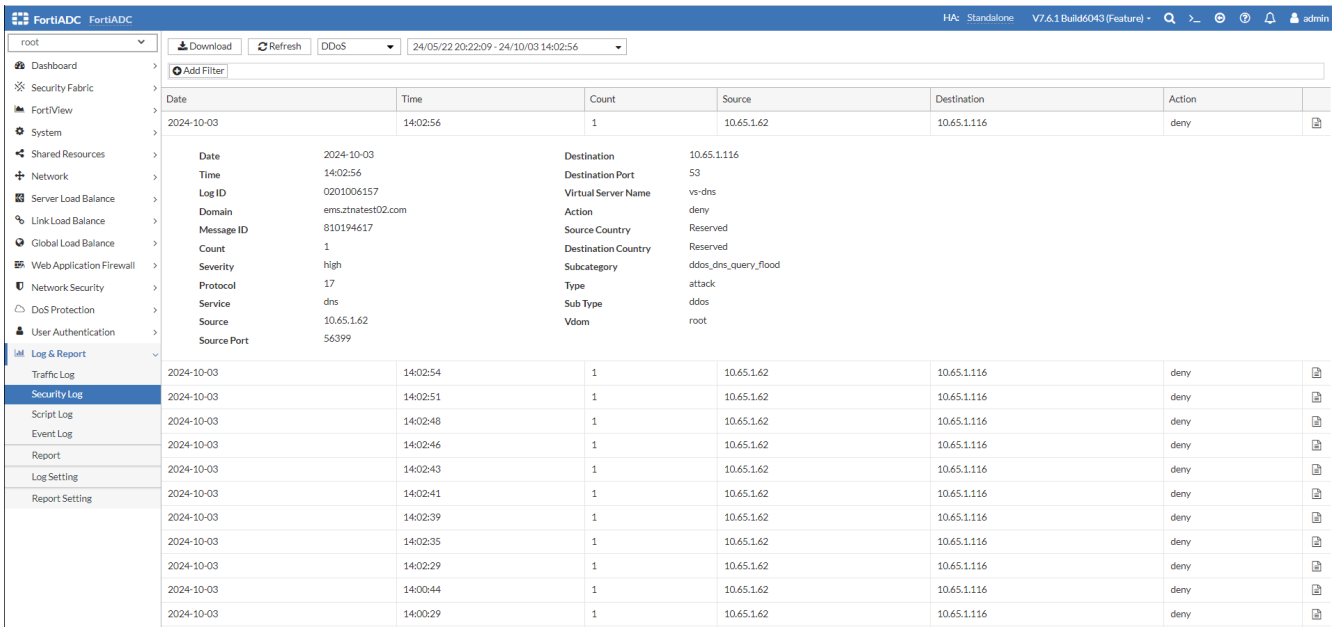
Step 4: Post-Attack Analysis and Traffic Logging

Maintaining a robust DNS infrastructure necessitates real-time monitoring and detailed post-attack analysis. This step focuses on the importance of traffic logging to capture DNS traffic patterns and attack vectors for forensic analysis. FortiADC will maintain comprehensive logs, enabling administrators to assess the effectiveness of mitigation strategies and refine policies accordingly. By monitoring security logs related to DNS Query Flood Protection and DNS Reverse Flood Protection, organizations can gain insights into attack trends and adjust their defenses proactively.

Traffic Logs:



Security Logs:



Global Server Load Balancing in Hybrid Environments

In hybrid IT architectures, balancing workloads across multiple data centers and cloud environments is critical for maintaining high availability and ensuring optimal performance. This use case examines the deployment of Global Server Load Balancing (GSLB) using FortiADC for HybridShop, a global e-commerce platform with infrastructure spanning both on-premises data centers and public cloud services. FortiADC's GSLB configuration enables dynamic traffic routing by leveraging real-time health checks, load conditions, and geographic proximity to direct user requests to the most responsive and available data centers. By integrating FortiADC for GSLB, HybridShop aims to achieve enhanced performance, disaster recovery readiness, and improved user experience through intelligent traffic management that adapts to changing conditions and mitigates the risk of service disruptions.

User Scenario: Implementing GSLB with FortiADC for a Hybrid E-Commerce Environment

HybridShop, a global e-commerce platform, utilizes a hybrid IT architecture that integrates on-premises data centers with multiple cloud providers to support its high-performance, geographically distributed operations. To optimize service availability, reduce latency, and ensure robust disaster recovery, HybridShop deploys FortiADC's Global Server Load Balancing (GSLB) solution. GSLB dynamically distributes client requests based on real-time metrics such as server health, response time, network latency, and geographic proximity. It uses DNS-based redirection and advanced traffic steering policies to direct traffic to the most suitable data center or cloud environment.

In this use case, FortiADC serves as the central control point for managing traffic across HybridShop's hybrid environment. It is strategically positioned to maintain high-bandwidth, low-latency connections to both on-premises data centers and cloud environments, ensuring seamless communication and comprehensive visibility into backend infrastructure. With active health checks and automated failover mechanisms, FortiADC effectively mitigates the risk of server overload and regional outages, enhancing overall fault tolerance and resource efficiency across HybridShop’s infrastructure.

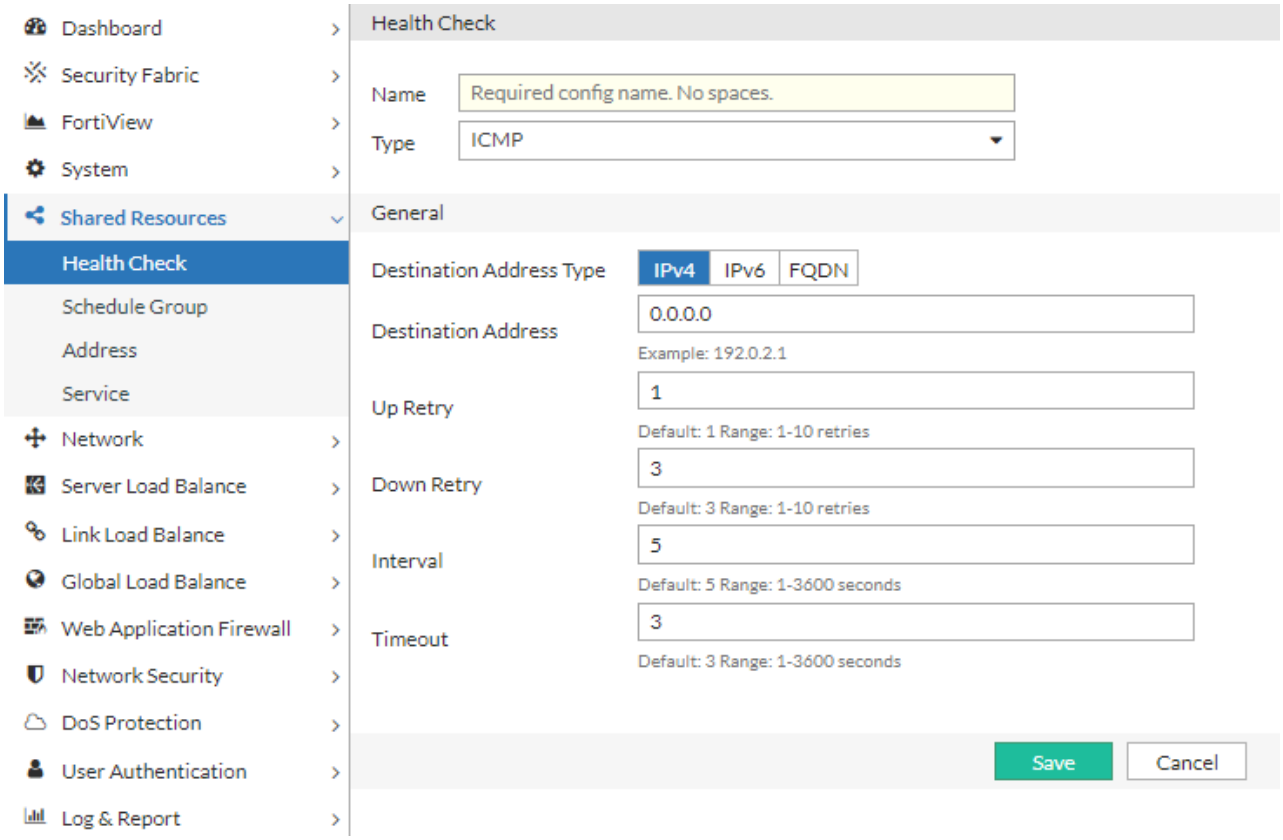
Step 1: Health Monitor Configuration

The first step in setting up GSLB with FortiADC is configuring robust health checks for all backend data centers and cloud environments. Health checks are crucial for monitoring the availability and performance of servers in real-time. FortiADC supports multiple protocols, such as HTTP(S), TCP, and ICMP, to ensure that only healthy and responsive data centers participate in load balancing. Setting appropriate thresholds for failed health checks and retries allows FortiADC to detect issues quickly and reroute traffic, ensuring continuous availability across the hybrid infrastructure.

To configure Health Check:

1. Navigate to **Shared Resources > Health Check**.
   The configuration page displays the **Health Check** tab.

2. Click **Create New** to display the configuration editor and configure the required settings for the selected **Type**.



For configuration details, see Configuring health checks on page 355.

3. Save the configuration.

Alternatively, predefined Health Check profiles can also be applied.

Step 2: Global Balancing Servers and Locations Settings

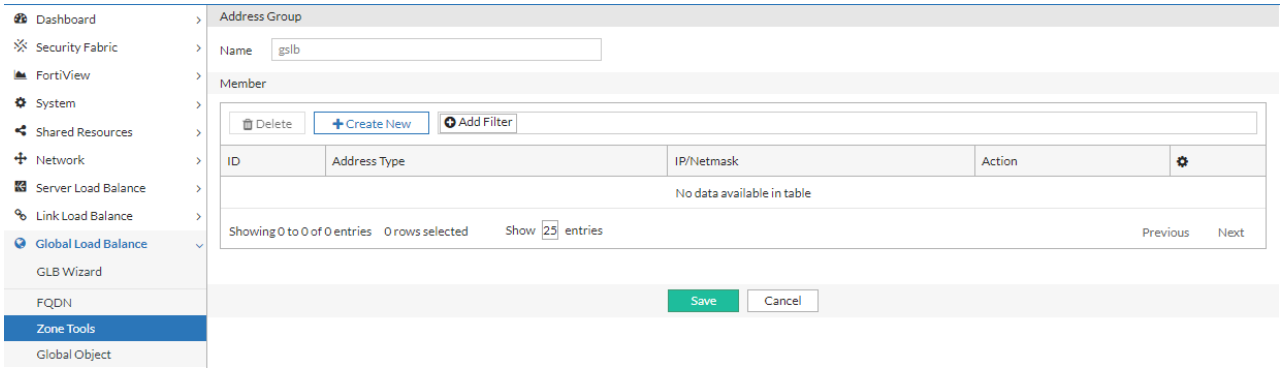
After completing the health check configurations, the next step involves defining the backend servers and their locations within the GSLB framework. This process includes configuring address groups to organize IP ranges, mapping physical and cloud-based data centers, and setting up remote DNS servers for traffic forwarding. Hosts are identified within each data center to participate in GSLB, and servers are defined using these hosts with associated IPs, ports, and health-check criteria. The virtual server pool is configured to establish load balancing and failover behaviors, while the Global DNS Policy governs traffic routing based on factors such as location and performance. Finally, general settings are fine-tuned to ensure the GSLB architecture operates efficiently within HybridShop's hybrid infrastructure.

Step 2.1: Configure Address Groups

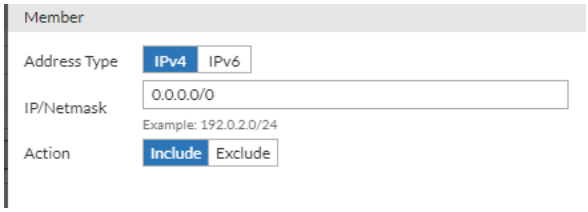
Define address groups to categorize IP ranges that represent various client locations or data center endpoints. Address groups play a critical role in HybridShop's GSLB policy by enabling DNS-based traffic steering decisions based on client proximity, network zones, or geographic distribution. FortiADC uses these groups to apply policy-based rules that ensure clients are directed to the optimal data center—whether on-premises or cloud-based—thereby enhancing response times and user experience across global markets. FortiADC provides two predefined groups: **any** for unrestricted matching and **none** for exclusions.

To configure an Address Group:

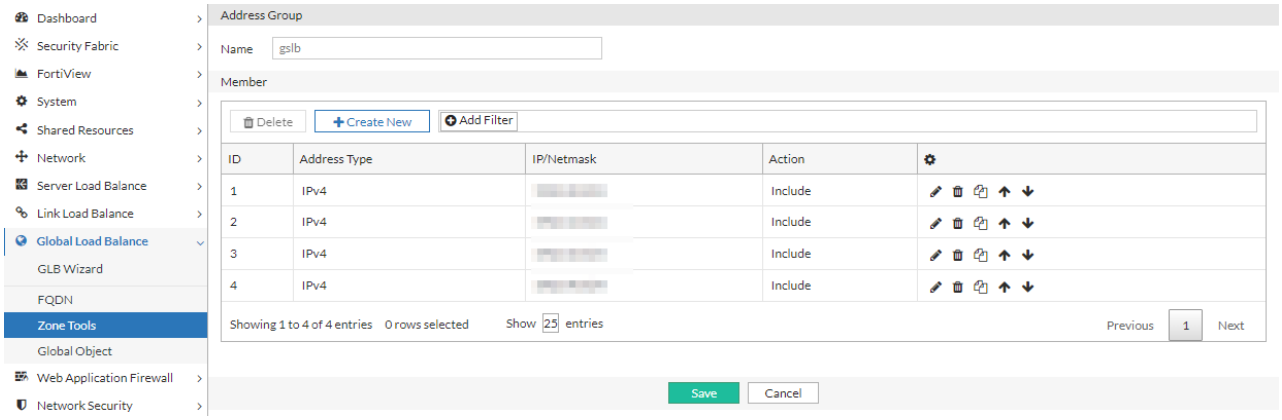
1. Navigate to **Global Load Balance > Zone Tools**. Click the **Address Group** tab.
2. Click **Create New** to display the configuration editor. Enter a name for the Address Group and click **Save** to enable configuration for the Member section.



3. Under the **Member** section, click **Create New** to add IP addresses. Save the configuration and repeat to create as many address objects as needed to include in the Address Group.



4. After you have added the desired number of IP addresses to the Address Group, click **Save** to commit the changes.



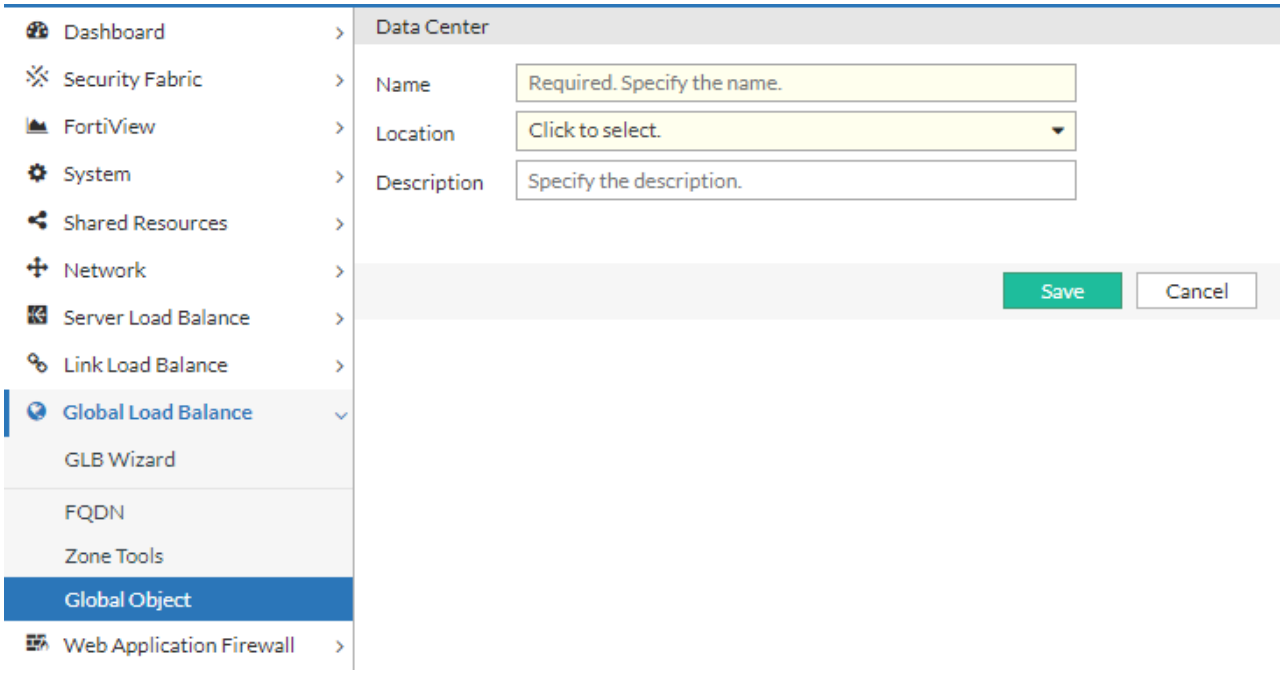
For configuration details, see Configuring an address group on page 651.

Step 2.2: Configure Data Centers

In HybridShop's architecture, data centers consist of on-premises facilities as well as cloud regions hosted by multiple providers. Define these data centers in FortiADC's GSLB configuration to establish strategic traffic distribution points. FortiADC monitors and manages traffic across all of these locations based on real-time metrics like latency and server load. By categorizing data centers geographically and functionally, HybridShop optimizes traffic flow, ensuring clients are routed to the nearest or most efficient center, improving service continuity and reducing latency.

To configure Data Center objects:

1. Navigate to **Global Load Balance > Global Object**. Click the **Data Center** tab.
2. Click **Create New** to display the configuration editor. Configure the required settings:



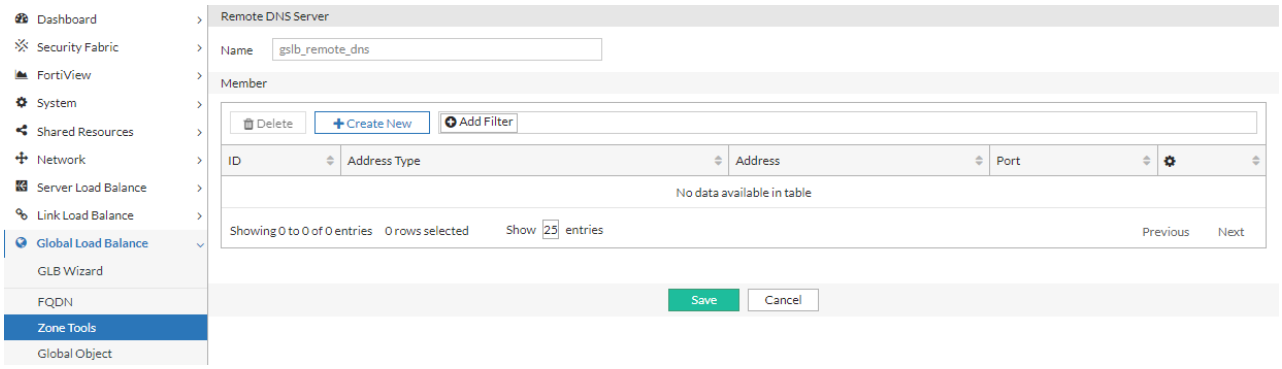
   For configuration details, see Configuring data centers on page 658.
3. Save the configuration.

Step 2.3: Configure Remote DNS Servers

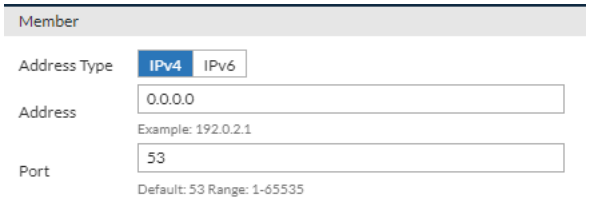
Set up remote DNS servers (also known as forwarders) to handle DNS queries that HybridShop's distributed architecture relies on. FortiADC uses these servers to redirect requests to the appropriate virtual servers, considering cloud regions and on-premises environments. This step is critical in ensuring that DNS resolution remains swift and accurate, enabling dynamic redirection based on real-time traffic conditions or network health. In addition, the DSSET list is configured to further secure the DNS traffic and authenticate the zones in HybridShop's architecture.

To configure a Remote DNS Server:

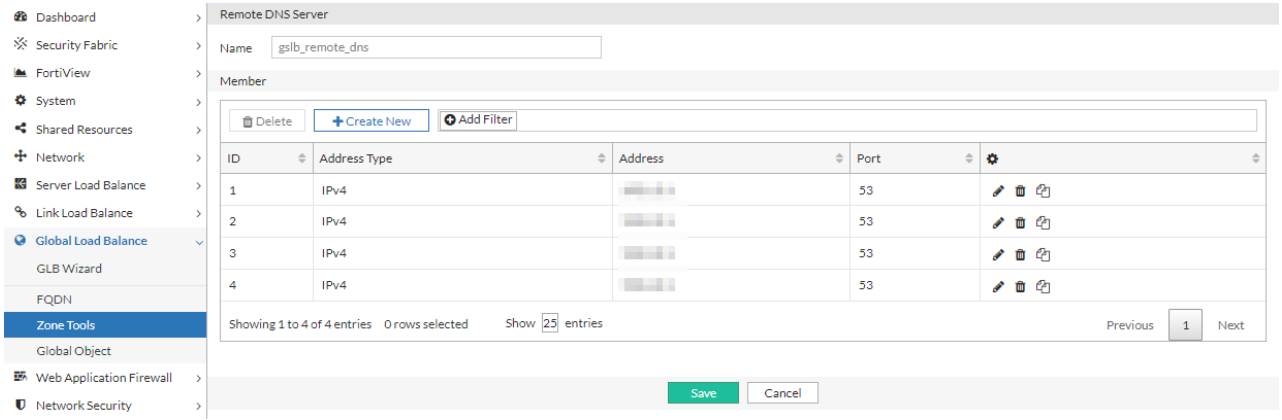
1. Navigate to **Global Load Balance > Zone Tools**. Click the **Remote DNS Server** tab.
2. Click **Create New** to display the configuration editor. Enter a name for the remote DNS Server and click **Save** to enable configuration for the Member section.



3. Under the **Member** section, click **Create New** to add the IP address of the remote DNS server. Save the configuration and repeat to create as many IP address objects as needed to include in the Remote DNS Server configuration.



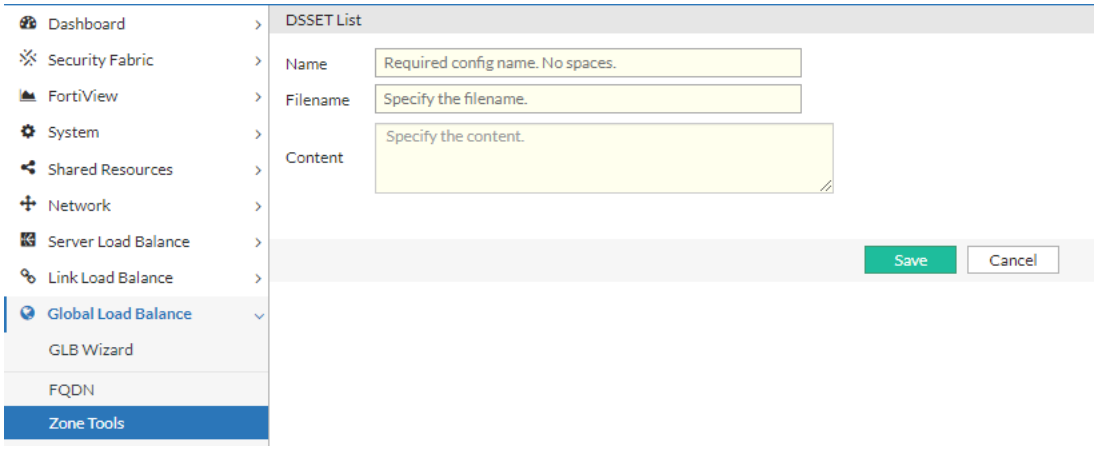
4. After you have added the desired number of IP addresses to the Remote DNS Server configuration, click **Save** to commit the changes.



For configuration details, see Configuring remote DNS servers on page 652.

To configure a DSSET List:

1. Navigate to **Global Load Balance > Zone Tools**. Click the **DSSET List** tab.
2. Click **Create New** to display the configuration editor. Configure the required settings:



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: `dns.example.com. IN DS 13447 5 1 A5AD9EFB6840F58CF817F3CC7C24A7ED2DD5559C`

For configuration details, see Configuring the DSSET list on page 650.

3. Save the configuration.

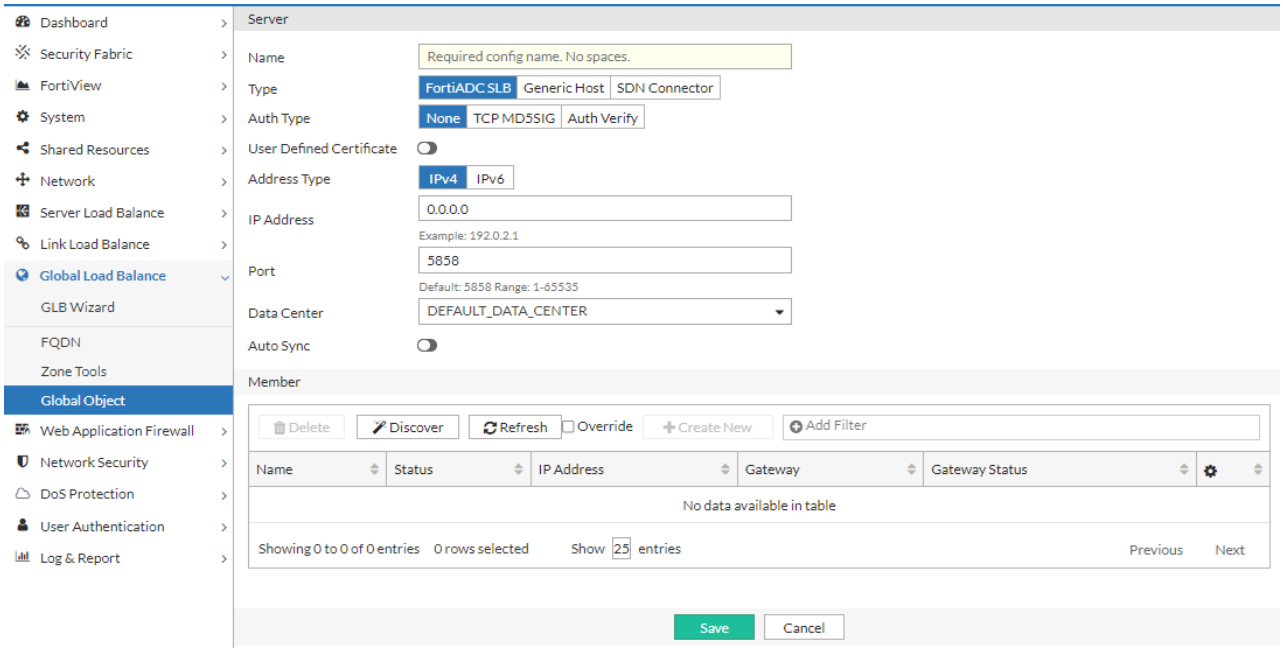
Step 2.4: Configure Servers and Server Pools

Hosts are grouped into server pools within FortiADC, which form the basis for load balancing decisions. Prior to configuring Hosts, configure the servers and server pools. For HybridShop, these server pools consist of resources from both on-premises and cloud infrastructures, ensuring that traffic can be balanced across different environments seamlessly. FortiADC allows HybridShop to define load balancing methods—such as round-robin or least connections—to distribute client requests evenly or based on specific performance metrics. Failover response settings ensure that when a server becomes unreachable, the system automatically redirects traffic to an alternative pool.

To configure a GLB Server:

1. Navigate to Global Load Balance > Global Object.
   The configuration page displays the Server tab.

2. Click Create New to display the configuration editor. Configure the required settings:



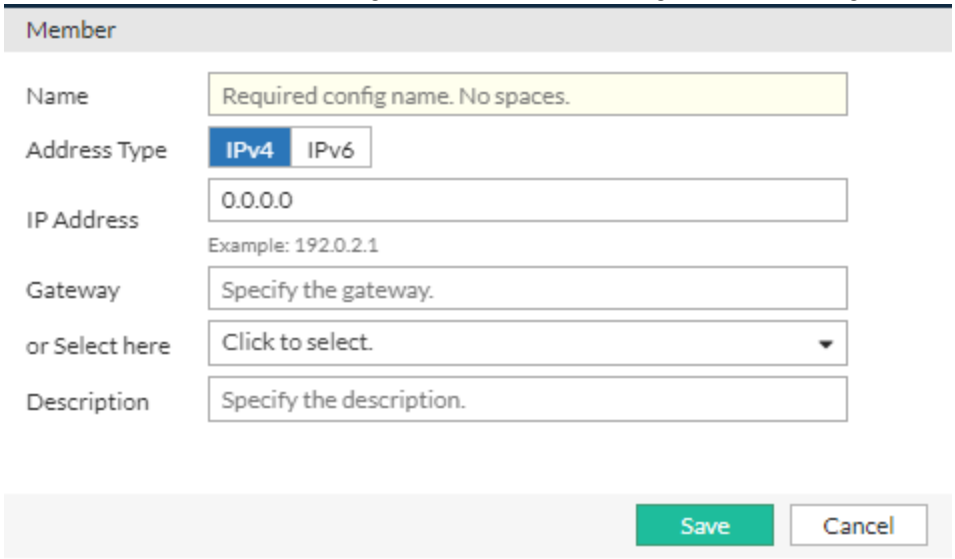
Setting	Description
Name	Configuration name. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No spaces. You reference this name in the virtual server pool configuration. **Note**: After you initially save the configuration, you cannot edit the name.
Type	Select the remote server to use for global server load balancing: • FortiADC SLB — use a FortiADC instance. • Generic Host — use a third party ADC or server. • SDN Connector — use an existing external connector that is connected to the FortiADC Security Fabric. **Note**: Currently, the SDN Connector option only supports AWS Connectors.
Auth Type	The **Auth Type** option is available if **Type** is **FortiADC SLB**. Select the authentication type: • None — No password. • TCP MD5SIG — With password, but cannot 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	The **Password** option is available if **Type** is **FortiADC SLB** and **Auth Type** is **TCP MD5SIG**. Enter the password to authenticate the key.

Setting	Description
	The password you enter here must match the password configured on the FortiADC appliance in a global sever load-balancing configuration.
User Defined Certificate	The **User Defined Certificate** option is available if **Type** is **FortiADC SLB**. Enable to use a self-defined certificate for authentication.
Certificate	The **Certificate** option is available if **Type** is **FortiADC SLB** and **User Defined Certificate** is **enabled**. Select the local certificate object to use for the GSLB server.
Address Type	The **Address Type** option is available if **Type** is **FortiADC SLB**. IPv4 or IPv6.
IP Address/IPv6 Address	The **IP Address** or **IPv6 Address** option is available if **Type** is **FortiADC SLB**. Specify the IPv4 or IPv6 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 responses.
Port	The **Port** option is available if **Type** is **FortiADC SLB**. Specify the port. Default: 5858 Range: 1-65535.
SDN Connector	The **SDN Connector** option is available if **Type** is **SDN Connector**. Select the SDN Connector to synchronize to the GSLB server. For public SDN type servers, GSLB can update the public IP dynamically. **Note**: Currently, only AWS connectors are supported.
Use SDN Private IP	The **Use SDN Private IP** option is available if **Type** is **SDN Connector**. Enable to use the SDN Private IP address.
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 automatic synchronization with the remote server. When enabled, Global load balancing will synchronize automatically with the server member. If auto-sync is enabled for SDN Connector type servers, all instances from the SDN connector will be added as server members. **Note**: When disabling auto-sync, the server member will be cleared and re-synced.
Health Check Control	The **Health Check Control** option is available if **Type** is **Generic Host** or **SDN Connector**. 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.

Setting	Description
	Note: Health checking is built-in, and you can optionally configure a gateway health check.
Health Check Relationship	The **Health Check Relationship** option is available if **Type** is **Generic Host** and **Health Check Control** is **enabled**. • 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	The **Health Check List** option is available if **Type** is **Generic Host** and **Health Check Control** is **enabled**. Select one or more health check configuration objects.

For configuration details, see Configuring servers on page 653.

3. Click **Save** to enable configuration for the Member section.
4. Under the **Member** section, configure the Member list configuration according to the GLB server type.



Setting	Description
FortiADC SLB	
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.
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.

Setting	Description
Name	Must match the virtual server configuration name on the local FortiADC.
Address Type	IPv4 or IPv6.
IP Address/IPv6 Address	Virtual server IPv4 or IPv6 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.
Generic Host	
Health Check Inherit	Enable to inherit the health check settings from the parent configuration. The Health Check Inherit option is enabled by default. Disable to specify health check settings in this member configuration.
Health Check Control	The **Health Check Control** option is available if **Health Check Inherit** is **disabled**. Enable health checking for the virtual server.
Health Check Relationship	The **Health Check Relationship** option is available if **Health Check Inherit** is **disabled** and **Health Check Control** is **enabled**. • 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	The **Health Check List** option is available if **Health Check Inherit** is **disabled** and **Health Check Control** is **enabled** Specify one or more health check configuration objects.
SDN Connector	
SDN Instance	Select an instance from the SDN's instance list.
Health Check Inherit	Enable to inherit the health check settings from the parent configuration. The Health Check Inherit option is enabled by default. Disable to specify health check settings in this member configuration.
Health Check Control	The **Health Check Control** option is available if **Health Check Inherit** is **disabled**. Enable health checking for the virtual server.
Health Check Relationship	The **Health Check Relationship** option is available if **Health Check Inherit** is **disabled** and **Health Check Control** is **enabled**. • 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	The **Health Check List** option is available if **Health Check Inherit** is **disabled** and **Health Check Control** is **enabled**

Setting	Description
	Specify one or more health check configuration objects.

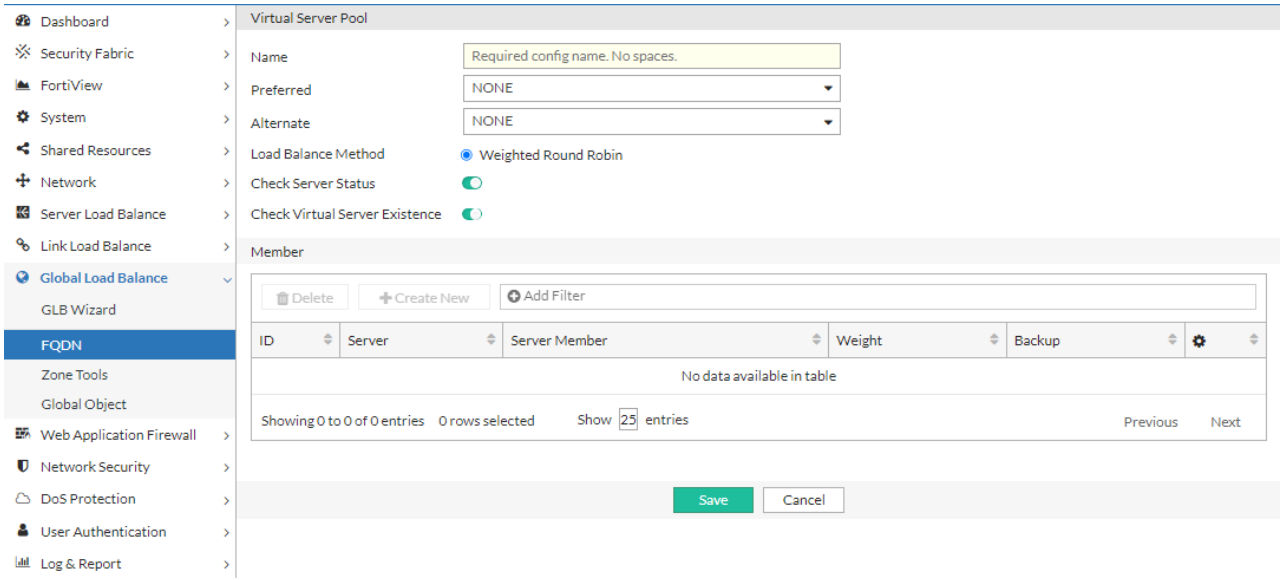
Note:

- If Health Check is disabled for SDN connector server members, their health check status will always appear as available since the status cannot be verified through health check.
- If the SDN instance is changed in the SDN connector, the SDN connector server member status will not be affected. For example, if the EC2 instance is terminated in AWS, that instance will still remain a server member in the SDN connector GLB server.

5. After you have added the desired number of Members to the Server configuration, click **Save** to commit the changes.

To configure a Virtual Server Pool:

1. Navigate to **Global Load Balance > FQDN**. Click the **Virtual Server Pool** tab.
2. Click **Create New** to display the configuration editor. Configure the required settings:

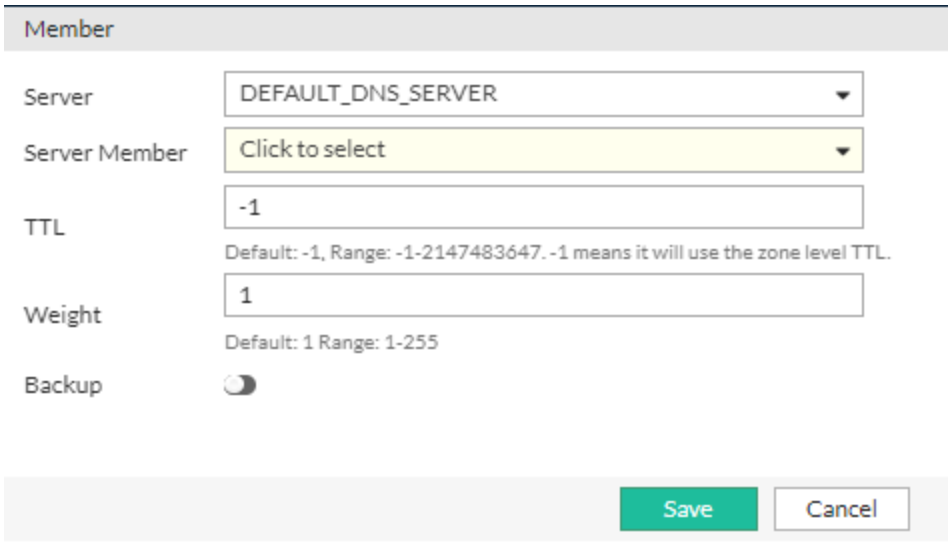


Settings	Guidelines
Name	Specify a unique name for the virtual server pool configuration. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No spaces. You reference this name in the host configuration. **Note**: After you initially save the configuration, you cannot edit the name.
Preferred	• None—No 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.

Settings	Guidelines
	• RTT—Virtual servers with the shortest latency link or closest to the data center will respond. • Least-Connnections—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. • Bytes-Per-Second—Virtual servers with the lowest traffic will respond. • Server-Performance—Virtual servers with better server-performancec in the CPU or Memory (whichever one you give more weight to) will respond.
Alternate	• 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-Connnections—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. • Bytes-Per-Second—Virtual servers with the lowest traffic will respond. • Server-Performance—Virtual servers with better server-performancec in the CPU or Memory (whichever one you give more weight to) will respond.
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.

For configuration details, see Configuring virtual server pools on page 621.

3. Save the configuration to enable configuration for the **Member** section.
4. Under the Member section, click **Create New** to display the configuration editor. Configure the required settings and save the configuration.



Settings	Guidelines
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.

5. After you have added the desired number of Members to the Virtual Server Pool configuration, click **Save** to commit the changes.

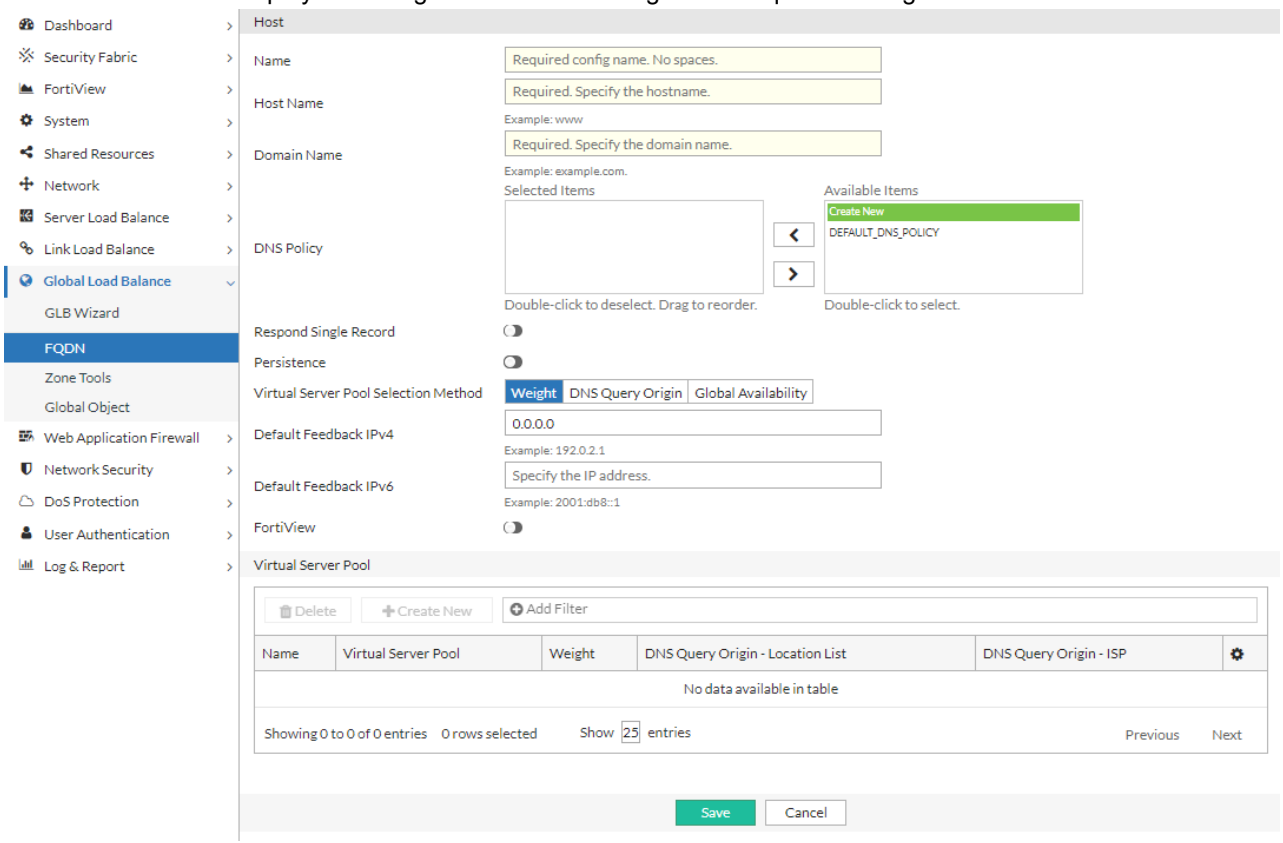
Step 2.5: Define Hosts (Physical and Cloud-Based Data Centers)

Hosts are the servers or virtual machines that handle incoming traffic within HybridShop's infrastructure. These hosts can reside either in local on-premises data centers or across cloud environments such as AWS or Azure. In this step, administrators must define each host by specifying its IP address, operational role, and health-check configuration. FortiADC's GSLB framework constantly monitors these hosts, ensuring that only healthy and responsive servers receive traffic, thus preventing client requests from reaching degraded or unavailable servers.

To configure a host:

1. Navigate to **Global Load Balance > FQDN**.
   The configuration page displays the **Host** tab.

2. Click **Create New** to display the configuration editor. Configure the required settings:

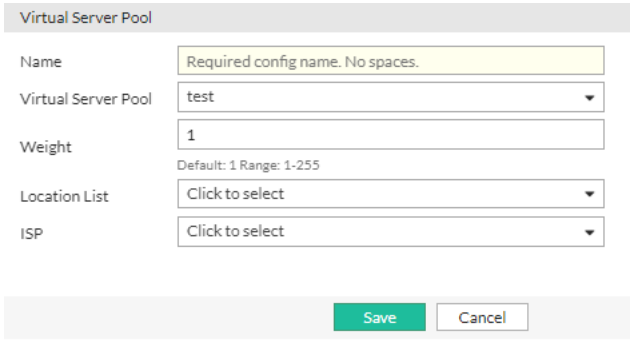


Settings	Guidelines
Name	Configuration name. Valid characters are `A-Z`, `a-z`, `0-9`, `_`, and `-`. No spaces. **Note**: After you initially save the configuration, you cannot edit the name.
Host Name	The hostname part of the FQDN, such as `www`. **Note**: You can specify the @ symbol to denote the zone root. The value substituted for @ is the preceding $ORIGIN directive.
Domain Name	The domain name must end with a period. For example: `example.com`.
DNS Policy	Select one or more DNS policies 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.

Settings	Guidelines
	• 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.

For configuration details, see Configuring hosts on page 619.

3. Save the configuration to enable configuration for the **Virtual Server Pool** section.
4. Under the Virtual Server Pool section, click **Create New** to display the configuration editor. Configure the required settings and save the configuration.



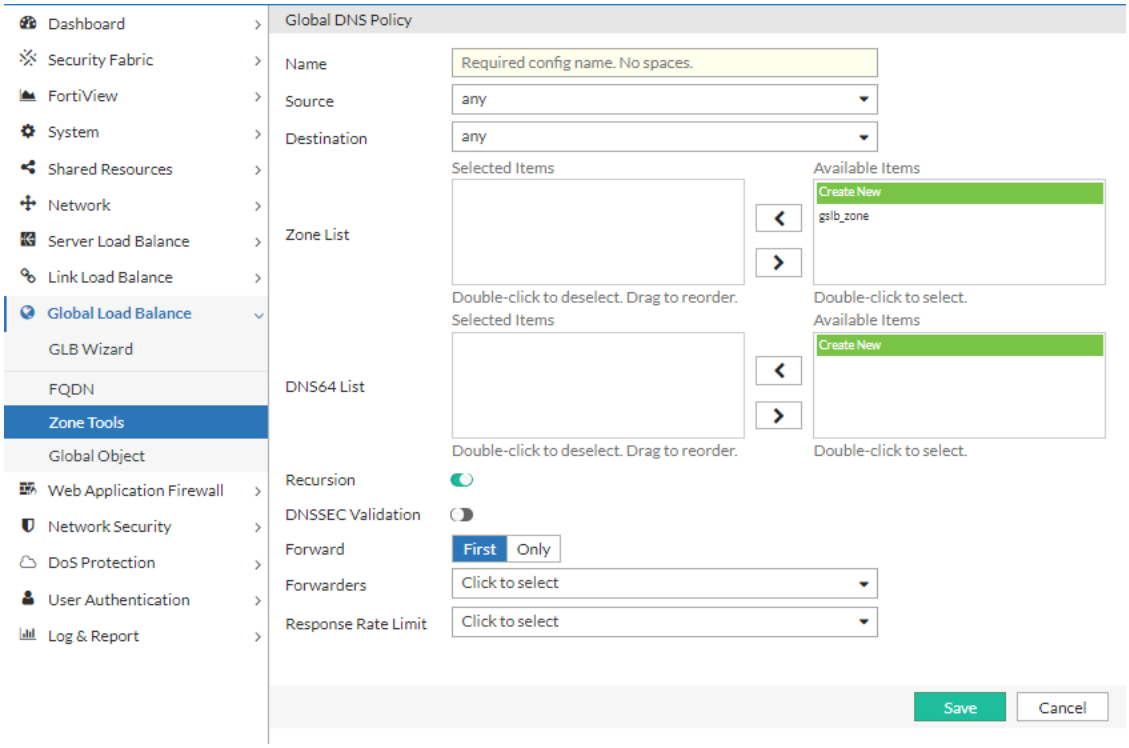
Settings	Guidelines
Name	Enter the mkey.
Virtual Server Pool	Select a virtual server pool from the list, or create a new one.
Weight	Assign a weight. Valid values range from 1 to 255.
Topology	Select a topology from the list, or create a new one.
ISP	Select an ISP from the list or create a new one.

Step 2.6: Configure Global DNS Policy

The Global DNS Policy is where HybridShop's traffic steering strategies are implemented. By mapping address groups to specific data centers or cloud regions, FortiADC's DNS policies enable precise control over how and where requests are routed. HybridShop leverages advanced traffic routing rules, such as geographic proximity and performance-based steering, to ensure that customers from different regions are directed to the data center best suited to handle their requests, optimizing both speed and resource utilization.

To the Global DNS Policy:

1. Navigate to **Global Load Balance > Zone Tools**. Click the **Global DNS Policy** tab.
2. Click **Create New** to display the configuration editor. Configure the required settings:



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.
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 Validation	Enables/disables DNSSEC validation.
Forward	• First—The DNS server queries the forwarders list before performing its own DNS lookup. • Only—Only queries the forwarders list. Does not perform its own DNS lookups.

Settings	Guidelines
	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.

For configuration details, see Configuring a Global DNS policy on page 626.

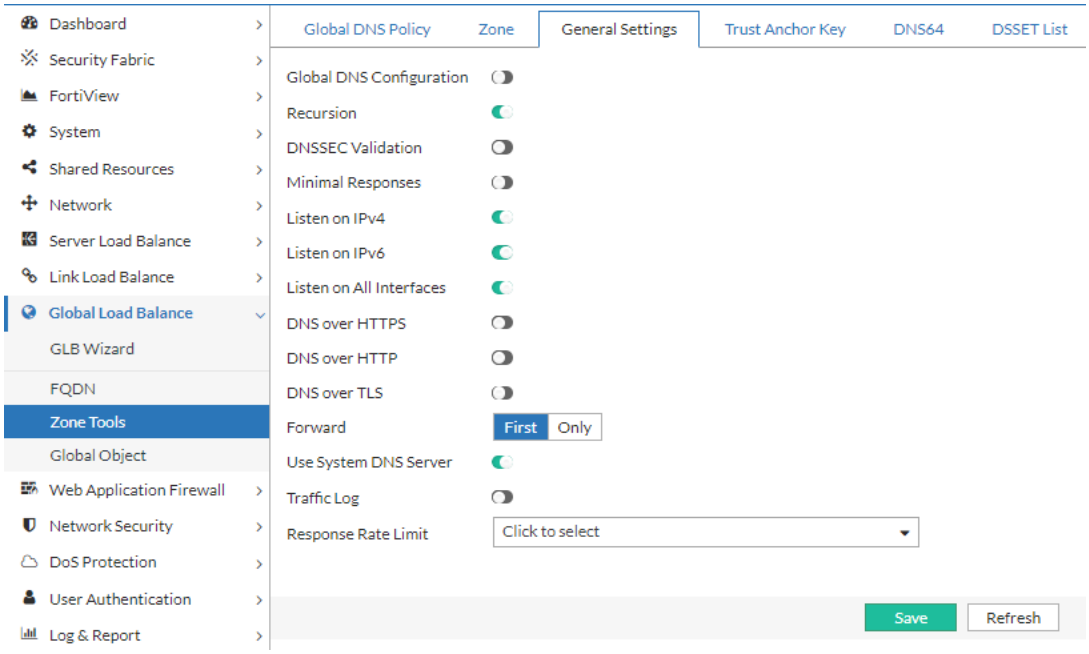
3. Save the configuration.

Step 2.7: Configure General Settings

Configure general DNS settings that apply when no specific policy match is found. For HybridShop, this ensures default behavior that directs client traffic to the nearest available data center, maintaining optimal performance even when more complex policies are not triggered. Additionally, DNSSEC settings may be configured to ensure the integrity and security of DNS queries within HybridShop's globally distributed environment, mitigating potential threats and ensuring data authenticity.

To configure general DNS settings:

1. Navigate to **Global Load Balance > Zone Tools**. Click the **General Settings** tab.
2. Configure the settings:



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 Validation	Enables/disables DNSSEC validation.
Minimal Responses	Enables/disables Minimal Responses to hide the **Authority Section** and **Additional Section** of DNS queries. When the DNS query only shows the minimal response, it can significantly increase the performance of the FortiADC DNS service by increasing the QPS.
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.
Interface List	The **Interface List** option is available if **Listen on All Interface** is disabled. If not listening to all interfaces, select one or more ports to listen on.
DNS over HTTPS	Enables/disables DNS over HTTPS to encrypt DNS queries using the HTTPS protocol.
DNS over HTTPS Port	The **DNS over HTTPS Port** option is available if **DNS over HTTPS** is enabled. Specify the port to listen on DNS over HTTPS. Default: 443 Range: 1-65535.
DNS over HTTPS Interface List	The **DNS over HTTPS Interface List** option is available if **DNS over HTTPS** is enabled. Select the interface(s) to listen on for DNS over HTTPS.
DNS over HTTP	Enables/disables DNS over HTTP to encrypt DNS queries using the HTTP protocol.
DNS over HTTP Port	The **DNS over HTTP Port** option is available if **DNS over HTTP** is enabled. Specify the port to listen on DNS over HTTP. Default: 80 Range: 1-65535.
DNS over HTTP Interface List	The **DNS over HTTP Interface List** option is available if **DNS over HTTP** is enabled. Select the interface(s) to listen on for DNS over HTTP.
DNS over TLS	Enables/disables DNS over TLS to encrypt DNS queries using the TLS protocol.
DNS over TLS Port	The **DNS over TLS Port** option is available if **DNS over TLS** is enabled. Specify the port to listen on DNS over TLS. Default: 853 Range: 1-65535.
DNS over TLS Interface List	The **DNS over TLS Interface List** option is available if **DNS over TLS** is enabled. Select the interface(s) to listen on for DNS queries for DNS over TLS.
Certificate	The **Certificate** option is available if **DNS over HTTPS** or **DNS over TLS** is enabled. Select the certificate object to apply for DNS over HTTPS or DNS over TLS. This certificate must refer to the DNS server domain or IP address. For details, see Configuring DNS over HTTPS and DNS over TLS on page 636.

Settings	Guidelines
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.

For configuration details, see Configuring general settings on page 634.

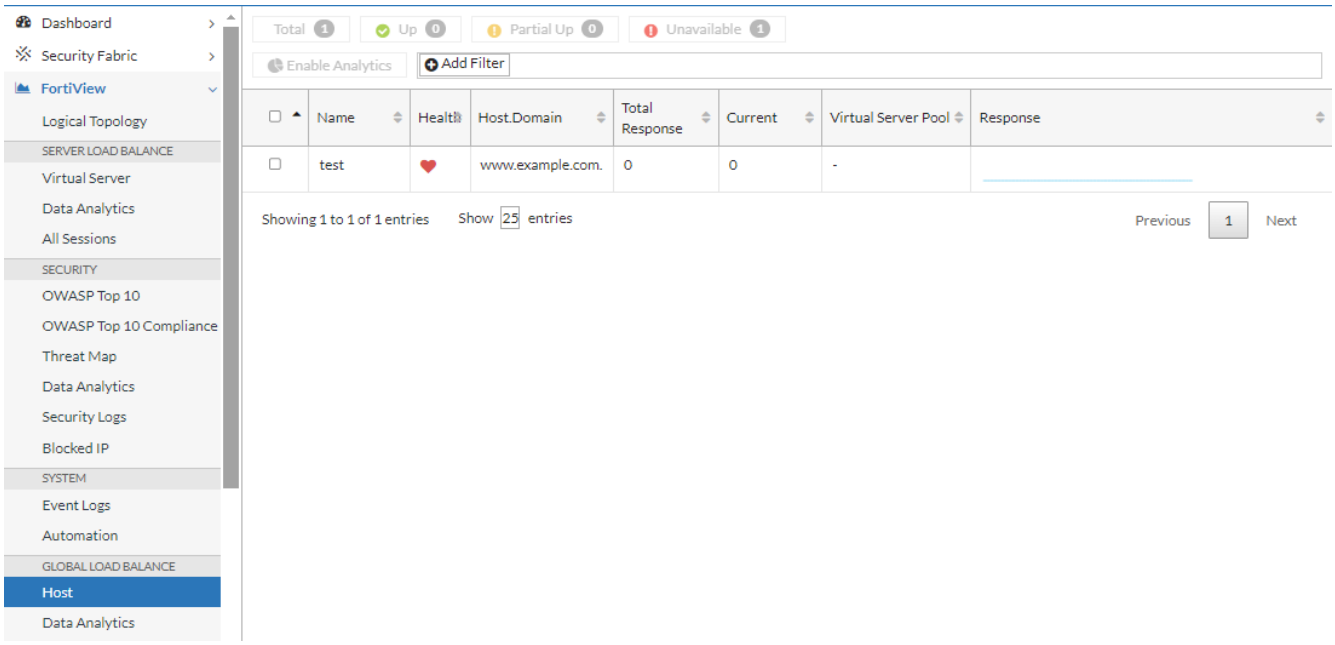
3. Save the configuration.

Step 3: Monitoring Traffic and Testing Failover

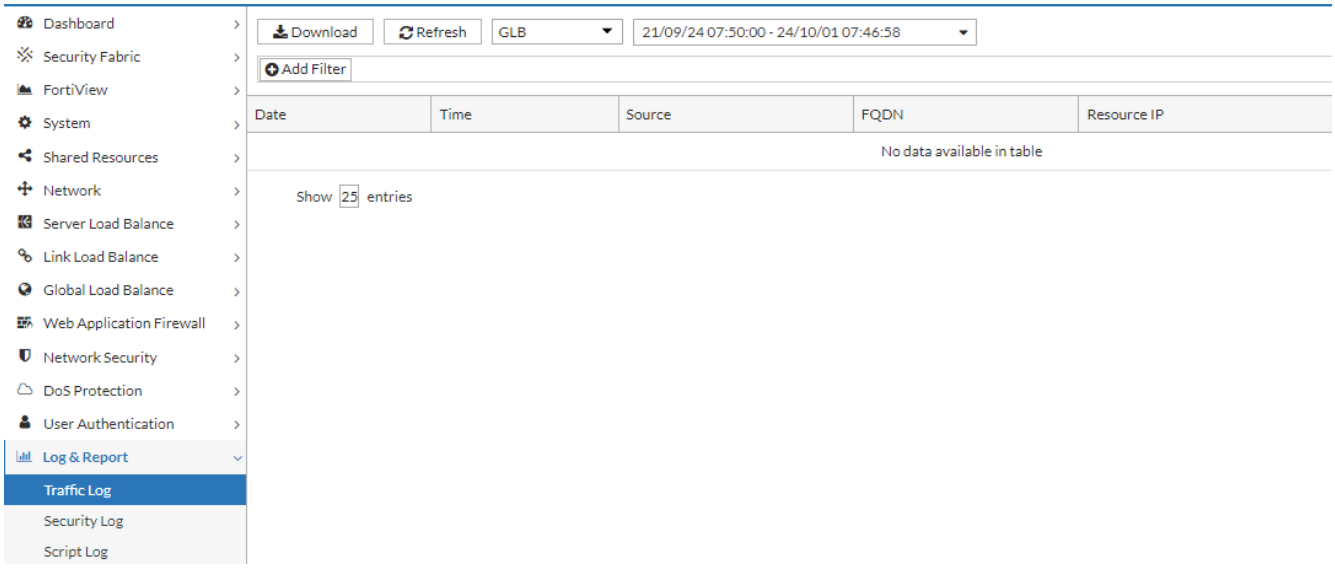
Once the configuration is complete, it is vital to monitor the traffic flow and test failover mechanisms. FortiADC's monitoring tools provide real-time insights into traffic distribution, server health, and load balancing performance. Simulating failover scenarios, such as data center outages or high traffic loads, helps verify that traffic is correctly rerouted to backup servers or data centers, ensuring the reliability and resilience of the GSLB setup. Testing these failover policies regularly guarantees that FortiADC can handle disaster recovery scenarios seamlessly, maintaining high availability for your hybrid environment.

Monitor GSLB traffic flow via FortiView and Traffic Log

FortiView:



Traffic Log:



Testing the Failover Mechanism

To verify the GSLB failover configuration, simulate a failover by causing a virtual server pool member to fail its health check. This should trigger the traffic to be seamlessly redirected to another pool member, ensuring proper failover functionality. Using tools such as "nslookup" or "dig", you can send a DNS query, and then collect and view its packet data using PCAP tools such as "wireshark".
In this example, we will use "nslookup".

1. Initiate a DNS query and record the IP address returned by the GSLB service.
2. Adjust the health check configuration for the designated virtual server pool member, setting its status to "down".
3. Inspect traffic logs to confirm that the virtual server pool member is marked as "down" and that the health check status reflects a failure.
4. Issue a subsequent DNS query and verify that the response returns a different IP address, indicating successful failover to another pool member.

Selecting the Optimal Mode: VDOM vs. ADOM

In FortiADC deployments, administrators often face the decision of whether to use Virtual Domains (VDOMs) or Administrative Domains (ADOMs) to manage and isolate network resources effectively. VDOMs allow the creation of multiple, independent virtual instances of FortiADC, each with its own configuration, routing, and policies, offering granular control over traffic segmentation and security enforcement. ADOMs, on the other hand, provide a logical separation of administrative tasks, allowing centralized management across multiple FortiADC instances without affecting their underlying configurations.

Choosing between VDOMs and ADOMs is crucial for optimizing resource allocation, maintaining security, and ensuring scalability in complex environments. This technical guide examines the core functionality of VDOMs and ADOMs in FortiADC, and provides recommendations on when to deploy each based on factors such as system architecture, administrative hierarchy, and traffic management requirements.

For details, navigate to the respective sections for Use-case for VDOM on page 1019 and Use-case for ADOM on page 1029.

Use-case for VDOM

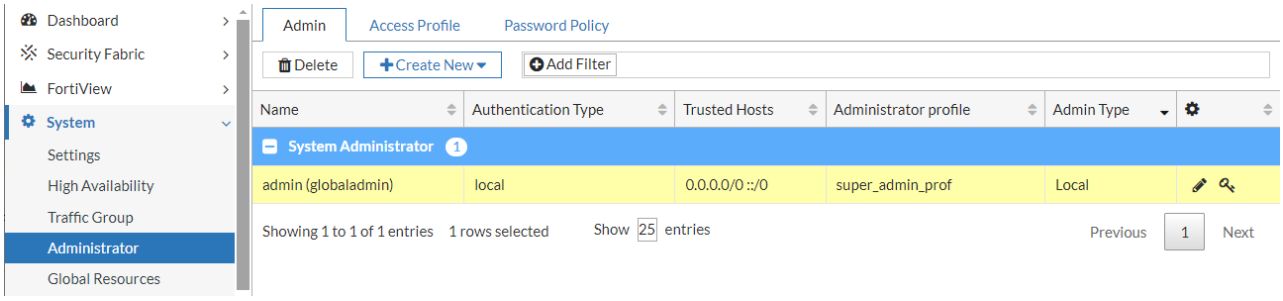
Virtual Domains (VDOMs) in FortiADC provide the ability to partition a single appliance into multiple independent virtual instances, ideal for environments requiring traffic segregation and resource isolation. Each VDOM operates as a separate entity, with its own routing, security policies, and configurations. This allows for greater flexibility and isolation in environments where multiple tenants or distinct security zones must be maintained. VDOMs are ideal for organizations needing to consolidate hardware while ensuring isolated management and traffic control across different networks or clients. For instance, a service provider hosting multiple customers with separate application environments can utilize VDOMs to isolate customer traffic and configurations from each other.

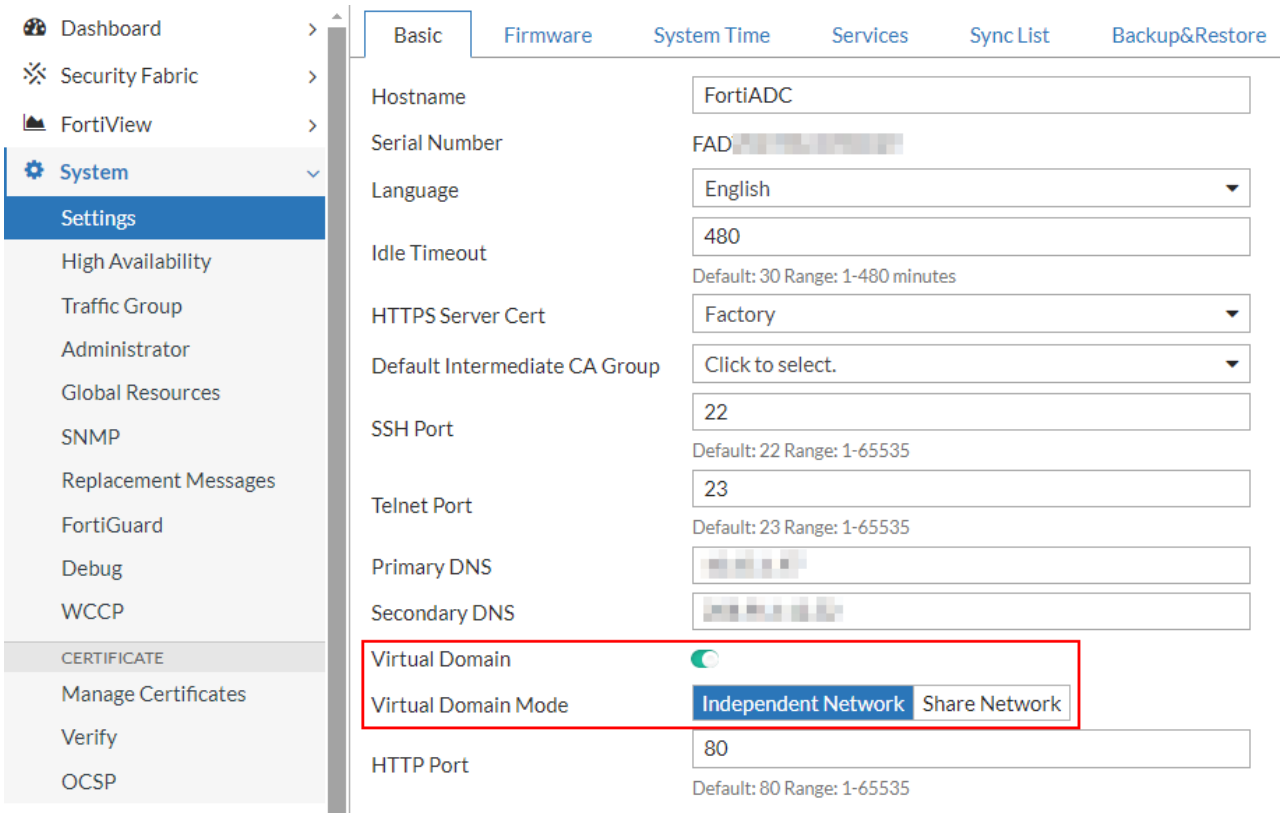
In this use-case, each VDOM is configured with its own interfaces, routing, and address spaces, ensuring that no traffic crosses between VDOMs. This method supports unique address spaces, allowing even duplicate IP addresses across different VDOMs without conflict.

Scenario: Isolating Multiple Customer Environments

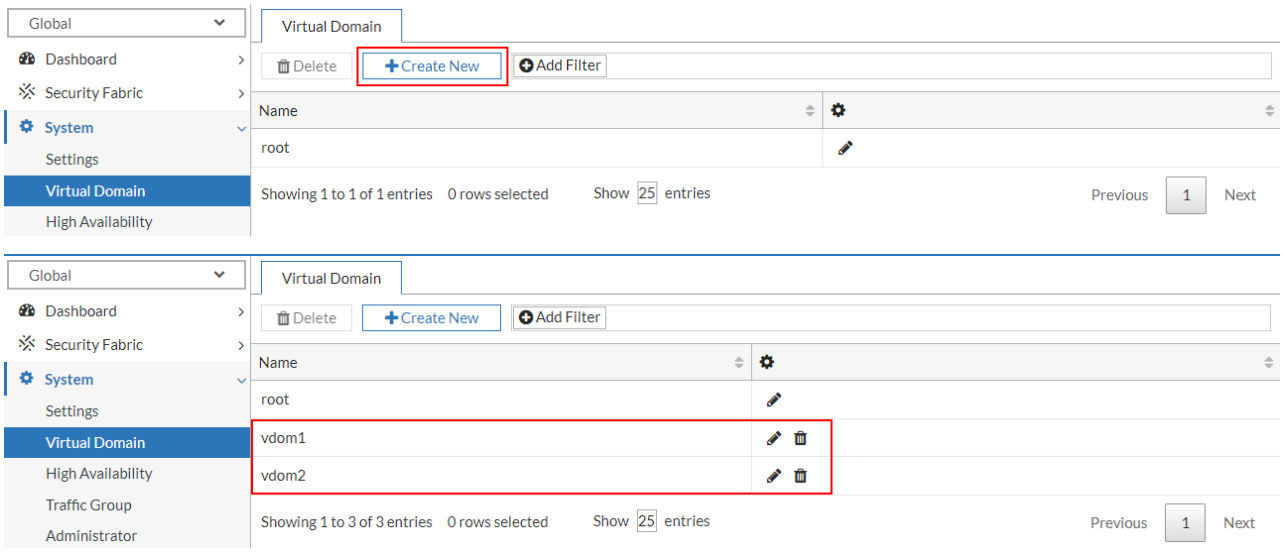
A service provider services multiple customers, where each customer deploys their own applications. To isolate customer traffic from each other, VDOMs can be configured as follows:

1. Log in as super admin and enable Virtual Domains in Independent Network mode.

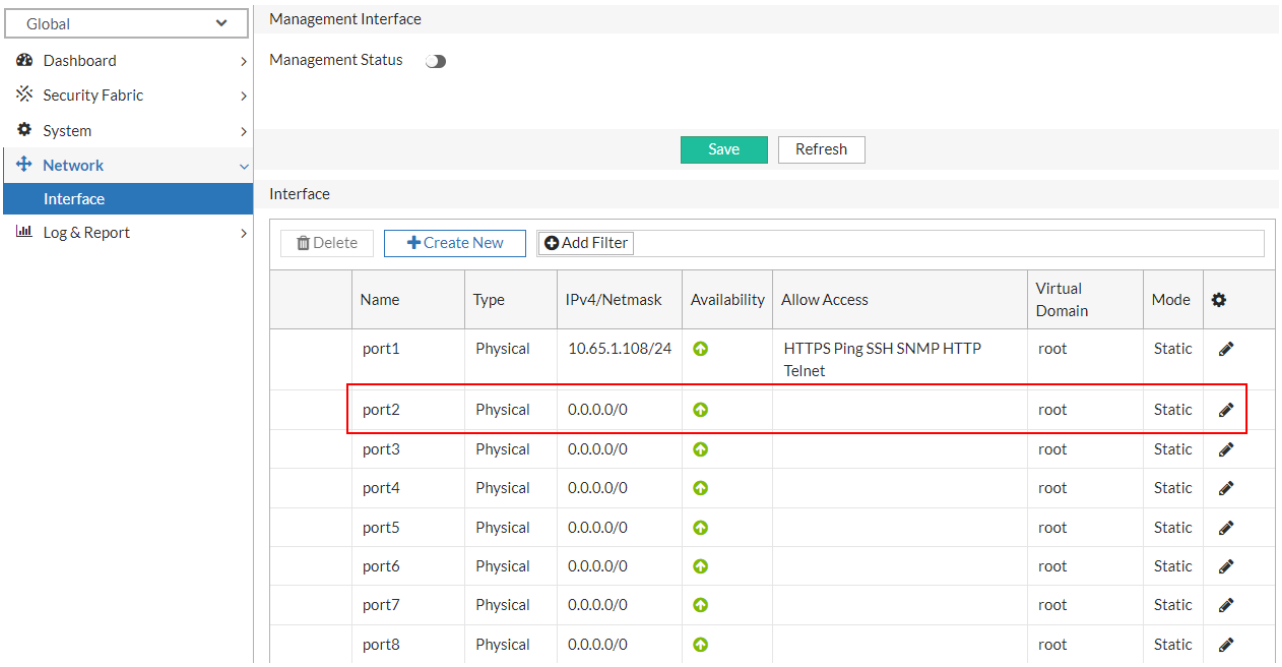


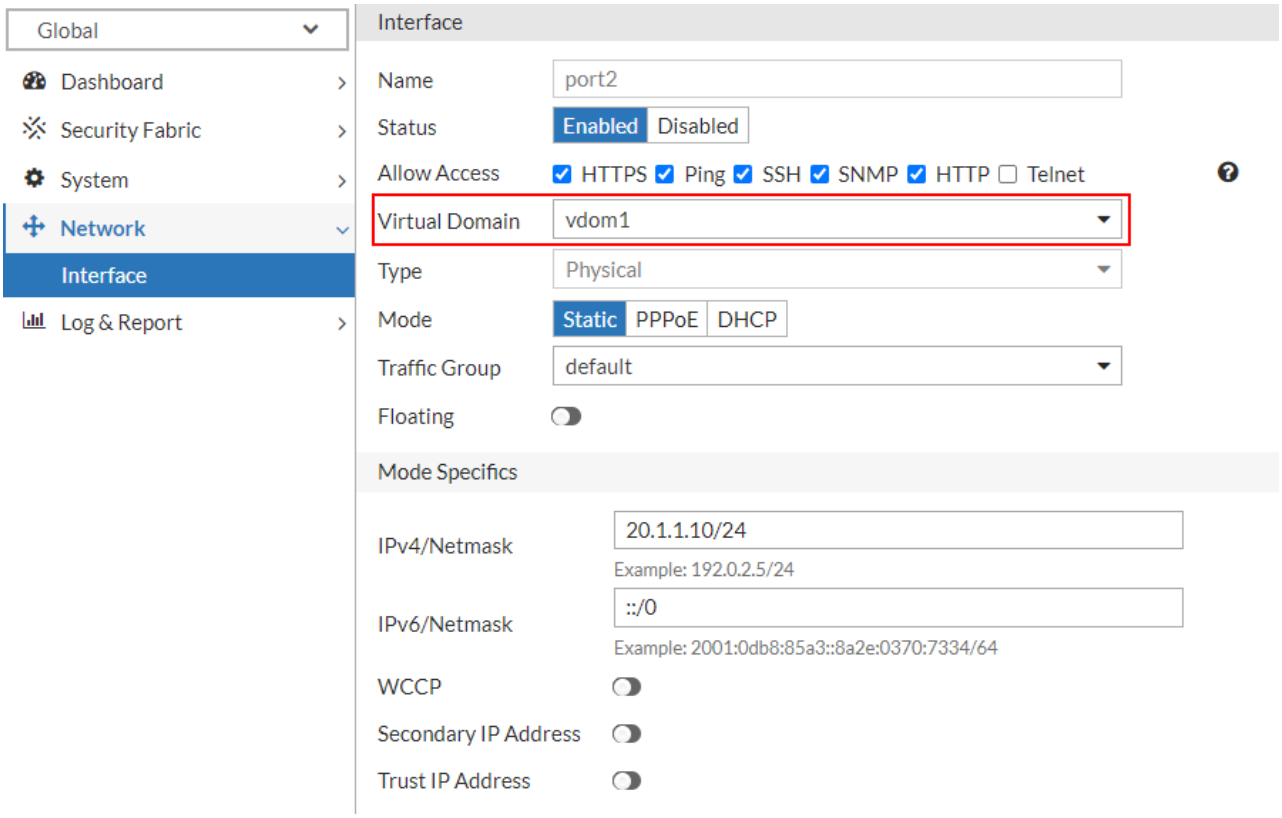


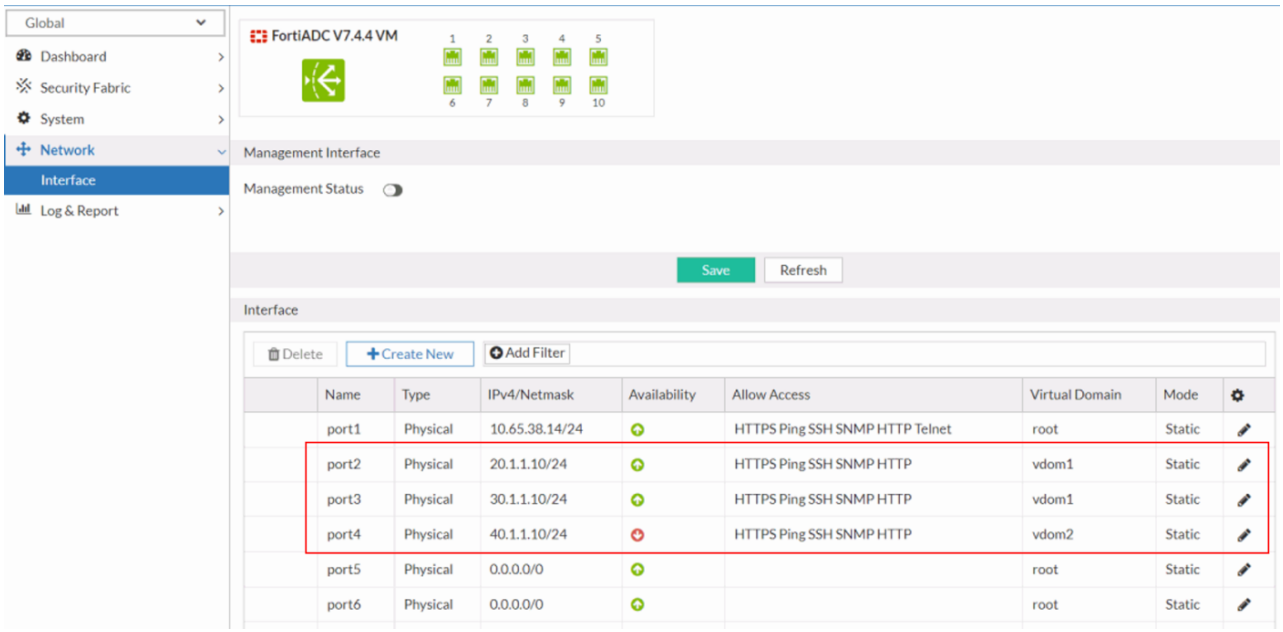
2. Provision separate VDOMs to host application environments for distinct customers. In this example, two VDOM instances—**vdom1** and **vdom2**—were configured to isolate and manage traffic for **customer1** and **customer2**, respectively.



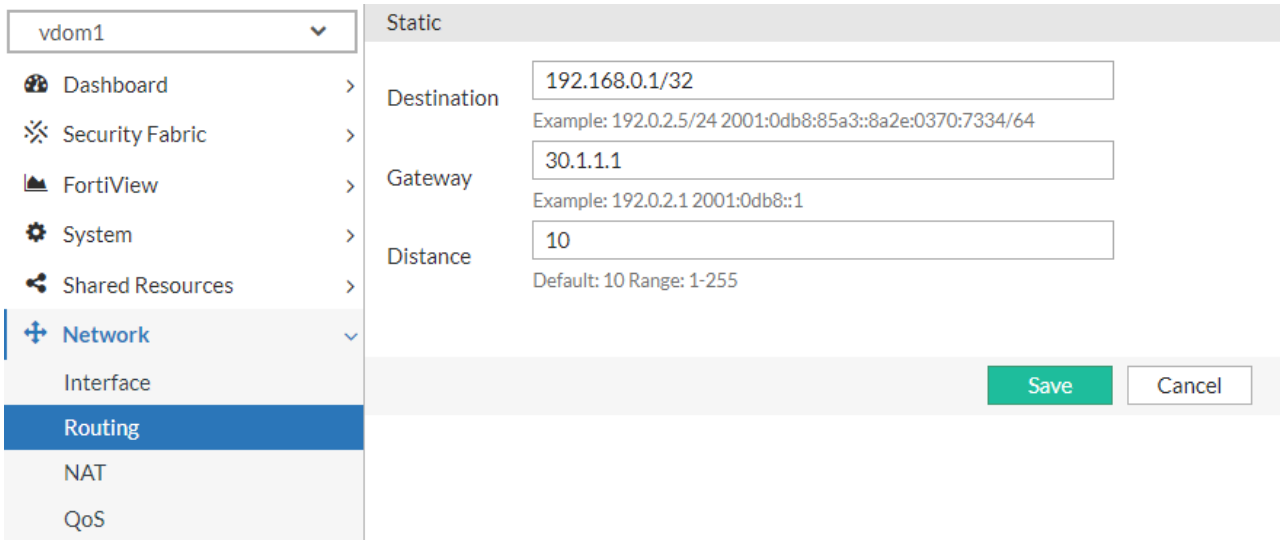
3. Allocate network interfaces to each VDOM and configure the corresponding IP addressing schemes. Ensure that each interface is bound to its respective VDOM with appropriate IP address assignments to facilitate traffic routing and network segmentation.

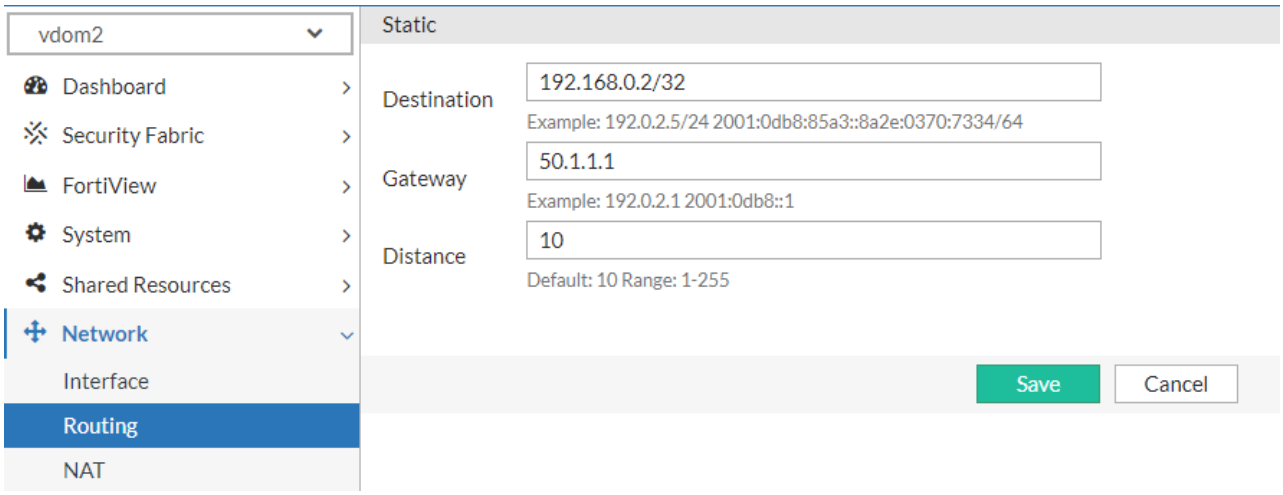

Global
Dashboard
Security Fabric
System
Network
Interface
Log & Report
Management Interface
Management Status
Save
Refresh
Interface
Delete
Create New
Add Filter
Name
Type
IPv4/Netmask
Availability
Allow Access
Virtual Domain
Mode
port1 Physical 10.65.1.108/24 HTTPS Ping SSH SNMP HTTP Telnet root Static
port2 Physical 0.0.0.0/0 root Static
port3 Physical 0.0.0.0/0 root Static
port4 Physical 0.0.0.0/0 root Static
port5 Physical 0.0.0.0/0 root Static
port6 Physical 0.0.0.0/0 root Static
port7 Physical 0.0.0.0/0 root Static
port8 Physical 0.0.0.0/0 root Static


Global
Dashboard
Security Fabric
System
Network
Interface
Log & Report
Interface
Name port2
Status Enabled Disabled
Allow Access HTTPS Ping SSH SNMP HTTP Telnet
Virtual Domain vdom1
Type Physical
Mode Static PPPoE DHCP
Traffic Group default
Floating
Mode Specifics
IPv4/Netmask 20.1.1.10/24
Example: 192.0.2.5/24
IPv6/Netmask ::/0
Example: 2001:0db8:85a3::8a2e:0370:7334/64
WCCP
Secondary IP Address
Trust IP Address

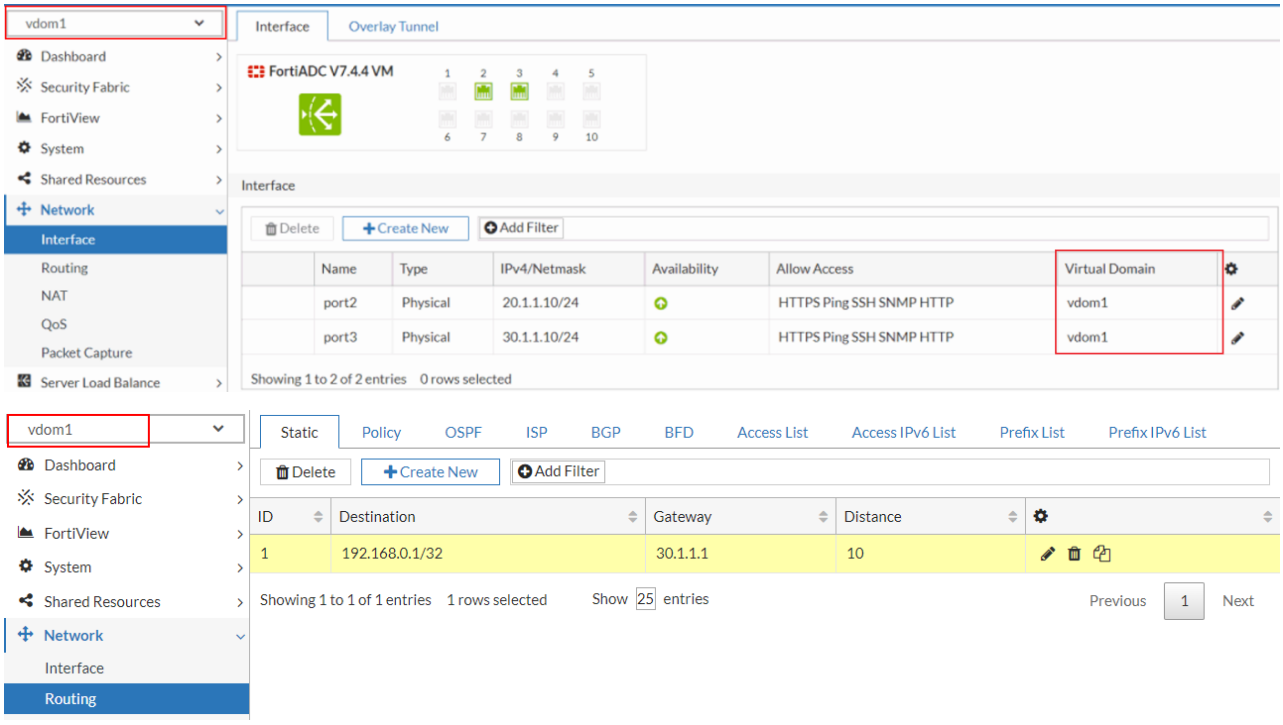


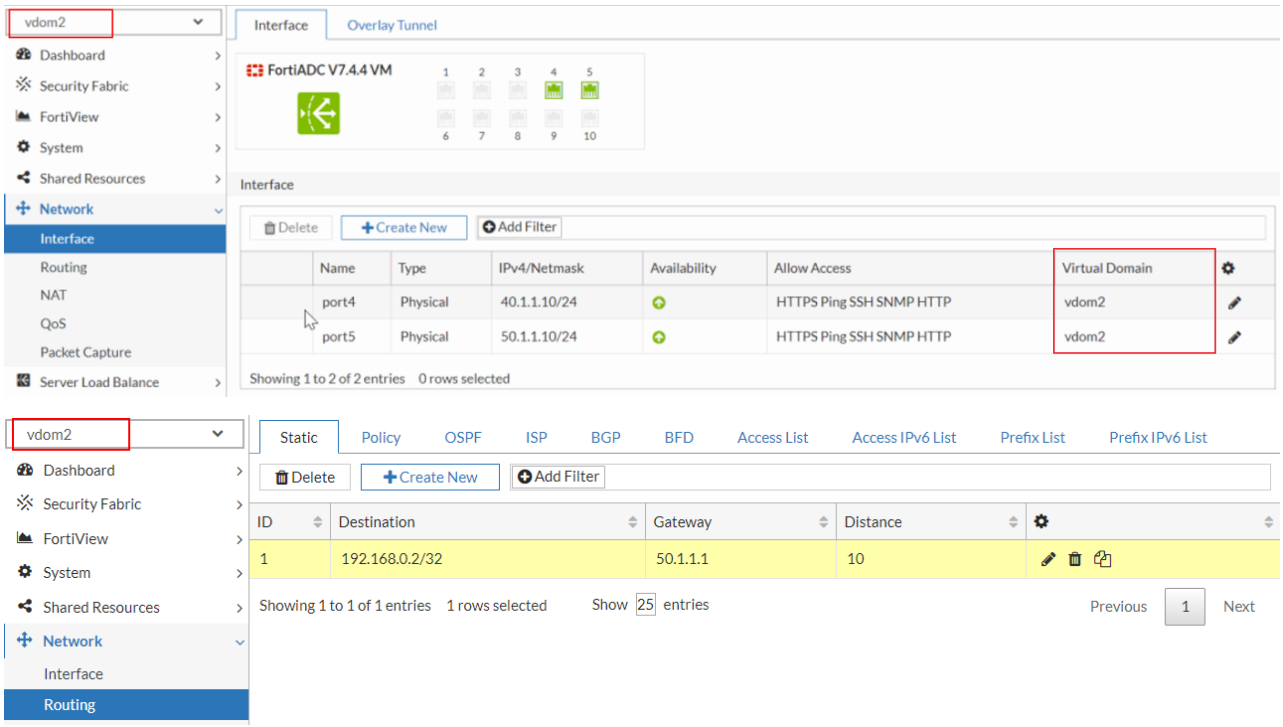
4. Within each VDOM, define a dedicated static route to ensure proper traffic forwarding to the corresponding backend server.
   In this example, configure a static route in **vdom1** to direct traffic toward **customer1's** real server, while **vdom2's** static route is set to forward traffic to **customer2's** designated server. Ensure each route is optimized for accurate destination reachability based on the configured next-hop and network topology.



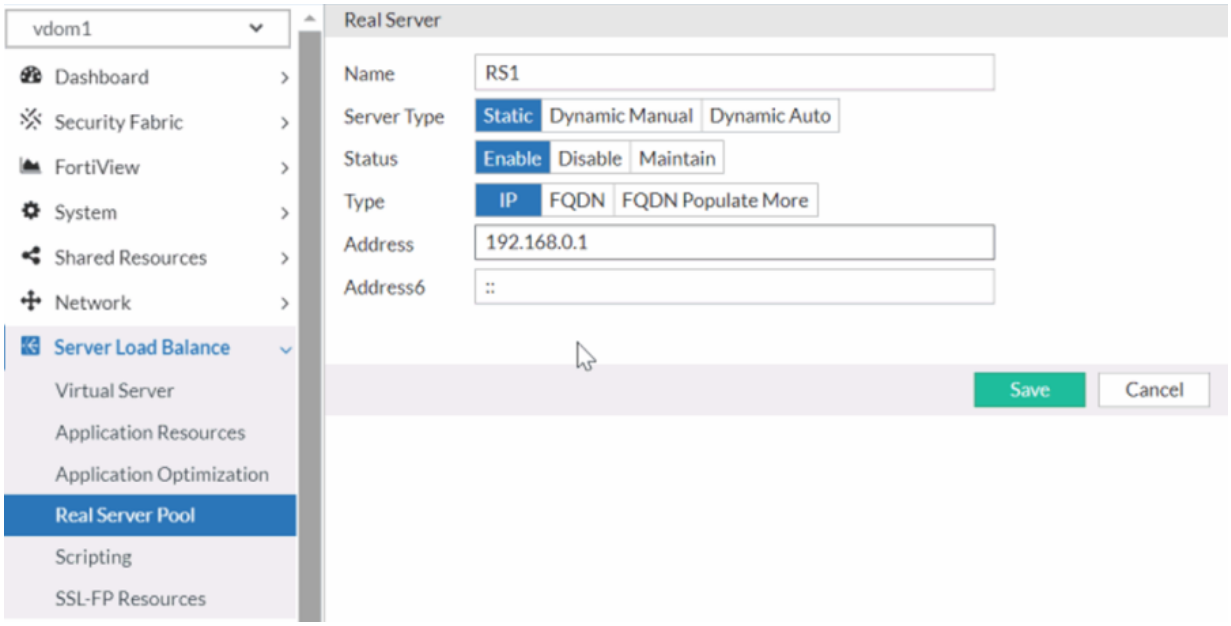


5. Validate that each VDOM has visibility restricted to its assigned network interfaces and routing table. Ensure that no cross-VDOM access to interfaces or route tables occurs, maintaining strict separation of resources and ensuring each VDOM operates within its own isolated network and routing domain.

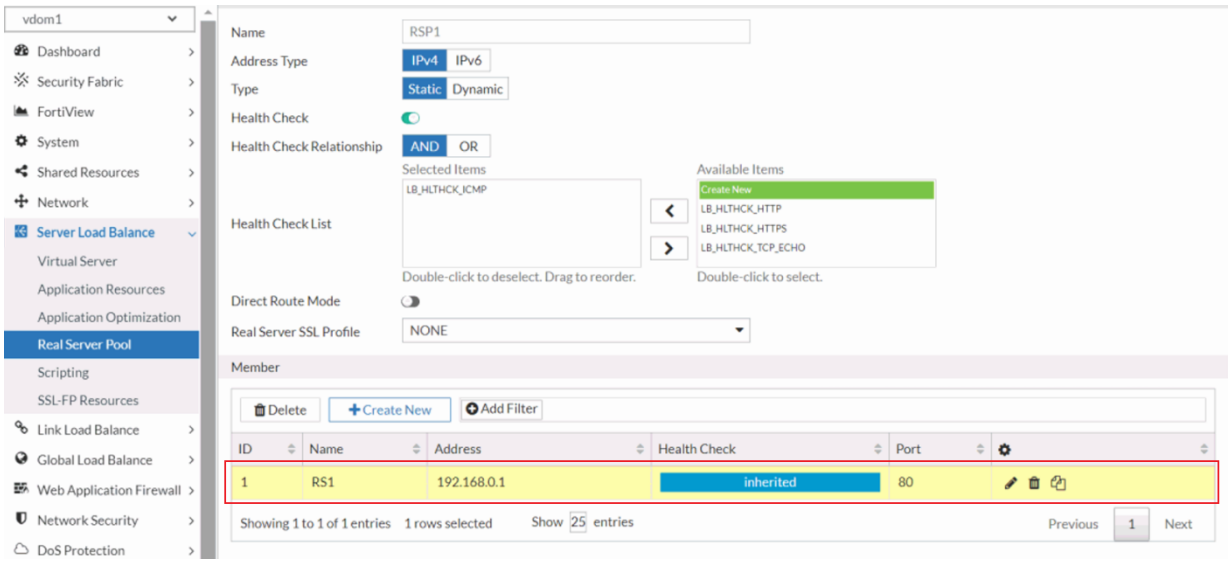




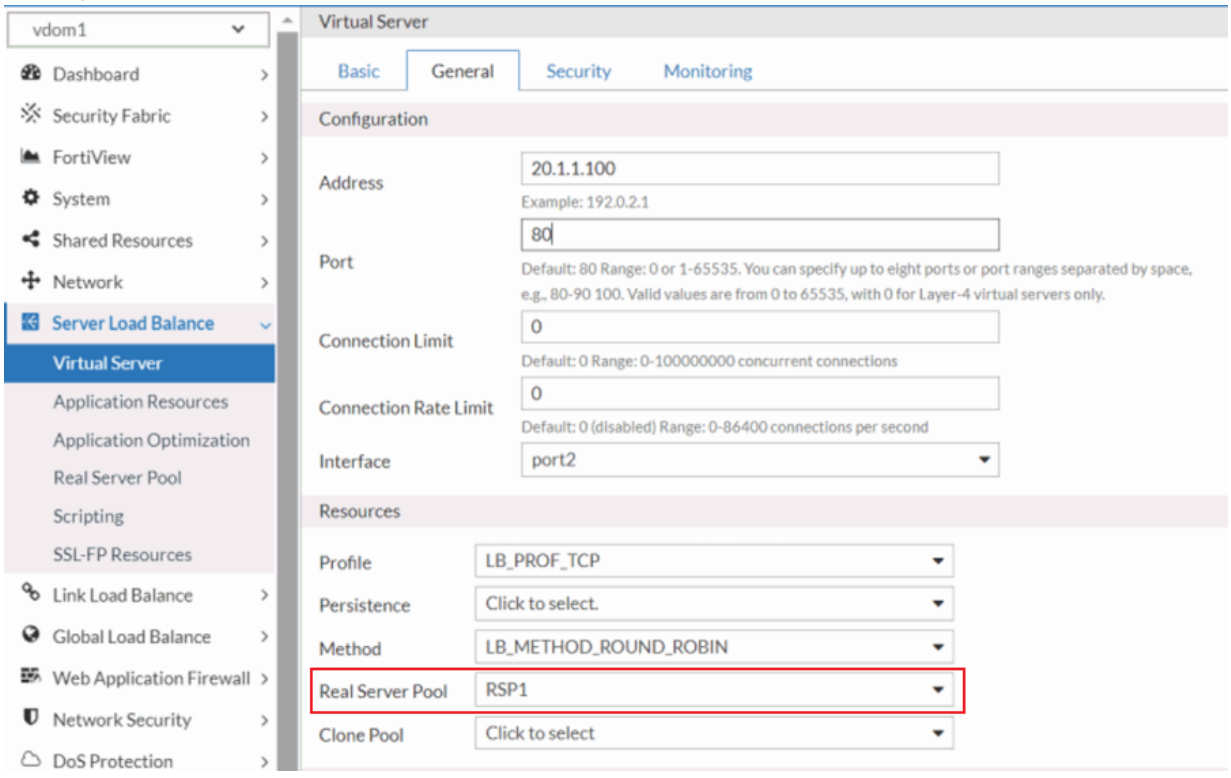
6. Configure server load balancing on a per-VDOM basis for each customer application, encompassing the complete setup of real servers, real server pools, and virtual server definitions.
   a. Create the real server to define and register the backend servers within each VDOM, specifying their IP addresses, ports, and any health check parameters required for monitoring server status.



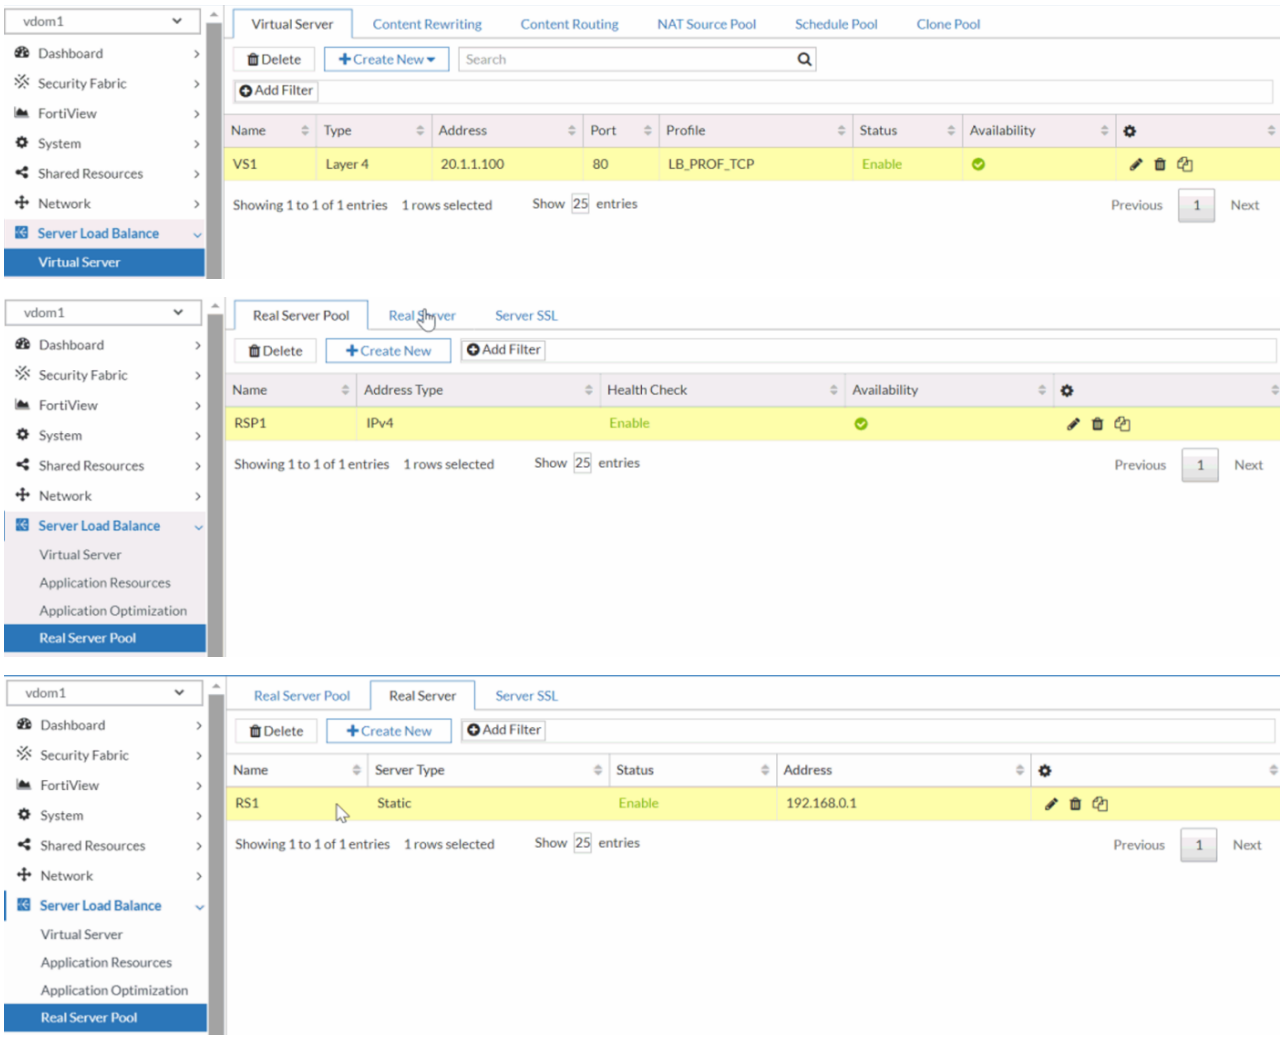
   b. Create and manage the real server pools within each VDOM, aggregating real servers into logical groups for load distribution. Configure load-balancing algorithms and session persistence settings tailored to the application's needs. In this example, the **RS1** real server is added as a member to the **RSP1** real server pool.



c. Create the virtual server within each VDOM, mapping them to the appropriate real server pools. Configure virtual servers with necessary IP addresses, ports, and access rules to handle incoming traffic and route it effectively to the designated real servers. In this example, the **VS1** virtual server is mapped to the **RSP1** real server pool.



7. Verify that each VDOM is configured to restrict visibility solely to its own virtual server, real server, and real server pool.



8. Verify that inter-VDOM traffic segmentation is properly enforced, ensuring complete isolation between traffic flows in each VDOM. Use the `execute ping` command to test network connectivity between each VDOM and its respective virtual servers.
   In this example, when **vdom1** pings **VS1** (its own virtual server), the packets should be successfully received. However, when vdom1 attempts to ping **VS2** (the virtual server in **vdom2**), no packets should be received, confirming isolation. Perform the same test from **vdom2**, where **VS2** should respond, but **VS1** should remain unreachable, validating proper segmentation and isolation between the VDOMs.
   a. Retrieve the IP address assigned to the virtual server in each VDOM by accessing the respective VDOM's configuration.
      The IP address assigned to **VS1**, the virtual server in **vdom1**, is **20.1.1.100**.

```
FortiADC (vdom1) # show load-balance virtual-server
config load-balance virtual-server
  edit "VS1"
    set interface port2
    set ip 20.1.1.100
    set load-balance-profile LB_PROF_TCP
    set load-balance-method LB_METHOD_ROUND_ROBIN
    set load-balance-pool RSP1
    set traffic-group default
  next
end
```

The IP address assigned to **VS2**, the virtual server in **vdom2**, is **40.1.1.100**.

```
FortiADC (vdom2) # show load-balance virtual-server
config load-balance virtual-server
  edit "VS2"
    set interface port4
    set ip 40.1.1.100
    set load-balance-profile LB_PROF_TCP
    set load-balance-method LB_METHOD_ROUND_ROBIN
    set load-balance-pool RSP2
    set traffic-group default
  next
end
```

b. From **vdom1**, use the execute ping command to test connectivity to both **VS1 (20.1.1.100)** and **VS2 (40.1.1.100)**. The ping to **VS1** should succeed, confirming proper connectivity within **vdom1**, while the ping to **VS2** should fail, verifying that inter-VDOM isolation is enforced.

```
FortiADC (vdom1) # exec ping -c 2 20.1.1.100
PING 20.1.1.100 (20.1.1.100): 56 data bytes
64 bytes from 20.1.1.100: icmp_seq=1 ttl=64 time=0.0 ms
64 bytes from 20.1.1.100: icmp_seq=2 ttl=64 time=0.1 ms

--- 20.1.1.100 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0.0/0.0/0.1 ms
```

```
FortiADC (vdom1) # exec ping -c 2 40.1.1.100
PING 40.1.1.100 (40.1.1.100): 56 data bytes
sendto: Network is unreachable
sendto: Network is unreachable
Timeout ...
Timeout ...
Timeout ...
Timeout ...
Timeout ...

--- 40.1.1.100 ping statistics ---
2 packets transmitted, 0 packets received, 100% packet loss
```

c. From **vdom2**, use the execute ping command to test connectivity to both **VS1 (20.1.1.100)** and **VS2 (40.1.1.100)**. The ping to **VS2** should succeed, confirming proper connectivity within **vdom2**, while the ping to **VS1** should fail, verifying that inter-VDOM isolation is enforced.

```
FortiADC (vdom2) # exec ping -c 2 40.1.1.100
PING 40.1.1.100 (40.1.1.100): 56 data bytes
64 bytes from 40.1.1.100: icmp_seq=1 ttl=64 time=0.1 ms
64 bytes from 40.1.1.100: icmp_seq=2 ttl=64 time=0.1 ms

--- 40.1.1.100 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0.1/0.1/0.1 ms
```

```
FortiADC (vdom2) # exec ping -c 2 20.1.1.100
PING 20.1.1.100 (20.1.1.100): 56 data bytes
sendto: Network is unreachable
sendto: Network is unreachable
Timeout ...
Timeout ...
Timeout ...
Timeout ...
Timeout ...

--- 20.1.1.100 ping statistics ---
2 packets transmitted, 0 packets received, 100% packet loss
```

9. Verify that the virtual server in each VDOM works properly by connecting to the virtual servers from the client.

```
[root@Client1 ~]# curl -k 20.1.1.100
Server1[root@Client1 ~]#
[root@Client1 ~]#
```

```
[root@Client2 ~]# curl -k 40.1.1.100
Server2[root@Client2 ~]#
[root@Client2 ~]#
```

By following these steps, you can efficiently deploy and manage isolated network environments for multiple customers, ensuring robust traffic separation and independent resource allocation.

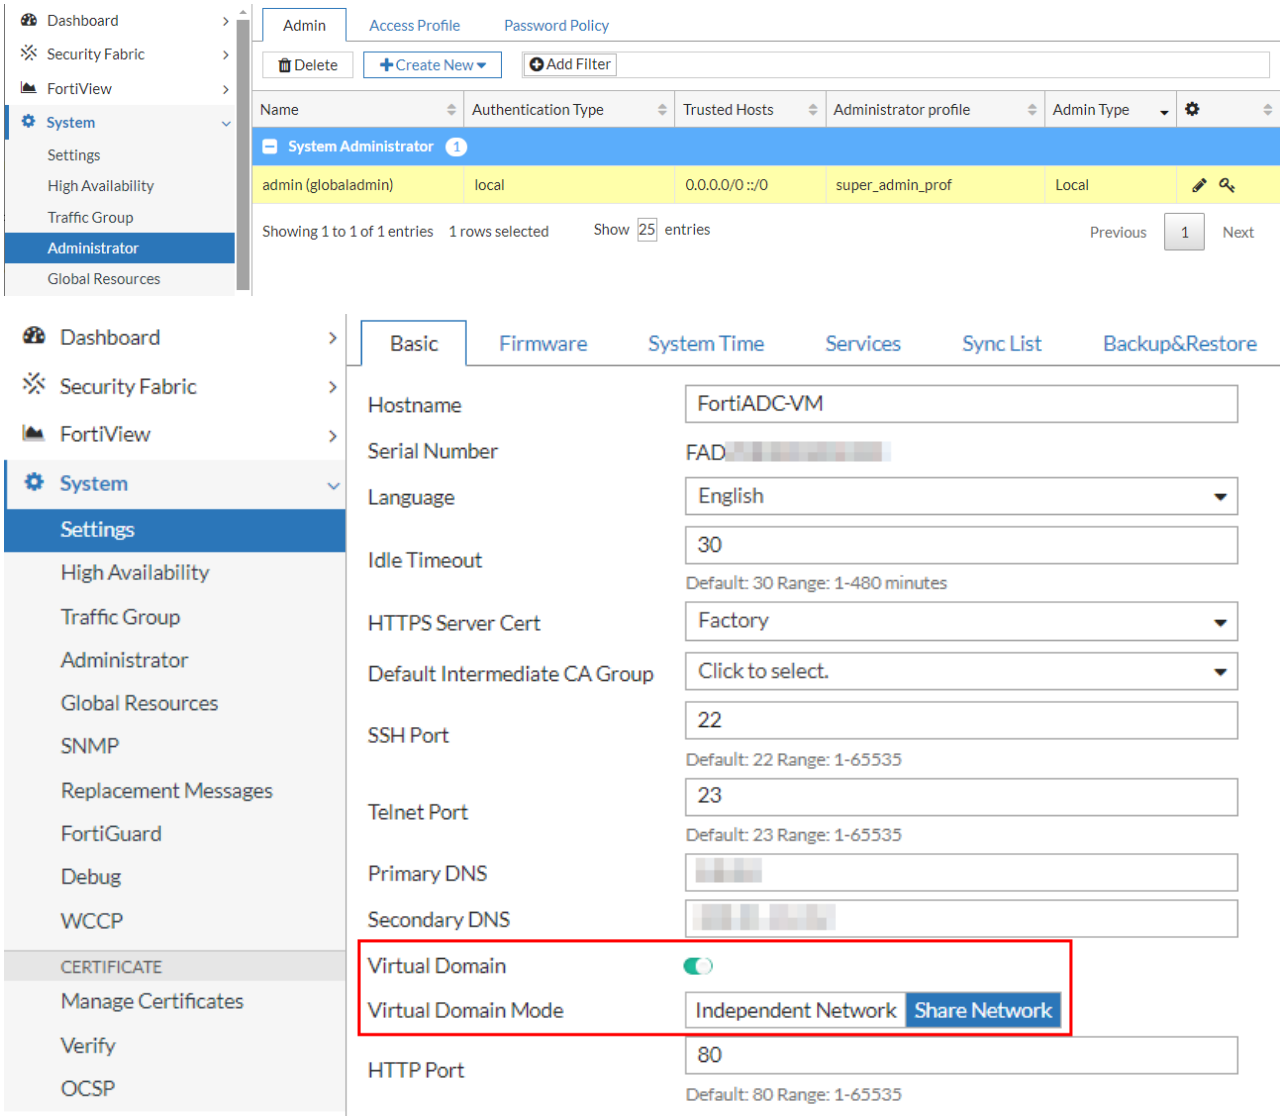
Use-case for ADOM

Administrative Domains (ADOMs) in FortiADC offer centralized management across multiple FortiADC instances while maintaining the integrity of each device's individual configuration. By leveraging ADOMs, administrators can streamline policy, certificate, and object management across distinct devices without the risk of cross-device configuration interference. ADOMs are particularly useful in multi-appliance environments where consistency in administrative operations is key. For instance, in environments where multiple administrators need to manage different server load-balancing configurations but without segregating the network into isolated domains, ADOMs provide an efficient solution. By using shared interfaces and routes, multiple administrators can work within the same infrastructure while being assigned specific resources and management zones.

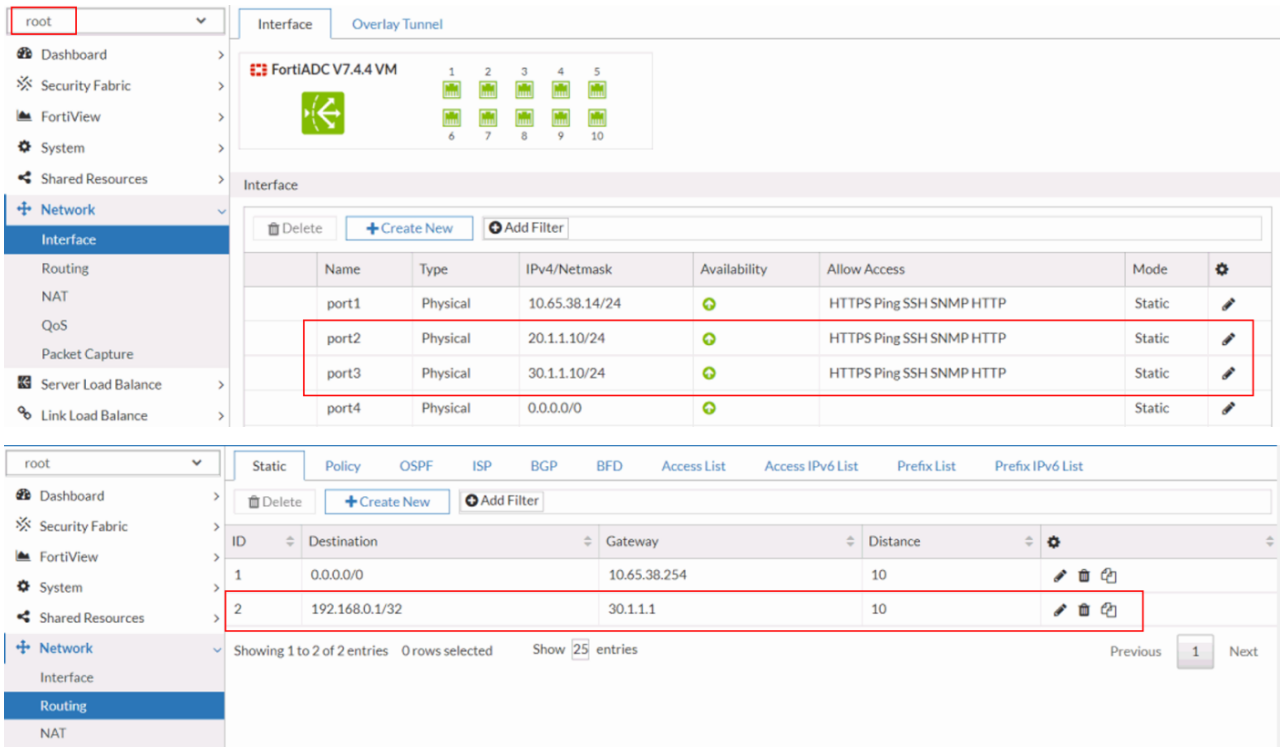
Scenario: Managing Server Load-Balancing with Multiple Administrators

For example, a customer wants to assign multiple administrators the responsibility to manage subsets of server load-balancing servers, but they do not want to break the network into multiple subdomains. In this case, ADOMs can be configured to share the same interfaces and routes while maintaining separate management zones for each admin.

1. Log in as super admin and enable Virtual Domains in Shared Network mode.

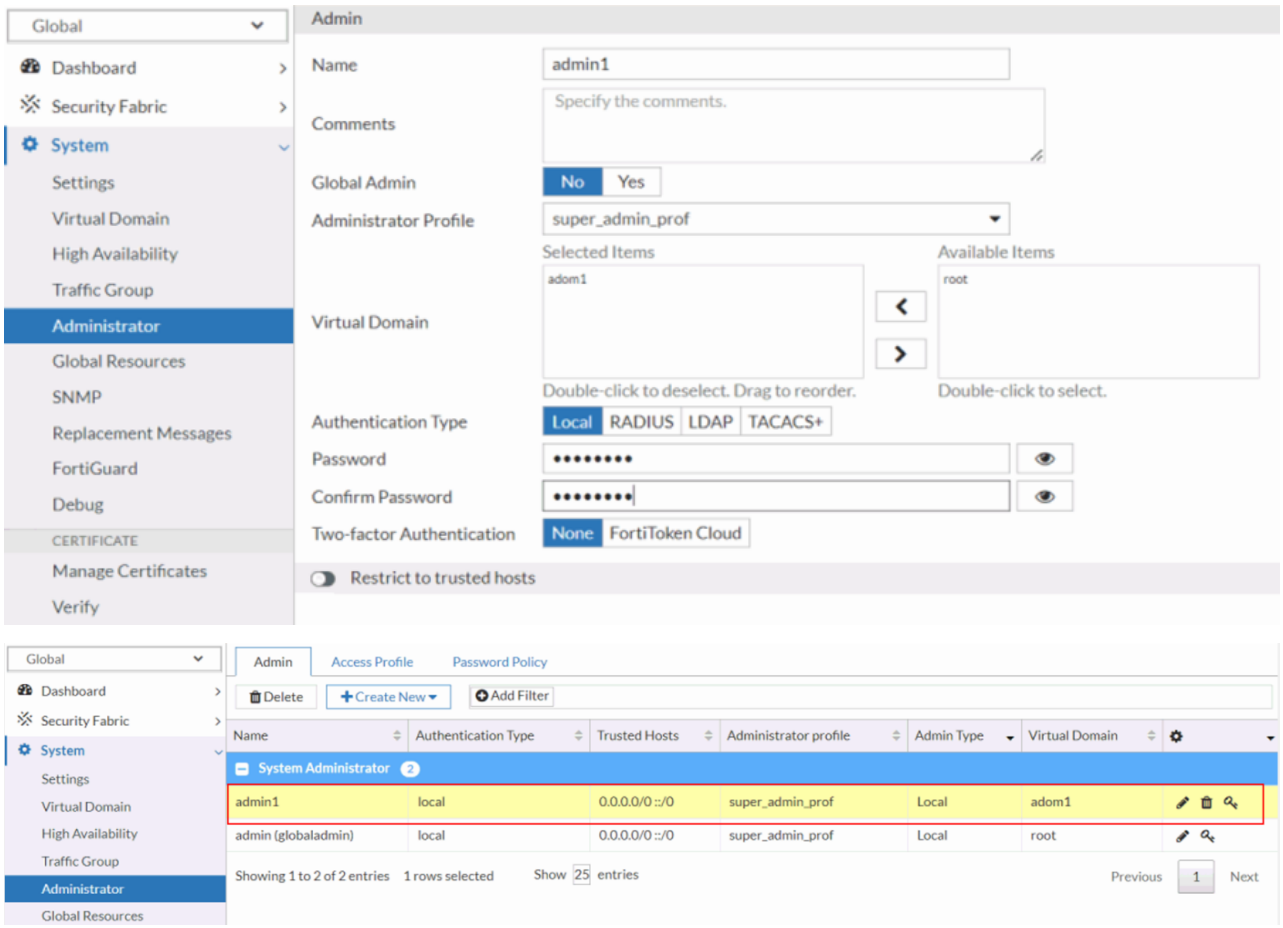


2. Configure the interfaces and establish network routing within the root ADOM to ensure proper communication across network segments. Assign the appropriate IP addresses and subnet masks to each interface, and define static or dynamic routing policies to direct traffic efficiently across the network infrastructure.



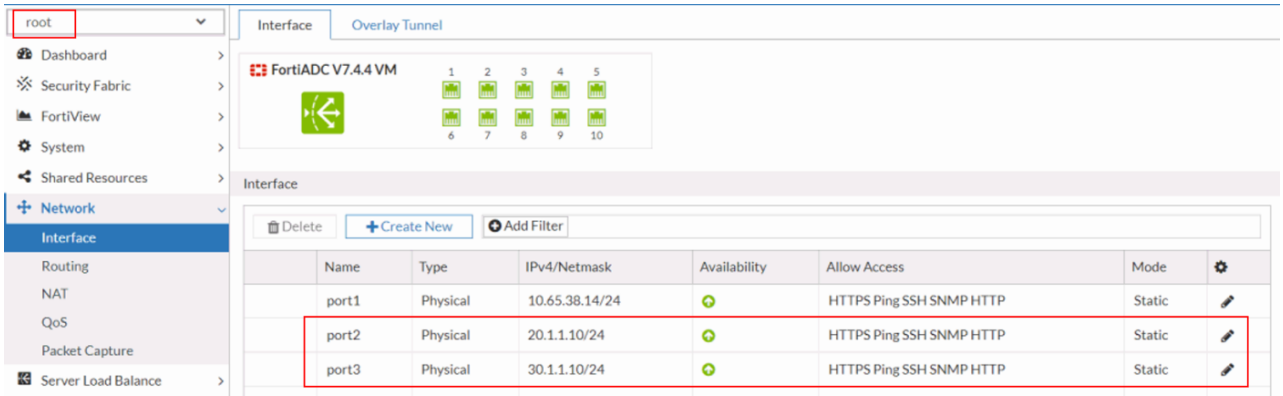
3. Create a non-root ADOM and assign a dedicated non-root ADOM administrator with the necessary permissions to manage and configure the ADOM. Ensure that role-based access control (RBAC) is properly applied to limit administrative privileges to the designated ADOM environment.

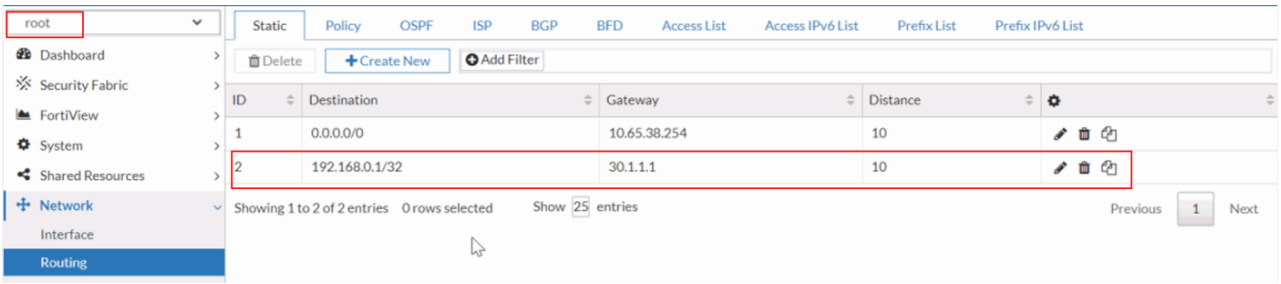




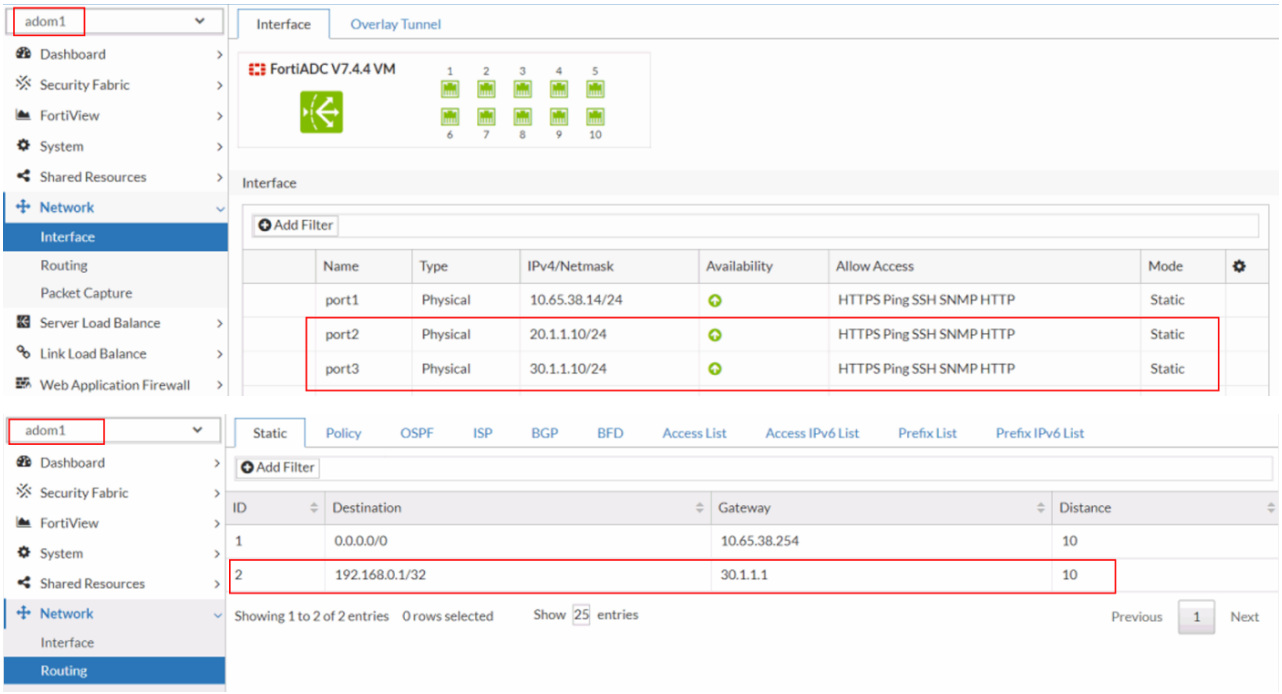
4. Verify that the interfaces and routes are accessible from both the root ADOM and the non-root ADOM (**adom1**). Ensure that the configurations are correctly propagated, allowing seamless visibility and management of the same network resources across both ADOM environments. Confirm that routing tables and interface assignments are consistent between the ADOMs, ensuring unified network operation.

root ADOM:

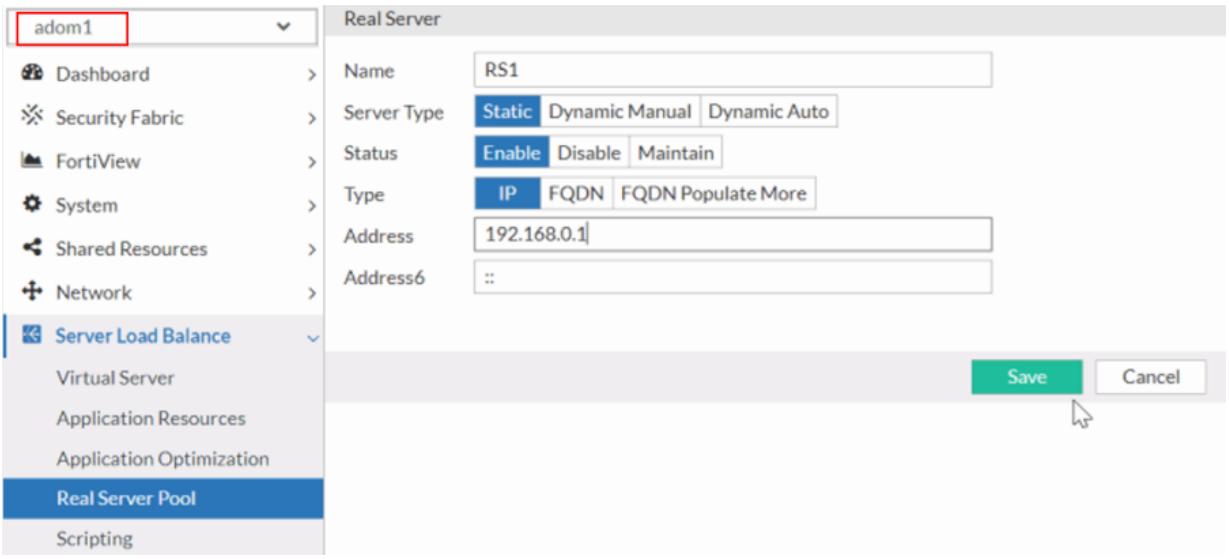




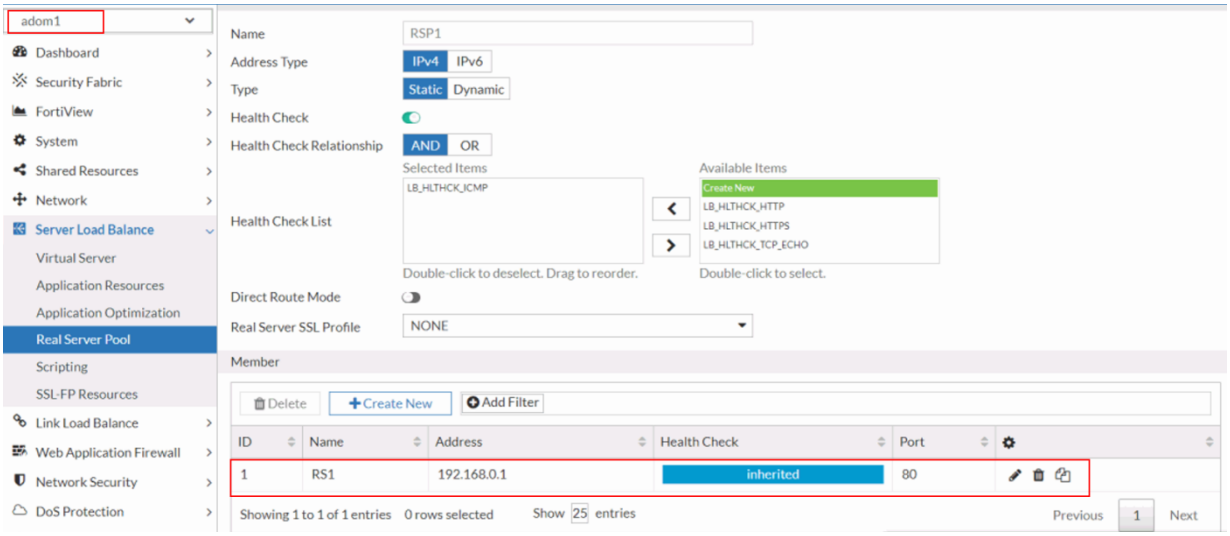
non-root ADOM (**adom1**):



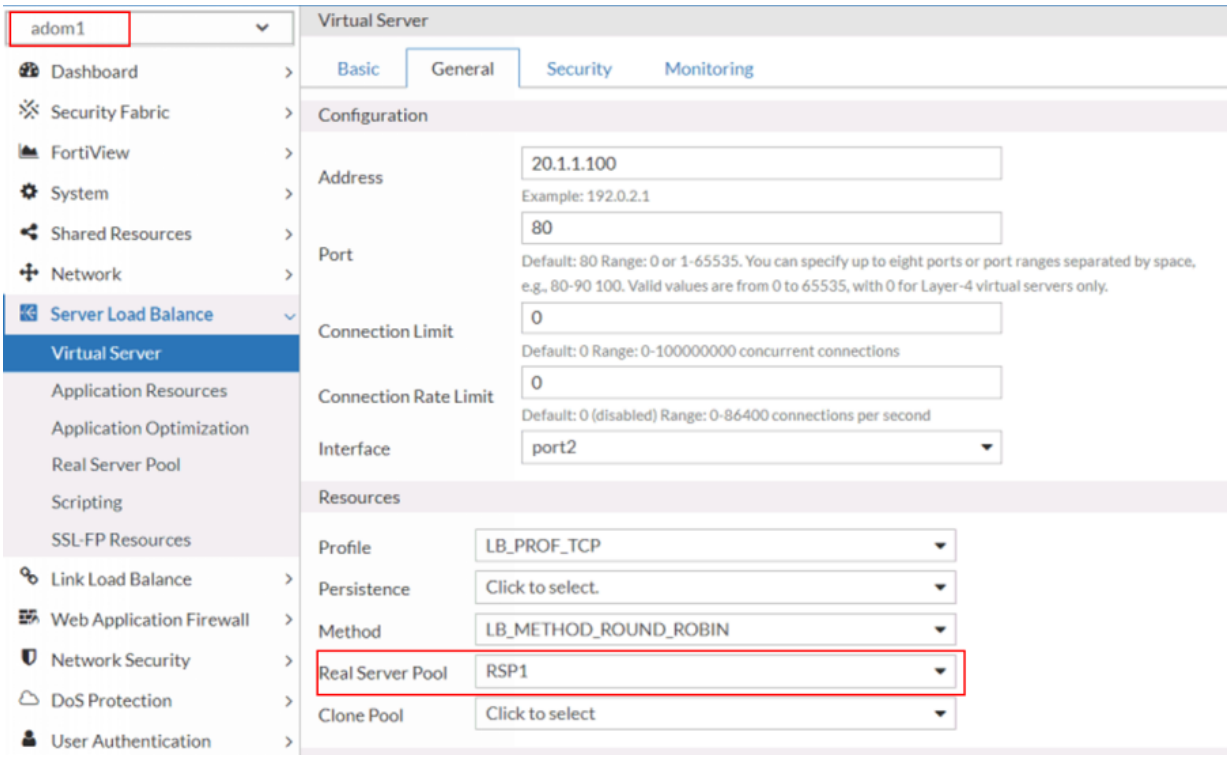
5. Create a virtual server within the non-root ADOM using the interfaces and routes inherited from the root ADOM. Ensure that the virtual server configuration leverages the correct network resources, including IP addresses, interfaces, and routing paths, to seamlessly integrate with the broader network architecture managed by the root ADOM.
   a. Create the real server by defining its IP address, port, and health check parameters. Ensure that the real server is properly configured to integrate with the load-balancing policies and is linked to the appropriate server pool for efficient traffic distribution.



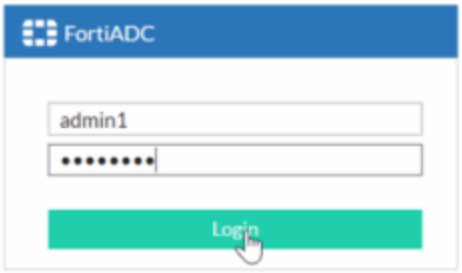
b. Create and manage the real server pools by aggregating real servers into logical groups for load distribution. Configure load-balancing algorithms and session persistence settings tailored to the application's needs. In this example, the RS1 real server is added as a member to the RSP1 real server pool.



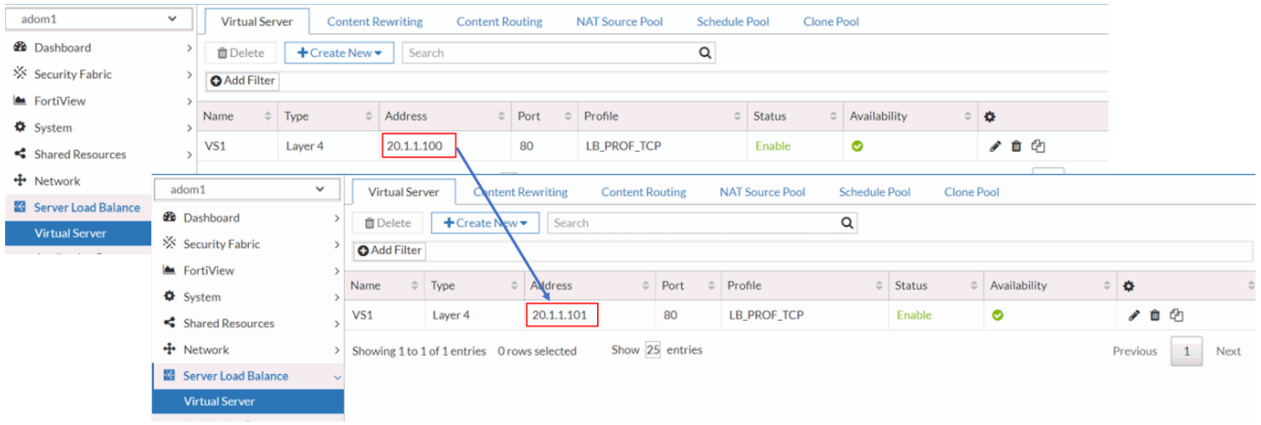
c. Create the virtual server in the ADOM, mapping them to the appropriate real server pools. Configure virtual servers with necessary IP addresses, ports, and access rules to handle incoming traffic and route it effectively to the designated real servers. In this example, the VS1 virtual server is mapped to the RSP1 real server pool.



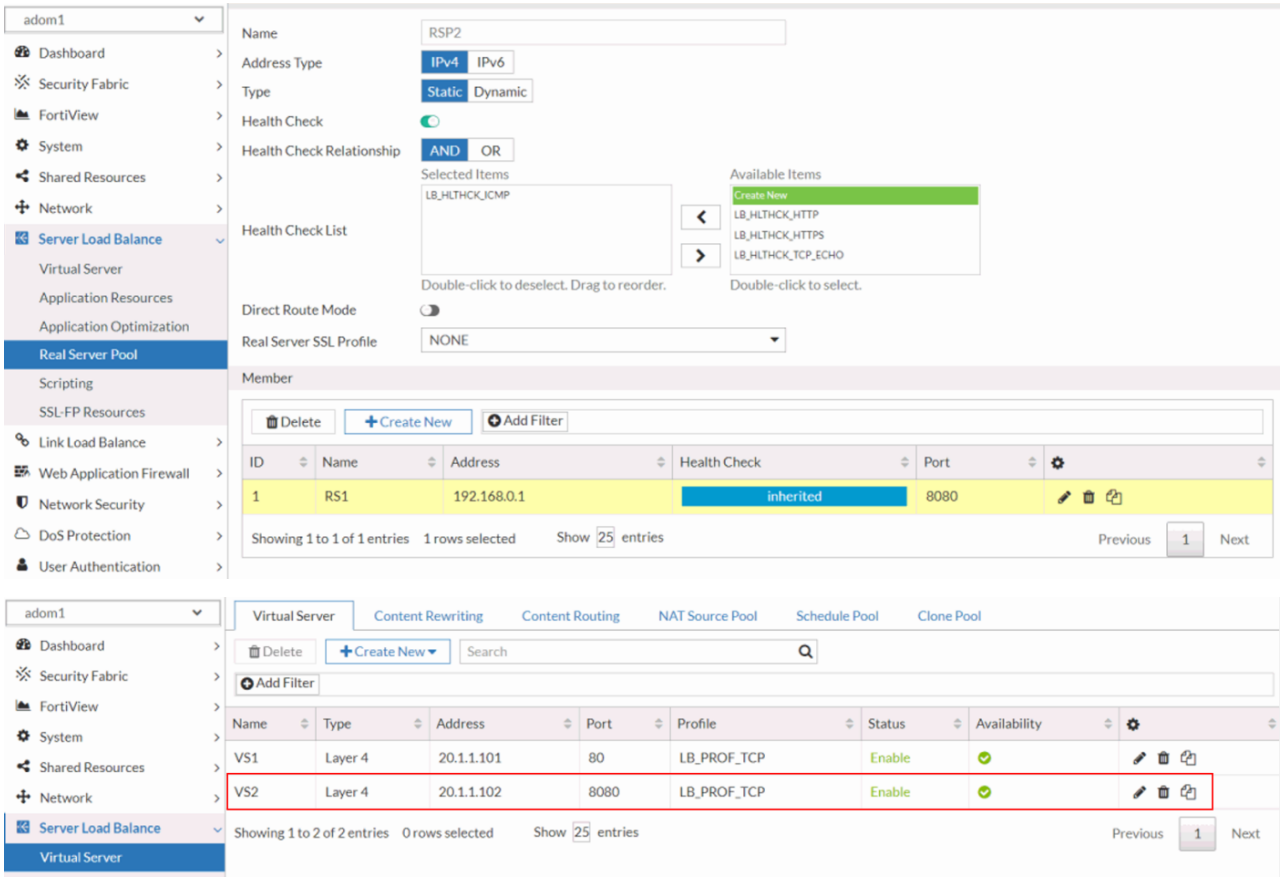
6. Log in as the non-root ADOM administrator and verify access to the root ADOM's interfaces and routes. Ensure that the non-root administrator can view and manage the network resources from the root ADOM, confirming that appropriate permissions and visibility settings are correctly configured.
   You should expect to see the same interface and routing settings in the root ADOM as those configured in the non-root ADOM (**adom1**).



7. Verify that the non-root ADOM administrator can effectively manage the virtual server created by the super administrator. Ensure that the non-root admin has the appropriate permissions to view, configure, and modify the virtual server settings as necessary. For example, make minor adjustments to the existing virtual server configuration to test whether the system permits these changes.



8. Verify that the non-root ADOM administrator can successfully create a new virtual server using the interfaces and routes inherited from the root ADOM. Ensure that the non-root admin can configure the virtual server with the appropriate network resources and routing parameters as defined in the root ADOM.



By leveraging ADOMs to streamline management, administrators can efficiently allocate resources and segregate responsibilities without compromising the unified network infrastructure. This approach simplifies operational tasks, enhances scalability, and ensures consistent management practices across large-scale deployments.

SAP Application Integration

For multinational corporations operating critical SAP applications, ensuring optimal performance and reliability is essential. These organizations face the challenge of managing high volumes of traffic across distributed environments while maintaining session persistence and securing sensitive data. FortiADC provides a comprehensive solution for these needs by offering advanced Server Load Balancing (SLB) capabilities and seamless integration with SAP environments. This use-case details how FortiADC can be effectively deployed to manage traffic distribution among SAP servers, ensuring high availability, load balancing efficiency, and robust data protection. By leveraging FortiADC's features, organizations can achieve improved application performance and resilience, even during peak traffic periods.

User Scenario: Enhancing SAP Application Delivery with the FortiADC SAP Connector

A multinational corporation specializing in high-end consumer products relies heavily on SAP applications for essential functions such as finance, inventory management, and customer relationship management. With a complex infrastructure spanning multiple data centers and cloud environments, the corporation faces challenges in efficiently managing traffic, maintaining session persistence, and ensuring data security. To address these challenges and enhance application performance, the company has decided to implement FortiADC, known for its advanced traffic management and security capabilities.

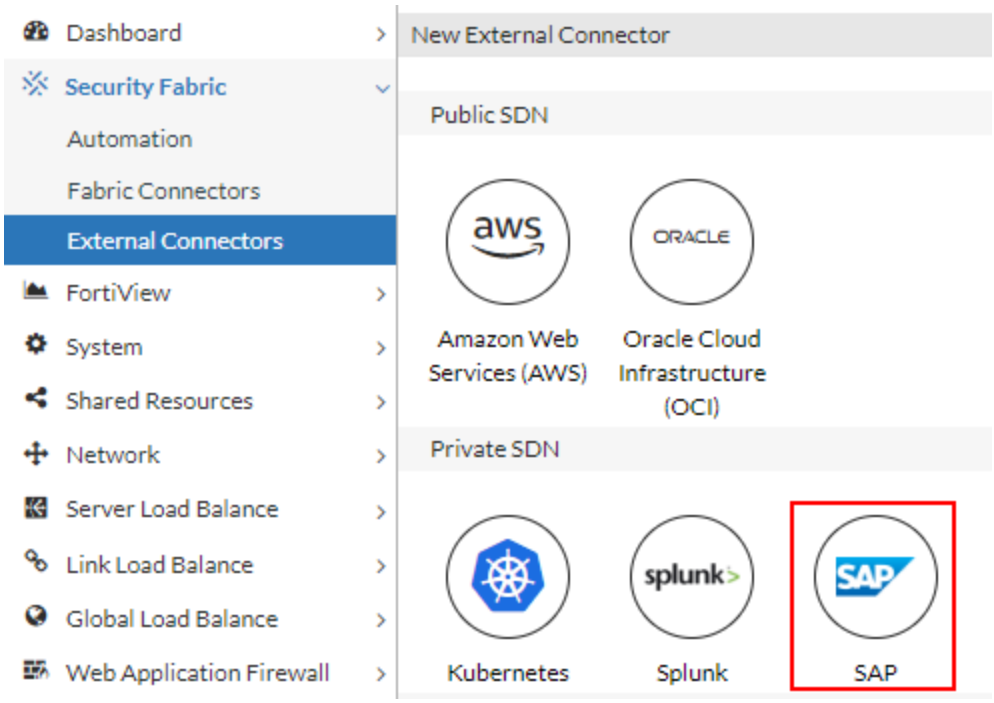
During a major business event, such as a product launch or financial reporting cycle, the corporation experiences a surge in traffic to its SAP applications. To maintain high availability and optimal performance, FortiADC is deployed to manage traffic distribution among multiple SAP servers. The FortiADC SAP connector ensures seamless session persistence and effective load balancing, preventing server overload and protecting sensitive data, thereby ensuring a reliable and secure user experience throughout the peak period.

Step 1: SAP Connector Setup

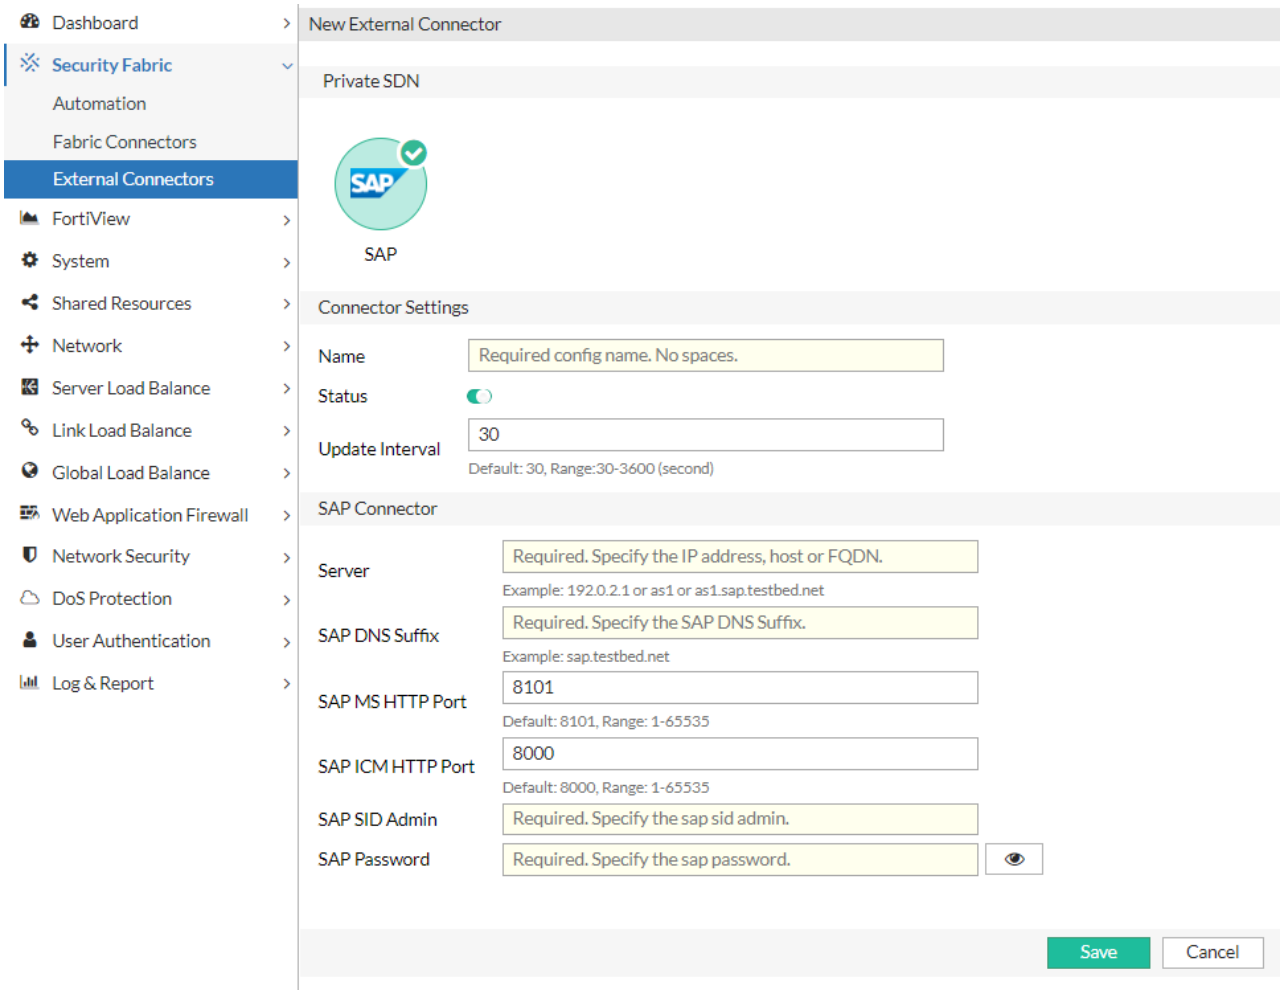
Configure the FortiADC SAP Connector to integrate seamlessly with your existing SAP infrastructure. The SAP Connector enables FortiADC to communicate with SAP servers, dynamically managing traffic flows based on server health and performance. This setup ensures that FortiADC is ready to manage SAP-specific traffic and distribute it across the available backend servers.

To create an SAP Connector:

1. Navigate to **Security Fabric > External Connectors**. Click **Create New** to display the configuration editor. Under the Private SDN section is the **SAP** connector option.

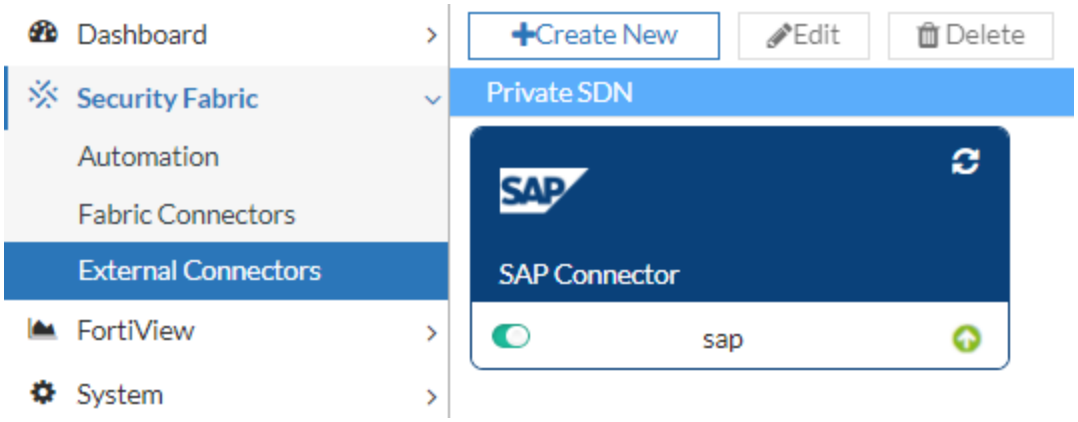

Dashboard
Security Fabric
Automation
Fabric Connectors
External Connectors
FortiView
System
Shared Resources
Network
Server Load Balance
Link Load Balance
Global Load Balance
Web Application Firewall
New External Connector
Public SDN
aws
Amazon Web Services (AWS)
ORACLE
Oracle Cloud Infrastructure (OCI)
Private SDN
Kubernetes
splunk>
Splunk
SAP
SAP

2. Click the SAP icon to display the configuration editor. Configure the required settings for the SAP connector.



For configuration details, see SAP Connector on page 105.

3. Click **Save** to commit the changes.
The newly created SAP connector will display on the External Connectors page.

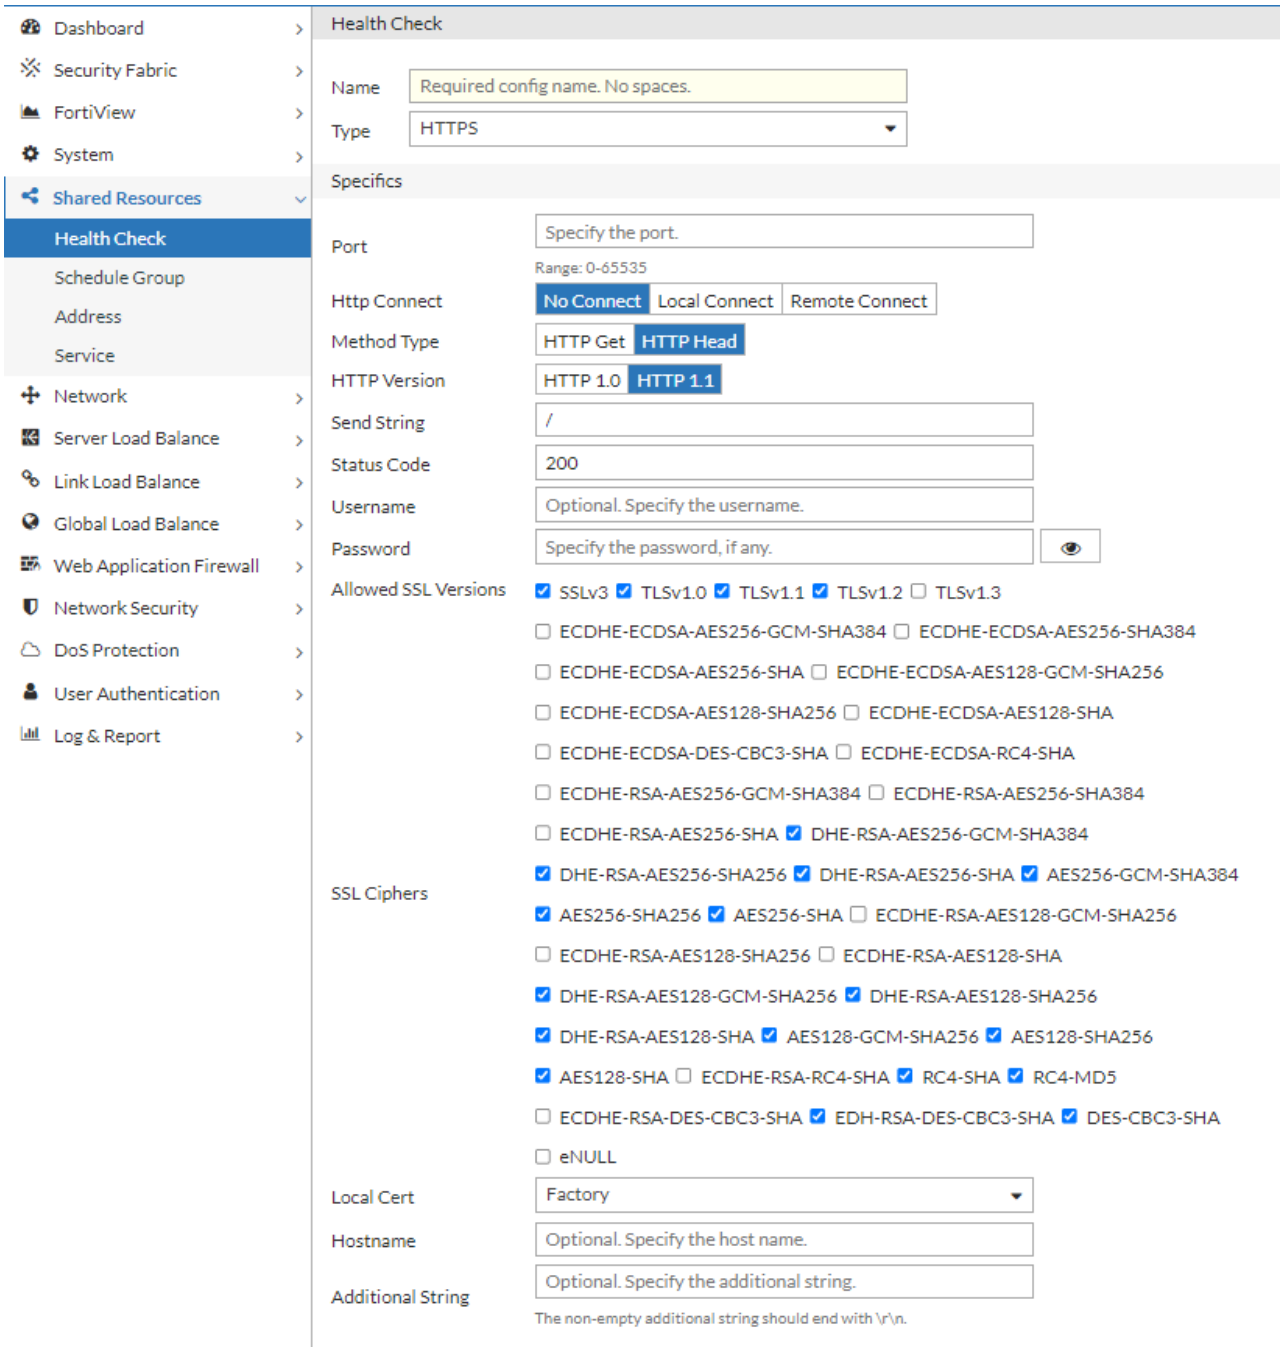


Step 2: Health Monitor Configuration

Once the SAP Connector is in place, configure core system settings to ensure stable traffic management. This includes setting up health monitoring for SAP servers, allowing FortiADC to check the real-time status and availability of each server in the pool. By actively monitoring server health, FortiADC ensures that traffic is directed only to servers that are operational, preventing potential bottlenecks or service outages.

To configure Health Check:

1. Navigate to **Shared Resources > Health Check**.
   The configuration page displays the **Health Check** tab.
2. Click **Create New** to display the configuration editor and configure the required settings.
   For this example, set the **Type** to **HTTPS**, as SAP applications typically communicate over the HTTPS protocol for secure data transmission.



For configuration details, see Configuring health checks on page 355.

3. Save the configuration.

Alternatively, predefined Health Check profiles can also be applied.

Step 3: Security and Data Protection

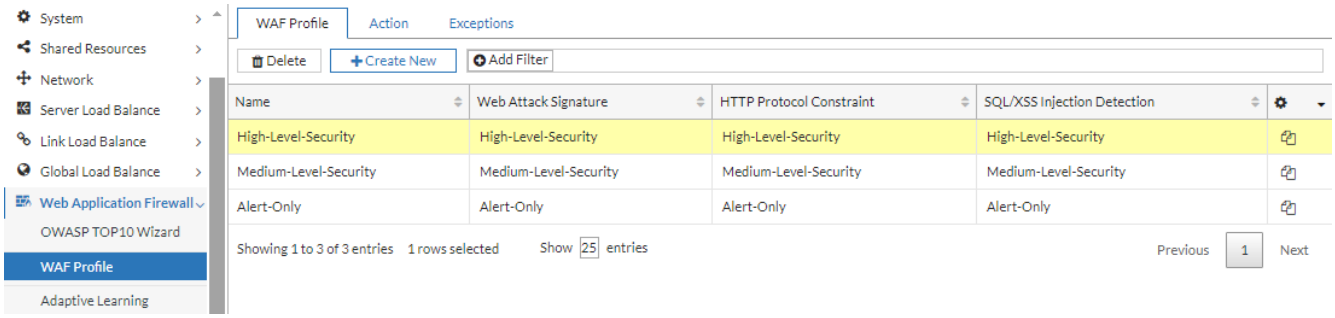
FortiADC plays a critical role in securing SAP traffic. Configure a Web Application Firewall (WAF) profile—either a custom profile or one of the predefined options—to safeguard your SAP applications from web-based threats and potential data breaches. This ensures that traffic is not only efficiently balanced but also securely routed, protecting

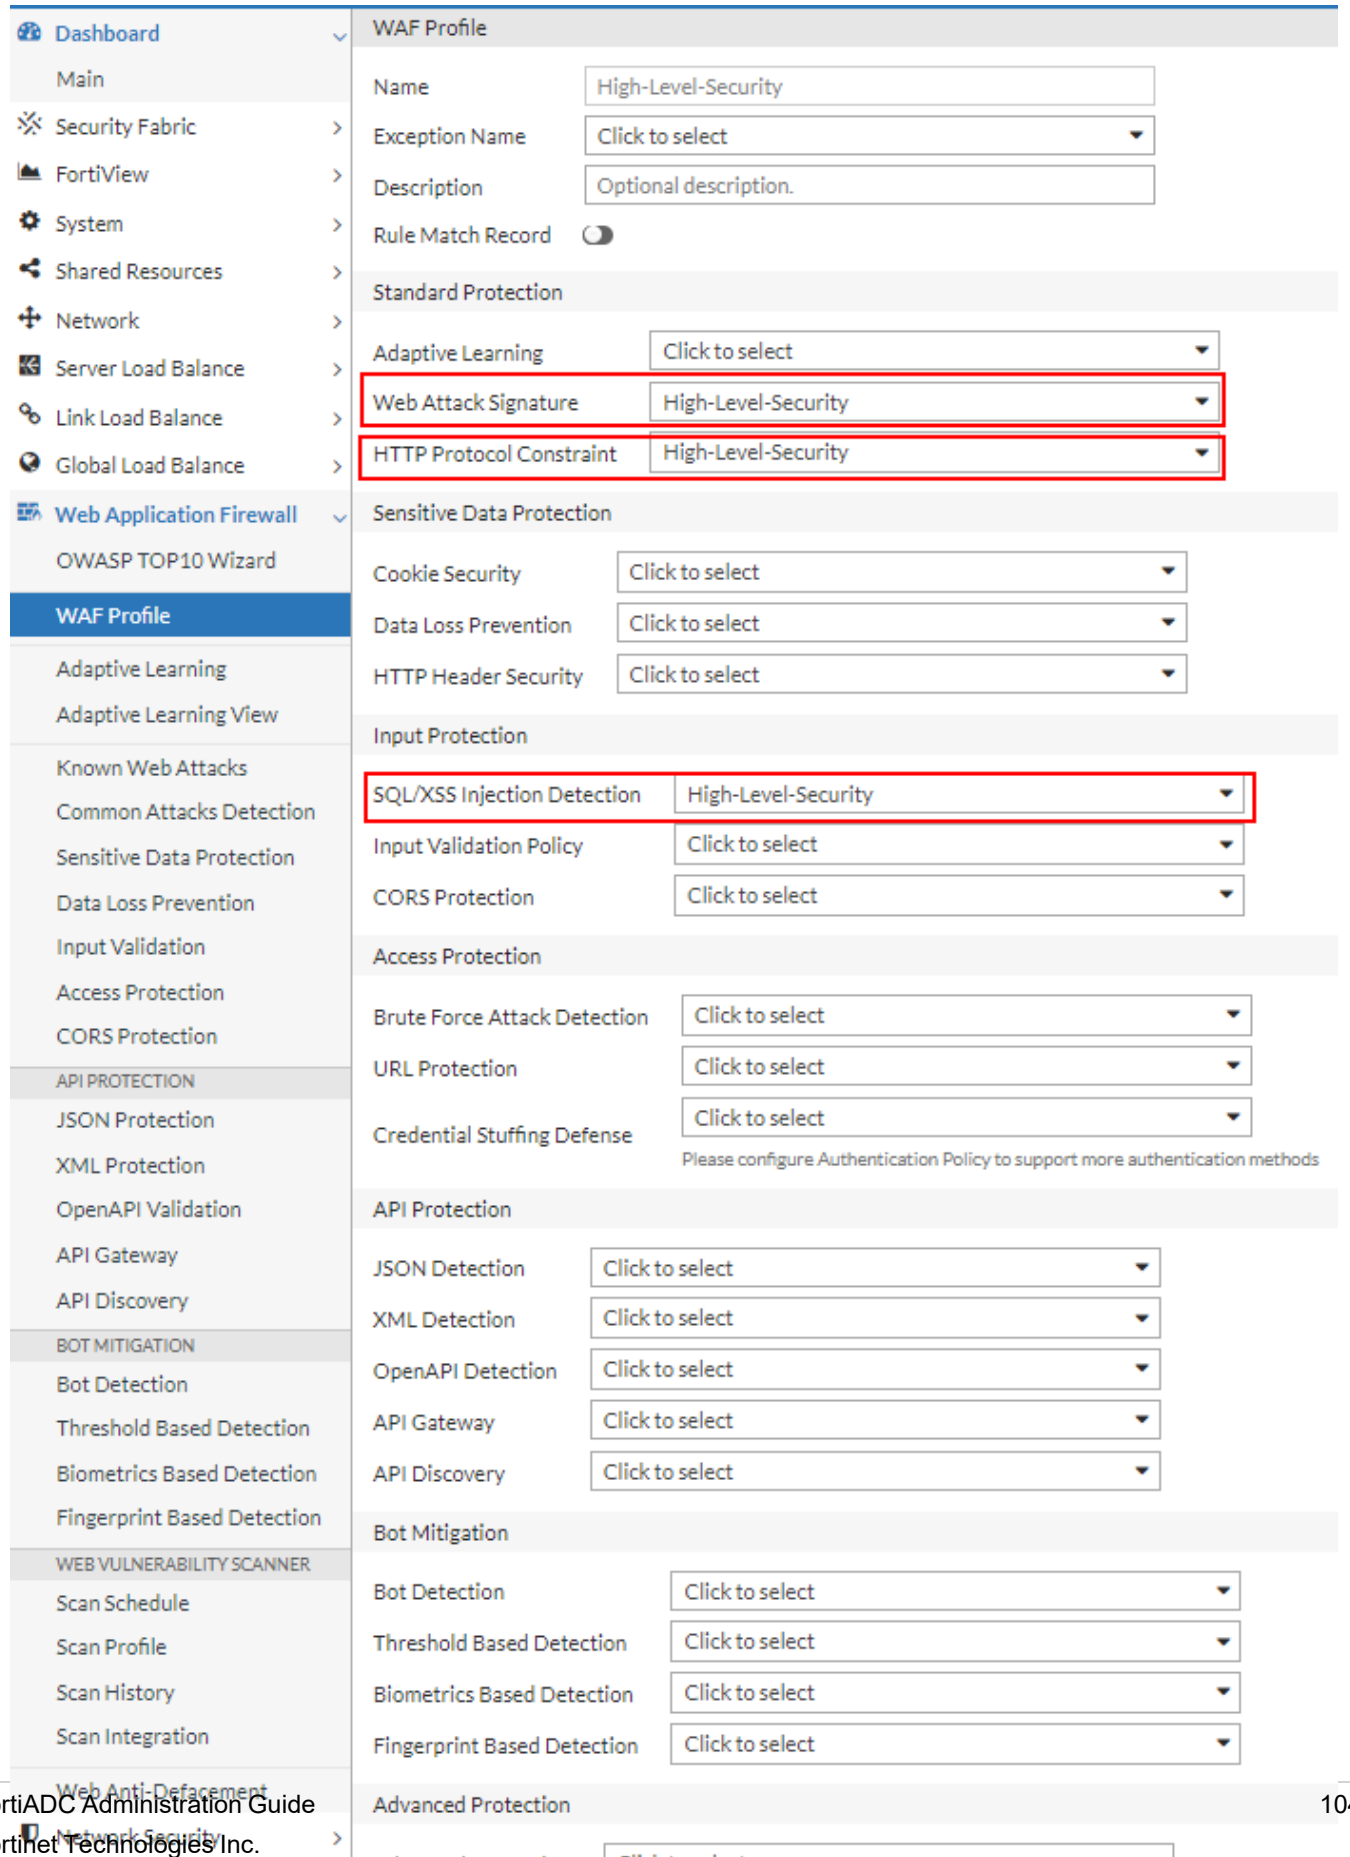
sensitive information such as financial records and customer data. Custom WAF profiles provide granular control over security policies, allowing for fine-tuned protection, while predefined profiles offer standardized settings optimized for common security requirements.

In this example, we recommend utilizing the **High-Level-Security** predefined WAF profile. This profile is specifically designed to provide robust, enterprise-grade protection by enabling a comprehensive set of advanced security features. It includes an extensive library of web attack signatures, strict HTTP protocol constraints, and precise SQL/XSS injection detection mechanisms. Together, these features fortify your SAP applications against sophisticated and evolving web-based threats, ensuring secure and uninterrupted access to critical business data and processes, such as financial transactions and customer information.

WAF Profile: High-Level-Security

The following describes in detail each of the WAF policy settings applied in the predefined WAF Profile named "High-Level-Security": Web Attack Signature, HTTP Protocol Constraint, and SQL/XSS Injection Detection.




Dashboard
Main
Security Fabric
FortiView
System
Shared Resources
Network
Server Load Balance
Link Load Balance
Global Load Balance
Web Application Firewall
OWASP TOP10 Wizard
WAF Profile
Adaptive Learning
Adaptive Learning View
Known Web Attacks
Common Attacks Detection
Sensitive Data Protection
Data Loss Prevention
Input Validation
Access Protection
CORS Protection
API PROTECTION
JSON Protection
XML Protection
OpenAPI Validation
API Gateway
API Discovery
BOT MITIGATION
Bot Detection
Threshold Based Detection
Biometrics Based Detection
Fingerprint Based Detection
WEB VULNERABILITY SCANNER
Scan Schedule
Scan Profile
Scan History
Scan Integration
Web Anti-Defacement
Network Security
DoS Protection
WAF Profile
Name
High-Level-Security
Exception Name
Click to select
Description
Optional description.
Rule Match Record
Standard Protection
Adaptive Learning
Click to select
Web Attack Signature
High-Level-Security
HTTP Protocol Constraint
High-Level-Security
Sensitive Data Protection
Cookie Security
Click to select
Data Loss Prevention
Click to select
HTTP Header Security
Click to select
Input Protection
SQL/XSS Injection Detection
High-Level-Security
Input Validation Policy
Click to select
CORS Protection
Click to select
Access Protection
Brute Force Attack Detection
Click to select
URL Protection
Click to select
Credential Stuffing Defense
Click to select
Please configure Authentication Policy to support more authentication methods
API Protection
JSON Detection
Click to select
XML Detection
Click to select
OpenAPI Detection
Click to select
API Gateway
Click to select
API Discovery
Click to select
Bot Mitigation
Bot Detection
Click to select
Threshold Based Detection
Click to select
Biometrics Based Detection
Click to select
Fingerprint Based Detection
Click to select
Advanced Protection
Advanced Protection
Click to select
CSRF Protection
Click to select

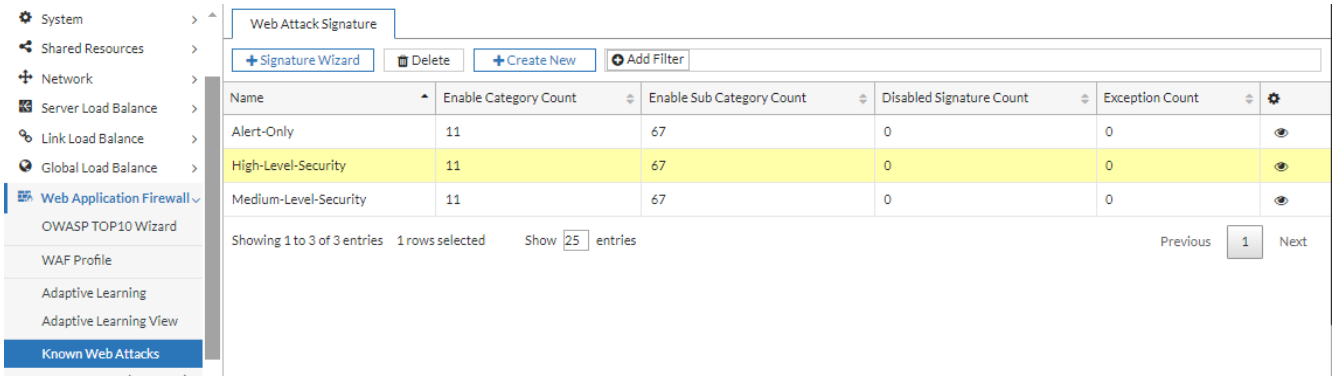
Web Attack Signature policy: High-Level-Security

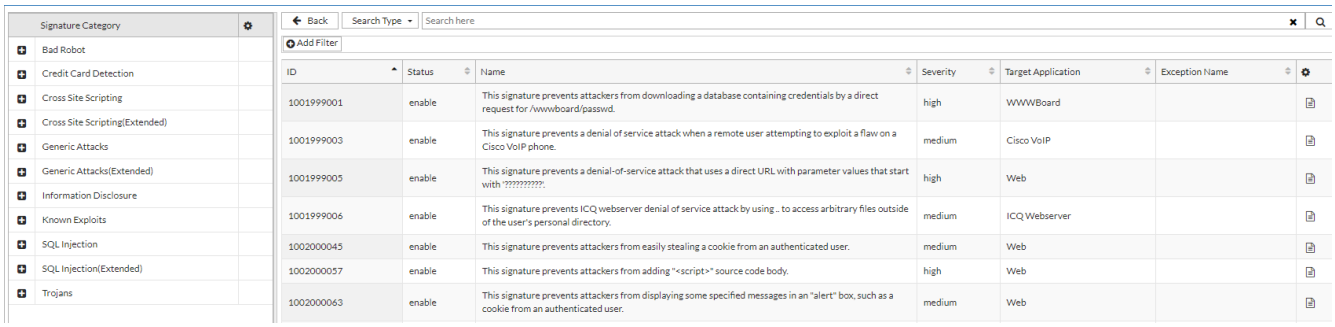
The FortiGuard Web Attack Signature service provides a database of attack signatures that is updated periodically to protect against new kinds of attacks.

The predefined Web Attack Signature policy "High-Level-Security" is applied in the predefined WAF Profile also named "High-Level-Security". This predefined Web Attack Signature policy contains the following settings for the class of scanpoints, the action when traffic matches the signatures, and the Web Attack Signature categories and subcategories of threats that are detected by the signatures.

Scanpoint	Description	Status
HTTP header	Scans traffic against HTTP header signatures.	Enabled
HTTP Request Body	Scans traffic against HTTP request body signatures.	Enabled
HTTP Response Body	Scans traffic against HTTP response body signatures.	Disabled

Event Severity	Action
High	Deny
Medium	Deny
Low	Alert





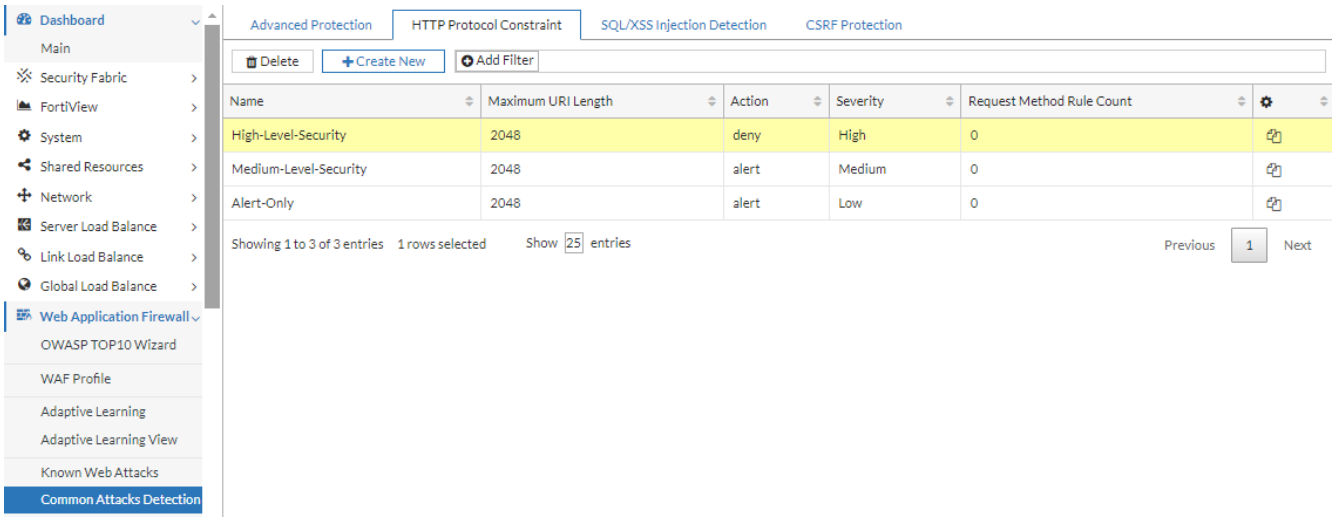
For more information, see the Configuring a Web Attack Signature policy on page 680.

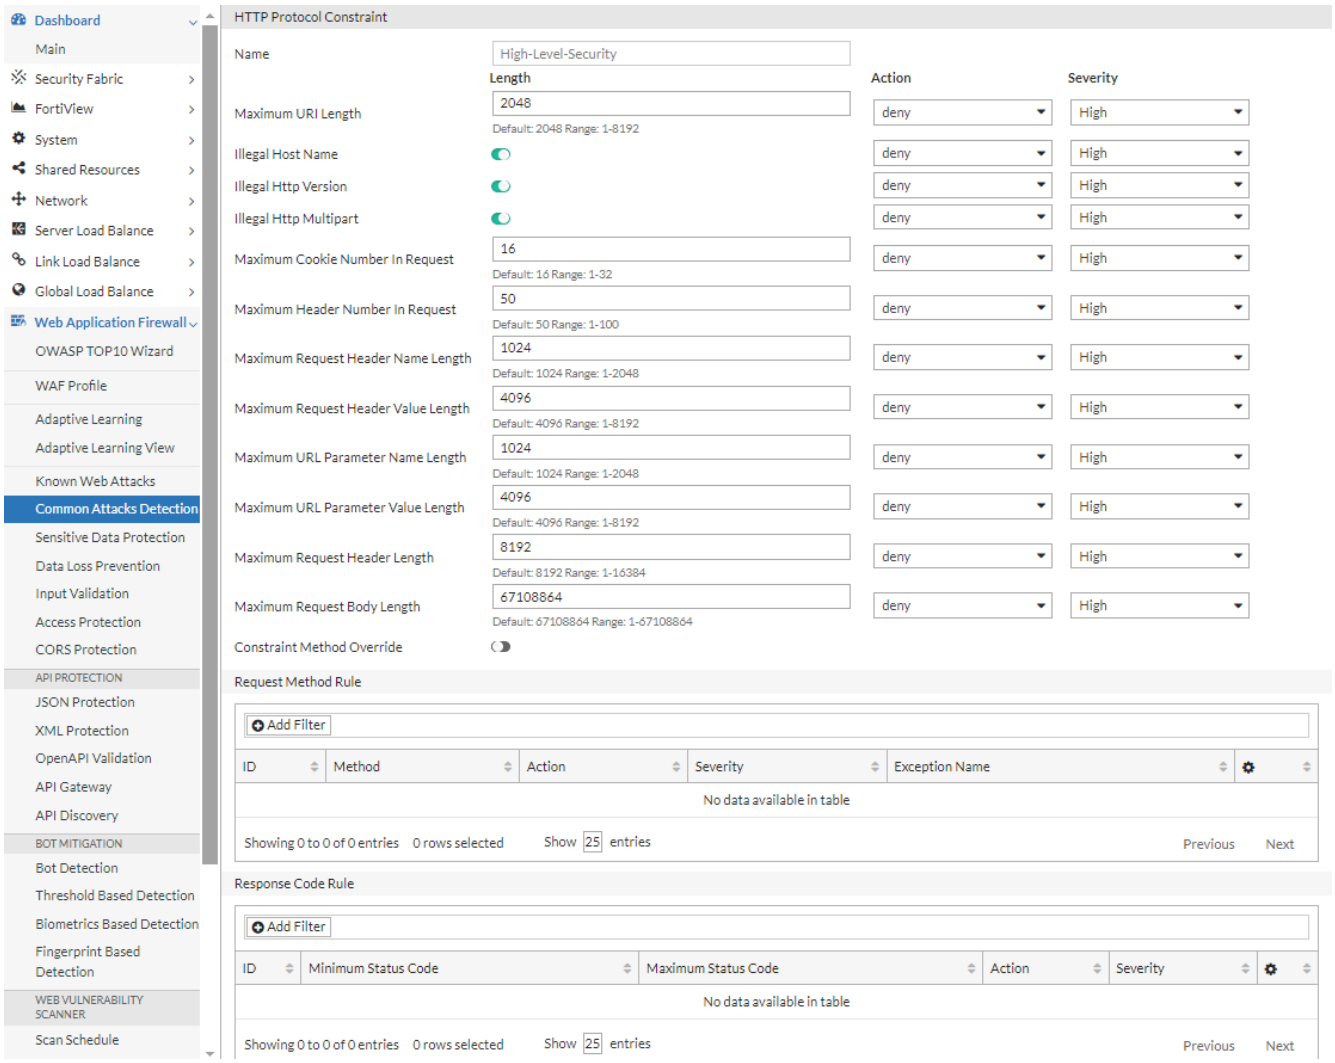
HTTP Protocol Constraint policy: High-Level-Security

As part of Common Attacks Detection, 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.

The predefined HTTP Protocol Constraint policy "High-Level-Security" is applied in the predefined WAF Profile also named "High-Level-Security", with the following policy settings:





For more information, see the Configuring an HTTP Protocol Constraint policy on page 690.

SQL/XSS Injection Detection policy: High-Level-Security

SQL/XSS Injection Detection policies are designed to identify and prevent SQL injection and cross-site scripting (XSS) attacks. These types of attacks occur when user-supplied data is improperly processed by an interpreter. In an SQL injection attack, malicious HTTP requests are crafted to manipulate SQL queries, allowing unauthorized access to a web application's database. XSS attacks, on the other hand, exploit vulnerabilities that enable the execution of malicious scripts within a user's web browser.

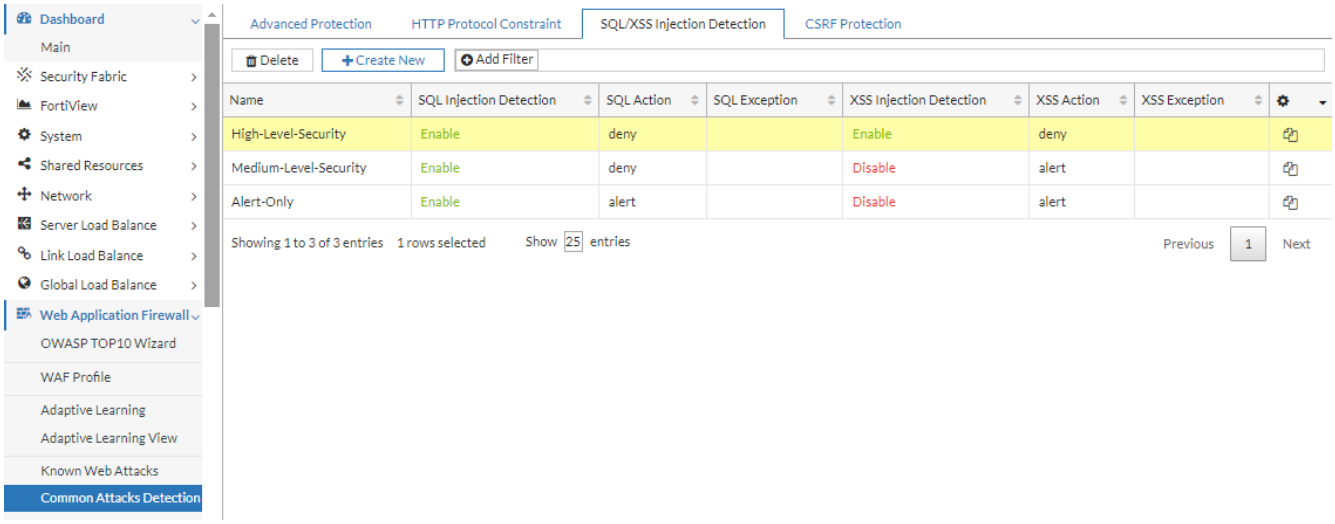
Unlike traditional signature-based detection methods, the Web Application Firewall (WAF) SQL and XSS injection detector module employs lexical analysis to identify these threats. This approach is both complementary and faster, offering enhanced detection capabilities.

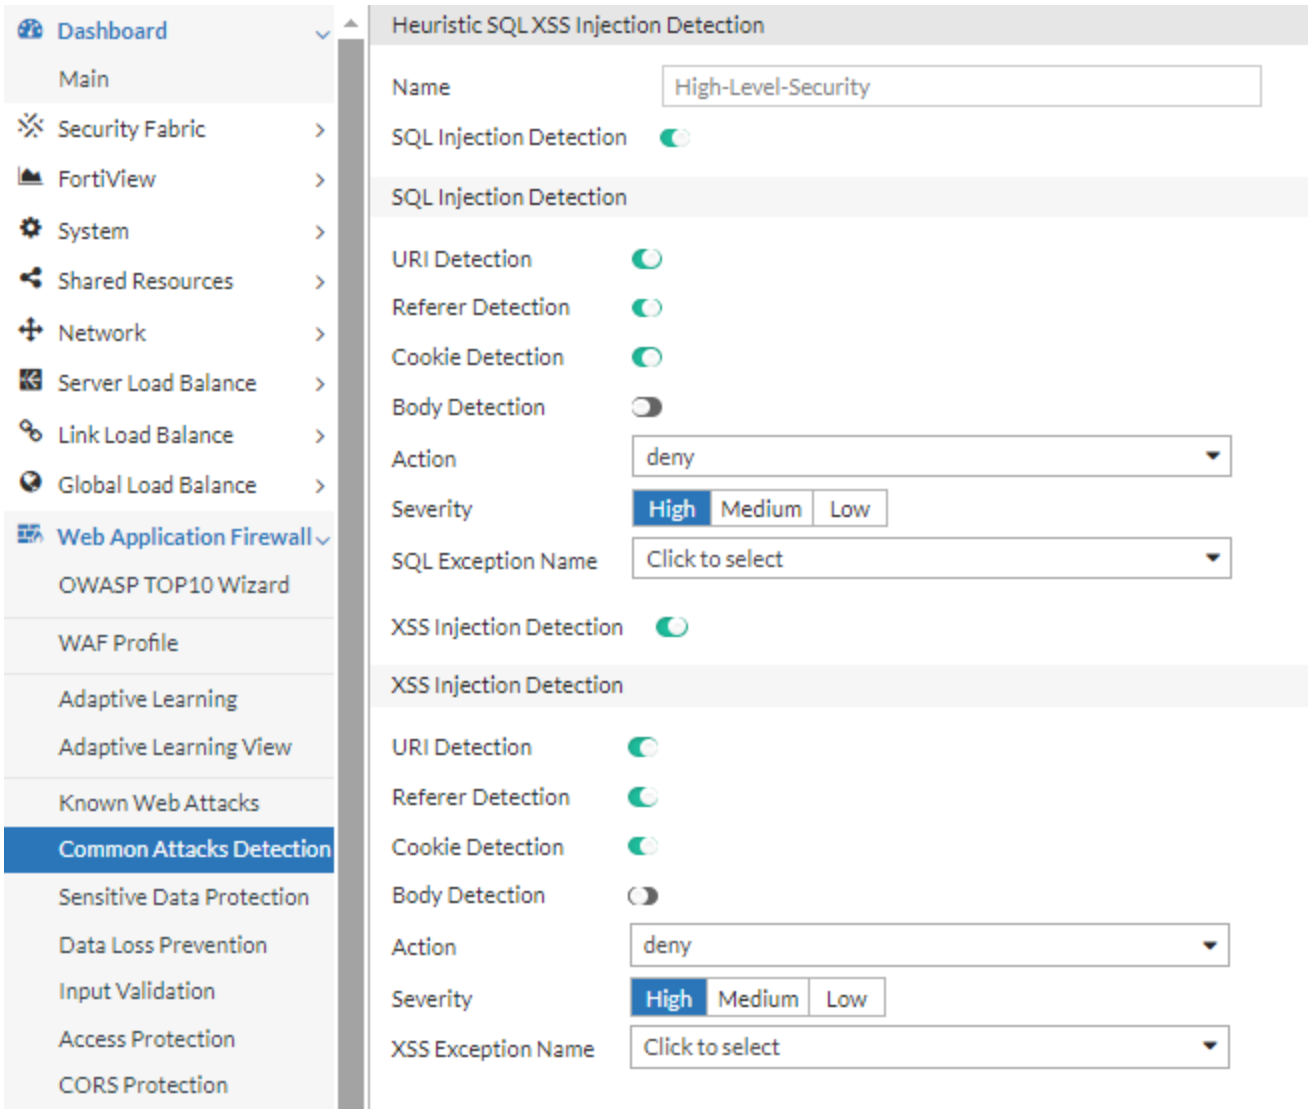
The policy configuration allows for enabling or disabling scanpoints, defining actions when traffic matches specific signatures, and setting the severity of events.

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.

The predefined SQL/XSS Injection Detection policy "High-Level-Security" is applied in the predefined WAF Profile also named "High-Level-Security", with the following policy settings:





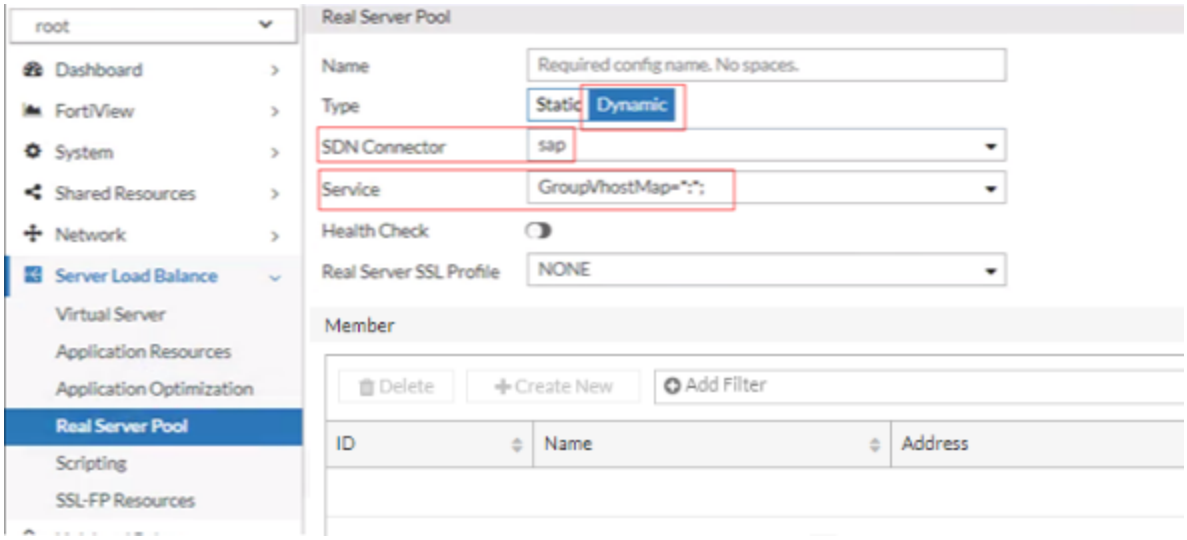
Step 4: Configuring Dynamic Traffic Distribution and Session Persistence

With the SAP Connector, health checks, and WAF profile configured, proceed to define load balancing policies and ensure session persistence. Set up dynamic real server pools populated from SAP instances and apply load balancing algorithms—such as round-robin or least connections—to efficiently distribute traffic across SAP servers, ensuring high availability and optimal resource utilization during peak traffic. Configure session persistence to ensure that user sessions are consistently routed to the same SAP server, which is critical for maintaining session continuity in processes like financial transactions and inventory management.

Configure the SAP real server pool and virtual server:

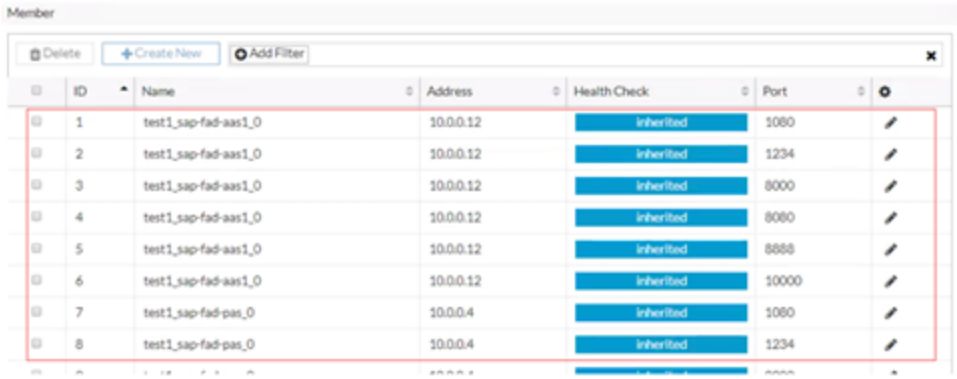
1. Create a dynamic real server pool that utilizes the SAP instances as real servers.
   a. Navigate to **Server Load Balance > Real Server Pool**.
   The configuration page displays the **Real Server Pool** tab.
   b. Click **Create New** to display the configuration editor. Configure the key settings:

Setting	Guideline
Name	Enter a unique name for the Real Server Pool.
Type	Set as Dynamic.
SDN Connector	Select the SAP connector previously created.
Service	Select the SAP service.
Health Check	Enable Health Check.
Health Check List	Select the Health Check profile configured as part of the initial configuration setup, or you can apply the predefined Health Check profile **LB_HLTHCK_HTTPS**.



For configuration details, see Using real server pools on page 552.

c. Save the configuration. After the real server pool configuration is saved initially, the Member section will automatically populate the real server pool with the SAP instances as real server pool members.



These real servers can also be viewed under the **Real Server** page.

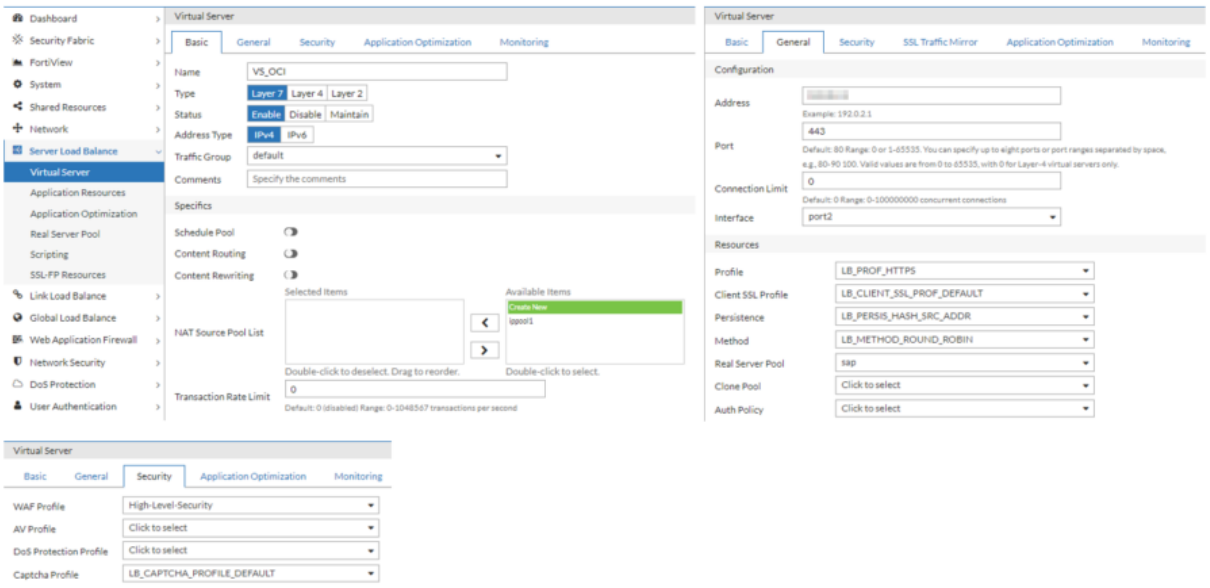
d. Save the Real Server Pool configuration to commit the changes for the Real Server Members.

2. Create the virtual server to handle load balancing for the SAP real server pool.

a. Navigate to **Server Load Balance > Virtual Server**.
The configuration page displays the **Virtual Server** tab.

b. Click **Create New** and select the **Advanced Mode** to display the configuration editor. Configure the key settings:

Setting	Guideline
Name	Enter a unique name for the Virtual Server.
Type	Select **Layer 7**.
Address	Enter the IP address for the Virtual Server.
Port	Enter **443** for the HTTPS common port
Interface	Select the interface to which you bind the VS.
Profile	Select the **LB_PROF_HTTPS** predefined Application Profile. For details, see Configuring Application profiles on page 451.
Client SSL Profile	Select the **LB_CLIENT_SSL_PROF_DEFAULT** predefined Client SSL Profile. For details, see Configuring Client SSL profiles on page 512.
Persistence	Select the **LB_PERSIS_HASH_SRC_ADDR** predefined Persistence profile. For details, see Configuring persistence rules on page 527.
Method	Select the **LB_METHOD_ROUND_ROBIN** predefined profile for the load-balancing method. The Round Robin method selects the next server in the series: server 1, then server 2, then server 3, and so on. For details, see Configuring load-balancing (LB) methods on page 526.
Real Server Pool	Select the SAP Real Server Pool configuration previously created.
WAF Profile	Select **High-Level-Security** (the recommended predefined WAF Profile). For details, see Configuring a WAF Profile on page 667.



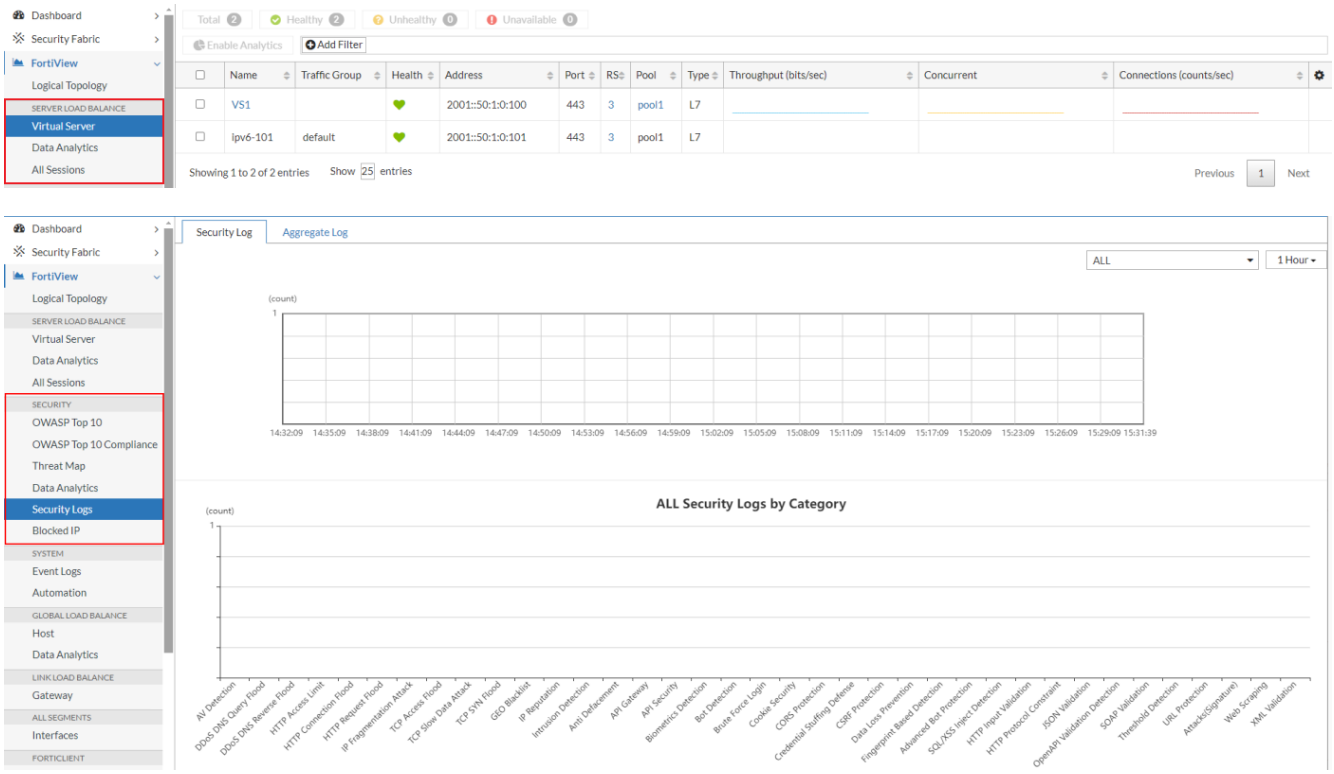
For configuration details, see Configuring virtual servers on page 384.

c. Save the configuration.

Step 5: Performance Monitoring and Optimization

Once the system is fully configured, set up performance monitoring tools within FortiADC to track traffic patterns, server performance, and overall network health. By regularly reviewing these metrics, you can fine-tune load balancing policies, optimize resource allocation, and ensure that SAP applications maintain peak performance even during high-demand periods like product launches or financial reporting cycles.

Monitoring tools such as **FortiView** enable real-time visibility and analysis of virtual server performance, providing detailed insights into server load balancing and security events.



Multi-Cloud Application Load Balancing

In multi-cloud architectures, where different applications are hosted across various cloud environments, optimizing load balancing becomes critical for maintaining application performance and reliability. FortiADC's Server Load Balancing (SLB) capabilities, augmented with external connectors, offer a technical solution for managing traffic across disparate platforms such as Amazon Web Services (AWS) and Oracle Cloud Infrastructure (OCI). This section explores how FortiADC's SLB dynamically allocates traffic based on real-time server metrics and load conditions. It will cover detailed SLB configuration, optimization strategies, and the benefits of using external connectors to achieve high availability and balanced resource utilization across multiple cloud-hosted applications.

User Scenario: Multi-Cloud Traffic Distribution with FortiADC

A leading global e-commerce enterprise operates its primary web application across multiple cloud environments—AWS and OCI—to achieve high availability, scalability, and enhanced performance. AWS is leveraged for its extensive global network, low-latency edge locations, and sophisticated services, while OCI is selected for its cost-effectiveness and superior computing capabilities.

To optimize traffic management across these heterogeneous cloud platforms, the company deploys FortiADC with integrated connectors for AWS and OCI. This configuration enables FortiADC to perform real-time traffic distribution based on server health metrics and load conditions, effectively mitigating the risk of overloading any single cloud environment and ensuring a consistent user experience.

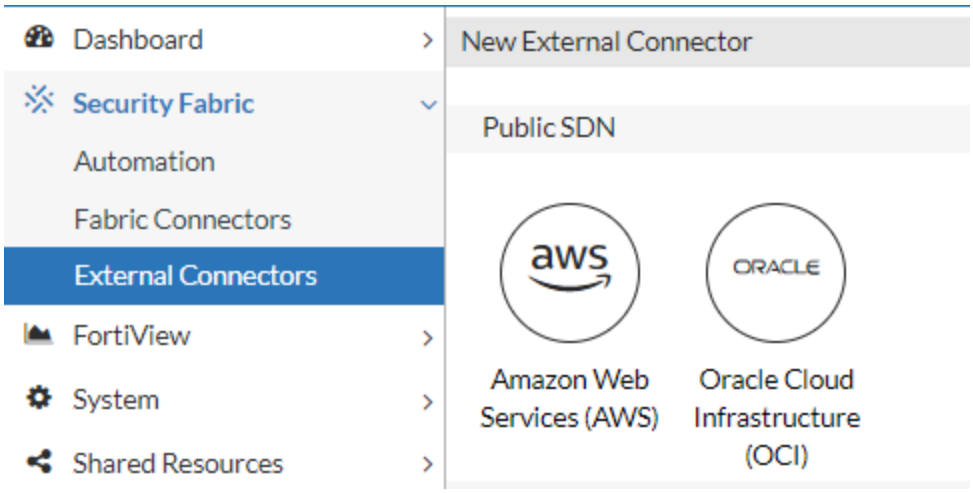
During a high-traffic global sales event, the e-commerce platform encounters a significant surge in user activity from diverse geographic regions. The IT team utilizes FortiADC to dynamically balance and route traffic between their AWS and OCI-based applications. By leveraging FortiADC's advanced load-balancing algorithms and real-time analytics, the team manages the influx of users efficiently, maintaining optimal performance and preventing any disruptions in service.

Step 1: Cloud Connector Setup

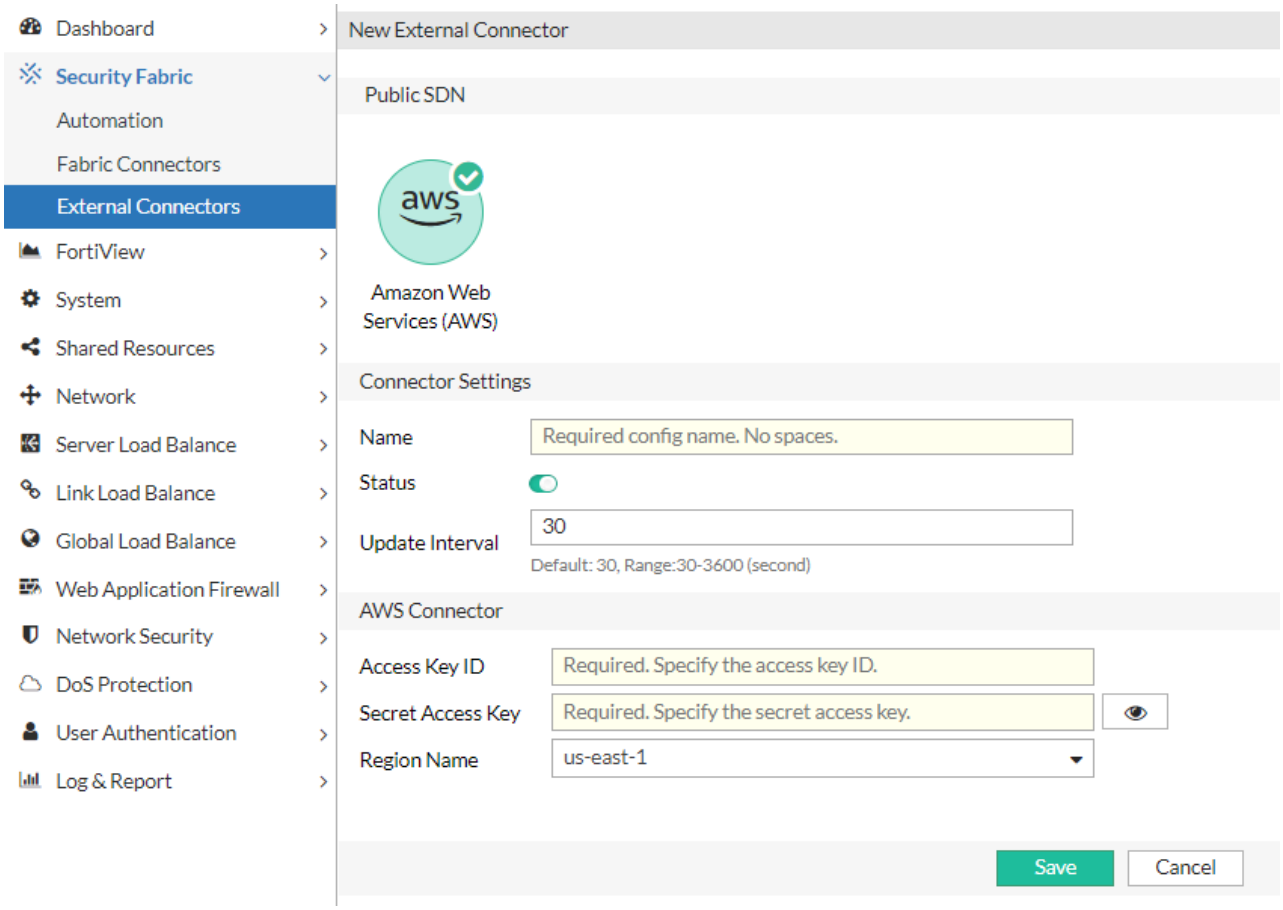
Configure the AWS and OCI cloud connectors on FortiADC to establish secure integration with both platforms. These connectors use API keys for authentication, allowing FortiADC to dynamically discover backend servers, such as AWS EC2 instances and OCI Compute instances. The connectors continuously monitor the availability and health of these servers, ensuring that FortiADC can route traffic to the most appropriate resource in real time.

To create an Amazon Web Services (AWS) connector:

1. Navigate to **Security Fabric > External Connectors**. Click **Create New** to display the configuration editor. Under the Public SDN section are the **Amazon Web Services (AWS)** and **Oracle Cloud Infrastructure (OCI)** connector options.

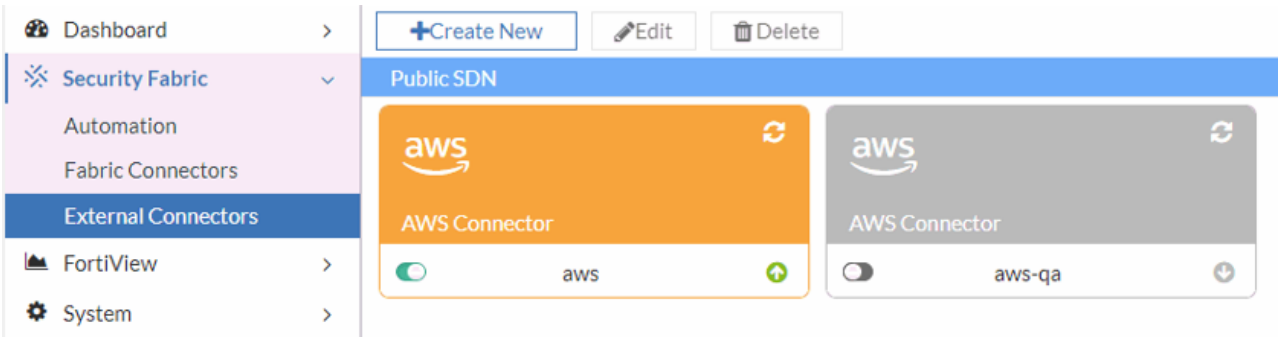


2. Click the Amazon Web Services (AWS) icon to display the configuration editor. Configure the required settings for the AWS connector.



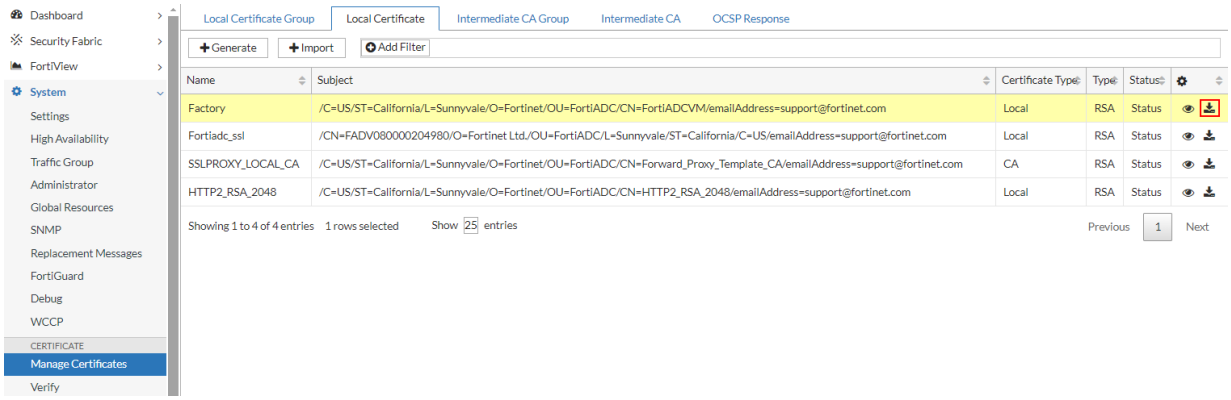
For configuration details, see Amazon Web Services (AWS) Connector on page 95.

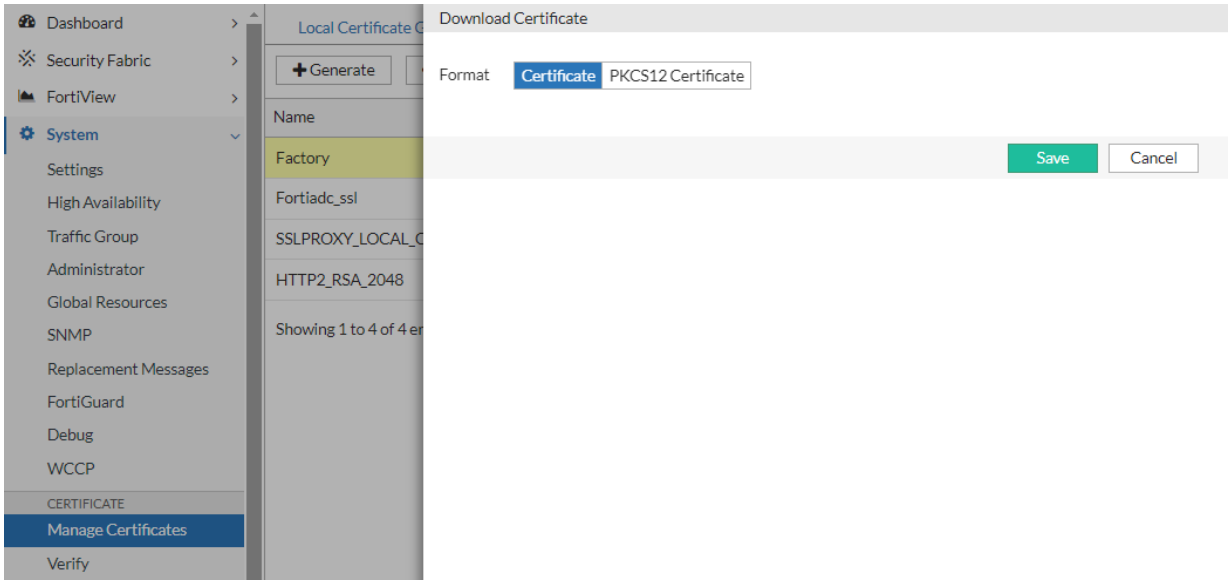
3. Click **Save** to commit the changes.
The newly created AWS connector will display on the External Connectors page.



To create an Oracle Cloud Infrastructure (OCI) connector:

1. Add an API key for the OCI user.
   This is required for secure access between FortiADC and OCI. The API key enables FortiADC to authenticate and interact with OCI services, allowing it to execute tasks such as querying resources, retrieving cloud-based data, and automating deployments.
   a. Download the **Factory Certificate** from **System > Certificates**.



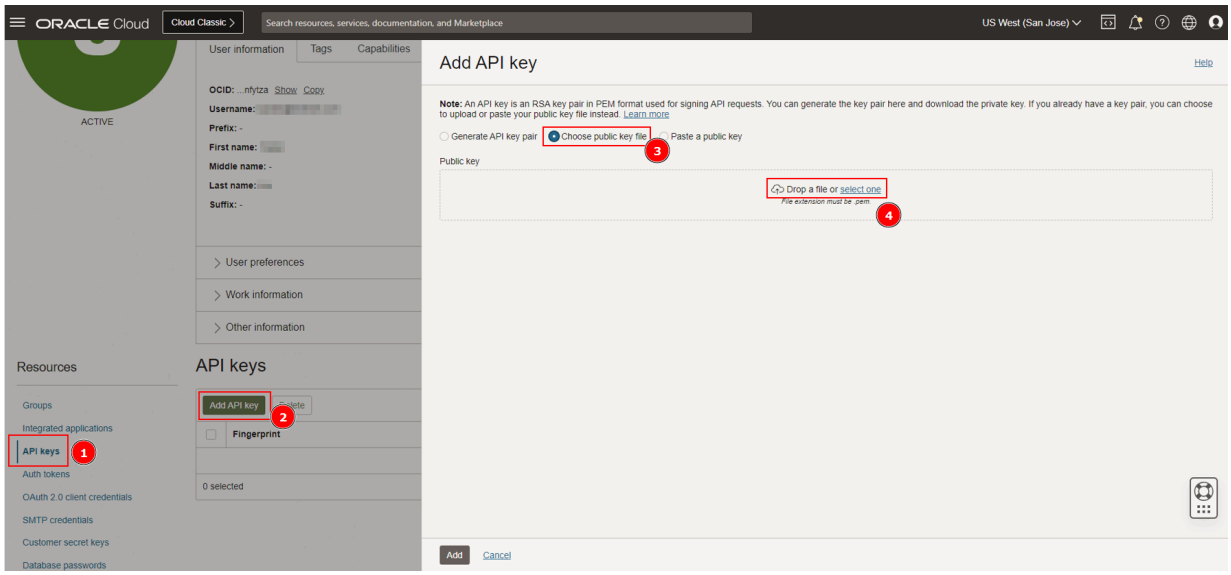


b. Convert the Factory Certificate to a public key file in PEM format using the following OpenSSL command.

```
openssl x509 -in <factory-certificate-file> -out <public-key-file>.pem
```

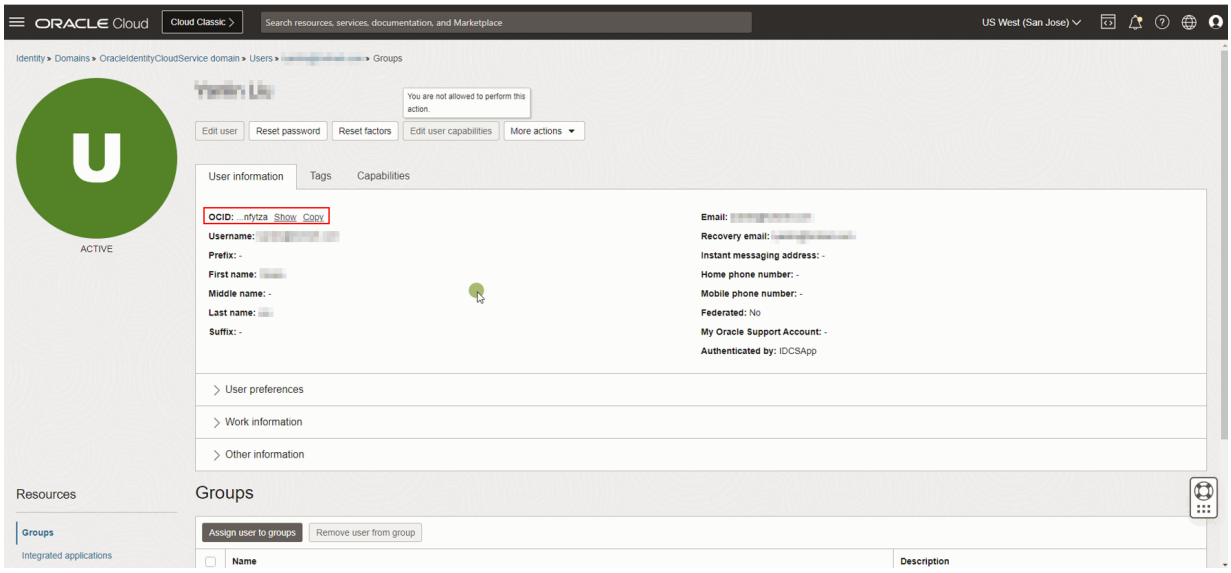
c. Upload the PEM File to Oracle Cloud.

From your Oracle Cloud Infrastructure (OCI) account, navigate to **Identify > Users** and select the user. Under **API Keys**, click Add **API Key** and upload the converted PEM file to register the API key.

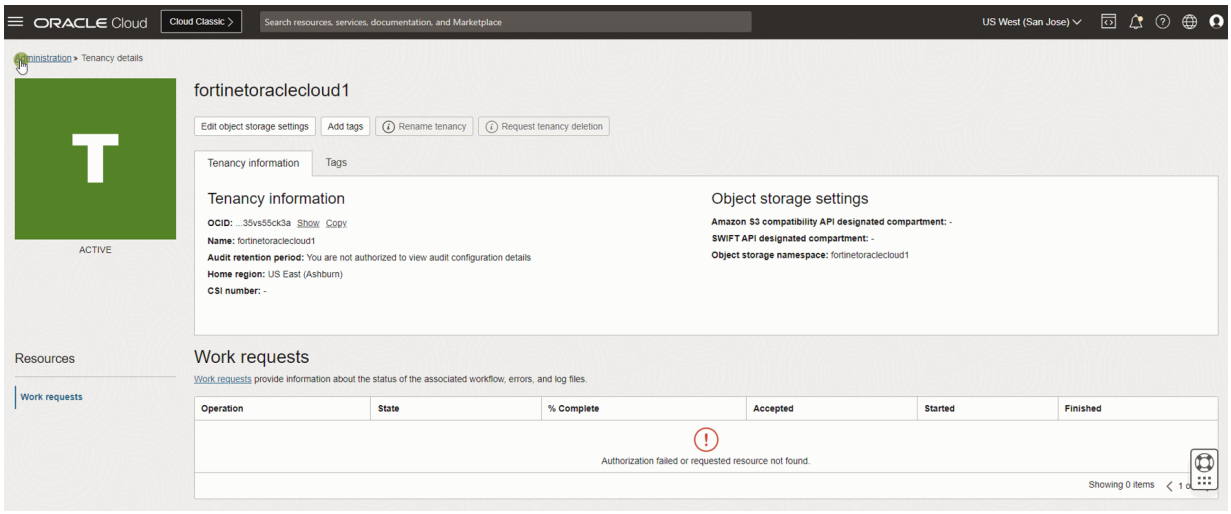


2. Retrieve the OCIDs from your Oracle Cloud Infrastructure (OCI) account required to create the OCI connector.

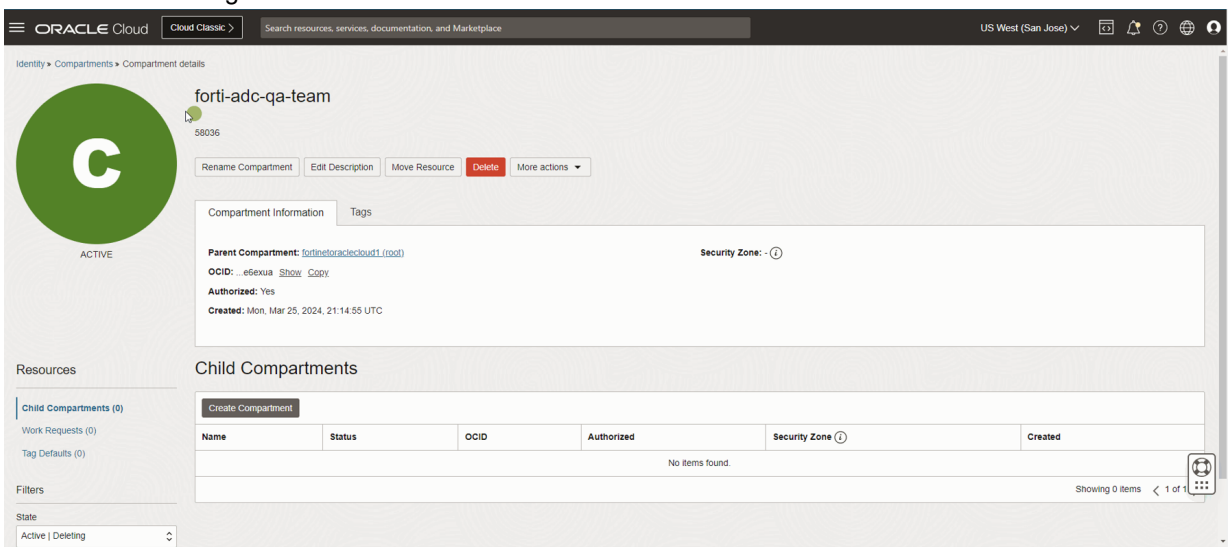
a. From your Oracle Cloud Infrastructure (OCI) account, navigate to **Identity > Users** and select the user. Copy the **OCID** from the **User information** to apply later in the OCI connector configuration.



b. Navigate to **Administration > Tenancy Details**. Copy the **OCID** from the **Tenancy information** to apply later in the OCI connector configuration.



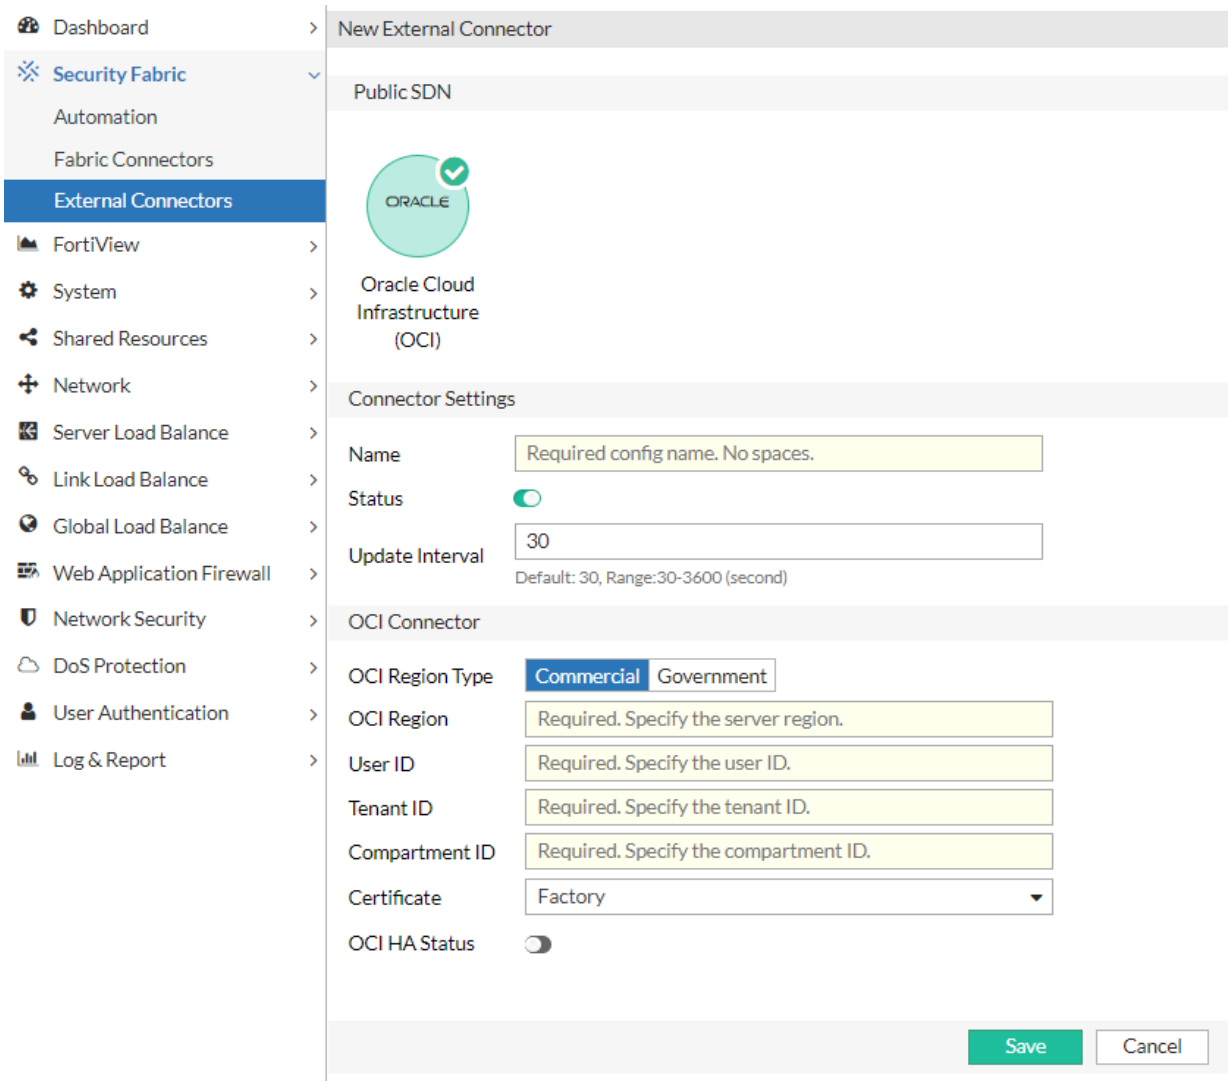
c. Navigate to Identity > Compartments. Copy the **OCID** from the **Compartment information** to apply later in the OCI connector configuration.



3. Create the a new Oracle Cloud Infrastructure (OCI) connector.

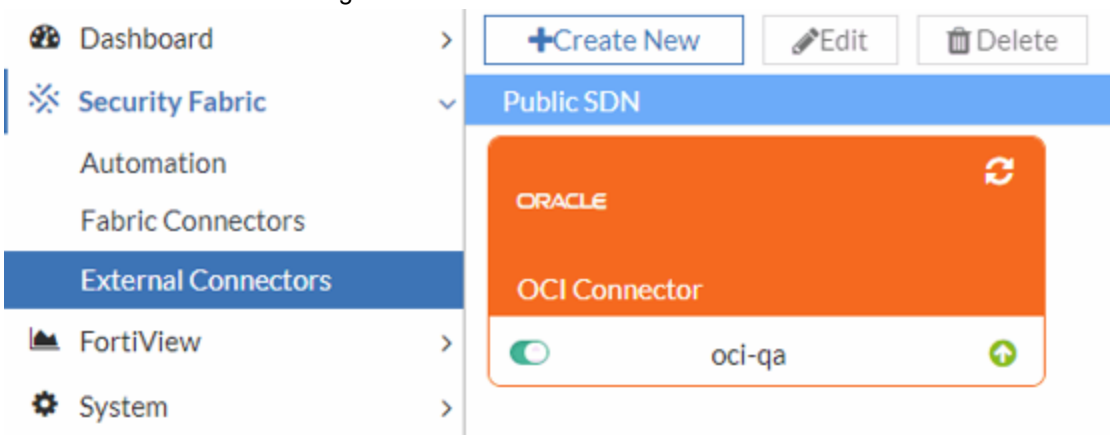
a. In FortiADC, go to **Security Fabric > External Connectors** and click **Create New** to configure a new Oracle Cloud Infrastructure (OCI) connector. Apply the OCIDs previously retrieved to the corresponding fields.

Parameter	Corresponding OCID
User ID	OCID from User information.
Tenant ID	OCID from Tenancy information.
Compartment ID	OCID from Compartment information.



For configuration details, see Oracle Cloud Infrastructure (OCI) Connector on page 98.

b. Click **Save** to commit the changes.

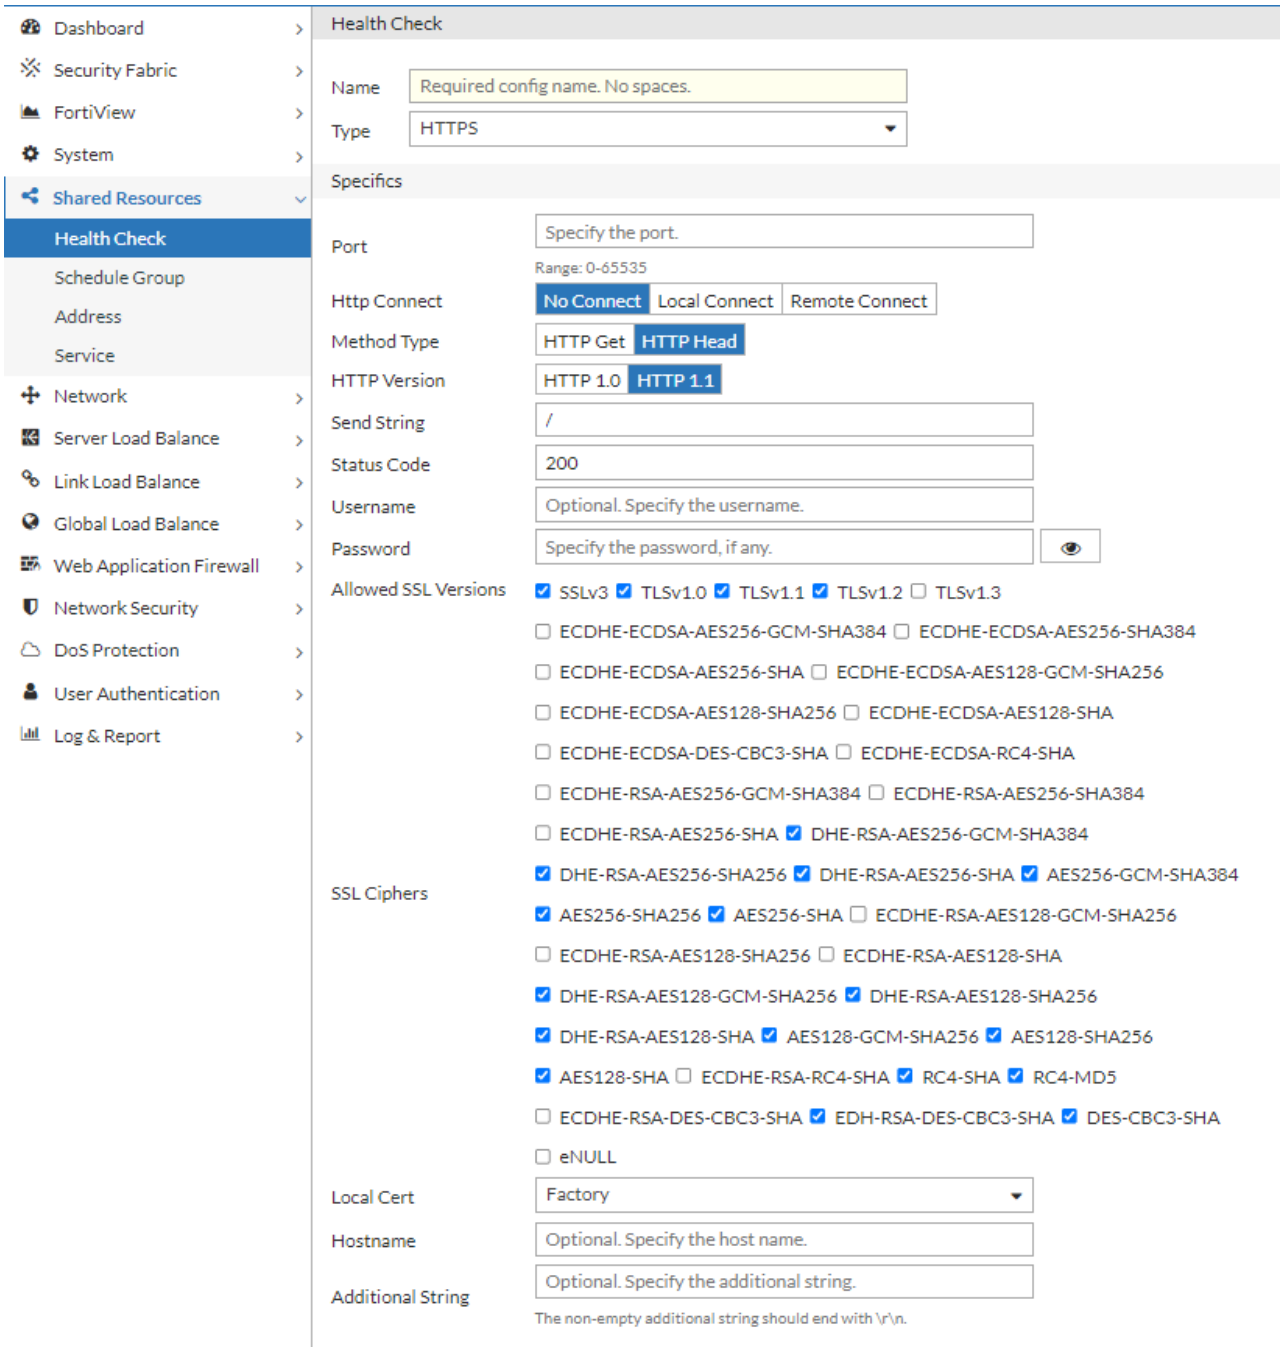


Step 2: Health Monitor Configuration

After setting up the cloud connectors, proceed with configuring FortiADC's core settings. This includes defining health checks for both AWS and OCI instances to ensure that only healthy servers receive traffic. By configuring detailed health check parameters, FortiADC can automatically remove any unhealthy instances from the load-balancing pool and redistribute traffic accordingly, maintaining application performance and availability.

To configure Health Check:

1. Navigate to **Shared Resources > Health Check**.
   The configuration page displays the **Health Check** tab.
2. Click **Create New** to display the configuration editor and configure the required settings.
   For this example, set the **Type** to **HTTPS**, as AWS and OCI applications typically communicate over the HTTPS protocol for secure data transmission.



For configuration details, see Configuring health checks on page 355.

3. Save the configuration.

Alternatively, predefined Health Check profiles can also be applied.

Step 3: Dynamic Traffic Distribution Configuration

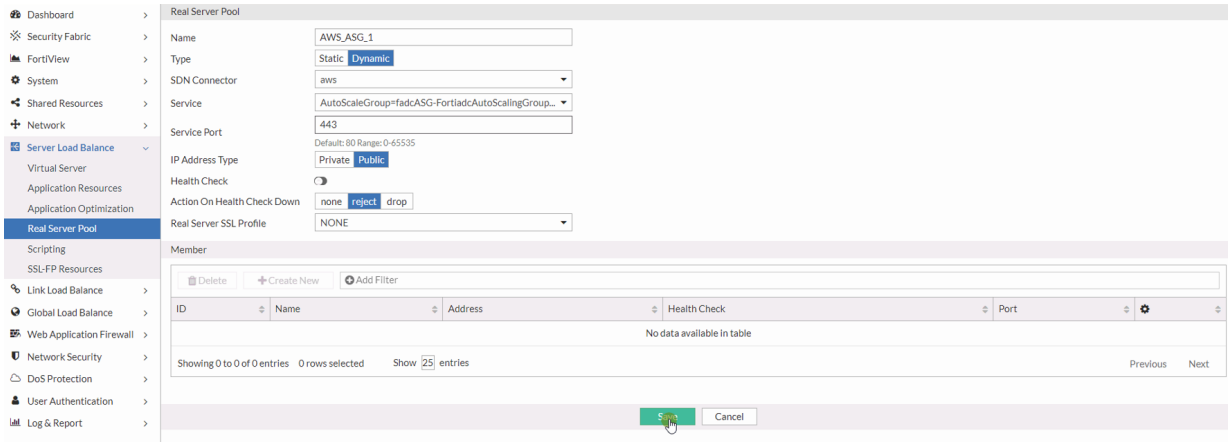
With health monitoring in place, FortiADC can now configure dynamic traffic distribution by defining real server pools populated from AWS and OCI instances, respectively. Separate virtual servers for each real server pool are then set up to apply routing policies based on server health, load, and user location. Selecting the appropriate load-balancing

algorithm—such as round-robin, least connections, or server capacity-based routing—is critical to optimizing traffic flow across the AWS or OCI real servers, ensuring efficient resource utilization and preventing any single real server from becoming a bottleneck.

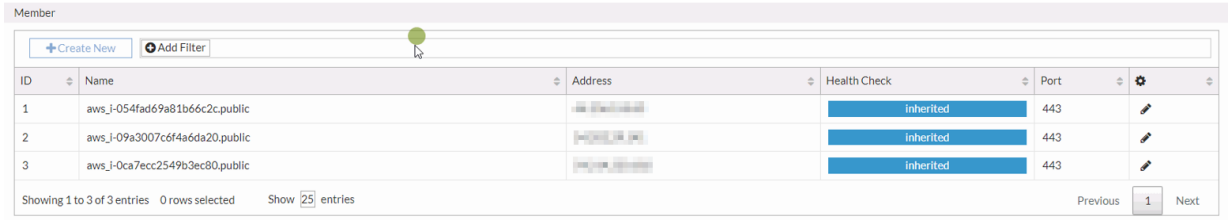
Configure the AWS real server pool and virtual server:

1. Create a dynamic real server pool that utilizes the AWS instances as real servers.
   a. Navigate to **Server Load Balance > Real Server Pool**.
   The configuration page displays the **Real Server Pool** tab.
   b. Click **Create New** to display the configuration editor. Configure the key settings:

Setting	Guideline
Name	Enter a unique name for the Real Server Pool.
Type	Set as Dynamic.
SDN Connector	Select the AWS connector previously created.
Service	Select the AWS service.
Service Port	Specify the port to use to access the AWS service.
Health Check	Enable Health Check.
Health Check List	Select the Health Check profile configured as part of the initial configuration setup, or you can apply the predefined Health Check profile **LB_HLTHCK_HTTPS**.



   For configuration details, see Using real server pools on page 552.
   c. Save the configuration. After the real server pool configuration is saved initially, the Member section will automatically populate the real server pool with the AWS instances as real server pool members.



These real servers can also be viewed under the **Real Server** page.

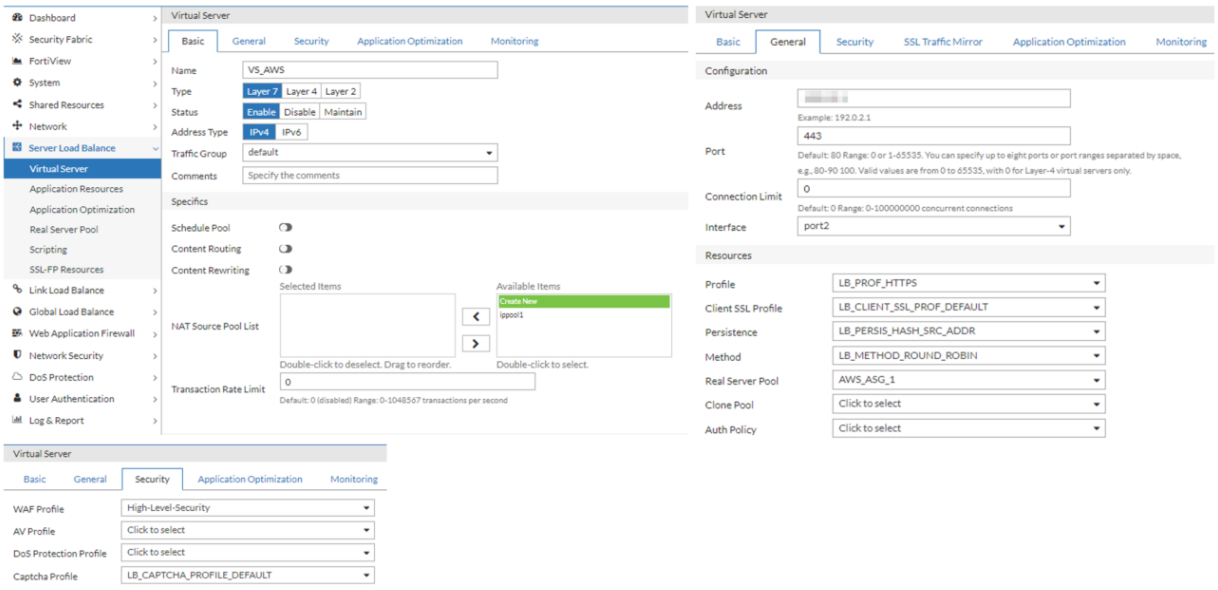
d. Save the Real Server Pool configuration to commit the changes for the Real Server Members.

2. Create the virtual server to handle load balancing for the AWS real server pool.

a. Navigate to **Server Load Balance > Virtual Server**.
The configuration page displays the **Virtual Server** tab.

b. Click **Create New** and select the **Advanced Mode** to display the configuration editor. Configure the key settings:

Setting	Guideline
Name	Enter a unique name for the Virtual Server.
Type	Select **Layer 7**.
Address	Enter the IP address for the Virtual Server.
Port	Enter **443** for the HTTPS common port
Interface	Select the interface to which you bind the VS.
Profile	Select the **LB_PROF_HTTPS** predefined Application Profile. For details, see Configuring Application profiles on page 451.
Client SSL Profile	Select the **LB_CLIENT_SSL_PROF_DEFAULT** predefined Client SSL Profile. For details, see Configuring Client SSL profiles on page 512.
Persistence	Select the **LB_PERSIS_HASH_SRC_ADDR** predefined Persistence profile. For details, see Configuring persistence rules on page 527.
Method	Select the **LB_METHOD_ROUND_ROBIN** predefined profile for the load-balancing method. The Round Robin method selects the next server in the series: server 1, then server 2, then server 3, and so on. For details, see Configuring load-balancing (LB) methods on page 526.
Real Server Pool	Select the AWS Real Server Pool configuration previously created.
WAF Profile	Select **High-Level-Security** (the recommended predefined WAF Profile). This profile is engineered to deliver robust, enterprise-grade protection by enabling a comprehensive suite of advanced security features. It includes an extensive library of Web Attack Signatures, stringent HTTP Protocol Constraints, and precise SQL/XSS Injection Detection mechanisms. These collectively fortify the environment against sophisticated and evolving web-based threats, ensuring unparalleled security for your virtual desktop infrastructure. For details, see Configuring a WAF Profile on page 667.



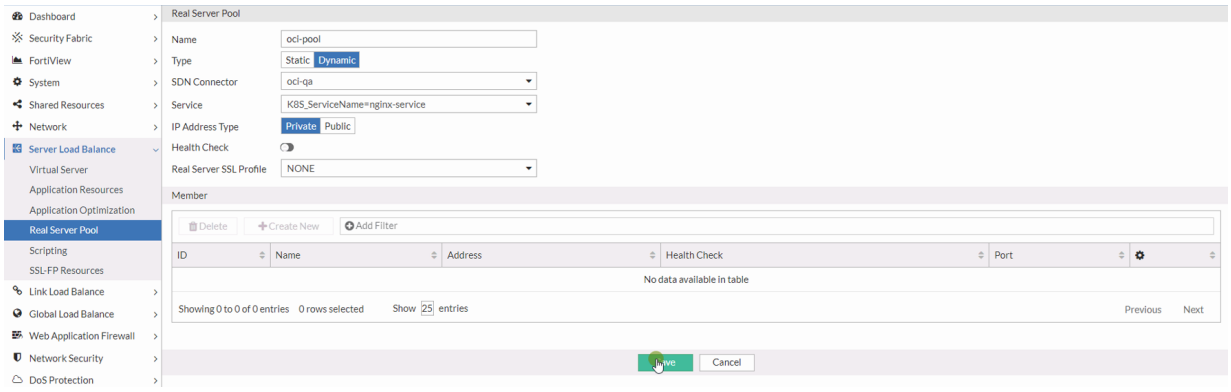
For configuration details, see Configuring virtual servers on page 384.

c. Save the configuration.

Configure the OCI real server pool and virtual server:

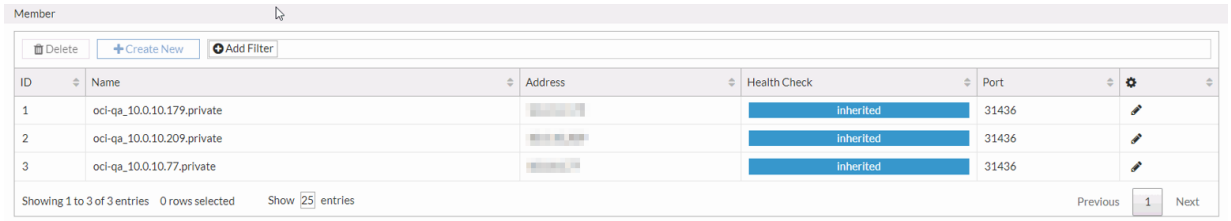
1. Create a dynamic real server pool that utilizes the OCI instances as real servers.
   a. Navigate to **Server Load Balance > Real Server Pool**.
   The configuration page displays the **Real Server Pool** tab.
   b. Click **Create New** to display the configuration editor. Configure the key settings:

Setting	Guideline
Name	Enter a unique name for the Real Server Pool.
Type	Set as Dynamic.
SDN Connector	Select the OCI connector previously created.
Service	Select the OCI service.
Health Check	Enable Health Check.
Health Check List	Select the Health Check profile configured as part of the initial configuration setup, or you can apply the predefined Health Check profile **LB_HLTHCK_HTTPS**.



For configuration details, see Using real server pools on page 552.

c. Save the configuration. After the real server pool configuration is saved initially, the Member section will automatically populate the real server pool with the OCI instances as real server pool members.



These real servers can also be viewed under the **Real Server** page.

d. Save the Real Server Pool configuration to commit the changes for the Real Server Members.

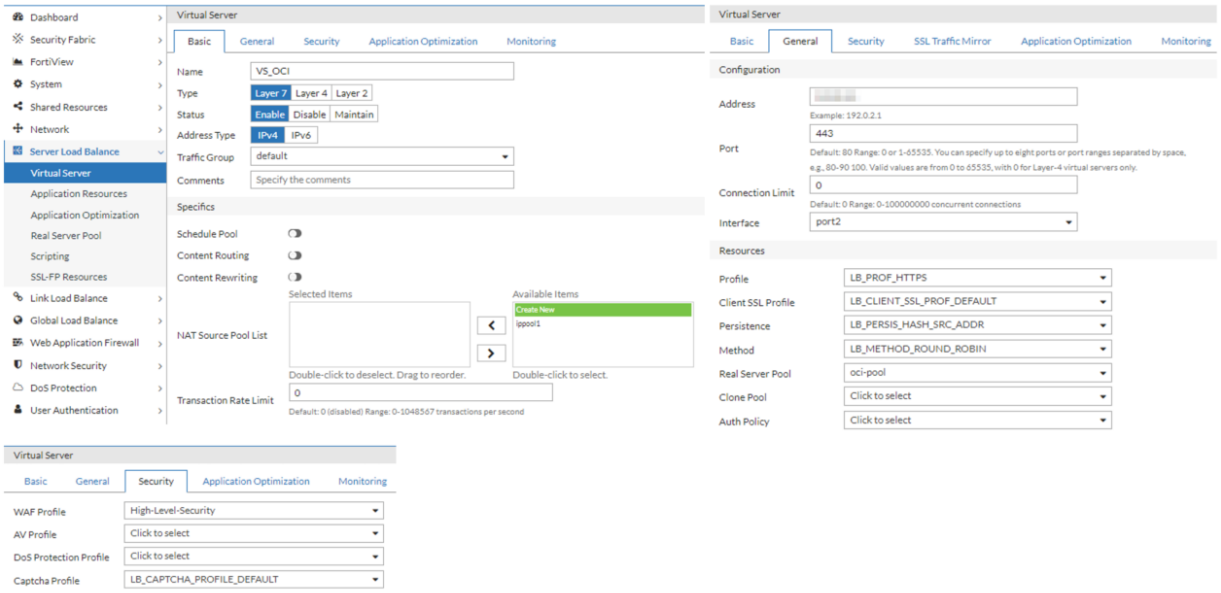
2. Create the virtual server to handle load balancing for the OCI real server pool.

a. Navigate to **Server Load Balance > Virtual Server**.
The configuration page displays the **Virtual Server** tab.

b. Click **Create New** and select the **Advanced Mode** to display the configuration editor. Configure the key settings:

Setting	Guideline
Name	Enter a unique name for the Virtual Server.
Type	Select **Layer 7**.
Address	Enter the IP address for the Virtual Server.
Port	Enter **443** for the HTTPS common port
Interface	Select the interface to which you bind the VS.
Profile	Select the **LB_PROF_HTTPS** predefined Application Profile. For details, see Configuring Application profiles on page 451.
Client SSL Profile	Select the **LB_CLIENT_SSL_PROF_DEFAULT** predefined Client SSL Profile. For details, see Configuring Client SSL profiles on page 512.
Persistence	Select the **LB_PERSIS_HASH_SRC_ADDR** predefined Persistence profile. For details, see Configuring persistence rules on page 527.

Setting	Guideline
Method	Select the **LB_METHOD_ROUND_ROBIN** predefined profile for the load-balancing method. The Round Robin method selects the next server in the series: server 1, then server 2, then server 3, and so on. For details, see Configuring load-balancing (LB) methods on page 526.
Real Server Pool	Select the OCI Real Server Pool configuration previously created.
WAF Profile	Select **High-Level-Security** (the recommended predefined WAF Profile). This profile is engineered to deliver robust, enterprise-grade protection by enabling a comprehensive suite of advanced security features. It includes an extensive library of Web Attack Signatures, stringent HTTP Protocol Constraints, and precise SQL/XSS Injection Detection mechanisms. These collectively fortify the environment against sophisticated and evolving web-based threats, ensuring unparalleled security for your virtual desktop infrastructure. For details, see Configuring a WAF Profile on page 667.



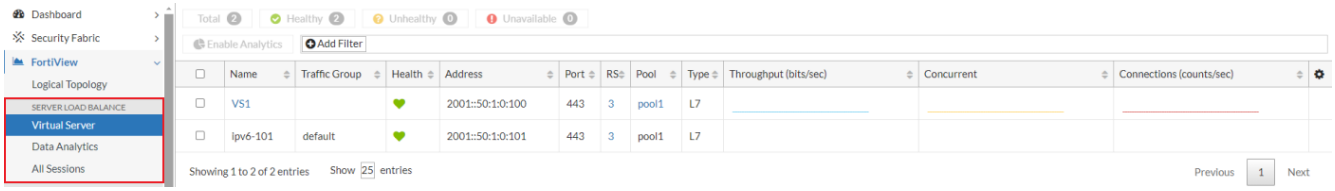
For configuration details, see Configuring virtual servers on page 384.

c. Save the configuration.

Step 4: Monitoring and Optimizing Performance

Once the traffic distribution and load-balancing algorithms are set, use FortiADC's built-in monitoring tools to track real-time performance metrics. This step focuses on analyzing traffic flow and server health of the AWS and OCI real server pools. Regular monitoring enables the IT team to fine-tune load-balancing settings, improve efficiency, and ensure that the e-commerce platform continues to provide a seamless experience during high-demand periods.

Monitoring tools such as **FortiView** enable real-time visibility and analysis of virtual server performance, providing detailed insights into server load balancing activity.



For more information, see the following topics about the FortiView Server Load Balance module:

- Virtual Servers on page 112
  - Virtual server details on page 113
  - Real server pool details on page 116
  - Real server pool member details on page 118
- Data Analytics on page 119
- All Sessions on page 121

Leveraging DevOps Tools (Ansible/Terraform) to Manage FortiADC

In today's fast-paced digital landscape, organizations are increasingly adopting DevOps practices to enhance operational efficiency, reduce deployment times, and minimize configuration errors. FortiADC, with its robust capabilities in application delivery and load balancing, integrates seamlessly with popular DevOps tools such as Ansible and Terraform. This section explores how these tools can be leveraged to automate the management and deployment of FortiADC configurations, ensuring consistent performance across multi-environment infrastructures while facilitating rapid and reliable changes.

User Scenario: Automating FortiADC Management with Ansible and Terraform

This user scenario focuses on "FinSecure," a financial services company that operates a complex multi-environment infrastructure featuring multiple FortiADC instances across development, staging, and production environments. To streamline operations and enhance efficiency while minimizing the risk of configuration errors, FinSecure seeks to automate the management and deployment of its FortiADC configurations using Ansible and Terraform. By implementing these DevOps tools, FinSecure aims to ensure consistent, rapid configuration changes that align with its operational requirements across all environments.

Role of Ansible

Ansible is an advanced automation framework designed for streamlined configuration management and application deployment. When applied to FortiADC, Ansible enables the automation of critical tasks such as the addition, modification, and deletion of FortiADC configurations. Leveraging the FortiADC Ansible Collection, which includes a set of purpose-built Ansible modules for managing FortiADC appliances, FinSecure can efficiently execute playbooks to enforce consistent configurations across its development, staging, and production environments. This automation framework minimizes manual intervention, ensuring rapid deployment and strict adherence to best practices and compliance policies, while enhancing operational reliability and configuration accuracy.

Step 1: Install FortiADC Ansible Collection

The first step in automating FortiADC management with Ansible is to install the FortiADC Ansible Collection. This collection provides a comprehensive suite of modules specifically tailored for configuring and managing FortiADC instances. By installing this collection, administrators can automate tasks such as adding, modifying, and deleting FortiADC configurations, streamlining operations across FinSecure's environments.

To install the FortiADC Ansible Collection, use the following command:

```
ansible-galaxy collection install fortinet.fortiadc
```

This command installs the latest version of the collection, enabling access to the necessary modules for FortiADC management. Before proceeding, ensure that the Ansible core version and Python environment meet the collection's version requirements to avoid compatibility issues. This setup establishes the foundation for automating configuration tasks efficiently within the infrastructure.

Step 2: Prepare Host Inventory

After installing the FortiADC Ansible Collection, the next step is to prepare the host inventory, a crucial component that defines the FortiADC instances Ansible will manage. The host inventory file specifies the target devices, allowing administrators to ensure that playbooks are executed on the correct FortiADC instances. Accurate host inventory preparation is essential for precise and consistent configuration management across multiple environments.

To prepare the host inventory, create or modify the inventory file located in the default path:

`/etc/ansible/hosts`

In this example, we create a host group named `[ADC1]` and provide the IP address of the FortiADC instance:

```
[ADC1]
44.215.84.44
```

By correctly defining the target hosts in the inventory file, you ensure that Ansible playbooks are applied to the designated FortiADC instances, enabling controlled and reliable automation across FinSecure's infrastructure.

Step 3: Prepare FortiADC Ansible Playbook

With the host inventory established, the next step is to create the FortiADC Ansible playbook. The playbook defines the specific tasks and configurations to be applied to FortiADC instances, allowing administrators to automate complex operations such as adding new servers, modifying existing configurations, or enforcing policy updates. A well-structured playbook ensures that configuration changes adhere to best practices, enhancing both the consistency and reliability of FortiADC management across multiple environments.

In this example, we will prepare a playbook to add a new real server to FortiADC. This is achieved using the `fadcos_real_server` module, which allows you to define and manage real server configurations.

```
---
- name:
  hosts: ADC1
  vars:
       ansible_network_os: fortinet.fortiadc.fadcos
       ansible_user: ansible
       ansible_passowrd: aa
       ansible_httpapi_use-ssl: yes
       ansible_httpapi_validate_certs: no
       ansible_httpapi_port: 443
  connection: httpapi
  gather_facts: false
  tasks:
    - name: add real server 1
      fadcos_real_server:
       action: add
       name: rs1
       ip: 6.6.6.1
       status: enable
       vdom: root

    - name: add real server 1
      fadcos_real_server:
       action: add
       name: rs2
       ip: 6.6.6.2
```

```
        status: enable
        vdom: root
```

For more complex configurations, consult the usage examples provided in the FortiADC Ansible Collection documentation for specific module options. By referencing the correct modules, you can automate various FortiADC configuration tasks, ensuring scalable, error-free management across your infrastructure.

Step 4: Run the Playbook

The final step in the Ansible workflow is to execute the prepared playbook, applying the defined configurations and tasks to the specified FortiADC instances. This critical action marks the transition from preparation to implementation, allowing administrators to execute configuration changes swiftly and efficiently while minimizing the potential for human error.

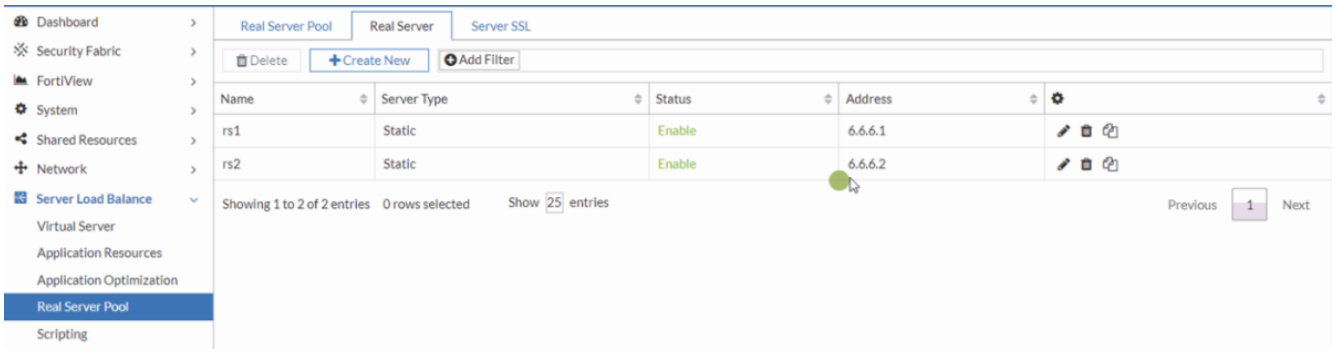
To run the playbook, execute the following command in the terminal:

```
ansible-playbook xxxx.yml -vvv
```

Replace `xxxx.yml` with the name of your specific playbook file. The `-vvv` option enables verbose output, providing detailed feedback on the playbook execution process, including task completion statuses and any encountered issues. This level of detail aids in troubleshooting and ensures that administrators can monitor the successful application of configurations to FortiADC instances, enhancing operational visibility and control over the automation process.



Access the FortiADC instance to confirm that the automated actions executed by Ansible have been successfully applied. In this example, two new real servers have been created with the IP addresses **6.6.6.1 (rs1)** and **6.6.6.2 (rs2)**.



Role of Terraform

Terraform is an Infrastructure as Code (IaC) tool engineered for provisioning and managing cloud infrastructure, providing an efficient framework for orchestrating the lifecycle of FortiADC instances within FinSecure's multi-environment architecture. Utilizing the FortiADC Terraform provider, FinSecure can declaratively specify infrastructure resources, enabling automated provisioning, modification, and decommissioning of FortiADC configurations. This approach ensures that the infrastructure adheres to predefined desired states, enhancing consistency and facilitating robust version control and change management practices. In this use case, the FortiADC Terraform provider is employed to interact with FortiADC-supported resources, allowing users to effectively manage and alter configurations using the requisite authentication credentials.

Step 1: Install FortiADC Terraform Provider

The initial step in utilizing Terraform to manage FortiADC is to install the FortiADC Terraform provider. This provider serves as a crucial bridge between Terraform and FortiADC, enabling administrators to programmatically interact with and manage FortiADC resources. By completing this installation, FinSecure establishes the capability to define and provision FortiADC configurations through Infrastructure as Code (IaC).

To install the FortiADC Terraform provider, incorporate the following configuration block into your Terraform configuration file:

```
terraform {
  required_providers {
    fortiadc = {
      source = "fortinetdev/fortiadc"
      version = "1.2.0"
    }
  }
}

provider "fortiadc" {
  # Configuration options
}
```

After adding the configuration, run the following command to initialize the Terraform workspace:

```
terraform init
```

This command downloads the necessary provider binaries and prepares the environment for managing FortiADC resources, ensuring FinSecure is equipped for efficient infrastructure management.

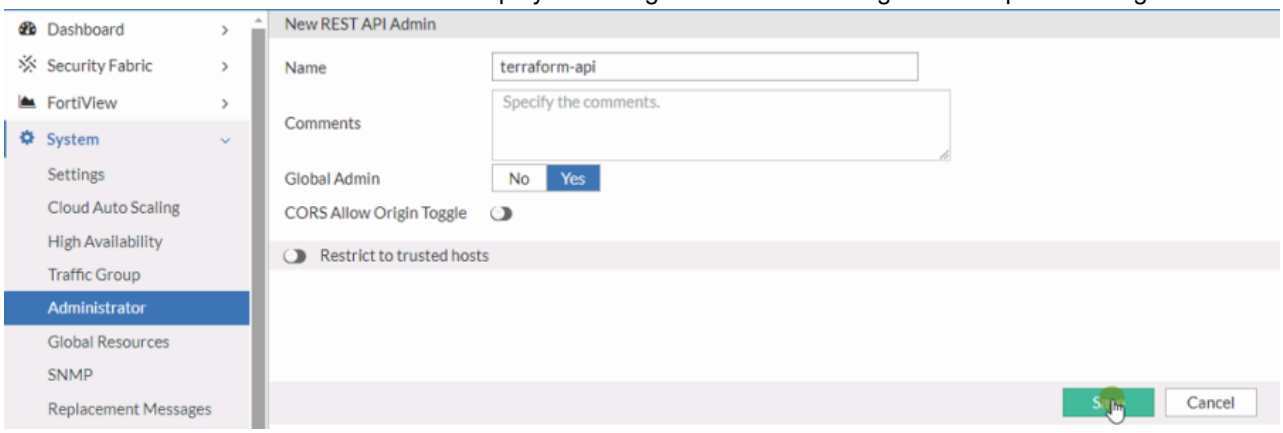
Step 2: Generate an API Token for FortiADC

After installing the Terraform provider, the next critical step is to generate an API token for FortiADC. This token facilitates secure authentication and communication between Terraform and the FortiADC instances. Proper configuration of this token is essential to ensure that Terraform has the necessary permissions to manage resources effectively while adhering to security protocols.

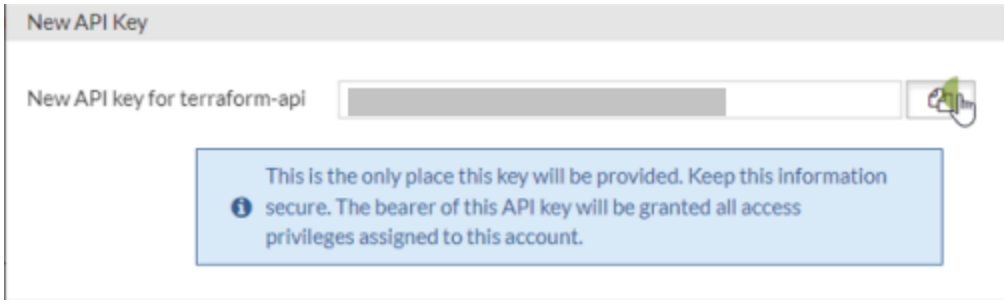
To generate an API token, create a new REST API administrator in FortiADC. Once the administrator account is created, the system will generate an API token that is displayed only once. It is crucial to record this token, as it cannot be retrieved later. This token will be required for subsequent configuration and operation within the tutorial. Ensure that the FortiADC provider is configured with the correct credentials, including the generated API token, before proceeding with any Terraform commands.

To create a REST API Administrator and obtain the API token:

1. In FortiADC, navigate to **System > Administrator**.
   The configuration page displays the **Admin** tab.
2. Click **Create New > REST API Admin** to display the configuration editor. Configure the required settings:



3. Once you successfully save the configuration, the **New API Key** dialog appears to display the newly generated API token. Record this API token for use later.



Step 3: Configure and Apply the Terraform Provider for FortiADC

With the API token in place, the next step involves configuring and applying the Terraform provider for FortiADC. This process includes defining the desired state of FortiADC instances within Terraform configuration files, which allows for automated provisioning and management of configurations.

To configure the FortiADC provider, include the API token obtained in Step 2 within the provider block. The configuration should resemble the following:

```
provider "fortiadc" {
  api_token = "YOUR_API_TOKEN"  # Replace with the API token from Step 2
}
```

After configuring the provider, you can utilize various FortiADC resources to define specific configurations. For instance, to add a new static router entry, use the following resource block:

```
resource "fortiadc_router_static" "example" {
  name = "ExampleRouter"
  # Additional configuration parameters go here
}
```

Refer to the FortiADC Terraform provider documentation for comprehensive examples and usage details related to different resource types. By executing the Terraform commands, FinSecure can provision and manage its FortiADC configurations efficiently, ensuring consistency and alignment with organizational requirements while leveraging the benefits of IaC.

Troubleshooting

This section includes the following topics:

- Logs on page 1072
- Tools on page 1072
- Debug on page 274
- Solutions by issue type on page 1076
- Resetting the configuration on page 1083
- Restoring firmware (“clean install”) on page 1083
- Additional resources on page 1086

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 dumpsystem
```

```
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 dumpsystem-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 dumpsystem-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.

Diff

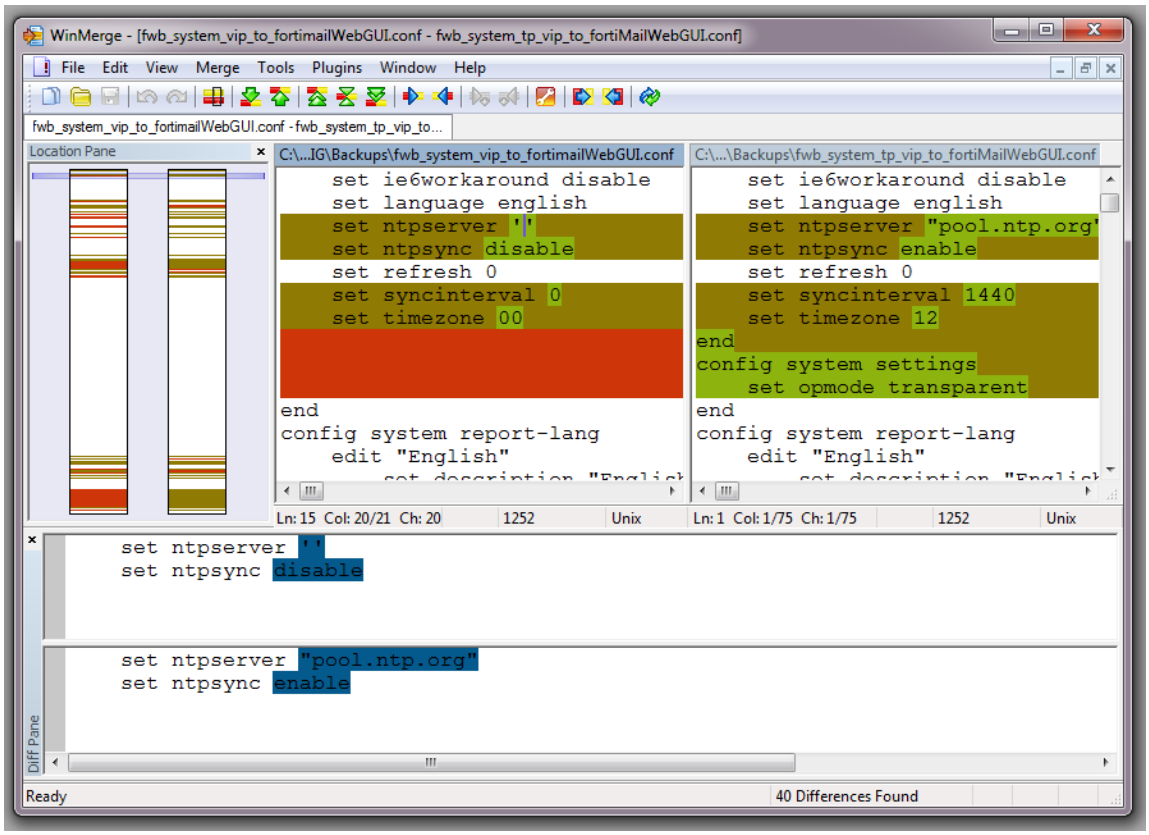
You can compare backups of the core configuration file with your current configuration. This can be useful if, for example:

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.

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 `traceroute` 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 `traceroute`. However, if the appliance does not respond, and there are no firewall policies that block it, ICMP type 0 (`ECHO_REPSPONSE`) might be effectively disabled.

To enable ping and traceroute 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 `traceroute` 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 blocklisting, 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-ottawatc_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.cgcil.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.

Appendices

This section includes the following FortiADC Appendix pages.

- Maximum Configuration Values on page 1088
- Supported image versions for EOS models on page 1096

Maximum Configuration Values

Maximum configuration objects - hardware models on page 1088 and Maximum configuration objects - virtual appliances on page 1093 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		100F/120F/200F/220F	300D/400D/300F/320F/400F/420F/ 1000F/1200F/2000F/2200F	4000F/4200F/5000F
System				
Administration	Administrative users	300	300	300
	Access profiles	16	64	64
	Virtual domains (VDOMs)	100F/120F/200F: 2 220F: 10	300D/400D/300F/1000F/ 2000F:10 400F: 20 320F: 25 420F: 45 1200F: 45 2200F: 60	90
Certificates	Any configuration object	256	256	256

Parameters		100F/120F/200F/220F	300D/400D/300F/320F/400F/420F/ 1000F/1200F/2000F/2200F	4000F/4200F/5000F
Shared Resources	Address	1024	2048	4096
	Address group	256	256	256
	Health checks	128	256	512
	ISP address book	32	32	32
	Schedule	256	256	256
	Schedule group	64	64	64
	Service	1024	2048	4096
	Service group	256	256	256
SNMP	SNMP community	16	16	16
	SNMP community Host	16	16	16
	SNMP user	16	16	16
Networking				
Interface	Physical network interfaces	100F/120F: 4 200F: 6 220F: 8	300D/400D: 4 300F: 8 320F: 10 400F: 10 420F: 12 1000F: 21 1200F: 24 2000F: 25 2200F 20	4000F: 15 4200F/5000F: 12
	VLAN interfaces	256	512	1024
Routing	ARP table	4096	4096	4096

Parameters		100F/120F/200F/220F	300D/400D/300F/320F/400F/420F/ 1000F/1200F/2000F/2200F	4000F/4200F/5000F
	entries (per VDOM)			
	Static routes	2048	4096	4096
	Policy routes	64	128	256
	ISP routes	32	32	32
NAT	Any configuration object	256	256	256
QoS	Any configuration object	256	256	256
Packet capture	Table	5	5	5
User				
Any configuration object		256	256	256
Server Load Balancing				
Virtual Servers	L4	1024	2048	4096
	L7	100F/120F/200F : 256 220F : 1024	300D/400D/300F/320F/400F/420F : 512 1000F : 1024 1200F/2000F/2200F : 2048	4096
	L7 HTTPs	100F/120F/200F : 128 220F: 1024	300D/400D/300F/400F : 256 1000F/2000F : 512 1200F/2200F: 2048	2048
Real Server Pool	Pools	1024	2048	4096
	Pool members	1024	2048	4096
	Real server SSL profiles	256	256	256

Parameters		100F/120F/200F/220F	300D/400D/300F/320F/400F/420F/ 1000F/1200F/2000F/2200F	4000F/4200F/5000F
Resources	Profiles	256	256	256
	Cache policies	256	256	256
	Compression policies	256	256	256
	Persistence policies	128	256	512
	Method policies	64	128	256
	Authentication policies	256	256	256
	Scripts	256	256	256
Content Rules	Content routing rules *Indicates maximum limits per VDOM. True maximums are calculated by ***PER_VDOM_LIMIT × MAX_VDOM***.	256*	512*	1024*
	Content rewriting rules	256	512	1024
Link Load Balancing				
Link Group	Gateway	1024	2048	4096
	Link group	512	1024	2048
	Link group member	1024	2048	4096
Virtual Tunnel Group	Virtual tunnel group	512	1024	2048

Parameters		100F/120F/200F/220F	300D/400D/300F/320F/400F/420F/ 1000F/1200F/2000F/2200F	4000F/4200F/5000F
	Virtual tunnel member	256	256	256
Policy	LLB policy rule	512	1024	1024
Global Load Balancing				
Zones		4096	4096	4096
FQDN hosts		4096	4096	4096
GLB policies		256	300D/300F/320F/400F: 256 420F/1000F/1200F/2000F/2200F: 1024	1024
Security				
Any configuration object		256	256	256
Log & Report				
Remote Syslog Servers		3	3	3

Maximum configuration values when HW SSL acceleration is enabled

Maximum configuration values - hardware models when HW SSL acceleration is enabled on page 1093 show the maximum number of configuration objects for hardware models able to support HW SSL acceleration.

The maximum number of processes that a virtual server is able to support can be increased by enabling HW SSL acceleration. This depends on whether the virtual server is enabled for alone mode, and which SSL hardware is supported by each model. For virtual servers enabled for alone mode, each will be handled by a separate httproxy process, whereas multiple virtual servers with alone mode disabled may share a single process. This may allow the number of virtual servers disabled for alone mode be unlimited for a virtual domain for when HW SSL acceleration is enabled.

For models that support HW SSL acceleration, they will either be compatible with the Cavium SSL or QAT SSL. For models using QAT SSL, the maximum number of processes depends on whether polling or epoll mode is enabled. Polling mode allows for four times the number of processes than epoll mode, however, epoll mode is higher performing. For models using Cavium SSL, there are no restrictions on the number of processes.

Note: Since the 6.2.0 release, the default mode for QAT SSL has been changed to polling.

Maximum configuration values - hardware models when HW SSL acceleration is enabled

Parameters	400D/400F/420F/1000F/ 1200F/2000F/2200F	4000F/4200F/5000F
Virtual Servers with alone mode enabled	2048	4096
HW SSL Process Number with QAT SSL in polling mode	400F: 64 420F: 128 1200F/2200F: 192	4200F/5000F: 192
HW SSL Process Number with QAT SSL in epoll mode	400F: 16 420F: 32 1200F/2200F: 48	4200F/5000F: 48
HW SSL Process Number with Cavium SSL	400D/1000F/2000F: 30720	4000F: 61440

Maximum configuration objects - virtual appliances

Parameters		VM01	VM02	VM04	VM08	VM16	VM32
System							
Administration	Administrative users	300	300	300	300	300	300
	Access profiles	8	16	64	64	64	64
	Virtual domains (VDOMs)	10	10	10	10	15	20
Certificate	Any configuration object	256	256	256	256	256	256

Parameters		VM01	VM02	VM04	VM08	VM16	VM32
Shared Resources	Address	512	1024	2048	4096	4096	4096
	Address group	256	256	256	256	256	256
	Health checks	64	128	256	512	512	512
	ISP address book	32	32	32	32	32	32
	Schedule	256	256	256	256	256	256
	Schedule group	64	64	64	64	64	64
	Service	512	1024	2048	4096	4096	4096
	Service group	256	256	256	256	256	256
SNMP	SNMP community	16	16	16	16	16	16
	SNMP community host	16	16	16	16	16	16
	SNMP user	16	16	16	16	16	16
Networking							
Interfaces	Physical network interfaces	10	10	10	10	10	10
	VLAN interfaces	128	256	512	1024	1024	1024
Routing	ARP table entries (per VDOM)	4096	4096	4096	4096	4096	4096
	Static routes	1024	2048	4096	4096	4096	4096
	Policy routes	32	64	128	256	256	256
	ISP routes	32	32	32	32	32	32
NAT	Any configuration object	256	256	256	256	256	256
QoS	Any configuration object	256	256	256	256	256	256
Packet Capture	Table	5	5	5	5	5	5
User							
Any configuration object		256	256	256	256	256	256
Server Load Balancing							
Virtual Servers	L4	512	1024	2048	4096	4096	4096
	L7	128	256	512	1024	1024	1024
	L7 HTTPs	64	128	256	512	512	512

Parameters		VM01	VM02	VM04	VM08	VM16	VM32
Real Server Pool	Pools	512	1024	2048	4096	4096	4096
	Pool members	512	1024	2048	4096	4096	4096
	Real server SSL profile	256	256	256	256	256	256
Resources	Profiles	256	256	256	256	256	256
	Cache policies	256	256	256	256	256	256
	Compression policies	256	256	256	256	256	256
	Persistence policies	128	128	128	256	256	256
	Method policies	32	64	128	256	256	256
	Authentication policies	256	256	256	256	256	256
	Scripts	256	256	256	256	256	256
Content Rules	Content routing rules *Indicates maximum limits per VDOM. True maximums are calculated by ***PER_VDOM_LIMIT* × *MAX_VDOM***.	128*	256*	512*	1024*	1024*	1024*
	Content rewriting rules	128	256	512	1024	1024	1024
Link Load Balancing							
Link Group	Gateway	512	1024	2048	4096	4096	4096
	Link group	256	512	1024	2048	2048	2048
	Link group member	512	1024	2048	4096	4096	4096
Virtual Tunnel	Virtual tunnel	256	512	1024	2048	2048	2048
	Virtual tunnel member	256	256	256	256	256	256
Policy	LLB policy rule	256	512	1024	2048	2048	2048
Global Load Balancing							
Zones		4096	4096	4096	4096	4096	4096
FQDN hosts		4096	4096	4096	4096	4096	4096
GLB policies		256	256	256	1024	1024	1024
Security							

Parameters	VM01	VM02	VM04	VM08	VM16	VM32
Any configuration object	256	256	256	256	256	256
Log & Report						
Remote Syslog Servers	3	3	3	3	3	3

Supported image versions for EOS models

For FortiADC models that are still under maintenance, you can install or upgrade to the latest image versions. However, for models that have reached their End of Support (EOS) date, the supported image versions are limited. In the following table, the EOS models and their latest supported versions are listed.

FortiADC model	Latest Mature Version
FortiADC 200D	7.2.1 Build 0220
FortiADC 700D	7.2.1 Build 0220
FortiADC 1500D	7.2.1 Build 0220
FortiADC 2000D	7.2.1 Build 0220
FortiADC 4000D	7.2.1 Build 0220
FortiADC 400D	7.4.6 Build 0377



www.fortinet.com

Copyright© 2025 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., 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 Chief Legal Officer, 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. 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.